

UDP25G-IP reference design

Rev1.1 9-Mar-23

1 Introduction

Compared to the TCP protocol, the UDP protocol minimizes protocol mechanisms when sending data. There is no handshake and no data recovery process to confirm that the receiver accepts all data correctly. However, like TCP, UDP provides checksums for data integrity and port numbers for addressing different functions at the source and destination in networks.

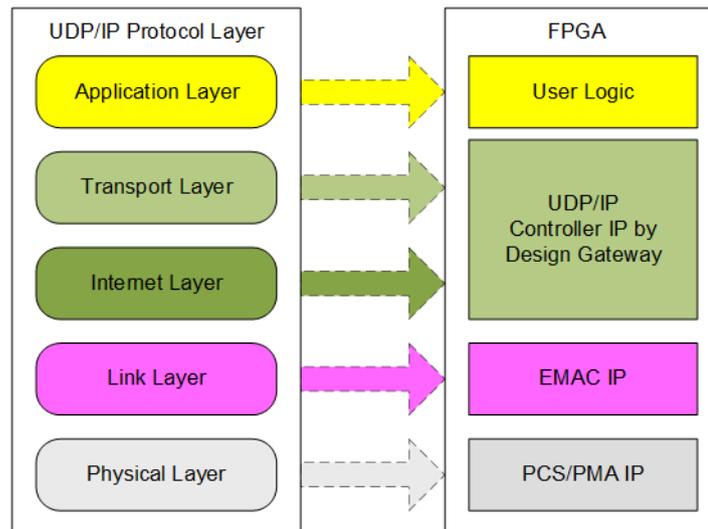


Figure 1-1 UDP/IP protocol layer

UDP25G-IP implements the Transport and Internet layer of the UDP/IP Protocol to build Ethernet packets from user data (UDP payload data) to EMAC. If the UDP payload data size is larger than a packet size, UDP25G-IP splits the user data into smaller size to fit in one packet. After that, the payload data is appended by UDP/IP header. Conversely, when the EMAC receives the Ethernet packet, UDP25G-IP extracts the data from the packet and verifies the header. If the header is valid, UDP payload data is forwarded to the user logic. Otherwise, the packet is rejected.

The lower layer protocols are implemented by EMAC-IP and PCS/PMA-IP. PCS/PMA-IP is provided by Xilinx FPGA while EMAC-IP can be implemented by DG 10G25GEMAC-IP or Xilinx 10/25G Ethernet Subsystem.

The reference design includes a simple user logic to transfer data using UDP25G-IP. UDP25G-IP transfers data with a PC or another UDP25G-IP on another FPGA board. To transfer data with a PC, the test application, `udpdatatest`, is called on the PC to send and verify UDP payload data via Ethernet connection at a high-speed rate. One application is called for transferring data in one direction, while two test applications are called for full-duplex test to send and receive data simultaneously.

To allow the user to control the test parameters and the operation of the UDP25G-IP demo via UART/JTAGUART, the CPU system is included. The user can easily set the test parameters and monitor the current status on the console. The firmware on CPU is built using bare-metal OS. More details of the demo are described below.

2 Hardware overview

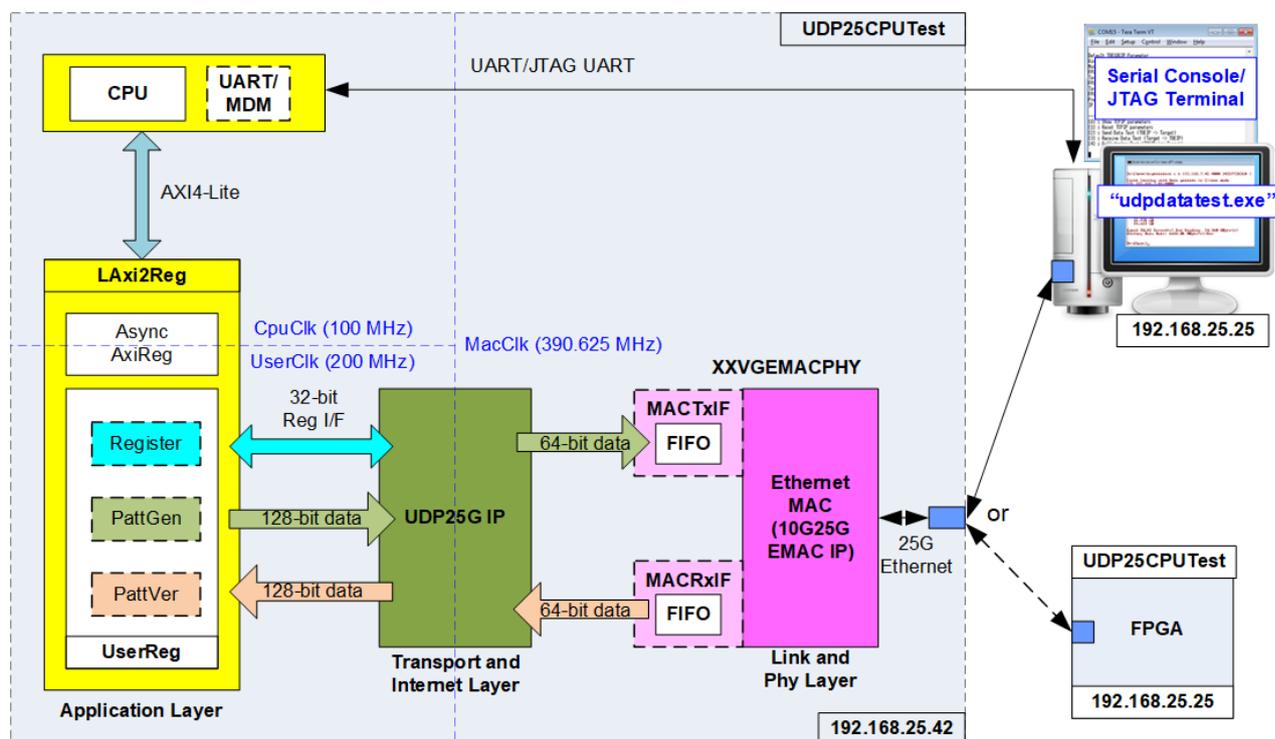


Figure 2-1 Demo block diagram

During testing, two devices are used to transfer data over a 25Gb Ethernet connection. The first device operates in Client mode and the second device operates in Server mode. In the demonstration, the Client device is the UDP25G-IP on an FPGA board, while the Server device can be either the UDP25G-IP on an FPGA board or a PC, as shown in Figure 2-1. If a PC is used, the test application “udpdatatest” must be executed on the PC to transfer data with the UDP25G-IP within the FPGA.

In the FPGA system, the UDP25G-IP is connected to 25G Ethernet System to implement all UDP/IP layers. If the DG 10G25GEMAC-IP and Xilinx PCS/PMA-IP are used to implement the 2G Ethernet System, the UDP25G-IP can be directly connected to it. However, if the Xilinx 25G Ethernet (MAC) Subsystem IP is employed, the adapter logic (MACTxIF and MACRxIF) must be added to server as an interface module between the UDP25G-IP and the Ethernet system.

The user interface of the UDP25G-IP is connected to UserReg within AsyncAxiReg, which includes a Register file for interfacing with the Register interface, PattGen for sending test data via Tx FIFO interface, and PattVer for verifying test data via Rx FIFO interface. The Register file of UserReg is controlled by CPU firmware through the AXI4-Lite bus.

The test design uses three clock domains, CpuClk, which is used for the CPU system, MacClk, which is the clock output from 10/25Gb Ethernet PCS/PMA, and UserClk, which is used for the user logic of the UDP25G-IP. Therefore, AsyncAxiReg is designed to support asynchronous signals between CpuClk and UserClk. More information about each module inside the UDP25CPUtest is described below.

Note: 1. UserClk can be modified to use the same clock as CpuClk for reducing clock resource.
2. The valid frequency range for UserClk of UDP25G-IP is 195.3125 – 390.625 MHz.

2.1 25G Ethernet System (25G BASE-SR)

The 25G Ethernet System comprises the MAC layer, PCS layer, and PMA layer to interface with the external device by 25G BASE-SR. To connect with UDP25G-IP, the user interface of 25G Ethernet System is 64-bit AXI4-stream, operated at 390.625 MHz. Three solutions of 25G Ethernet System are shown in Figure 2-2.

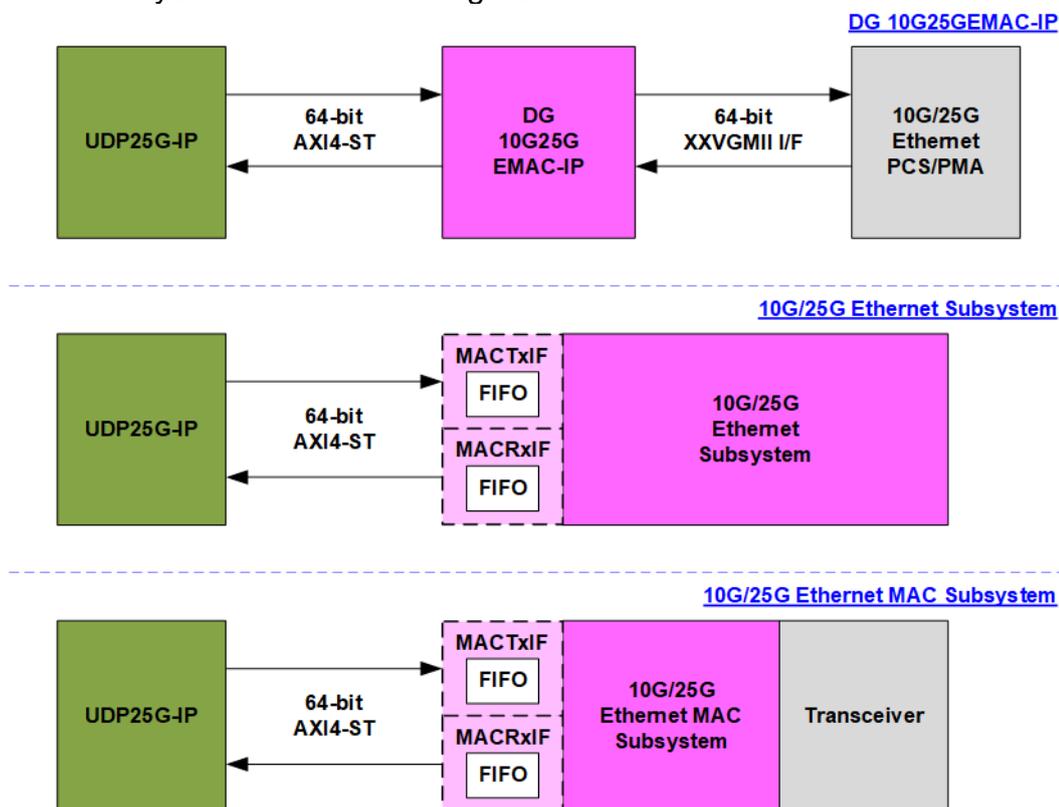


Figure 2-2 UDP25G-IP with three Ethernet system solutions

The first solution combines the DG 10G25GEMAC-IP and Xilinx PCS/PMA module to optimize IP resource and reduce latency time. The UDP25G-IP and 10G25GEMAC-IP can be connected directly. However, the Xilinx PCS/PMA module is free only when the RS-FEC feature is not enabled. To enable RS-FEC feature, Design Gateway provides another IP, 25GMEAC/PCS+RSFEC-IP. Further details can be found on our website.

