CMOS VLSI Design

Lab 1: Gate Design

Harris

The only way to become a good chip designer is to design chips. This is the first of five labs in which you will use the Electric VLSI Design System to design the 8-bit MIPS microprocessor described in the CMOS VLSI Design book. The labs will guide you through mastering schematic entry, layout, simulation, and verification of a complex system. Chip design is a very time-consuming activity, so certain portions of the processor will be provided for you to reduce the tedious work and to illustrate good design style. The design you do in each lab will form a component of the microprocessor, so careful work on each lab will save you time at the end.

Begin by reviewing the MIPS processor architecture and multicycle microarchitecure covered in Section 1.7 of CMOS VLSI Design. You may also wish to refer to Section 5.5 of Patterson & Hennessy, Computer Organization and Design (3rd Edition).

This lab guides you through the design of a 2-input NAND gate. You will draw and simulate schematics. Then you will draw the layout and verify that it satisfies design rules and matches the layout. Using your NAND gate and an inverter, you'll design a 2-input AND gate. Finally, you'll design your own 2-input NOR and OR gates.

In Lab 2, you will design a full adder schematic and layout to learn to handle a more complicated cell. In Lab 3, you will build an arithmetic/logic unit from your AND, OR, and full adder. You will attach it to a bitslice to produce the MIPS datapath. In Lab 4, you will modify the Verilog code for the MIPS controller to implement the ADDI instruction, then synthesize and place & route it. Finally, in Lab 5, you will put the entire processor together, test it, and generate files from which it could be manufactured. At the end of these five labs, you will be prepared to independently design a chip of your own. On average, previous students have taken about 6 hours to complete Lab 1 and 10 hours each for Labs 2-5.

1. The Electric VLSI Design System

Integrated circuits have become sufficiently complex that Computer-Aided Design (CAD) tools are essential; nobody could design a 100-million transistor chip by hand on a schedule that would complete the chip while it was still interesting. We will use CAD tools for schematic entry, layout, simulation, and verification. Unfortunately, no CAD tools are perfect.

The leading industry-standard tool is made by Cadence. It normally sells for six figures per seat, but is available at extremely generous discounts to universities. However, it runs only on Unix and is nearly a full-time job to setup and maintain. Mentor Graphics also produces professional CAD tools of a similar nature. The Tanner tools run on NT and are easier to support, but cost more than Cadence for universities! The freely available Magic and Sue tools are popular at many universities, but Magic again is limited to Unix and has a primitive, albeit powerful, user interface. All of these tools have their fair share of bugs.

The Electric VLSI Design System is a computer-aided design tool developed by Steve Rubin. It has powerful notions of connectivity, runs on Windows, Unix, and Macintosh, and is very well integrated. It is also GNU-licensed so you may freely download your own personal copy and even modify the source code. The drawback is that many features of Electric are still in development, so the tool crashes more often than you would like and sometimes does unintuitive things. These labs are written for Electric 6.08ac (Feb. 24, 2003), which is available precompiled for Windows on the web page. You may also download newer versions and versions for different operating systems from www.staticfreesoft.com, but the tool changes quickly enough that you are likely to encounter incompatibilities. Electric has been rewritten in Java since these labs were developed. The new version has many improvements but is not compatible with these labs.

Be sure to save often so you do not lose too much work. While Electric rarely corrupts libraries, it is wise to periodically create backups of your library just in case.

2. Getting Started

Copy the mipsparts.elib library from the web page to your account and rename it to lab1_xx where xx are your initials. This library contains many parts of the MIPS processor that are provided to you. You will add your new designs to the library as you work through the labs. Do not modify the parts that were already provided lest you break them in ways that will not become apparent until you put the entire system together in Lab 5.

Double-click on Electric to start the program. Dismiss the splash screen. Choose Info • See Manual from the menu to bring up the Electric manual in a web browser. The manual is also available online at www.staticfreesoft.com. Skim through the following sections:

      Chapter I: 1-2, 6-9
      Chapter II: 1-6
      Chapter III: 1-12

Don't worry if it doesn't all make sense yet. After you complete this lab, go back and skim over the sections that you initially found confusing. Refer back to the other chapters of the manual as you need help with specific features of Electric.

