

Version 1.1 2/1/2004 9:48 PM

This lab can be started during Lab 4 and completed during Lab 5, if necessary.

## Goals

1. Add a user designed peripheral to a basic MicroBlaze system.
2. Demonstrate the required structure necessary for interfacing user-designed cores to the Xilinx cores.

## Preparation

1. Review the handout outlining the EDK project structure provided in Lab 1. We will be focusing on the pcores subdirectory for this lab.
2. If you are unfamiliar with the profiling of code, read the manual page for “gprof”, which is available on the ugsparc with the “man” command.

## Background

To this point, you have only been adding cores from the existing library. If you cannot find a core with the required functionality, you will have to add your own.

This lab adds a simple core that can be used to profile the code running on a MicroBlaze. This has similar functionality to the “gprof” utility for profiling that is available in the GNU tools. However, it is much more accurate. You may find it useful for your project.

## Step-by-step

1. Copy your working lab1 project into a new directory.
2. Add a pcores subdirectory to your new project directory.

Using an ssh window, login to your ugsparc account. The files you will need for this lab can be found in `~pc/labs/User_Area/lab4.5`. Copy: 1) `system.c` to the code directory of your project; 2) the `snoopy_v1_00_a` directory into the pcores directory; and 3) the `sst` script and `example_results.txt` to the root of the project directory. Note: You may find the “`cp -r`” command useful.

3. Take time to look through the directory structure of the snoopy core. The naming structure is essential for XPS to be able to detect a user’s peripheral. All user cores must be located in the pcores subdirectory or in a globally specified path to a peripheral repository. User cores can be defined using

either VHDL or Verilog, but the tools do not support simulation of mixed designs. Therefore, it may be preferable to write the designs in VHDL

Along with the hdl files used to implement the core, the user must also include a .pao file (Peripheral Analyze Order) and an .mpd file (Microprocessor Peripheral Description). The .pao file lists the order in which files in your design should be synthesized to resolve component architectures. The .mpd file describes the external interface of the core to a system. For more information on these files and their structure, go to the Embedded System Tools Guide.

The easiest method for including user IP into an EDK project, is to follow an example. When you develop your own cores for your project, you should use this as a guide. The cores provided by Xilinx in the EDK6.1\hw\XilinxProcessorIPLib\pcores may also be used as a reference. There is also a document describing how to add user cores to Embedded Systems (the User Core Template Reference Guide) but it is meant for relatively complex cores, providing the user an IP Interface (IPIF) to simplify user core connections to system buses.

4. The snoopy core is a snooping profiler that is able to profile software running on a softcore processor in real time. The counters calculate the exact number of clock cycles spent executing contiguous address ranges. The user specifies the number of counters and the lower and upper bounds for each counter before synthesis. This information can be used by embedded system designers to determine which, if any, sections of the software should be moved to hardware to achieve the required design specifications.
5. Open the Add/Edit cores menu. The snoopy core should appear in the list of peripherals you can add to your design. Add the snoopy core and connect it to the slave opb bus. The core requires 0x100 bytes and a 0x100 byte alignment. Since the tools resolve connections to the opb and lmb buses based on the address of a peripheral, we suggest address range 0xfffff00 to 0xfffffff to guarantee the peripheral resides on the opb.
6. For the snoopy core to work, it must be interfaced with the system clock, and the PC\_EX and valid\_instr ports on the Microblaze core. Connect the clock from the snoopy core (OPB\_Clk) to the system clock. Don't forget to check the net name and the scope under the Ports tab. Unfortunately, the PC\_EX and valid\_instr ports on the snoopy core are not visible through the Add/Edit core interface. Therefore, they will have to be added manually to the hardware description.
7. Close the XPS GUI. Go to your project directory and open the system.mhs file. You are going to edit this file by hand. It is important to remember that

this file contains the project description used by XPS to generate your MicroBlaze system. Therefore, it is safest to only edit the mhs file when XPS is not running.

8. You are going to add two lines of code to the MicroBlaze and snoopy module descriptions. Each module description begins with a BEGIN <module\_type>. Look for the MicroBlaze core and the snoopy core and add the following two lines just before both END statements:

```
PORT PC_EX = PC_EX  
PORT valid_instr = valid_instr
```

9. Save and close the mhs file. Reopen XPS and your current project. Go to the Add/Edit cores menu. If you look under the ports tab, you should now see that the valid\_instr and PC\_EX ports of the processor have been connected to snoopy.
10. The core lets you set a reset address for clearing the counters, the number of counters you want to use (maximum of 16) and the upper and lower bounds for the instruction addresses. Choose an address that is within the address range assigned to the core. (Hints: if you used the suggested address range for the core, the default reset address, 0xfffffe4, will be fine. Remember this is found under the Parameters tab.)
11. Remove the old system source file and add the new source file you copied into the code directory, system.c. Compile the source program with ut optimizations and including debug flags to see the generated executable. If you get an error, go to the compiler options menu and verify that the executable is being placed in the appropriate directory. Disassemble your executable into a file.
12. Open the disassembled file to determine the address ranges you will profile. You will be selecting contiguous address ranges to profile based on function calls. Edit the parameters for the snoopy core to profile the following functions:

Counter	Function
0	_start
1	exit
2	_crtinit
3	main
4	_exception_handler
5	_interrupt_handler
6	_program_clean
7	_program_init

8	print
9	putnum
10	outbyte
11	XUartLite_SendByte
12	XUartLite_RecvByte
13	All Functions (complete program)

The lowerbound should be the starting address of the function and the upperbound should be the address of the last instruction in the function. Change the default value for the NUM\_COUNTERS parameter to 14

13. Now you can generate the bitstream to download onto the FPGA.
14. After the bitstream has been downloaded onto the FPGA, start up the XMD window. Connect to the XMD stub and then download the executable onto the board using the XMD download command **dow** (*Remember that when you start up the XMD window you will be in the project root directory and the executable.elf file is in the ./microblaze\_0/code/ subdirectory*).
15. The sst script you copied to your project directory can be used to read the counters. To reset them, use the memory write command in the XMD window: *mwr reset\_address <value>*, where the reset\_address is one of the core parameters. The counters are designed to reset independent of the written value. Type "source sst" and a file res.out will be generated. Open the file res.out to view the values stored in each of the 64-bit counters. The most significant portion is stored in 0x00 and the least significant portion in the 0x04 address. As you have just reset the counters, all the values should be zero.
16. After you have run your application, you can check the new counter values again using the sst script. Open res.out to see what the values for each counter are. Are they all non-zero values? Why or why not? (*Hint: look at the disassembled code to understand the values.*)
17. Now copy res.out to res.bak. Reset the counter values as previously instructed and rerun the program. Run the sst script again and look at the values of the counters. Are they all the same as the values in your res.bak file. Why or why not? You should compare your results to those found in the file example\_results.txt to verify the counter values.
18. You can also use the counters to determine how many clock cycles are required to execute a single instruction by setting both the upper and lower bounds to the same address. If the instruction is executed only once, the number of clock cycles should be the same as what is specified in the MicroBlaze processor manual. You can use any remaining lab time to

familiarize yourself with the available XMD commands or change the core parameters to better understand how your application runs on the MicroBlaze processor.