https://dgway.com/UDP-IP_X_E.html

DG 10G25GEMAC-IP

https://dgway.com/products/IP/GEMAC-IP/dg_10g25gemacip_data_sheet_xilinx.pdf

The second solution employs the Xilinx 10G/25G Ethernet Subsystem, which implements Ethernet MAC and PCS/PMA function. However, the user interface of 10G/25G Ethernet Subsystem is incompatible with the UDP25G-IP's MAC interface. The UDP25G-IP requires to transfer each packet continuously, while the 10G/25G Ethernet Subsystem may de-assert data ready or data valid signal to pause data transmission before the end of the packet. Therefore, adapter logics (MACTxIF and MACRxIF) with small FIFO are used to interface between UDP25G-IP and 10G/25G Ethernet Subsystem. Please visit Xilinx website for more information about this solution.

<https://www.xilinx.com/products/intellectual-property/ef-di-25gemac.html>

The third solution involves using the Ethernet MAC Subsystem, a Hard IP available on the Versal device, to implement the Ethernet MAC layer and PCS layer. The IP must be connected to the Transceiver module to process the PMA layer. The user interface of the IP can be configured to several modes. In this reference design, the 64-bit non-segmented mode with independent clock is applied. The minimum user clock frequency of this mode is 390.625 MHz. Adapter logics are also required to interface between UDP25G-IP and Ethernet MAC Subsystem. More information about this solution can be found on the Xilinx website.

<https://www.xilinx.com/products/intellectual-property/mrmac.html>

Note: The default reference design disables the RS-FEC feature, so the network equipment of the test environment must disable RS-FEC feature. Please contact us for the demo system that enables RS-FEC feature.

The details of the adapter logics, MAC25GTxIF and MAC25GRxIF, are described below.

MAC25GTxIF

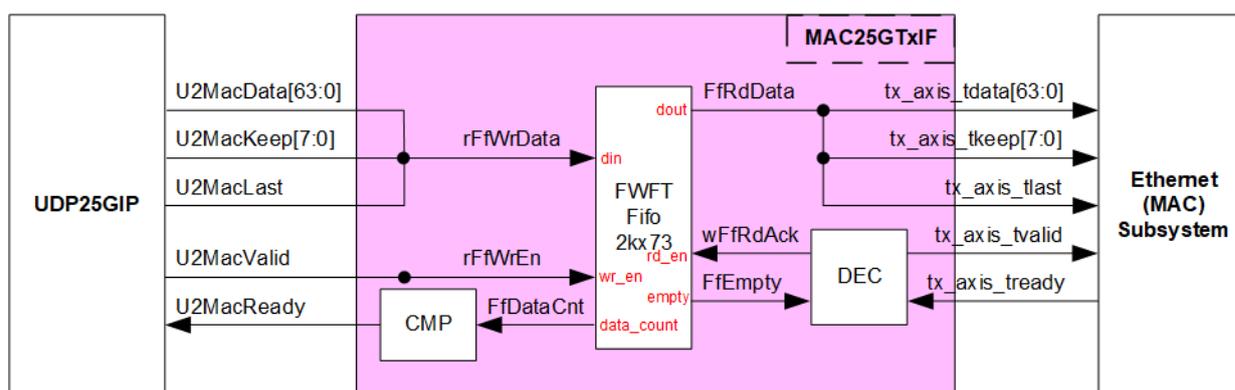


Figure 2-3 MAC25GTxIF Timing diagram

Tx interface timing diagram of Ethernet (MAC) Subsystem and UDP25G-IP are different. UDP25G-IP needs to send each packet continuously, but the Xilinx Ethernet (MAC) Subsystem does not support this. The EMAC may pause receiving data by de-asserting tx_axis_tready signal before the end of the packet.

To solve this issue, the MAC25GTxIF is designed to store transmitted data from UDP25G-IP when EMAC is not ready to receive new data. The FIFO depth is 2048 to store at least one data packet during pausing time. Since the maximum packet size of UDP25G-IP reference design is 8960 bytes or 1120 of 64-bit data, the FIFO depth of 2048 is sufficient. The FIFO is First-Word Fall-Through (FWFT) FIFO, so the read data is valid for reading at the same time as asserting read enable to 1b.

The operation of MAC25GTxIF is divided into two parts: Writing FIFO and Reading FIFO. Timing diagrams of each part are illustrated in Figure 2-4 and Figure 2-5.

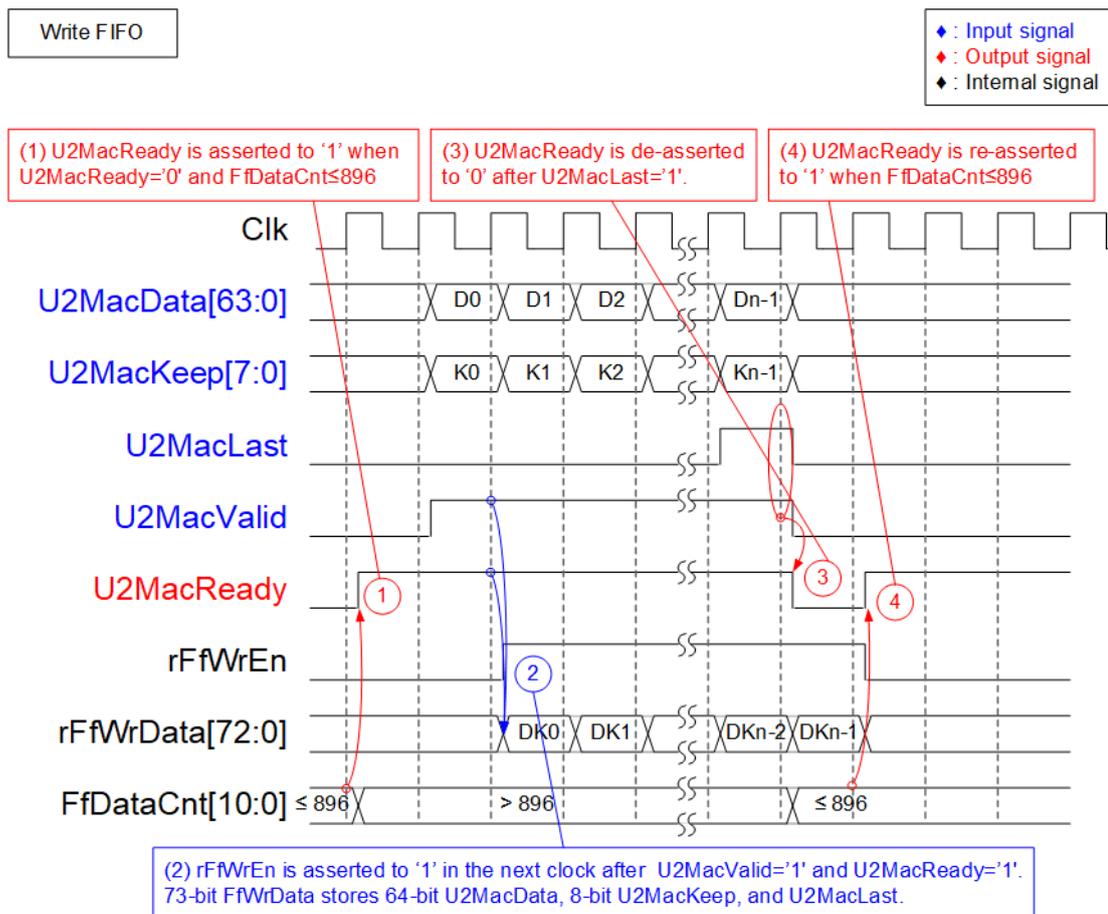


Figure 2-4 Timing diagram for transferring data from UDP25G-IP to FIFO

- 1) Before asserting U2MacReady to 1b to receive a new packet from user, two conditions must be satisfied. First, it must be enough free space in FIFO to store a maximum packet size, 9014 bytes. To verify this, the upper bit of FfDataCnt is read to ensure that the amount of data in FIFO is not greater than 896 (indicating that there is at least 1151 of 64-bit free space). Second, the previous packet must have been completely transferred, as monitored by U2MacReady being set to 0b.
- 2) The user begins transmitting a packet by asserting U2MacValid to 1b. The input signals from user (U2MacData, U2MacKeep, and U2MacLast) are considered valid and stored in the FIFO when both U2MacValid and U2MacReady are asserted to 1b. Following this, the inputs are stored in the FIFO by asserting rFfWrEn to 1b. The 73-bit Write data to the FIFO includes 64-bit data (U2MacData), 8-bit empty byte (U2MacKeep), and the end flag (U2MacLast).
- 3) Once the final data of a packet is received (U2MacLast=1b and U2MacValid=1b), U2MacReady is de-asserted to 0b to pause data transmission, allowing for FfDataCnt to be read.
- 4) If FfDataCnt indicates that there is sufficient free space in the FIFO, U2MacReady will be re-asserted to 1b in the next cycle.