3. Schematic Entry

Your first task is to create a schematic for a 2-input NAND gate. Each design is kept in a *facet*; for example, your schematic will be in the nand2{sch} facet, while your layout will eventually go in the nand2{lay} facet and your AND gate will go in the and2{sch} facet. Choose File • Open Library to open your lab1_xx library. Choose Facet • Edit Facet to bring up the Facet dialog. Click New Facet. Enter nand2 as the facet name and schematic as the view. A new editing window will appear with the title lab1_xx:nand2{sch} indicating the library, facet name, and view.

Electric defines various technologies for schematics and layout. To draw transistor-level schematics, you will need to select the Analog Schematic technology by choosing Technology • Change Current Technology and selecting the *schematic, analog* technology. This technology file contains basic circuit elements such as transistors, resistors, capacitors, and power and ground.

Your goal is to draw a gate like the one shown in Figure 1. Choose Windows • Toggle Grid to turn on a grid to help you align objects. Left-click on an nMOS transistor symbol in the components menu on the left side of the screen. Left-click on your schematic window to drop the transistor into your layout. Repeat until you have two nMOS transistors, two pMOS transistors, the circular power symbol, and the triangular ground symbol arranged on the page. You may move the objects around by left-clicking and dragging. The transistors default to a width/length value of 2/2. Double-click on the pMOS transistor and change its width to 12. Recall that nMOS transistors are roughly twice as strong as pMOS transistors. So a single nMOS transistor would only have to be 6 wide. However, because the nMOS transistors are in series, they should also be 12 wide.


Figure 1: nand2{sch}

Now, make the connections. Left-click on a port such as the gate, source, or drain of a transistor. Right-click on another port to create a wire connecting the ports. Continue until all the blue wiring is completed.

Finally, you will need to provide *exports* defining inputs and outputs of the facet. Left click on the end of a wire where you need to create the export for input *a*. You should see a small square box highlighted at the end of the wire. If the entire wire is highlighted, you clicked on the middle of the wire instead of the end, so try again. Once you have selected the end of the wire, choose Export • Create Export. Give your export the name *a*. Give it the characteristic Input. Repeat with the other input *b*. Export *y* as an Output.

Use File • Save (*Ctrl-s*) to save your library. Get into the habit of saving often because Electric crashes fairly often. Also, learn the keyboard shortcuts for the commands you use frequently.

4. Switch-Level Simulation

Our next step is to simulate the schematic to ensure it is correct. Electric has two built-in *switch-level simulators*: IRSIM and ALS. ALS is buggy, so we will use IRSIM. IRSIM treats transistors as switches that may be ON or OFF; it also understands resistance and capacitance based on transistor sizes and crudely estimates switching delays.

First select Tools • Simulation (Built-in) • Simulation Options. Make sure that IRSIM is selected for Simulation engine. Check that the scmos0.3.prm parameter file is selected; this contains resistance and capacitance data for the technology we are using. Start the simulation by selecting Tools • Simulation (Built-in) • Simulate. A waveform window will appear listing each of the *nets* in your design. If you created your exports properly, you will see nets for *a*, *b*, and *y*, as shown in Figure 2. You will also see an unnamed nets for the node between the series nMOS transistors; in the figure it is named net3.

The simulator has two vertical white cursors. The primary cursor with no x is used to create stimulus. Click and drag the cursor near the left edge (near time 0). Click on the *a* input in the simulation window and press *h* or 1 to drive the input high. Drag the cursor to some later time. Click on the *b* input and press h to drive it high. Use the *l* or 0 key sometime later to drive *a* and *b* back low, as shown in the figure. Check that the *y* output matches the behavior you would expect for a NAND2 gate. If it does not, fix the bug in your schematic and resimulate.

The secondary cursor with an x is only used to measure delays relative to the first cursor. You will find that unloaded gate delays are under 100 ps. Gate delay depends on the capacitance being driven, so attaching a realistic load will slow the gate.

Use the Windows • Adjust Position • Tile Vertically command to arrange the windows so that you can see both the simulation window and your schematic at the same time. Watch the color-coding of the wires in the schematic change as you drag the primary cursor back and forth. On the schematic, magenta indicates high, blue indicates low, and

