Introducton to CMOS VLSI Design (E158)

Lab 1: MuddLib

Harris

In this class, we will be building and using MuddLib, a standard cell library. The library is defined with strict rules so that cells snap together gracefully. In this lab, you will add a cell to MuddLib.

I. Library Overview

The HMC MuddLib contains a broad range of standard cells optimized for both datapath and control circuits. MuddLib is found in the E158/Labs directory on Charlie. Some key features of the library include:

- Many power levels for each gate
- 90 λ cell height
- Inputs and outputs on metal2 in consistent locations
- Vertical metal2 wires on 8 λ pitch
- Horizontal metal3 wires on 10 λ pitch (with room for 7 to pass over each cell)
- Datapath cell flavors for minimum area
- Control cell flavors containing local inverters to simplify interfacing
- Wide range of flip-flops with scan, reset, set, and enable features

An nMOS transistor of width 6 λ is defined to have unit drive. We assume that pMOS transistors have half the drive of comparable nMOS transistors. Cell names include the power level. For example, nand2_1x, nand2_1_5x, nand2_2x, and nand2_4x are four flavors of 2-input NAND gates with 1x, 1.5x, 2x, and 4x unit drive. The higher power level cells have proportionally larger input capacitance. Cells also come in a _lp low power level that use minimum size (4 λ) transistors throughout. These cells tend to have poor and unequal delays but save power on noncritical paths.

Cells are sized to achieve minimum average delay rather than equal rise and fall delays. This has a further benefit of reducing area and average power consumption. However, the rising and falling resistances may be different than the average resistance. Appendix A describes how the cell sizes are computed to achieve these goals.

Some cells come in both _dp and a _c flavors to optimize for both datapaths and control circuits. The datapath flavors factor out all unnecessary inverters. For example, a mux2_dp accepts both s and s_b select signals. The complementary select signal can be produced using a large inverter in the zipper and can be driven to all the multiplexers in the datapath. A mux2_c accepts a single select signal s. It is intended for use in

synthesized controllers where the inverter is best integrated into the multiplexer to simplify routing.

II. Schematic guidelines

Names are case-sensitive. For ordinary gates, export the inputs as a, b, c, … Input a should be the outer input (closer to GND for series nMOS stacks, or closer to VDD for series pMOS stacks). Name the output y. Inverted inputs should be followed with a _b suffix (e.g. en_b, clk_b). Outputs do not have a _b unless both true and complementary forms are produced.

Draw neat schematics. Place series transistors one grid unit apart. Place the nMOS and pMOS transistors two grid units apart. Avoid four-way junctions except when necessary.

Noninverting gates should be drawn as a flat schematic rather than hierarchically composed from smaller inverting gates. This way, the layouts can be fully compacted.

All schematics should be exhaustively simulated with IRSIM to confirm functionality. Weird cells such as SRAMs and critical circuits such as flip-flops should be simulated with HSPICE to guarantee functionality.

III. Symbol Guidelines

Symbols should be automatically generated from the schematic so that exports appear neatly. The body of the symbol should be deleted and redrawn with the appropriate shape for the gate. A typical logic gate should be about 6 grid units square; complex gates should be slightly larger. Draw inversion bubbles wherever they are appropriate, including on inverted inputs such as en_b.

IV. Layout Guidelines

The goal of the layout guidelines is to produce standard cells that snap together like Legos and that are easy to wire up. The cells use metal1 horizontal, metal2 vertical, and metal3 horizontal. Transistors should all be oriented with vertical polysilicon.

Use the templates wherever possible. The cells should have VDD and GND on a 90 λ center-to-center pitch. These power and ground busses should be 8 λ wide metal1. Each cell should have exactly one vdd and one gnd export. Exports are case sensitive.

The cells should conform to the wiring grids shown below. Well and substrate contacts should be placed on an 8 λ pitch centered under the VDD and GND busses. The busses should extend 2 λ beyond the contacts (and the cell contents) on both sides. Nothing in the cell should extend into this buffer zone lest it cause conflicts in wiring to other cells. All exports should be placed on metal 2. Usually exports will be placed on a metal1-metal2 via, but the export can also be placed on a metal2 pin on a longer stretch of metal2. No two exports should be placed in the same column.

Design the cell to use as few columns as possible. Sometimes this involves a clever layout. For example,

When multiple sizes of the same cell exist, exports should be placed in the same location on all cells whenever possible so that one size may easily be changed with another. However, certain high-drive versions of cells may require extra columns and hence cannot be interchanged with smaller cells. It is usually most efficient to start with the highest-drive version that fits in the minimal number of columns. Do the layout and place the pins. Then size down the transistors while keeping the pins fixed. Finally, do layouts of any high-drive versions that require additional columns.

Series transistors should be compacted to 3 λ apart to minimize diffusion capacitance. Contacted transistors should use a contact of the same size as the transistor whenever possible.

All geometry should be on a 1 λ grid. Polysilicon wires should be 2 λ wide. All other internal wires should be 4 λ wide. Metal2 should be used within the cell only when absolutely necessary, and only in vertical lines on the routing grid. Metal3 should never be used within the cell.

Providing enough space for inputs can be difficult on large gates. Polysilicon wires can be jogged where necessary to make room for contacts. Vias can also be stacked on top of contacts where necessary.

The widest transistors that can fit in a cell are 37 λ for pMOS and 27 λ for nMOS. Because of pin and wiring constraints, not all cells can actually be this big.

The layout should pass DRC, NCC, and ERC.

V. Characterization Guidelines

You do not need to characterize the cells in this lab. These note apply to the team building the Databook Generator.