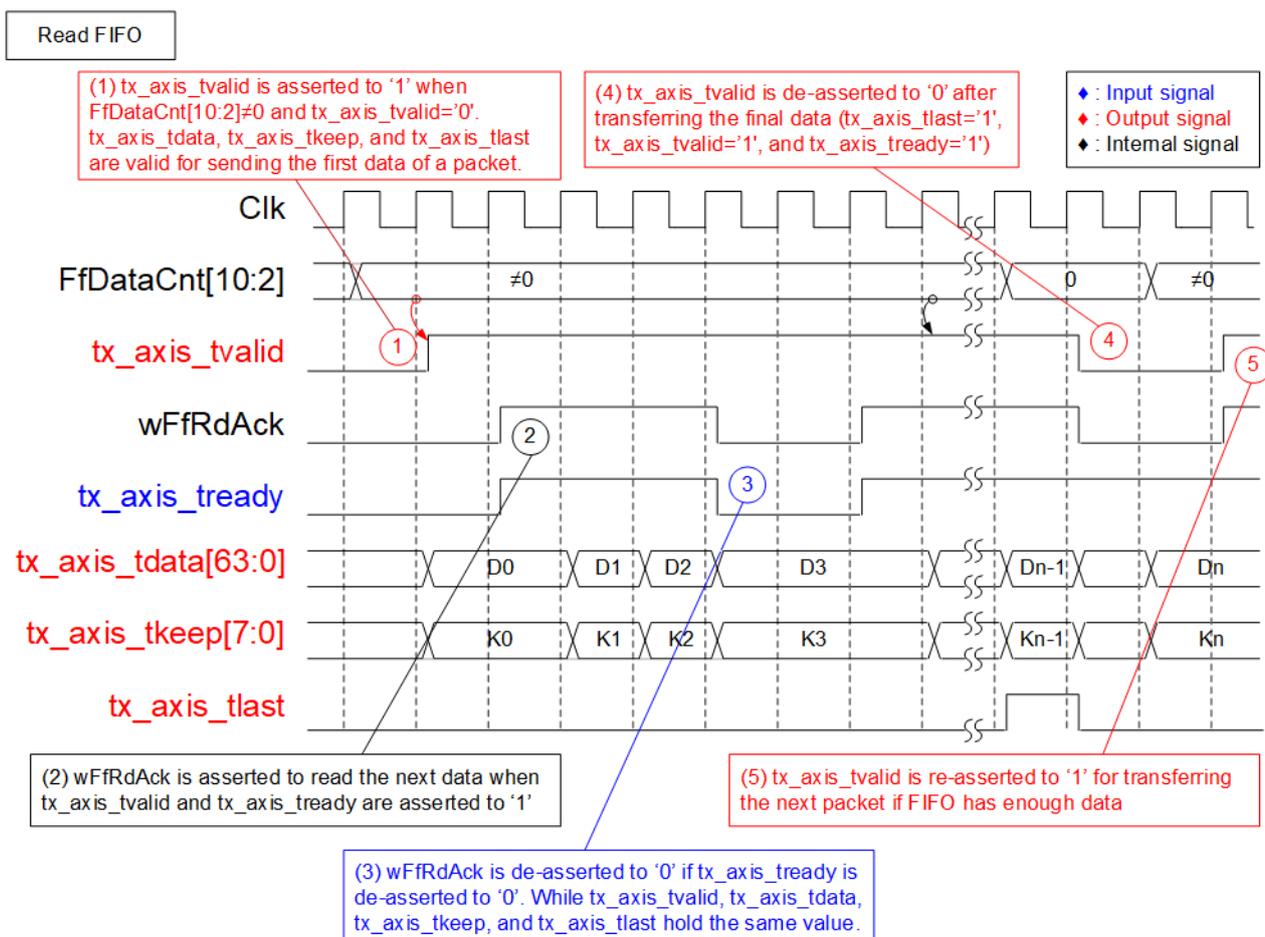


Figure 2-5 Timing diagram for transferring data from FIFO to EMAC

- 1) The transmission of a new packet begins only when the FIFO has data stored ($FfDataCnt[10:2] \neq 0$) and there is no packet transmitting ($tx_axis_tvalid = 0b$). To initiate the transmission, the tx_axis_tvalid signal is asserted to 1b, along with the valid output signals to EMAC which include 64-bit tx_axis_tdata , 8-bit tx_axis_tkeep , and tx_axis_tlast .
- 2) Upon complete transmission of data to EMAC ($tx_axis_tvalid = 1b$ and $tx_axis_tready = 1b$), $wFfRdAck$ is asserted to 1b to read the next set of data from FIFO.
- 3) If tx_axis_tready is de-asserted to 0b, $wFfRdAck$ is de-asserted to 0b to pause the reading of new data from FIFO. This means that all output signals sent to EMAC will hold the same value until EMAC re-asserts tx_axis_tready to 1b.
- 4) After the final data of a packet is fully transferred ($tx_axis_tlast = 1b$ and $tx_axis_tready = 1b$), tx_axis_tvalid is de-asserted to 0b to pause the data transmission, and allowing the checking of data size in FIFO for transferring the next packet.
- 5) The transmission of the next packet begins when the FIFO has sufficient data. The process returns to step 1 to transmit the new packet.

MAC25GRxIF

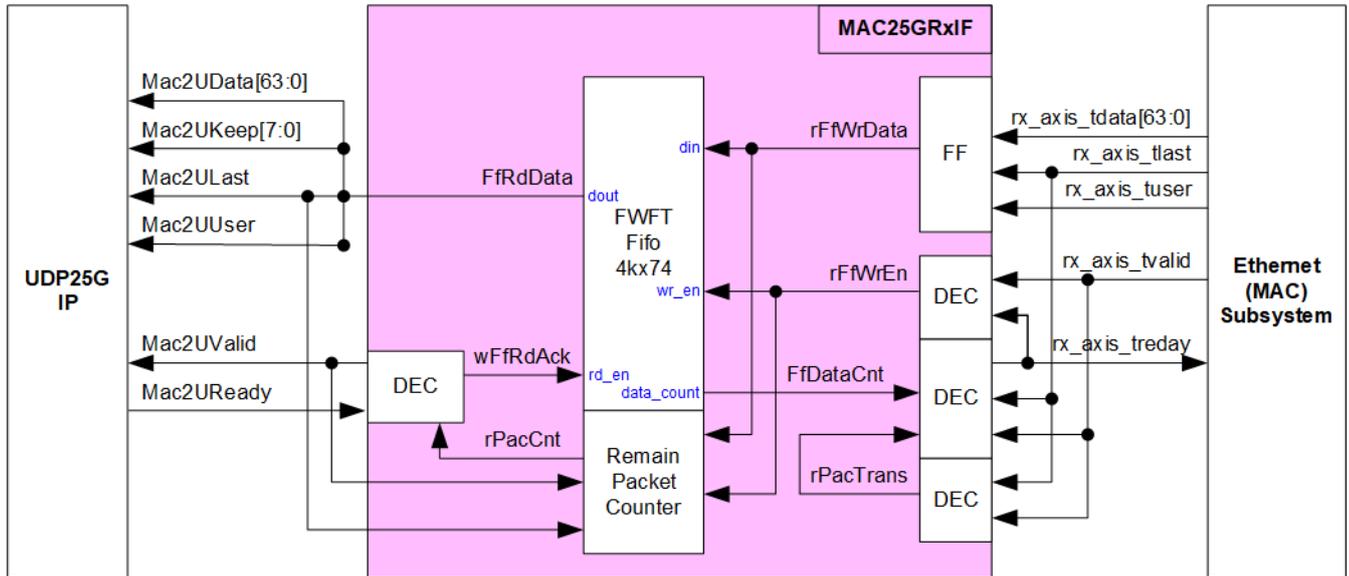


Figure 2-6 MAC25GRxIF Timing diagram

The EMAC interface of UDP25G-IP requires to receive the packets continuously, so Mac2UValid must remain set to 1b from the beginning to the end of each packet. However, the Xilinx Ethernet (MAC) Subsystem may pause data transmission by de-asserting the valid signal (rx_axis_tvalid) before the packet is complete.

To solve this issue, the MAC25GRxIF is designed to store all packet data transmitted from EMAC before forwarding it to UDP25G-IP. This guarantees that all data is available for transfer until the end of the packet. The FIFO has a depth of 4096, which can store multiple Ethernet packets, and operates as a First-Word Fall-Through FIFO.

The Remain packet counter keeps track of the number of packets stored in the FIFO. When a new packet is received from EMAC, the counter is increased, and it decreased when the packet is completely forwarded to UDP25G-IP.

The logic of MAC25GRxIF is split into two groups: Writing FIFO and Reading FIFO. The timing diagrams for each group are displayed in Figure 2-7 and Figure 2-8, respectively.

