

White-Rabbit NIC Gateway

21 Jun. 2012

Contents

1	Introduction	2
2	Gateway elements	3
2.1	DIO core	4
3	Configuration & Control	5
3.1	DIO core utilization	6
3.2	Interrupt handling	6
4	Application & examples	7
5	Troubleshooting	8

1 Introduction

The White-Rabbit Network Interface Card (WR-NIC) project is concerned with the development of gateway and software to make the combination of a SPEC and a DIO mezzanine behave as a Network Interface Card (NIC) under Linux. Basic demo uses two SPEC boards, one configured as grandmaster and one as slave. Different simple use cases will be provided as basic demo. This document focus on the description of the project gateway. This manual is part of the associated hardware project, hosted at <http://www.ohwr.org/projects/wr-nic/repository>, whose git repository hosts the latest version.

NIC project elements

Application
GN4124 drivers

Software
(Linux PC)



FPGA gateway

Hardware
(2 SPEC + 2 FMC-DIO)




Figure 1: Basic wr-nic project elements

Note that the WR-NIC project inherits many code and working methodology of many other project. Specially important to highly the following ones:

1. White-Rabbit core collection: <http://www.ohwr.org/projects/wr-cores> (wishbonized branch). Look at the WR PTP Core. It is a black-box standalone WR-PTP protocol core, incorporating a CPU, WR MAC and PLLs. It is also convenient to look at TxTSU and NIC project. For further details, search for its related wbgen2 files (extension .wb).
2. Software for White-Rabbit PTP core: <http://www.ohwr.org/projects/wrpc-sw> (a subproject of the previous one)
3. Gennum GN4124 core: <http://www.ohwr.org/projects/gn4124-core>
4. The platform independent core collection: <http://www.ohwr.org/projects/general-cores>. An

important one is the Wishbone crossbar with is download at the `DOWNLOAD_PATH/wr-cores/ip_cores/general-cores/modules/wishbone`.

2 Gateway elements

The main block modules of the architecture are described on next figure.

Basic on-chip architecture




Figure 2: Main HDL blocks of the wr-nic gateway

Here's a quick description of each block:

- The **DIO core** allows configuration of each one of the 5 channels of the DIO mezzanine as input or output. For inputs, it provides an accurate time stamp (using time information from the WRPC, not shown in the diagram) and optionally a host (PCIe) interrupt via the IRQ Gen block. For outputs, it allows the user to schedule the generation of a pulse at a given future time, or to generate it immediately.
- The **IRQ Gen** block receives one-tick-long pulses from other blocks and generates interrupt requests to the GN4124 core. It also includes interrupt source and mask registers.
- The **WB intercon** block ensures seamless interconnection of Wishbone masters and slaves using a crossbar topology. The **GN4124 core** is a bridge between the GN4124 PCIe interface chip and the internal Wishbone bus, allowing communication with the host and interrupts → pipelined version!

- The **WRPC (White Rabbit PTP Core)** communicates with the outside world through the SFP socket in the SPEC, typically using fiber optics. It deals with the WR PTP using an internal LM32 CPU running a portable PTP stack. It forwards/receives non-PTP frames to/from the NIC block, using two pipelined Wishbone interfaces (master and slave for forwarding and receiving respectively). It also provides time information to other cores (not represented in the diagram), and time-tags for transmitted and received frames that can be read through Wishbone for diagnostics purposes. Future versions will include the PPSi library instead of the current PTP stack.
- The **NIC core** ensures communication between the host and the WRPC. More precisely, it interrupts the host and provides a descriptor that the host can use to fetch incoming frames. For outgoing frames, it receives a descriptor from the host, fetches the frame using PCIe DMA via the GN4124 core and sends it to the WRPC using a pipelined Wishbone interface.
- The **TxTSU module** collect timestamps with associated frame identifiers and puts them in a shared FIFO (port identifier is also included although not required for the SPEC card because only one Ethernet port is available). A IRQ is triggered when FIFO is not empty so drivers could read TX timestamps and frame/port identifiers.

2.1 DIO core

The **DIO Core** block is the HDL block that access to the fmc-dio-5chls mezzanine card. Their main elements are shown on the next figure.




Figure 3: Main DIO Core block elements

The different submodules description are:

- **GPIO**: It allows to enable/disable the output drivers and the termination resistors.
- **I2C**: It allows to set the threshold of the ADCMP604 fast LVDS comparator and to access to write/read data to the EEPROM memory (24AA64).
- **Onewire**: It is used for temperature acquisition
- Modules to generate or stamping pulses:
 1. **Pulse generator**: It produces a 1-tick-long pulse in its output when the time passed to it through a vector equals a pre-programmed time.
 2. **Pulse stamper**: It associates a time-tag with an asynchronous input pulse
- Additional Wishbone slave core generated elements are:
 1. Trigger registers (time counters for pulse generation)
 2. FIFOs to store the timestamps of input signals
 3. Interrupt control registers that allow to configure the interrupts generated when there are data in the FIFOs.
 4. Monostable register, which generates a single clock cycle-long

3 Configuration & Control

Accessing to the different system elements is as simple as doing a read/write memory access. The memory map of the different elements of the board are:




Figure 4: Memory map

Any address within this memory space may be addressed by the PC to configure corresponding module (WRPTEP, NIC, TxTSU, etc...). Many of them are already well known so we will skip their descriptions. In next paragraphs we supposed that the utilization of these elements are already known and we will just focus on the DIO core which is the new block added to the OHWR.