Ideally, each cell will be characterized. A datasheet would be generated, either automatically, or manually. The datasheet should include at least the following information on each cell:

- The logic function and symbol
- Cell width
- Input capacitance of each pin (for power consumption)

- Internal switching capacitance (for power consumption)
- Delay characterization (on a per-pin basis)
    - Input capacitance
    - Parasitic delay (rising and falling) (ps)
    - Delay / load (ps / pF)
    - g_u, g_d, g_{avg}: Logical effort (rising and falling)
    - p_u, p_d, p_{avg}: Parasitic delay (rising and falling, normalized)
- Power characterization
    - Input capacitance of each pin (note, different than for delay)
    - Internal switching capacitance
- Schematic and layout of at least one power level

Most of these entries will have one column for each power level, along with an average across all power levels.

VI. Notes on Cells

Datapath multiplexers come in both inverting and noninverting flavors. An M-input datapath multiplexer consists of N tristates inverters in parallel followed by an optional output inverter. The datapath multiplexers minimize the delay from the data input at the expense of delay decoding the select inputs in the zipper. Control multiplexers come in noninverting flavors only because the inverter provides useful amplification before driving long control wires. They accept a minimum number of encoded select signals and use one or two levels of transmission gates between the input and output inverters.

Flip-flops are all scannable except flop_noscan. Datapath flops receive both clk and clk_b. Control flops receive only clk, and generate both local clocks using one and two inverters.

Invbuf cells contain chains of one and two inverters to produce true and complementary versions of control signals. They are intended for use in zippers. Clkbuf cells are similar, but are designed for equal rise and fall times to avoid distorting the duty cycle of the clock.

VII. Adding Cells

Copy MuddLib into your directory and name it MuddLib_xx, where _xx are your initials.

Cells have been randomly assigned to students. Refer to the spreadsheet to find which one you own.

Your task in this lab is to add new cell layouts to MuddLib. All of the cells in the library have a schematic and symbol for one of the sizes. Refer to the MuddLib Standard Cell Library spreadsheet to determine which sizes should exist. Copy the schematic and symbol and change the names and sizes (including the text on the icon) to reflect the new sizes. Sometimes a cell has many dimensions listed for a given drive strength; these refer

to the various stages of a gate. For example, a 4x mux2_c is given as 9, 6 + 18, 12 + 37, 27. This means that pMOS and nMOS should be 9 and 6 λ wide for the input inverters, 18 and 12 for the transistors in the body of the multiplexer, and 37 and 27 for the output inverter. Similarly, the xor3_1x is listed as 30,21 + 9,6 and the xor3_4x is listed as 9,6 + 18,12 + 37,27. The input inverters are 9,6 in both cases. For the _1x cell, the body of the gate is 30, 21. For the _4x cell, the body couldn't be 4 times larger. Instead, it is 18, 12, followed by a 4x (37,27) inverter. The polarity is exchanged so that the output still performs XOR. If you have questions, ask.

Create a layout for the biggest size cell that can be done without folding. Where possible, you can steal and modify layout from similar cells that are already complete. This especially applies to flip-flops. It is usually best to sketch a stick diagram first before starting the time-consuming layout. Print a color copy of your layout on a transparency and bring it to class on January 25. Color transparencies are available for the manual feed tray of C3PO in the ECF. Be sure that you conform to all of the design guidelines. The class will review all of the cells during lecture. (Note that the largest cell sizes are just estimates of what will fit. You may discover that it is infeasible to do a layout of one of the larger cells; for example, there may not be enough room for the wires. If the largest cell is infeasible, inform Prof. Harris and skip down to the next largest cell.)

Based on any feedback from the class, make any improvements. Then duplicate the layout for the other drive strengths and change the transistor sizes while attempting to keep the ports in the same location. Email your final MuddLib_xx to Prof. Harris before class on January 30; this is all that you have to turn in for this lab.

Library cell design is done to an unusually high standard. You will receive an A on this lab if Prof. Harris would make no change to your submitted circuits. You will receive a B if minor changes would be recommended but the cell is usable as completed. You will receive a C or lower if the cell needs to be corrected by somebody else before it is usable.

A. Cell Sizing Principles

Define a 6-λ wide nMOS transistor to have unit resistance and unit capacitance. Assume pMOS transistors have half the mobility, so a 12-λ pMOS has unit resistance.

A traditional 1x cell would be designed to have unit resistance pulling up and down. Hence, a unit inverter would be 12/6, a unit NAND2 would be 12/12, and a unit NOR would be 24/6. In the muddlib, cells are designed to minimize average delay rather than to have equal rising and falling resistances. This leads to a P/N ratio that is the square root of the ratio that gives equal delays (see *Logical Effort* by Sutherland, Sproull, and Harris). Such a P/N ratio also reduces cell sizes and capacitance as well as slightly improving delay. However, the rise and fall times are no longer equal.

For example, consider a unit inverter. It should have a P/N ratio of $\sqrt{2}$:1 and an average resistance of 1. If the nMOS transistor width is W, the pMOS must be $W\sqrt{2}$. The average resistance is:

$$R = \frac{\frac{6}{W} + \frac{12}{W\sqrt{2}}}{2}$$

For unit resistance, this can be solved for W = 7.24, leading to 10.24/7.24 transistor widths. We round to the nearest integer: 10/7. A _2x inverter uses twice these sizes: 20/14. A _4x inverter should use 40/28 transistors. However, the largest that fit without folding are 37/27. This is close enough, so we use these sizes rather than folding the transistors.

The NAND has a 1:1 P/N ratio, so the 1x cell still has 12/12 transistors. The NOR has a $\sqrt{4}$:1 P/N ratio, so it uses 16/8 transistors in the 1x cell.