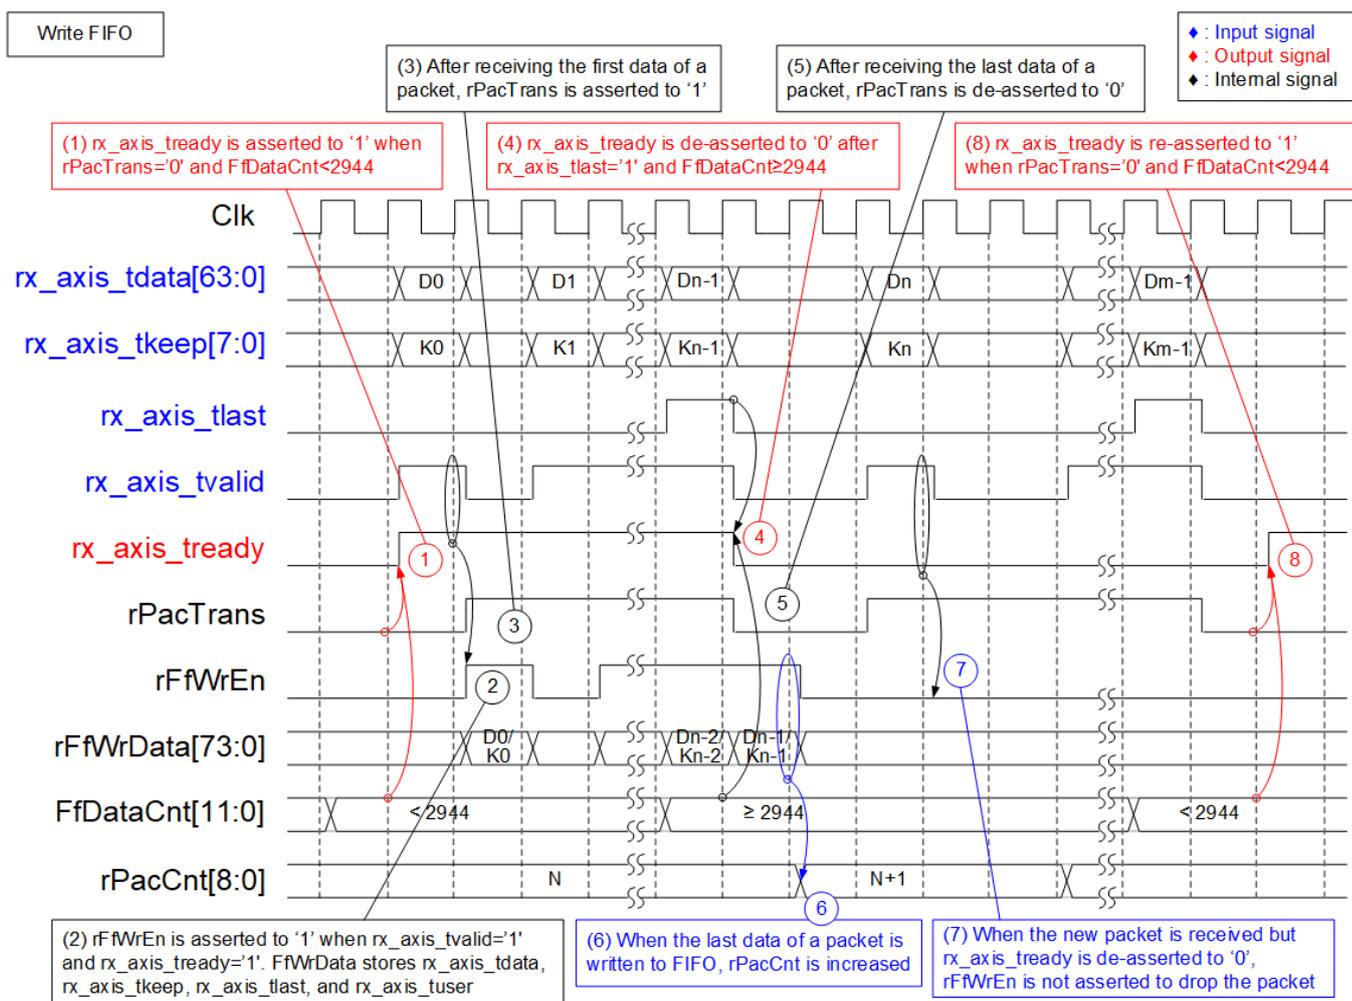


Figure 2-7 Timing diagram for transferring data from EMAC to FIFO

- 1) The free space size in FIFO is checked by reading FfDataCnt. If it is less than 2944 (the free space in FIFO is more than 1153), this is sufficient to store the maximum packet size (9014 bytes). Also, rPacTrans must be equal to 0 to confirm that the packet is not currently being transmitted. Once these conditions are met, rx_axis_tready is asserted to 1b to begin receiving data from EMAC.
- 2) When a new packet is ready to be transferred (rx_axis_tvalid is set to 1b), the data and control signals from EMAC are stored in the FIFO, including 64-bit rx_axis_tdata, 8-bit rx_axis_tkeep, rx_axis_tlast, and rx_axis_tuser. rFfWrEn is set to 1b to write the 74-bit data to the FIFO.
- 3) After the first data of a packet is received, rPacTrans is set to 1b and remains until the end of packet is received. This allows to use rPacTrans for monitoring the packet transmission status.
- 4) If the final data of a packet is received and there is not enough free space in the FIFO (FfDataCnt≥2944), rx_axis_tready is de-asserted to 0b to pause data reception.
- 5) After the final data of a packet is received, rPacTrans is de-asserted to 0b to change the packet transmission status from Busy to Idle.
- 6) Once the final data of a packet has been stored in the FIFO (rFfWrEn=1b and rFfWrData[72] which is last flag =1b), the packet counter (rPacCnt) is increased by 1 to indicate the total number of packets stored in the FIFO.

- 7) If a new packet is received but rx_axis_tready is still de-asserted to 0b, the received packet will be dropped and not be stored in the FIFO.
- 8) When there is no packet currently being transmitted and there is enough free space in the FIFO, rx_axis_tready is re-asserted to 1b to resume data reception.

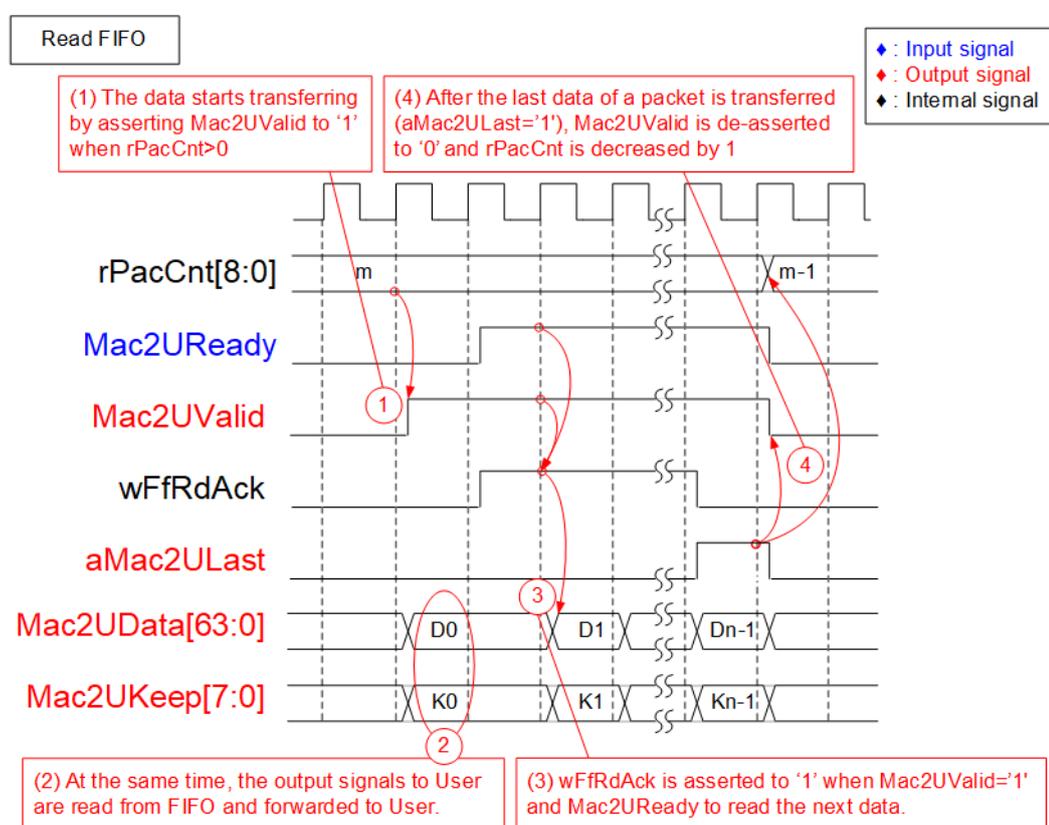


Figure 2-8 Timing diagram of data transferring from FIFO to UDP25G-IP

- 1) Data transmission from FIFO to UDP25G-IP begins only when there is at least one packet stored in FIFO (indicated by rPacCnt not being equal to 0). To start the transmission, Mac2UValid is asserted to 1b.
- 2) The data and control signals read from FIFO are sent to UDP25G-IP, including 64-bit Mac2UData, 8-bit Mac2UKeep, Mac2ULast, and Mac2UUser. Mac2UValid remains asserted to 1b until the end of packet to ensure continuous data transmission.
- 3) Once the data is completely transferred to UDP25G-IP (Mac2UValid=1b and Mac2UReady=1b), wFfRdAck is asserted to 1b to read the next data.
- 4) After the final data of a packet is transferred (aMac2ULast=1b and Mac2UValid=1b), Mac2UValid and wFfRdAck are de-asserted to 0b to pause a packet transmission. Also, rPacCnt is decreased by 1 after the completion of one packet transfer.

2.2 UDP25G-IP

UDP25G-IP implements UDP/IP stack and fully offload engine without requiring the CPU and the external memory. User interface has two signal groups - control signals and data signals. Control and status signals use Single-port RAM interface for write/read register access. Data signals use FIFO interface for transferring data stream in both directions. More information can be found from the datasheet.

https://dgway.com/products/IP/UDP25G-IP/dg_udp25gip_data_sheet_xilinx.pdf

2.3 CPU and Peripherals

The 32-bit AXI4-Lite is used for CPU access to peripherals such as Timer and UART in the test system. Control and status signals are connected to be registers for CPU access as a peripheral through the 32-bit AXI4-Lite bus. The CPU assigns a different base address and address range to each peripheral, allowing access to one peripheral at a time.

In the reference design, the test hardware is connected to the CPU system as a peripheral with a specified base address and range. Therefore, the LAXi2Reg module that interfaces with the CPU must support the AXI4-Lite bus standard for CPU writing and reading, as shown in Figure 2-9.

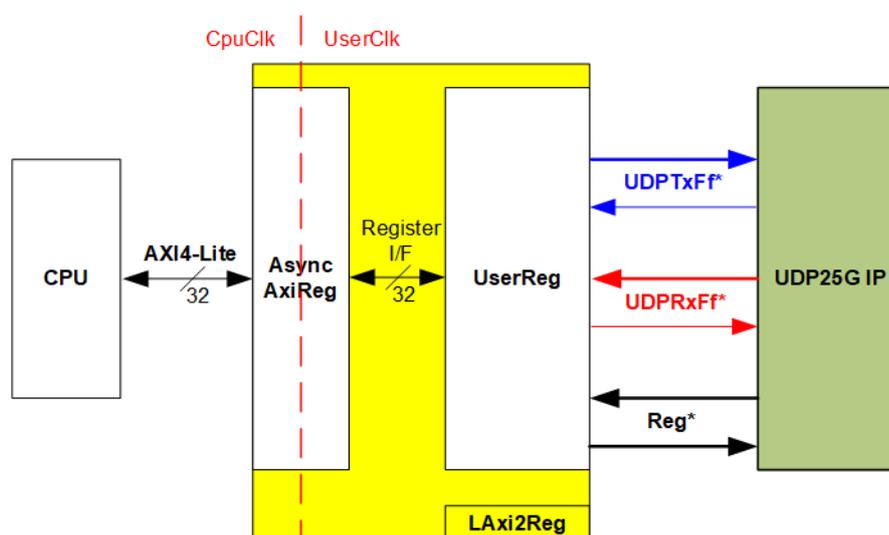


Figure 2-9 LAXi2Reg block diagram

The LAXi2Reg module includes two parts: AsyncAxiReg and UserReg. AsyncAxiReg is designed to convert the AXI4-Lite signals to a simple Register interface with a 32-bit data bus size, similar to AXI4-Lite data bus size. In addition, AsyncAxiReg includes asynchronous logic to support clock domain crossing between the CpuClk domain and UserClk domain.