black indicates x, an illegal value. When a signal is selected in the waveform window, the corresponding net will be highlighted on the schematic. Similarly, when a net is selected in the schematic then its waveform representation will be highlighted. Watching the voltage levels change on the schematic is helpful for debugging problems. Study your simulation and determine why the node between the two nMOS transistors behaves as it does. A thick purple bar means that the state is floating, undefined, or could not be made into a definite logic level.
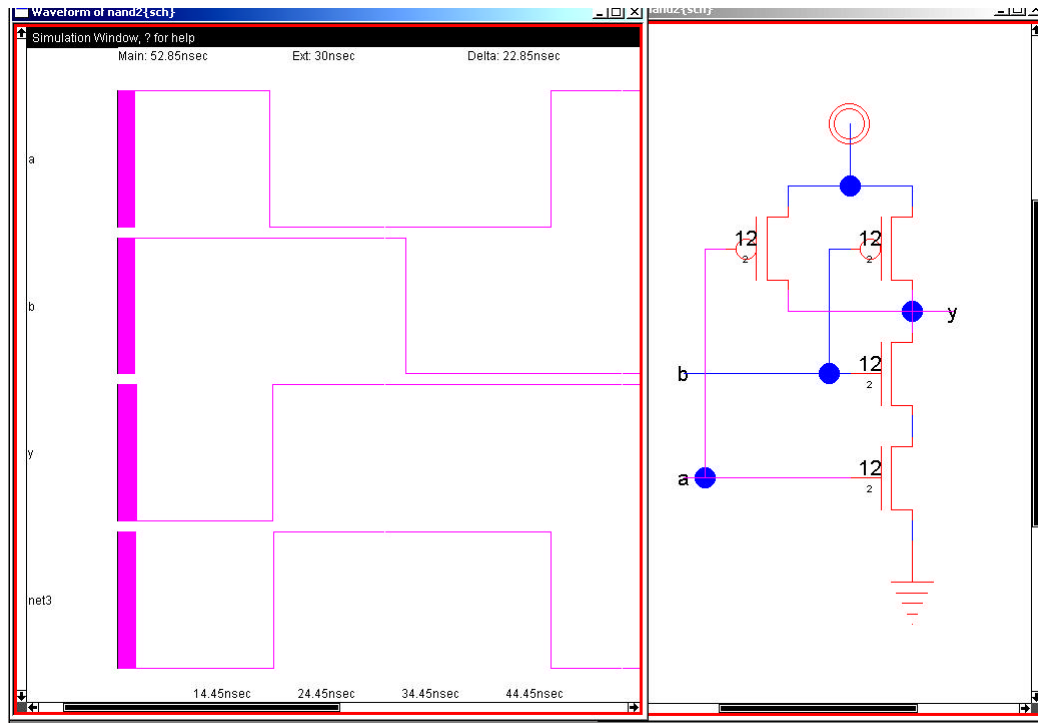

Figure 2: Simulation of nand2{sch}

Make a habit of simulating each facet after you draw it so you catch errors while the design is fresh in your mind.

Electric cannot print waveforms directly. To print your simulation results, press the P ("Preserve") key while in the waveform viewer. This will create a facet named nand2{sim} containing the waveforms. Open that facet and print it.

5. Layout

Now that you have a schematic simulating correctly, it is time to draw the layout. Choose Facet • Edit Facet to bring up the Facet dialog. Click New Facet. Enter nand2 as the facet name and layout as the view. We will be targeting the AMI 0.5 μm process but using MOSIS submicron scalable rules so we could easily adapt to the AMI 1.5 μm process or others. To setup the technology file, choose Technology • Change Current Technology and select *mocmos*, the MOSIS CMOS technology. Then choose Technology

• Change Units and set Lambda for mocmos to 600 half-milimicrons, i.e. 0.3 μm.[1] Use the Arc • New Arc Options to set the default width for Metal-1 to 4 and for Metal-2 to 4 for convenience of layout. Finally, choose Technology • Technology Options and select 3 metal layers, submicron rules, and alternate active and poly contact rules. The submicron rules are documented on the MOSIS web page. Alternate Active and Poly contact rules are an older, cleaner set of design rules that don't involve half-lambda dimensions.

Your goal is to draw a layout like the one shown in Figure 3. It is important to choose a consistent layout style so that various cells can "snap together." In this project's style, power and ground run horizontally in metal2 at the top and bottom of the cell, respectively. The spacing between power and ground is 80λ center to center. No other metal2 is used in the cell, allowing the designer to connect cells with metal 2 over the top later on. nMOS transistors occupy the bottom half of the cell and pMOS transistors occupy the top half. Each cell has at least one well and substrate contact. Inputs and outputs are given metal1 ports within the cell.

You may find it convenient to have another sample of layout visible on the screen while you draw your gate. Use the Edit • New Facet Instance command and select inv{lay}, then click to drop this inverter in the layout window. Highlight the inverter and use the Facet • Expand Facet Instances • All the way command to view the contents of the inverter. Study the inverter until you understand what each piece represents.

