EECS 151/251A FPGA Lab
Lab 5: Serial I/O - UART

Prof. John Wawrzynek, Nicholas Weaver
TAs: Arya Reais-Parsi, Taehwan Kim
Department of Electrical Engineering and Computer Sciences
College of Engineering, University of California, Berkeley

Contents

1 Before You Start This Lab **1**

2 Lab Setup **2**

3 Serial Device **2**
- 3.1 Framing 2
- 3.2 Transmitting 3
- 3.3 Receiving 3
- 3.4 Putting It All Together 3
- 3.5 Simulation 4

4 Echo **4**
- 4.1 PMOD serial interface 5
- 4.2 Checkoff 5

1 Before You Start This Lab

Before you proceed with the contents of this lab, we suggest that you get acquainted with what a ready/valid handshake is. Read `resources/Verilog/ready_valid_interface.pdf` in the `fpga_labs_sp18` folder. It may be helpful to draw out the timing diagrams for the ready/valid handshake to gain understanding. If you are having trouble, ask a TA. Please note that the serial line itself is not a ready/valid interface. Rather, it is the modules you will work with in this lab (`UATransmit` and `UAReceive`) that have the ready/valid handshake for interfacing with other modules on the FPGA.

In this lab you will implement a UART (Universal Asynchronous Receiver / Transmitter) device, otherwise known as a serial interface. Your working UART from this lab will be used in your project to talk to your workstation (desktop) over a serial line.

2 Lab Setup

Run `git pull` in the `fpga_labs_sp18` folder to fetch the latest skeleton files. You will need to copy over your `debouncer.v`, `edge_detector.v`, and `synchronizer.v` from Lab 4 to your sources directory (overwriting the skeleton files). (You could also just copy the contents.)

3 Serial Device

You are responsible only for implementing the **transmit** side of the UART for this lab. As you should have inferred from reading the ready/valid tutorial, the UART transmit and receive modules use a ready/valid interface to communicate with other modules on the FPGA.

Both the UARTs receive and transmit modules will have their own separate set of ready/valid interfaces connected appropriately to external modules.

3.1 Framing

On the `Pynq-Z1` board, the physical signaling aspects (such as voltage level) of the serial connection will be taken care of by off-FPGA devices. From the FPGA's perspective, there are two signals, `FPGA_SERIAL_RX` and `FPGA_SERIAL_TX`, which correspond to the receive-side and transmit-side pins of the serial port. The FPGA's job is to correctly frame characters going back and forth across the serial connection. Figure 1 below shows a single character frame being transmitted and will be extremely useful in understanding the protocol.

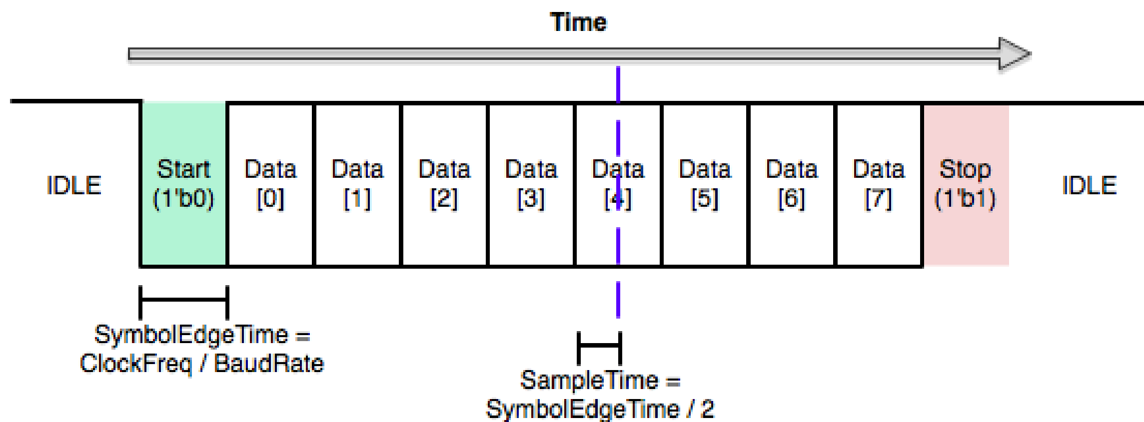

Figure 1: UART Frame Structure

In the idle state the serial line is held high. When the TX side is ready to send a character, it pulls the line low. This is called the start bit. Because UART is an asynchronous protocol, all timing within the frame is relative to when the start bit is first sent (or detected, on the receive side).

The frame is divided up in to 10 uniformly sized bits: the start bit, 8 data bits, and then the stop bit. The width of a bit in cycles of the system clock is then naturally given by the system clock

frequency divided by the baudrate. The baudrate is the number of bits sent per second; in this lab the baudrate will be 115200. Notice that both sides must agree on a baudrate for this scheme to be feasible.

3.2 Transmitting

Let us first think about sending a character using this scheme. Once we have a character that we want to send out, transmitting it is simply a matter of shifting each bit of the character, plus the start and stop bits, out of a shift register on to the serial line.

Remember, the serial baudrate is much slower than the system clock, so we must wait $SymbolEdgeTime = \frac{ClockFreq}{BaudRate}$ cycles between changing the character we're putting on the serial line. After we have shifted all 10 bits out of the shift register, we are done unless we see another transmission immediately after.