UserReg includes the Register file for the parameters and the status signals of the test logics. Both the data interface and control interface of UDP25G-IP are connected to UserReg. Further details of AsyncAxiReg and UserReg are described below.

2.3.1 AsyncAxiReg

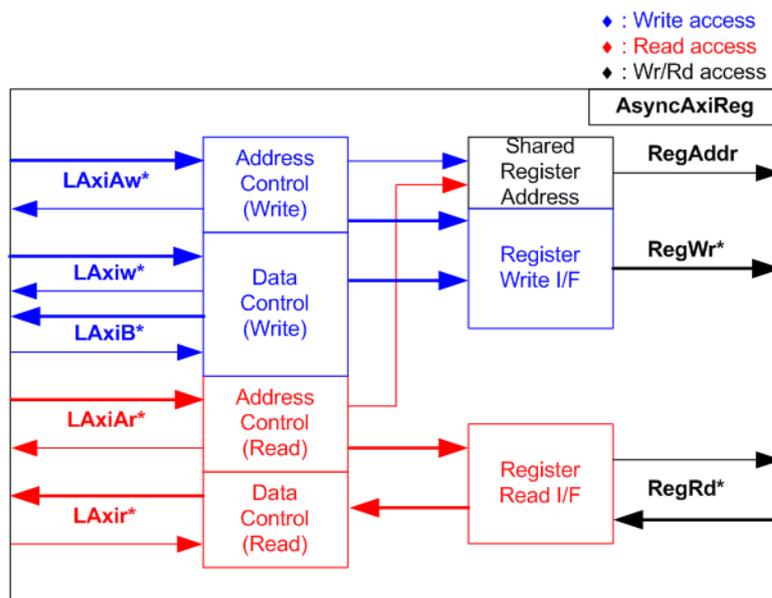


Figure 2-10 AsyncAxiReg Interface

The AXI4-Lite bus interface signals are categorized into five groups: LAXiAw* (Write address channel), LAXiw* (Write data channel), LAXiB* (Write response channel), LAXiAr* (Read address channel), and LAXir* (Read data channel). More information on creating custom logic for the AXI4-Lite bus can be found in the following document.

https://github.com/Architech-Silica/Designing-a-Custom-AXI-Slave-Peripheral/blob/master/designing_a_custom_axi_slave_rev1.pdf

According to the AXI4-Lite standard, the write channel and read channels operate independently for both control and data interfaces. Therefore, the logic in the AsyncAxiReg module that interfaces with the AXI4-Lite bus is divided into four groups: Write control logic, Write data logic, Read control logic, and Read data logic, as shown on the left side of Figure 2-10. The Write control I/F and Write data I/F of the AXI4-Lite bus are latched and transferred to become the Write register interface with clock domain crossing registers. Similarly, the Read control I/F of the AXI4-Lite bus is latched and transferred to the Read register interface, while Read data is returned from the Register interface to the AXI4-Lite bus via clock domain crossing registers. In the Register interface, RegAddr is a shared signal for write and read access, loading the value from LAXiAw for write access or LAXiAr for read access.

The Register interface is compatible with a single-port RAM interface for write transaction. However, the read transaction of the Register interface has been slightly modified from the RAM interface by adding the RdReq and RdValid signals to control read latency time. Since the address of the Register interface is shared for both write and read transactions, the user cannot write and read the register simultaneously. The timing diagram of the Register interface is shown in Figure 2-11.

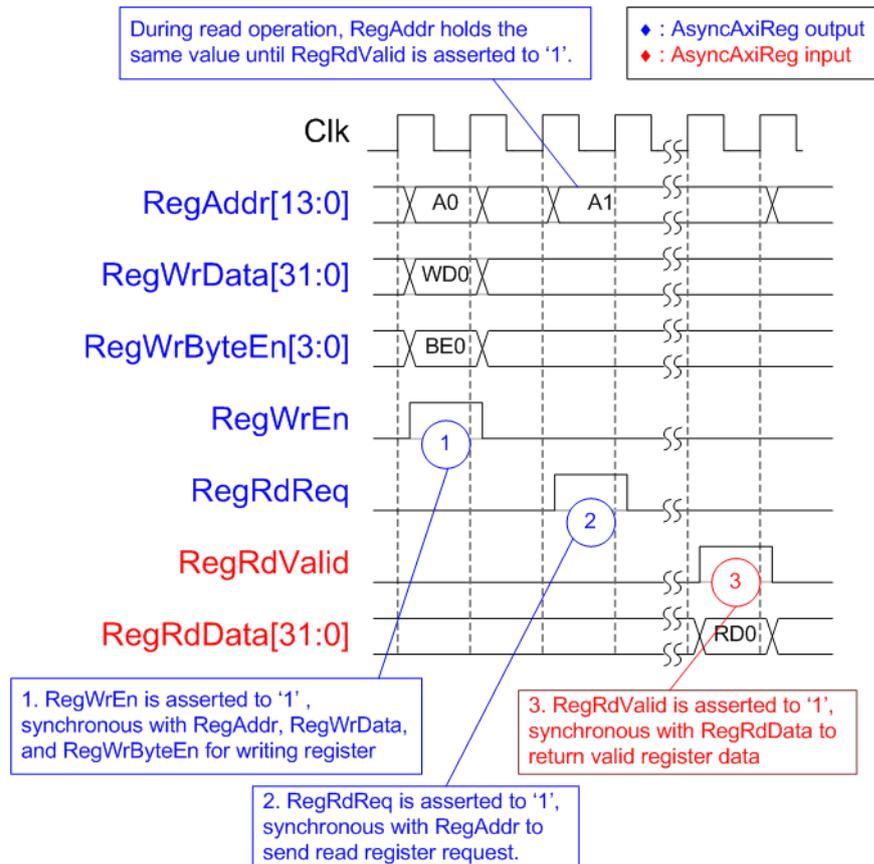


Figure 2-11 Register interface timing diagram

- 1) Timing diagram to write register is similar to that of a single-port RAM. The RegWrEn signal is set to 1b, along with a valid RegAddr (Register address in 32-bit units), RegWrData (write data for the register), and RegWrByteEn (write byte enable). The byte enable consists of four bits that indicate the validity of the byte data. For example, bit[0], [1], [2], and [3] are set to 1b when RegWrData[7:0], [15:8], [23:16], and [31:24] are valid, respectively.
- 2) To read register, AsyncAxiReg sets the RegRdReq signal to 1b with a valid value for RegAddr. The 32-bit data is returned after the read request is received. The slave detects the RegRdReq signal being set to start the read transaction. In the read operation, the address value (RegAddr) remains unchanged until RegRdValid is set to 1b. The address can then be used to select the returned data using multiple layers of multiplexers.
- 3) The slave returns the read data on RegRdData bus by setting the RegRdValid signal to 1b. After that, AsyncAxiReg forwards the read value to the LAXir* interface.

2.3.2 UserReg

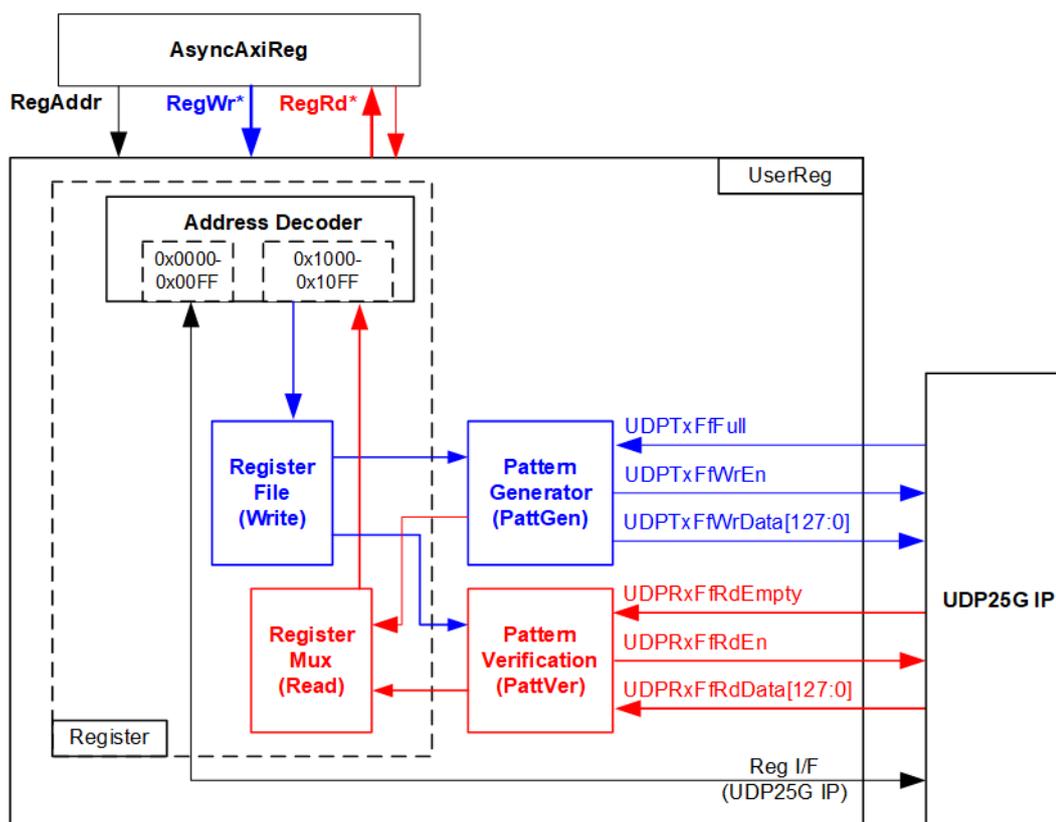


Figure 2-12 UserReg block diagram

The UserReg module includes three functions: Register, Pattern generator (PattGen), and Pattern verification (PattVer). The Register block decodes the requested address from AsyncAxiReg and selects the active register for a write or read transaction. The PattGen block is designed to send 128-bit test data to UDP25G-IP following FIFO interface standard, while the PattVer block is designed to read and verify 128-bit data from UDP25G-IP following FIFO interface standard.