---

[1] Lambda is normally defined as half of the minimum drawn gate length. Therefore, the minimum drawn gate length (polysilicon width) is 0.6 μm even though the vendor describes the process is 0.5 μm.

Figure 3: nand2{lay}

Start by drawing your nMOS transistors. Recall that an nMOS transistor is formed when polysilicon crosses N-diffusion. N-diffusion is represented in Electric as green diffusion surrounded by a dotted yellow N-select layer all within a hashed black P-well background. This set of layers is conveniently provided as a 3-terminal transistor node in Electric. Move the mouse to the components menu on the left side of the screen. As you move the mouse over various objects, the node name will appear on the status line next to the word NODE near the bottom left corner of the screen. Left click on the N-Transistor, shown in Figure 4, and click again in the layout window to drop the transistor in place. Use the Edit • Rotate • 90 Degrees Counterclockwise command to rotate the transistor so that the red polysilicon gate is oriented vertically. There are two nMOS transistors in series in a 2-input NAND gate, so we would like to make each wider to compensate. Double-click on the transistor. In the node information dialog, adjust the width to 12.

Figure 4: nMOS transistor before and after rotation and sizing

We need two transistors in series, so copy and paste the transistor you have drawn or use the Edit ● Duplicate command. Drag the two transistors along side each other so they are not quite touching. Left click the diffusion (source/drain) of one of the transistors and right click on the diffusion of the other transistor to connect the two. Then drag the two transistors until the polysilicon gates are 3λ apart, looking like they do in Figure 3. You will probably find it helpful to turn on the grid using the Windows ● Toggle Grid command. The grid defaults to small dots every lambda and large dots every 10λ. You can change this with the Windows ● Grid Options command. Also by default, objects snap to a 1-λ grid. This can be changed by using Windows ● Alignment Options. You can also move objects around with the arrow keys on the keyboard. Pressing the *h* or *f* keys cause the arrows to move objects by a half or full lambda, respectively. You will avoid messy problems by keeping your layout on a lambda grid as much as possible. Inevitably, though, you will create structures that are an odd number of lambda in width and thus will have either centers or edges on a half-lambda boundary.

Electric has an interactive *design rule checker* (DRC). If you place elements too closely together, it will report errors in the Electric Messages window. Try dragging one of the transistors until its gate is only 2λ from the other. Observe the DRC error. Then drag the transistors back to proper spacing. When you are in doubt about spacing, use the Tools ● DRC ● Hierarchical Check command to ask Electric to recheck the entire facet and any subfacets it might contain. Use the > key to sequence through and highlight errors.

An Aside on Design Rules

The definitive design rules are available on the MOSIS web page. MOSIS is a service that collects small orders for designs and combines them into an order large enough to interest a semiconductor manufacturing facility. For best density, one could optimize for a particular fabrication process and design in units of microns rather than lambda. However, the layout would require redrawing when ported to a different process because the design rules and feature size will have changed. MOSIS has developed a set of *Scalable CMOS* design rules that are sufficiently conservative to work for virtually all processes. The original rules, *SCMOS*, apply to older

processes. We will be using the *SUBM* submicron rules that are even more conservative and suffice for more advanced processes. They are adequate for both the AMI 0.5 and 1.5 micron processes. Finally, the *DEEP* deep submicron rules are slightly more conservative and are necessary for best performance in the 0.25 μm and below processes. For historical reasons, all three sets of rules are selected in Electric by choosing the scmossub process, then using the Technology Options dialog. Tables 3 and 4 on the MOSIS web page list the differences between these design rules.

http://www.mosis.org/Technical/Designrules/scmos/scmos-main.html

We will be fabricating using the SCNE technology code defined in Table 5 of the web page. SCNE means Scalable CMOS with N-wells and an Electrode layer, i.e. polysilcon 2. Click on each of the layers in the table to see the design rules. For example, the Metal1 rules are shown in Figure 5 below. We will follow the SUBM rule set, requiring metal width and spacing of 3λ.