3.3 Receiving

The receive side is a bit more complicated. Fortunately, we will provide the receiver module. Open `lab5/lab5.srcs/sources_1/new/uart_receiver.v` so you can see the explanation below implemented.

Like the transmit side, the receive side of the serial device is essentially just a shift register, but this time we are shifting bits from the serial line into the shift register. However, care must be taken into determining when to shift bits in. If we attempt to sample the serial signal directly on the edge between two symbols, we are exceedingly likely to sample on the wrong side of the edge (or worse, when the signal is transitioning) and get the wrong value for that bit. The correct solution is to wait halfway into a cycle (until `SampleTime` on the diagram) before reading a bit in to the shift register.

One other subtlety of the receive side is correctly implementing the ready/valid interface. Once we have received a full character over the serial port, we want to hold the valid signal high until the ready signal goes high, after which the valid signal will be driven low until we receive another character.

This requires using an extra flip-flop (the `has_byte` reg in `uart_receiver.v`) that is set when the last character is shifted in to the shift register and cleared when the ready signal is asserted. This allows us to correctly implement the ready/valid handshake.

3.4 Putting It All Together

Although the receive side and transmit side of the UART you will be building are essentially orthogonal, we are packaging them into one UART module to keep things tidy. If you look at `uart.v`, you will see that this module consists of instantiations of `uart_receiver` and `uart_transmitter`, but there are also two `iob` registers that the serial lines are fed through. What are these for? The `iob` directive tells the synthesis tool to pack those registers into a special block called an `IOB`, which

is used to drive and sense from the IO pins. Using an `IOB` helps ensure that you will have a nice, clean, well-behaved off-chip signal to use as an input or output to your serial modules.

The diagram below shows the entire setup:

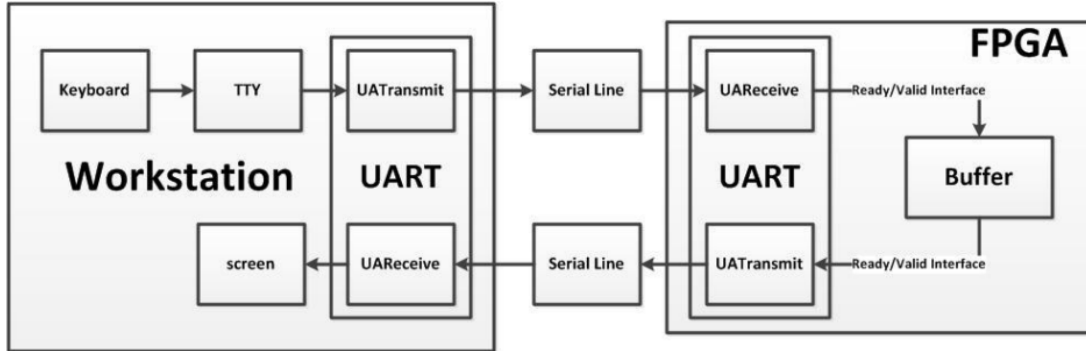

Figure 2: High Level Diagram

3.5 Simulation

We have provided a simple testbench, called `uart_testbench` that will run some basic tests on two instantiations of the UART module with their RX and TX signals crossed so that they can talk to each other. There is also a `.do` file that will run the test. You should note that this test bench reporting success is **not** by itself a good indication that your UART is working properly. The testbench does not attempt to test back to back UART transmissions so you will have to add that test in yourself. Due to the way x's are treated by Modelsim if many signals in your design are undefined the testbench may erroneously pass. Make sure to look at the waveform to see that everything appears to be working properly and that you adequately purge your simulation of high Z and undefined X signals.

If the testbench prints out `# <EOF>` it means that it timed out; this indicates that the testbench was stuck waiting for a condition that never became true. Inspect the waveform and match it up to the testbench code to see where it hangs and why. You shouldn't need to increase any of the timeouts in the `.do` files.

4 Echo

Your UART will eventually be used to interact with your CPU from your workstation. However, since you don't have a CPU yet, you need some other way to test that your UART works on the board.

We have provided this for you. The provided `z1top` contains a very simple finite state machine that does the following continuously:

- Pulls a received character from the `uart_receiver` using ready/valid

- If the received character is an ASCII letter (A-Z or a-z), its case is inverted (lower to upper case or upper or lower case)
- If the received character isn't an ASCII letter, it is unmodified
- The possibly modified character is sent to the `uart_transmitter` using ready/valid to be sent over the serial line one bit at a time

Check using the provided `echo_testbench.v` testbench that everything works as it should in simulation. This testbench is heavily commented to help you understand the communication between the 2 UARTs and the communication over the ready/valid interface. The file often refers to the UART on the workstation as the off-chip UART and the UART on the FPGA as the on-chip UART.

4.1 PMOD serial interface

The Pynq-Z1 does not have an RS-232 serial interface! Well, ok, it does, kind of. The USB interface you use for programming the board (and for RS-232 communication with the board's ARM-based system) is actually only connected to the Processor Subsystem, not the Programmable Logic. We can't use it from our design directly. So, we need to add a serial interface: this is usually a chip wired to connect a device as a serial host/client to another device, with appropriate voltage level shifting to meet the electrical specification (e.g. RS-232).

The PMOD expansion modules you need to interface with your workstation will be availble in the next lab. You're done for now.

4.2 Checkoff

Walk your TA through the simulation results and show that your UART behaves as expected.