Register Block

The address range is split into two areas: UDP25G-IP register (0x0000-0x00FF) and UserReg register (0x1000-0x10FF). The Address decoder decodes the upper bits of RegAddr to select the active hardware. Since the Register file inside UserReg is 32-bit bus size, Write byte enable (RegWrByteEn) is not used. To write hardware registers, the CPU must use a 32-bit pointer to place a 32-bit valid value on the write data bus.

For reading a register, a multiplexer selects the data to return to CPU by using the address. The lower bits of RegAddr are applied to select the active data within each Register area. While the upper bits are used to select the returned data from each Register area. The total latency time of read data is equal to one clock cycle, and RegRdValid is created by RegRdReq by asserting a D Flip-flop. More details of the address mapping within the UserReg module are shown in Table 2-1

Table 2-1 Register map Definition

Address	Register Name	Description
Wr/Rd	(Label in the "udp25gtest.c")	
BA+0x0000 – BA+0x00FF: UDP25G-IP Register Area More details of each register are described in UDP25G-IP datasheet.		
BA+0x0000	UDP_RST_INTREG	Mapped to RST register within UDP25G-IP
BA+0x0004	UDP_CMD_INTREG	Mapped to CMD register within UDP25G-IP
BA+0x0008	UDP_SML_INTREG	Mapped to SML register within UDP25G-IP
BA+0x000C	UDP_SMH_INTREG	Mapped to SMH register within UDP25G-IP
BA+0x0010	UDP_DIP_INTREG	Mapped to DIP register within UDP25G-IP
BA+0x0014	UDP_SIP_INTREG	Mapped to SIP register within UDP25G-IP
BA+0x0018	UDP_DPN_INTREG	Mapped to DPN register within UDP25G-IP
BA+0x001C	UDP_SPN_INTREG	Mapped to SPN register within UDP25G-IP
BA+0x0020	UDP_TDL_INTREG	Mapped to TDL register within UDP25G-IP
BA+0x0024	UDP_TMO_INTREG	Mapped to TMO register within UDP25G-IP
BA+0x0028	UDP_PKL_INTREG	Mapped to PKL register within UDP25G-IP
BA+0x0034	UDP_SRV_INTREG	Mapped to TDH register within UDP25G-IP
BA+0x0038	UDP_RST_INTREG	Mapped to SRV register within UDP25G-IP
BA+0x003C	UDP_VER_INTREG	Mapped to VER register within UDP25G-IP
BA+0x0040	UDP_DML_INTREG	Mapped to DML register within UDP25G-IP
BA+0x0044	UDP_DMH_INTREG	Mapped to DMH register within UDP25G-IP
BA+0x1000 – BA+0x10FF: UserReg control/status		
BA+0x1000	Total transmit length (Low)	Wr [31:0] – 32 lower bits of 44-bit total transmit size in 128-bit unit.
Wr/Rd	(USER_TXLENL_INTREG)	Valid from 1-0xFFFF_FFFF_FFFF. Rd [31:0] – 32 lower bits of 44-bit current transmit size in 128-bit unit. The value is cleared to 0 when USER_CMD_INTREG is written by user.
BA+0x1004	Total transmit length (High)	Wr [11:0] – 12 upper bits of 44-bit total transmit size in 128-bit unit.
Wr/Rd	(USER_TXLENH_INTREG)	Rd [11:0] – 12 upper bits of 44-bit current transmit size in 128-bit unit.
BA+0x1008	User Command	Wr [0] – Start transmitting. Set 0b to start transmitting.
Wr/Rd	(USER_CMD_INTREG)	[1] – Data verification enable (0b: Disable data verification, 1b: Enable data verification) Rd [0] – PattGen Busy (0b: Idle, 1b: PattGen is busy) [1] – Data verification error (0b: Normal, 1b: Error) This bit is auto-cleared when user starts new operation or reset.
BA+0x100C	User Reset	Wr [0] – Reset signal. Set 1b to reset the logic.
Wr/Rd	(USER_RST_INTREG)	This bit is auto-cleared to 0b. [8] – Set 1b to clear read value of USER_RST_INTREG[8] to 0b Rd [8] – Latched value of IntOut, the output from IP (0b: Normal, 1b: IntOut has been asserted) This flag can be cleared by system reset condition or setting USER_RST_INTREG[8]=1b. [16] – Ethernet linkup status from Ethernet MAC (0b: Link down, 1b: Link up)
BA+0x1010	FIFO status	Rd[3:0] - Mapped to UDPRxFfLastRdCnt signal of UDP25G-IP
Rd	(USER_FFSTS_INTREG)	[15:4] - Mapped to UDPRxFfRdCnt signal of UDP25G-IP [24] - Mapped to UDPTxFfFull signal of UDP25G-IP
BA+0x1014	Total receive length (Low)	Rd[31:0] – 32 lower bits of 44-bit current receive size in 128-bit unit
Rd	(USER_RXLENL_INTREG)	The value is cleared to 0 when USER_CMD_INTREG is written by user.
BA+0x1018	Total receive length (High)	Rd[11:0] – 12 upper bits of 44-bit current receive size in 128-bit unit
Rd	(USER_RXLENH_INTREG)	
BA+0x1080	EMAC IP version	Rd[31:0] – Mapped to IPVersion output from DG 10G25GEMAC-IP when the system integrates DG 10G25GEMAC-IP.
Rd	(EMAC_VER_INTREG)	

Pattern Generator

The logic diagram and timing diagram of Pattern Generator (PattGen) are illustrated in Figure 2-13 and Figure 2-14, respectively.

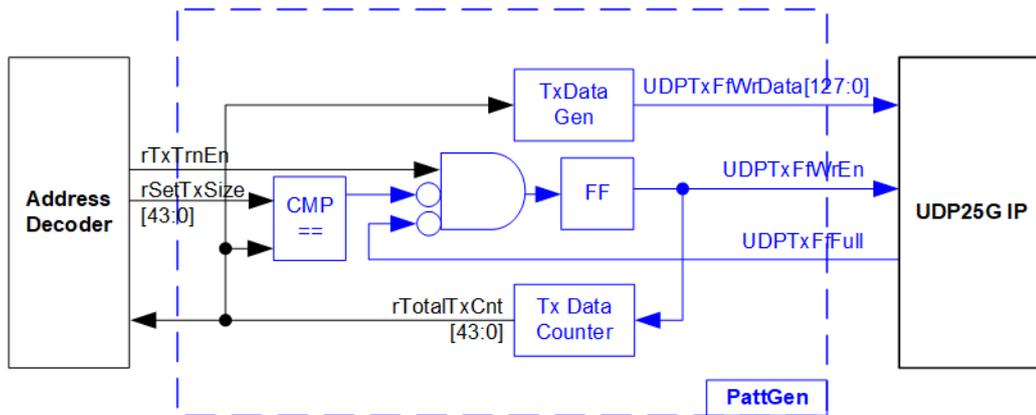


Figure 2-13 PattGen block

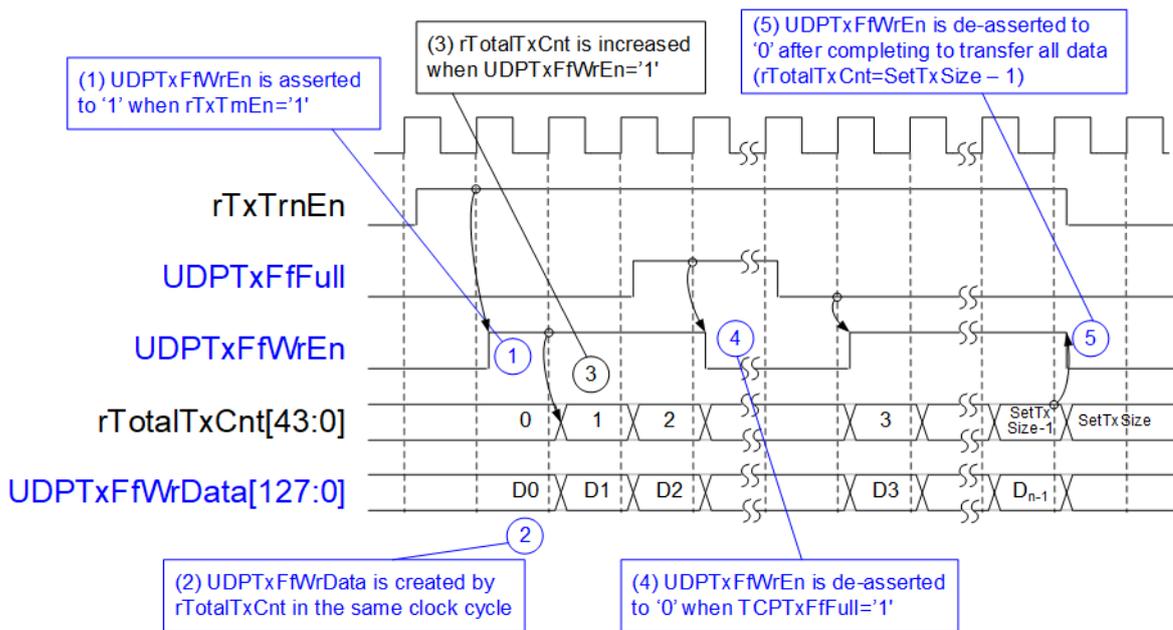


Figure 2-14 PattGen timing diagram

When USER_CMD_INTREG[0] is set to 0b, PattGen initiates the operation of generating test data by setting rTxTrnEn to 1b. While rTxTrnEn remains set to 1b, UDPTxFWrEn is controlled by UDPTxFfFull. If UDPTxFfFull is 1b, UDPTxFWrEn is de-asserted to 0b. The data counter, rTotalTxCnt, checks the total amount of data sent to UDP25G-IP. The lower bits of rTotalTxCnt generate 32-bit incremental data for the UDPTxFWrData signal. Once all data has been transferred, equal to rSetTxSize, rTxTrnEn is de-asserted to 0b.

Pattern Verification

The logic diagram and timing diagram of Pattern Verification (PattVer) are illustrated in Figure 2-15 and Figure 2-16, respectively. The verification feature is executed when the verification flag (rVerifyEn) is enabled.