SCMOS Layout Rules - Metal1

| Rule	Description	Lambda		
		SCMOS	SUBM	DEEP
7.1	Minimum width	3	3	3
7.2	Minimum spacing	2	3	3
7.3	Minimum overlap of any contact	1	1	1
7.4	Minimum spacing when either metal line is wider than 10 lambda	4	6	6

Figure 5: SCMOS Metal1 Rules

Next we will create the contacts from the N-diffusion to metal1. Diffusion is also refered to as *active area*. Drop a square of Metal-1-N-Active-Contact in the layout window and double-click to change the properties to a Y size of 12. You will need a second contact for the other end of the series stack of nMOS transistors, so duplicate the contact you

have drawn. Move the contacts near each end of the transistor stack and draw diffusion lines to connect the transistors. Then move the contacts even closer; you only need a gap of 1λ between the metal and polysilicon. Use the design rule checker to ensure you are as close as possible but no closer. Using similar steps, draw two pMOS transistors in parallel and create contacts from the P-diffusion to metal1. At this point, your layout should look something like Figure 6:


Figure 6: Contacted transistors for nand2 layout

Draw wires to connect the polysilicon gates, forming inputs *a* and *b*, and the metal1 output node *y*. Then add metal2 power and ground lines. You can add a line by selecting a pin such as metal-2-pin from the palette then right-clicking nearby to draw the line. Use the grid to ensure they are 80λ apart from the center of each line. This is the same spacing as the power/ground lines of the inverter. A metal1-metal2-contact, also known as a via, is required to connect the metal1 to the metal2 lines. Left-click on an active contact. Right-click on the ground line. Electric will automatically create the necessary via for you while making the connection. Complete the other connections to power and ground. Let power and ground extend 2 lambda beyond the contents of the cell (excluding wells) on either side so that cells may "snap together" with their contents separated by 4λ so design rules are satisfied.

Recall that well contacts are required to keep the diodes between the wells and source/drain diffusion reverse biased. We will place an N-well contact and a P-well contact in each cell. It is often easiest to drop the Metal1-N-Well-Con near the desired destination (near VDD), then right click on the power line to create the via. Then drag the contact until it overlaps the via to form a stack of N+ diffusion, the diffusion to metal1 contact, metal1, the metal1-metal2 via, and metal2. Repeat with the P-well.

In our datapath design style, we will be connecting gates with horizontal metal2 lines. Metal2 cannot connect directly to the polysilicon gates. Therefore, we will add contacts from the polysilicon gate inputs to metal1 to facilitate connections later in our design. Place a metal-1-polysilicon-1-contact near the left polysilicon gate. Connect it to the polysilicon gate and drag it near the gate. You will find a 3λ separation requirement from

the metal1 in the contact to the metal forming the output *y*. Add a short strip of metal1 near the contact to give yourself a landing pad for a via later in the design. You may find Electric wants to draw your strip from the contact in polysilicon rather than metal1. To tell Electric explicitly which layer you want, move the mouse over the palette until it is over the blue Metal-1 arc square and click. Then draw your wire.

Electric is agnostic about the polarity of well and substrate; it generates both n- and p-well layers. In our process that has a p-substrate already, the p-well, indicated by black slanting lines, will be ignored. The n-well, indicated by small black dots, will define the well on the chip. Electric only generates enough well to surround the n and p diffusion regions of the chip. It is a good idea to create rectangles of well to entirely cover each cell so that when you abut multiple cells you don't end up with awkward gaps between wells that cause design rule errors. To do this, choose Edit • New Pure Layer Node. Create an N-Well-Node and a P-Well-Node. Double-click on each to change the N-Well-Node size so that it entirely covers the existing n-well. Similarly, change the P-Well-Node. You will find the pure layer nodes are annoying because you will tend to select them when you really wanted to select a transistor or wire. To avoid this problem, shift-click to select both pure layer well nodes and use the Edit • Selection • Make Selected Hard to make the layers hard to select. You will then need to hold the Alt key while clicking to select a well. You can use the Make Selected Easy command if you want to restore a well to be easily selected. Electric also provides the Edit • New Special Object • Coverage Implants command that automatically creates hard-to-select pure layer nodes for N and P wells that is convenient for simple geometries that the tool understands.

Create exports for the cell. When you use the cell in another design, the exports define the locations that you can connect to the cell. Click near the end of the short metal1 input lines that you just drew on the left gate. You will see a small white box highlighted, corresponding to the pin at the end of the cell. If you accidentally selected the entire line instead, click elsewhere in space to deselect the line, then try again to find the pin. You may also try holding the *ctrl* key while clicking to cycle through selections. Add an input export called *a*. Repeat for input *b*. Export output *y* from the metal line connecting the nMOS and pMOS transistors. You may have to place an extra pin and connect it to the output line to give yourself a pin to export as *y*. Also export *vdd* and *gnd* from the metal 2 lines; these should be of type power and ground, respectively. Electric recognizes vdd and gnd as special names, so be sure to use these names.