3.1 DIO core utilization

The DIO core, according to its architecture already shown, it allows to read input data of each of the 5 channel with precise time-tag information provided by the WRPTP core. It is also possible to program output at a precise time or we could just generate an monostable output immediately. In addition, it is also possible to configure different boards elements as terminator resistors or reference voltage Level using the DAC. All this elements could be controlled independently for each of the 5 channels. More information about the different board configuration elements is available at: <http://www.ohwr.org/projects/fmc-dio-5chttla>

In order to use input/output channels as previously described, the following actions are required:

- Standard GPIO output generation is selected by default. In order to use monostable output (time-programmed or immediate), each channel should be properly configured written to the `di_out_mode` register. A value of 1 indicates that channel will be used for monostable output, otherwise (0 by default), channel will use the values assigned by the GPIO logic block.
- Programmed pulse generation: Generate a programmed input at any time at channel X (X between 0 and 4 identifies the required card channel). For this purpose you need to perform the following actions:
 - Set the required time. This means to provide the 40 bits for the time value and the number of cycles (28 bits). This requires to write the registers `dio_trigX_seconds`, `dio_trighX_seconds` (high part of the time value) and `dio_cyc0_cyc`.
 - Checking if the the board is ready for accepting new triggers. A 1 will be found at each bit of `dio_trig_rdy` register only when required channel is ready to accepts new request.
 - Arming the trigger. You need to write a 1 a the corresponding bit of the `dio_latch_time_chX` bit field.
After these operations, a monostable output will be presented on the desired channel at the requested time.
- Immediate pulse generation: A immediate pulse is generated a the output of each of the card channels just by writing a corresponding 1 at the bit field `dio_puls_inmed` when output mode is set to monostable outputs. No erase is required.
- Input time-tagging: for each of the 5 inputs, if a 1 is detected at this channel, a precise time information is stored on logic FIFOs including the 40 bits time counters and 28 bits more for the cycles. For accessing this information you need to read `dio_tsfX_tag_tai` (32 low bits), `dio_tsfX_tag_taih` (high bits), `dio_tsfX_tag_cycles`. Each time the time tag of any channel is stored, the `fifo not empty` flag generates an interruption to the PC. In the next section we will describe this mechanisms.

A detailed information about the memory maps and related registers names are available by generating html documentation of <http://www.ohwr.org/projects/wr-nic/repository/revisions/master/changes/modules> file. Download the file and generate the HTML documentation using `wbgen2` tool (for instance `wbgen2 -D diocore.htm wr_nic.wb`).

3.2 Interrupt handling

The VIC module block is in charge of handling the different interrupts and provide proper registers to inform of the source of each interruption. The main interrupt signal is communicated to the PC

using the gn4124 core and the `irq_req_p1_i` signal and proper GPIO pins needs to be assigned at the hardware level.

The base address, as shown on the memory map figure is 0x00060000. It handles the following interrupts sources:

- TxTSU interrupts -> at source 0.
- WRSW-NIC interrupts -> at source 1.
- DIO-core interrupts -> at source 2.

Low sources have the highest priority. In order to check the register layout, get the HTML help from the `wb_vic.wb` file. Information about VIC control and configuration registers are provided there. Because VIC module is done to cooperate with `wbgen2` peripherals Embedded Interrupts controllers (EICs), the related information should also be checked. Basically, VIC inform about the main interrupt source and then we need to check the wishbone peripheral interrupt register to complete the interrupt information. Please note that each peripheral generating interrupts has own interrupts registers so a proper configuration of them is required to set-up the interrupt operation.

For instance for the DIO core, please check the status of interrupt registers by looking at `wr_nic.wb` as previously described.

4 Application & examples

[JAVIER: THIS IS MUCH MORE FOR APPROPRIATED IN A GLOBAL DOCUMENT OR ONE ABOUT DEMO EXAMPLES / USE CASES. PERHAPS IT SHOULD BE DELETED FOR FUTURE VERSIONS]

This project could be used as starting demo with White-Rabbit technology. Between others, example applications are:

- Simple transmission of PPS from the master to the slave, with nothing hooked to the external inputs of the boards.
 - The master is free-running. The master host reads system time and schedules a pulse output on the next second. Then it gets an interrupt and from then on it schedules a pulse on each second.
 - The slave host does the same. Looking at the outputs on a scope we should see them perfectly aligned.
- Transmitting an external frequency in the 100 Hz range.
 - The user supplies a ~ 100 Hz square wave on one of the inputs of the master card. The master host reads the time value of the rising edge of the external pulse upon IRQ. Then it adds a constant time (something like 1 ms) and sends a frame with that value to the slave.
 - The slave schedules a pulse to be produced at that time. On the scope we should see a constant time offset between the two pulses.
- Network latency measurements. This is interesting if we connect a switch between the SPEC cards. By using the timestamps on Ethernet frames we could get the measurement of the network latency, verify if it is constant or how traffic affect this parameter.

-
- Many other options are possible. For instance, the measurement of the DIO card latency (we know when we generate a pulse and we could measure when it is received). Please feel free to propose new experiments!

5 Troubleshooting

Properly setting of interrupts registers or wrong memory maps are the typical errors at this stage (in addition to HDL bugs!).

Further information will be provided.

T.B.D.