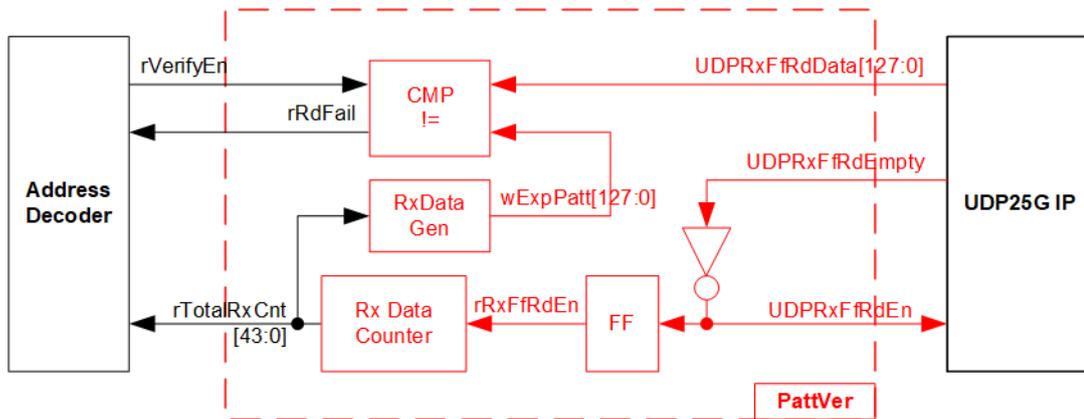


Figure 2-15 PattVer block

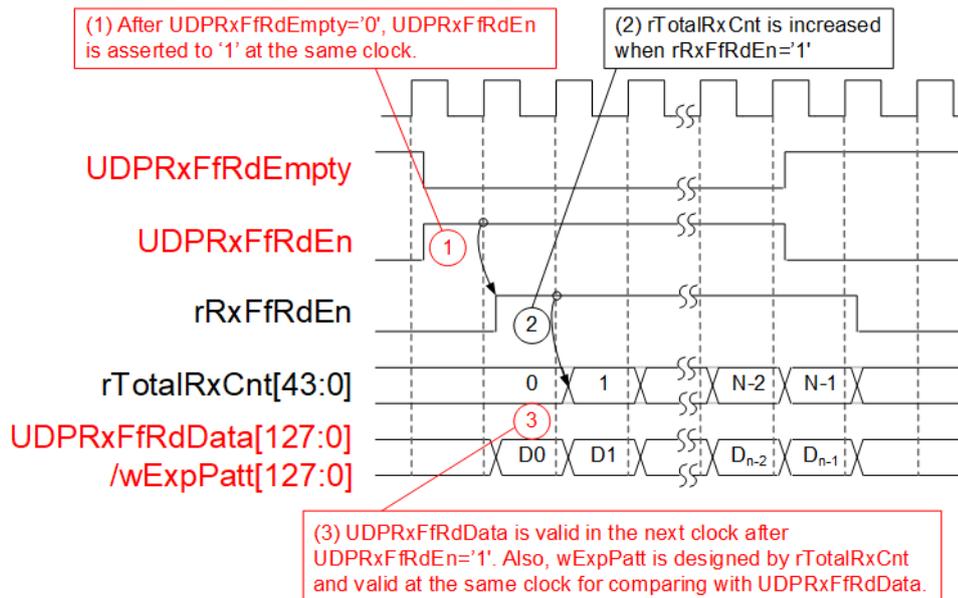


Figure 2-16 PattVer Timing diagram

When rVerifyEn is set to 1b, the verification logic is processed. It compares the received data (UDPRxFfRdData) with the expected data (wExpPatt). If comparison fails, rRdFail is asserted to 1b. The UDPRxFfRdEn signal is created by applying NOT logic to UDPRxFfRdEmpty. The data for comparison, UDPRxFfRdData, becomes valid in the next clock cycle. To count the total size of received data, rTotalRxCnt is enabled by rRxFfRdEn, which is delayed by one clock cycle from UDPRxFfRdEn. The lower bits of rTotalRxCnt are applied to generate wExpPatt for comparison with UDPRxFfRdData. Therefore, UDPRxFfRdData and wExpPatt are valid in the same clock cycle and can be compared using rRxFfRdEn signal.

3 CPU Firmware on FPGA

The reference design uses a bare-metal OS for the CPU firmware operating, which facilitates hardware handling. When executing the test system, the first step is to initialize the hardware, described in more details below.

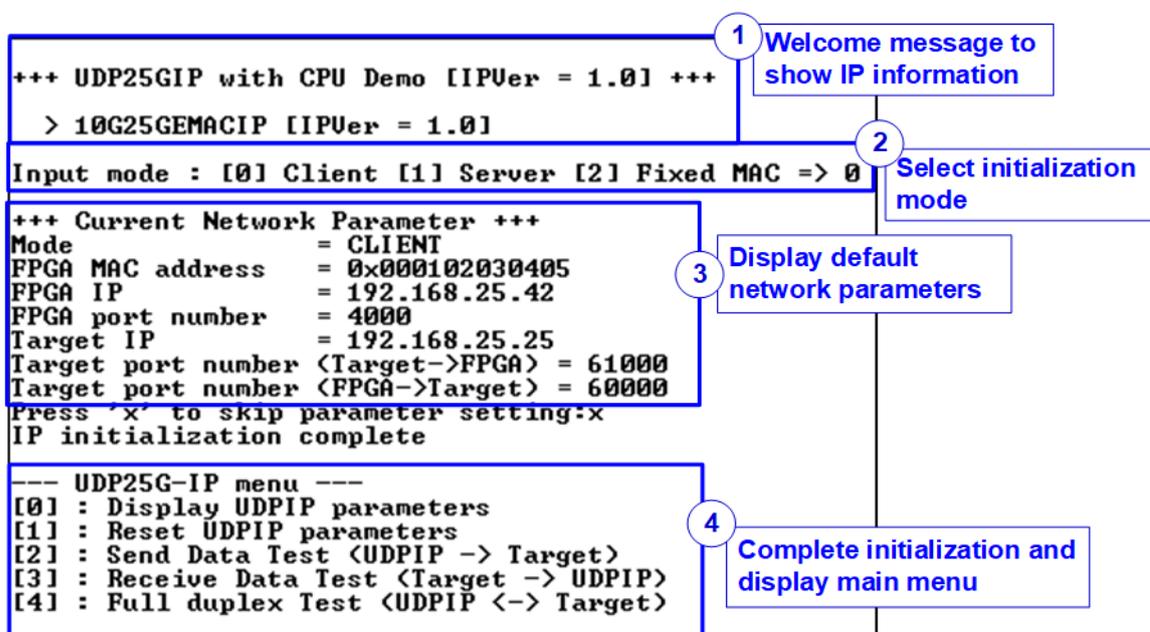


Figure 3-1 System initialization in Client mode by using default parameters

Figure 3-1 illustrates the four-step process for hardware initialization, which is described below.

- 1) Upon FPGA boot-up, the firmware polls the status of the 25G Ethernet link (USER_RST_INTREG[16]). The CPU waits until the link is up, and then displays a welcome message to show IP information.
- 2) The menu to select the initialization mode of UDP25G-IP is displayed, allowing the user to choose the Client, Server, or Fixed-MAC mode.

Note:

- When running in Client mode, UDP25G-IP sends an ARP request to obtain the MAC address of the target device from the ARP reply. When running in Server mode, UDP25G-IP waits until an ARP request is received to decode the MAC address and return an ARP reply. When running in Fixed-MAC mode, the user needs to know MAC address of the target device for setting to UDP25G-IP.
- When running the test environment with one FPGA board and Test PC, it is recommended to set the FPGA to run as Client mode.
- When using two FPGA boards in a test environment, there are three options to establish the connection between them. The first option is to set on board as the Client and the other as the Server. The second option is to configure both boards in Fixed-MAC mode. The third option is to set one board to Fixed-MAC mode and the other board to act as the Client.

- 3) The CPU displays the default values of the network parameters including the initialization mode, FPGA MAC address, FPGA IP address, FPGA port number, Target IP address, and Target port number. The firmware has two default parameter sets for the operation mode: Server parameter set (used for Server mode only) and Client parameter set (used for both Client and Fixed-MAC mode). When setting to Fixed-MAC mode, an extra parameter, Target MAC address, is also displayed. The user can select to complete the initialization process using the default parameters or by updating some parameters. The details of how to change the parameter are provided in Reset parameters menu (topic 3.2).
- 4) The CPU waits until the IP completes the initialization process by checking if busy status (UDP_CMD_INTREG[0]) is equal to 0b. After that, "IP initialization complete" is displayed with the main menu. There are five test operations in the main menu, and more details of each menu are described below.

3.1 Display parameters

This menu displays the current value of all UDP25G-IP parameters. The following steps are executed to display parameters.

- 1) Read the initialization mode.
- 2) Read all network parameters from each variable in the firmware following the initialization mode, i.e., source (FPGA) MAC address, source (FPGA) IP address, source (FPGA) port number, Target MAC address (only displayed in fixed MAC mode), Target IP address, and Target port number.

Note: The source parameters are the FPGA parameters set to UDP25G-IP, while the Target parameters are the parameters of a PC or another FPGA.

- 3) Print out each variable.

3.2 Reset parameters

This menu is used to change some UDP25G-IP parameters, such as IP address and source port number. After setting the updated values to UDP25G-IP, the CPU resets the IP to re-initialize the process using new parameters. Finally, the CPU waits until the initialization is completed. The following steps are executed to reset the parameters.

- 1) Display all parameters on the console, similar to topic 3.1 (Display parameters).
- 2) If the user uses the default value, skip to the next step. Otherwise, display the menu to set all parameters.
 - i) Receive the initialization mode from the user. If the initialization mode is changed, display the latest parameter set of new mode on the console.
 - ii) Receive the remaining parameters from the user and verify all inputs. If the input is invalid, the parameter is not updated.
- 3) Force reset to UDP25G-IP by setting UDP_RST_INTREG[0]=1b.
- 4) Set all parameters to UDP25G-IP register, such as UDP_SML_INTREG and UDP_DIP_INTREG.
- 5) De-assert UDP25G-IP reset by setting UDP_RST_INTREG[0]=0b to initiate the initialization process.
- 6) Clear PattGen and PattVer logic by sending a reset to user logic (USER_RST_INTREG[0]=1b).
- 7) Monitor the UDP25G-IP busy flag (UDP_CMD_INTREG[0]) until the initialization process is finished (busy flag is de-asserted to 0b).

3.3 Send data test

This menu allows the user to execute Send data test. The user can set the parameters such as total transmit length. If all inputs are valid, the data is transferred by sending 32-bit incremental test data. The operation is completed when all data is transferred. The following are the steps to send data.