A *facet center* looks like a small cross and indicates a reference point for your cell. Having a facet center avoids some off-grid weirdnesses in Electric. If there is not one already, place a facet center in the lower left corner[2] of the icon (on top of the *gnd* export) using the Edit • New Special Object • Facet Center command. Drag the facet center where you want it. If you deselect it, then try to select it again, you'll find you can't pick it up. The facet center is by default *hard to select*. To select hard to select nodes on a Windows machine, hold the Alt key down while clicking. On other machines, see the Special Select key combination in section 1.8 of the Electric manual. Using facet

---

[2] It is not strictly necessary for the facet center to be in the true center of the facet. Electric just uses the center to provide a reference point within the facet. This helps Electric keep cells on a clean grid.

centers on all your icons and layout facets will avoid nasty off-grid errors later in your work.

Your design should now resemble Figure 3 and should pass DRC. When you are done drawing the nand2 layout, click on the inverter and press delete to remove it from the facet. Save the library.

Electric has a fun feature to show a 3D rendition of your layout. Use the Windows • 3D Display • View in 3 Dimensions command. Click and drag with the mouse to rotate the layout. You will see the layer stack from diffusion on the bottom through polysilicon, metal1 and metal2. Vias are shown in white and contacts from metal1 to poly or diffusion in black. Does the 3-D visualization match your mental picture of the layout? When you are done, use Windows • 2D Display • View in 2 Dimensions to restore normal viewing.

6. Layout Verification

Layout verification involves more checks than schematic verification. The checks include *Design Rule Checks* (DRC), *Electrical Rule Checks* (ERC), *Network Consistency Checking* (NCC), and simulation.

You will likely find yourself invoking these menu options often and may wish to give yourself keyboard shortcuts. For example, you might use the Info • User Interface • Quick Key Options dialog to map the Tools • DRC • Hierarchical Check command to the F1 function key, Tools • Electrical Rules • Analyze Wells to F2, and Tools • Network • NCC Control and Options to F3.

First, be sure the design satisfies the layout design rules by running a hierarchical DRC. This checks the layout and any facets it might incorporate; in this case the nand2 is a *leaf* cell and has no subfacets. Correct any DRC violations that might remain.

Next, run Tools • Electrical Rules • Analyze Wells. This check ensures that every N-well has a contact to VDD and every P-well has a contact to GND reasonably close by. ERC will report the number of transistors found. Check that this matches your expectation; it should be 2 nMOS and 2 pMOS for the nand2. It also reports the farthest distance of any part of the layout from a well or substrate contact; if this is greater than 100, consider adding more contacts to avoid latchup risks. If any errors are reported, fix them. A common cause of ERC violations is to not export power and ground or to give them the wrong names or the wrong type of export.

Next, simulate the layout. Apply a complete set of stimulus to the two inputs to convince yourself that the gate is working correctly.

Electric includes a powerful tool called a network consistency checker that checks that the schematic and layout are equivalent. This is especially valuable when your design is too complex to verify completely through simulation. The checker relies upon graph

theory algorithms. If your layout and schematic match, you have much greater confidence that a bug hasn't crept into your design. Unfortunately, if the two don't match, the tool becomes very confused and provides few hints. Also, NCC uses a rather complex algorithm and has been the source of quite a few bugs in Electric, so don't trust it blindly.

To use NCC, first use Facets • Edit Facet to be sure both the nand2{sch} and nand2{lay} facets are open. Close all other facets; NCC assumes the two windows open are the schematic and layout views of the cell being verified. Choose Tools • Network • NCC Control and Options. In the dialog box, set for all facets to Recurse through hierarchy. Check the Check export names, Check component sizes, and Ignore Power and Ground boxes. Then click Do NCC. If your design is correct, you should get a message of:

Comparing facet nand2{lay} with facet nand2{sch}
- Flattening hierarchy; Ignoring Power and Ground nets
- Parallel components not merged; Series transistors not merged
- Checking export names and component sizes
- No facet overrides
Extracting networks from nand2{lay}...
Extracting networks from nand2{sch}...
Facet nand2{lay} has 4 components, 4 nets; facet nand2{sch} has 4 components, 4 nets
Facets nand2{lay} and nand2{sch} are equivalent (0.11 seconds)

This message means that the layout and schematic each have four components, i.e. the two nMOS and two pMOS transistors. They also each have four nets: A, B, Y, and the wire between the series nMOS transistors. Note that power and ground are ignored and therefore do not appear in this count. NCC finds that the facets are equivalent.

If your design is not equivalent in any nontrivial way, NCC will likely mark too much mismatching to be useful. Identifying the error is part of the art of using NCC. If possible, simulate both the layout and schematic carefully to eliminate obvious errors. Carefully examine your design by hand to look for errors. Here are some common errors:

- The layout and schematic being compared should be the only two windows open.
- Port names must agree and be in the same order. For example, the order of series transistors in the nMOS stack on the NAND gate matters. One input is closer to GND than the other. If the inputs a and b are in the opposite order in the layout than in the schematic, you will see the message:
    ******* Export differences have been found! (0.01 seconds)
  If you press the > key repeatedly to show each error until no more errors exist, you will see the inputs highlighted and see the following messages printed:
    NCC: Export names 'nand2{sch}:b' and 'nand2{lay}:a' do not match
    NCC: Export names 'nand2{sch}:a' and 'nand2{lay}:b' do not match
- Transistor sizes must agree.
- Forgetting to export power and ground in the layout will cause many errors.
- Be sure the *Ignore Power and Ground* network option is checked. Otherwise, you will get errors like:
    Note: facet flop{lay} has 1 power and 1 ground net(s),
    while facet flop{sch} has 0 power and 0 ground net(s)

Another powerful tool for tracking problems is NCC Preanalysis. This command lists all the components and networks in the design. You should expect equal numbers of connections in the layout and schematic, so by looking for discrepancies, you can find clues about your design error. Try running Preanalysis on a good design so that you know what the report looks like when the design is correct; that way you will have an easier time interpreting the results when the design is incorrect.

If desperate, you can turn on the Verbose NCC (text) or (graphics) options in the NCC options dialog to get hints about how NCC seeks to match components between the layout and schematic.

7. Hierarchical Design

Now that you have a 2-input NAND gate, you can use it and an inverter to construct a 2-input AND gate. Such hierarchical design is very important in the design of complex systems. You have found that the layout of an individual cell can be quite time consuming. It is very helpful to reuse cells wherever possible to avoid unnecessary drawing. Moreover, hierarchical design makes fixing errors much easier. For example, if you had a chip with a thousand nand gates and made an error in the nand design, you would prefer to only have to fix one nand cell so that all thousand instances of the nand inherit the correction than to need to fix each nand individually.

Each schematic has a corresponding symbol, called an *icon*, used to represent the cell in a higher-level schematic. For example, open the inv{sch} and inv{ic} to see the inverter schematic and icon provided. You will need to create an icon for your 2-input NAND gate. When creating your icon, it is a good idea to keep everything aligned to the 1-λ grid, this will make connecting icons simpler and cleaner when you need it for another facet.

Open your nand2{sch} and choose View • Make Icon. Electric will create a generic icon based on the exports looking something like Figure 7. It will drop the icon in the schematic for handy reference; drag the icon away from the transistors so it leaves the schematic readable.


Figure 7: nand2{ic} from Make Icon

A schematic is easier to read when familiar icons are used instead of generic boxes. Modify the icon to look like Figure 8. Pay attention to the dimensions of the icon; the overall design will look more readable if icons are of consistent sizes.


Figure 8: nand2{ic} final version

Click on the icon and choose Facets • Down Hierarchy to drop in to the icon. The technology will automatically change to *artwork*. A palette will appear with various shapes. Delete the generic red box but leave the input and output lines. Turn on the grid.

If there is not already a facet center in the facet, place a *facet center* in the center of the icon using the Edit • New Special Object • Facet Center command. Drag the facet center where you want it.

The body of the NAND is formed from an open C-shaped polygon, a semicircle, and a small circle. To form the semicircle, place an unfilled circle. Double-click to change its size to 6x6 and to span only 180 degrees of the circle. Use the rotate commands under the Edit menu to rotate the semicircle into place. Place another circle and adjust its size to 1x1. You will need to change the alignment options under the Windows menu to 0.5 to move the circle into place, then set alignment back to 1. Alternatively, you can press h and use the arrow keys to move objects by ½ grid increments, then press f to return to full grid movement.