- 1) Receive the transfer size and packet size from the user and verify that all inputs are valid. If any input is invalid, the operation is cancelled.
- 2) Set the UserReg registers - the transfer size (USER_TXLEN_INTREG), the Reset flag to clear the initial value of test pattern (USER_RST_INTREG[0]=1b), and the Command register to start the data pattern generator (USER_CMD_INTREG=0). The test pattern generator in UserReg starts to generate test data to UDP25G-IP.
- 3) Display the recommended parameters of the test application on PC by reading the current parameters in the system. Wait until the user presses any key to start the IP sending operation.
- 4) Set parameters to UDP25G-IP to start the operation. The packet size is set to UDP_PKL_INTREG, and the total size is set to UDP_TDL/H_INTREG. Finally, UDP_CMD_INTREG is set to 1b to start IP sending data.
- 5) Wait for UDP25G-IP to complete the operation by monitoring the busy flag of the IP (UDP_CMD_INTREG[0]=0b). While monitoring the busy flag, the CPU reads the current transfer size from the user logic (USER_TXLENL/H_INTREG) and displays it on the console every second.
- 6) Once the operation is completed, the CPU calculates the performance and displays the test result on the console.

3.4 Receive data test

This menu allows the user to execute Receive data test. The user can set the parameters such as total receive length. If all inputs are valid, a 32-bit incremental test data is created for verification with the received data from PC/FPGA when the data verification is enabled. The following are the steps to receive data.

- 1) Receive the total transfer size and data verification mode from user and verify that all inputs are valid. The operation is cancelled if some inputs are invalid.
- 2) Set the UserReg registers, i.e., the Reset flag to clear the initial value of the test pattern (USER_RST_INTREG[0]=1b) and data verification mode (USER_CMD_INTREG[1]=0b/1b to enable/disable).
- 3) Display recommended parameter (similar to Step 3 of Send data test).
- 4) Wait until the total number of received data (USER_RXLENL/H_INTREG) is equal to the set value (complete condition), or the number of received data is not updated for 100 msec (timeout condition). During receiving data, the CPU displays the current number of received data on the console every second.
- 5) Stop the timer. Check the interrupt from the timeout (USER_RST_INTREG[8]) and data verification flag (USER_CMD_INTREG[1]) registers when the verification mode is applied. If some errors are found, the error message will be displayed.
- 6) Calculate performance and show the test result on the console.

3.5 Full duplex test

This menu enables full duplex testing by simultaneously transferring data between the FPGA and another device (PC/FPGA) in both directions. User-defined parameters, such as the total transfer length, are received to initiate the test. If all inputs are valid, the data transfer begins and completes when the data is completely transferred in both directions.

Note: When testing with a PC, the transfer size on the test application (udpdatatest) must match the transfer size set on the FPGA. Two “udpdatatest” are executed, one for sending data and another for receiving data using different port number. When testing with two FPGAs, the port number for sending and receiving data must be the same.

The steps to execute a full duplex test are as follows.

- 1) Receive the total data size (using the same size for both transfer directions), packet size, and data verification mode (enabled or disabled) from the user and verify that all inputs are valid. The operation is cancelled if some inputs are invalid.
- 2) Set UserReg registers including transfer size (USER_TXLENL/H_INTREG), the Reset flag to clear the initial value of the test pattern (USER_RST_INTREG[0]=1b), and the Command register to start data pattern generator with data verification mode (USER_CMD_INTREG=0 or 2).
- 3) Display the recommended parameters for the test application running on the PC by reading the current parameters in the system.
- 4) Set UDP25G-IP registers, including packet size (UDP_PKL_INTREG), total transfer size (UDP_TDL/H_INTREG), and Send command (UDP_CMD_INTREG=1). The IP begins sending data once the UDP_CMD_INTREG is set to 1b. For receiving data, the IP is always ready to receive data without any additional setting.
- 5) The CPU controls data flow of both directions simultaneously, with two tasks running during the test, as follows.
 - a) To send data, the CPU reads the busy flag (UDP_CMD_INTREG[0]) and waits until it is de-asserted to 0b. The busy flag is de-asserted to 0b when the Send command is finished.
 - b) To receive data, the CPU reads the total number of received data. The read process finishes when the total number of received data is equal to the set value (no data lost). If the total number of received data does not change for 100 msec (timeout), the read process is also finished.

If the data is not completely transferred, the current number of transmitted data size (USER_TXLENL/H_INTREG) and received data size (USER_RXLENL/H_INTREG) are read and displayed on the console every second.

- 6) Stop the timer and check the data verification status (USER_CMD_INTREG[1]). If a verification error is found, an error message is displayed.
- 7) Calculate performance and display the test result on the console.

3.6 Function list in User application

This topic describes the function list to run UDP25G-IP operation.

void init_param(void)	
Parameters	None
Return value	None
Description	Reset parameters following the description in topic 3.2. In the function, show_param and input_param function are called to display parameters and get parameters from user.

int input_param(void)	
Parameters	None
Return value	0: Valid input, -1: Invalid input
Description	Receive network parameters from user, i.e., the initialization mode, FPGA MAC address, FPGA IP address, FPGA port number, Target MAC address (when run in Fixed-MAC mode), Target IP address, and Target port numbers. If all inputs are valid, the parameters are updated. Otherwise, the value does not change. After receiving all parameters, calling show_param function to display parameters.

void show_cursize(void)	
Parameters	None
Return value	None
Description	Read current number of transmitted data and number of received data by reading USER_TXLENL/H_INTREG and USER_RXLENL/H_INTREG in Byte, KByte, or MByte unit.

void show_interrupt(void)	
Parameters	None
Return value	None
Description	Read interrupt status from UDP_TMO_INTREG and decode interrupt type to display the details of interrupt on the console.

void show_param(void)	
Parameters	None
Return value	None
Description	Display the parameters following the description in topic 3.1.

void show_result(void)	
Parameters	None
Return value	None
Description	Read USER_TXLENL/H_INTREG and USER_RXLENL/H_REG to display total transmitted data size and total received data size. Read the global parameters (timer_val and timer_upper_val) and calculate total time usage to display in usec, msec, or sec unit. Finally, transfer performance is calculated and displayed on MB/s unit.

int udp_rcv_test(void)	
Parameters	None
Return value	0: The operation is successful -1: Receive invalid input or error is found
Description	Run Receive data test following description in topic 3.4. It calls show_interrupt, show_cursize, and show_result function.

int udp_send_test(void)	
Parameters	None
Return value	0: The operation is successful -1: Receive invalid input or error is found
Description	Run Send data test following description in topic 3.3. It calls show_cursize and show_result function.

int udp_txrx_test(void)	
Parameters	None
Return value	0: The operation is successful -1: Receive invalid input or error is found
Description	Run Full duplex test following described in topic 3.5. It calls show_interrupt, show_cursize, and show_result function.

void wait_ethlink(void)	
Parameters	None
Return value	None
Description	Read USER_RST_REG[16] and wait until the Ethernet connection is linked up

4 Test Software on PC

```

Command Prompt

C:\SW>udpdatabest

[ERROR] The application requires 5 input parameters and 2 optional input parameter
*****
UDP Data Transfer Test Version 1.6
*****
udpdatabest [Dir] [FPGAIP] [FPGAPort] [PCPort] [ByteLen] <Pattern> <Timeout>

[Dir]          Transfer direction of PC
                t:Transmit data  r:Receive data
[FPGAIP]       FPGA IP Address
[FPGAPort]     FPGA Port number<0-65535>
[PCPort]       PC Port number<0-65535>
[ByteLen]      Transfer length(Byte)
<Pattern>     <Optional>Disable/Enable Data pattern in transferring
                0:Disable  1:Enable,  Default=1 (Enable)
<Timeout>     <Optional>Timeout in msec<50-65535>. Default=100.

[Example]  udpdatabest r 192.168.11.42 4000 60000 4294967295

```

Figure 4-1 “udpdatabest” application usage

The “udpdatabest” application is executed to send or receive UDP data on a PC. It requires five mandatory parameters and two optional parameters. It is important to ensure that the parameter inputs match the parameters set on the FPGA. The details of each parameter input are as follows.

Mandatory parameters

- 1) Dir: : t – when PC sends data to FPGA
: r – when PC receives data from FPGA
- 2) FPGAIP : IP address setting on FPGA (Default value is 192.168.25.42)
- 3) FPGAPort : Port number of FPGA (Default value in FPGA is 4000)
- 4) PCPort : PC port number for sending or receiving data
(Default is 61000 for PC to FPGA and 60000 for FPGA to PC)
- 5) ByteLen : Transfer length for sending or receiving in byte unit. This value must be aligned to 16 from UDP25G-IP limitation.

Optional parameters

- 1) Pattern (optional): Default value when user does not input this parameter is 1.
0 – Generate dummy data in transmit mode or disable data verification in receive mode.
1 – Generate incremental data in transmit mode or enable data verification in receive mode.
- 2) Timeout (optional): Timeout for receiving data in msec unit.
Default value when user does not input this parameter is 100.
The 100 ms is recommended value for running with UDP25G-IP.

Transmit data mode

The steps for running the test application in transmit mode are as follows.

- 1) Get the parameters from the user and verify that all inputs are valid.
- 2) Create a socket and configure the received buffer properties.
- 3) Set the IP address and port number based on the user parameter, and then establish a connection.
- 4) Populate the Send buffer with data for transmission. While the data is being sent, the application prints the total amount of data sent every second to the console.
 - a) If Pattern=1, the Send buffer is filled with a 32-bit incremental pattern.
 - b) If Pattern=0, the Send buffer is not filled and dummy data is used for the test.
- 5) After all data has been sent, the application displays the test results, including the total size of transmitted data and the performance.

Receive data mode

The steps when running the test application in receive mode are as follows.

- 1) Follow step (1)-(3) from the Transmit data mode.
- 2) Continuously read data until the total number of received data equals the set value. If there is no new data received before the timeout, the operation is cancelled. During the data reception, the application prints the total amount of received data every second.
 - a) If Pattern=1, the received data is verified using a 32-bit incremental pattern that increases every four bytes of received data.
 - b) If Pattern=0, the received data is not verified.
- 3) If the read loop finishes due to a timeout, the application displays a “Timeout” message with the total number of lost data and received data. The total time used is also reduced by timeout value.
- 4) After the operation is complete, the application displays the test results, including the performance and the total amount of received data.

5 Revision History

Revision	Date	Description
1.1	9-Mar-23	Update Ethernet system solution
1.0	2-Jun-21	Initial version release

Copyright: 2021 Design Gateway Co,Ltd.