The opened-polygon shown in Figure 9 can be used to form the C-shaped body. Drop an opened-polygon object. Select it and choose Edit • Special Function • Outline Edit to enter outline edit mode. In this mode, you can use the left button to select and move points and the right button to create points. You should be able to form the shape with four clicks of the right button to define the four vertices. Outline edit mode is not entirely intuitive at first, but you will master it with practice. Choose Edit • Special Function • Exit Outline Edit when you are done. If your shape is incorrect, delete it, drop another opened-polygon, and try again.

Figure 9: Opened-Polygon

Electric is finicky about moving the lines with inputs or outputs. If you left-click and drag to select the line along with the input, everything moves as expected. If you try to move only the export name, it won't move as you might expect. Therefore, make a habit of

moving both the line and export simultaneously when editing icons. Note that the line is just an open-polygon and can be shortened if desired by entering Outline Edit mode.

Use the Text item in the artwork palette to place a label "nand2" in the center of the icon.

Now that you have an icon with three exports, create a new schematic called and2. Change the technology back to *schematic, analog* because you are drawing a schematic again now. Use Edit • New Facet Instance to create ("*instantiate*") the nand2{ic} and an inv{ic}. Wire the two together and create exports on inputs *a* and *b* and output *y*. Double-click on the wire between the two gates and give it a name like *yb* so you know what you are looking at in simulation. It is good practice to label every net in a design if practical. When you are done, your and2 schematic should look like Figure 10. If the line between the gates is black rather than blue, you neglected to return to the *schematic, analog* technology and were still drawing using the *artwork* palette.


Figure 10: and2{sch}

Simulate your and2 gate to ensure it works. You can use the Tools • Simulation • Down Hierarchy command to descend into the nand gate and see its internal node. In the schematic editor, you can double-click on each gate and assign it a name to help differentiate cells when you simulate.

Next, create a new layout called and2. Change the technology to mocmos. Instantiate the nand2{lay} and inv{lay} layouts. ALWAYS use the New Facet Instance to create layout from preexisting facets. NEVER build a cell by cutting and pasting entire existing cells. If you do, then make a correction to the original cell, your correction will not propagate to the new layout. However, this does mean that you will need to make all the connections to existing ports in your instantiated cells.

Initially the layouts appear as black boxes with ports. Select both and use the Facet • Expand Facet Instances • All the Way command to view the contents of each layout. Wire together power and ground. Move the cells together as closely as possible without violating design rules. You may need to place large blobs of pure layer nodes over the n-well and p-well to avoid introducing well-related errors from notches in the wells. Connect the output of the nand2 to the input of the inv using metal1. Remember that connections may only occur between the ports of the two cells. Also connect the power and ground lines of the cells using metal2. Export the two inputs, the output, and power and ground. The final and gate should resemble Figure 11. Run DRC, ERC, and simulation to verify your design.

Figure 11: and2{lay}

Finally, do a network compare. NCC has three modes in the NCC Options: Expand hierarchy, Recurse through hierarchy, and neither. Expand hierarchy smashes your entire design into one giant netlist for its internal comparison. Recurse through hierarchy checks each subfacet on each hierarchical level. NCC is a relatively new feature that has received many bug fixes, so it is wise to try both Expand hexarchy and Recurse through hierarchy to catch any problems that one mode misses. Do not check both Recurse through hierarchy and Expand hierarchy, as it will give meaningless results. In each case, be sure to check Check export names, Check component sizes, and Ignore Power and Ground. Do not set any individual facet overrides. After you check a facet, Electric marks it as clean and thus not in need of rechecking.

If NCC reports a problem, check that your inputs are in the correct order. If *a* and *b* are reversed between the layout and schematic, you will get an error that nodes are wired differently. If VDD or GND are not exported, you will also get many errors and your simulations will produce invalid levels. Tracking NCC mismatches is very difficult, so you are best off ensuring your design is correct by means of careful drawing and simulation rather than drawing sloppily and hoping to catch problems with the checker. Keep a list in your notebook of any design errors you made that caused NCC problems and note the symptoms reported by NCC. In the future, you can use this experience to help you find problems in complicated designs more quickly.

Now that your and2 gate is complete, use the Info • Check and Repair Libraries command to look for any inconsistencies in your library. You shouldn't expect any, but Electric has been known to do strange things from time to time, especially in the hands of novice users. Therefore, you may wish to check your library every few hours.

8. NOR / OR Gate Design

Your MIPS processor supports the OR instruction as well as the AND instruction. Therefore, you will need to create a 2-input NOR and a 2-input OR gate. Call these nor2 and or2. The nMOS transistors should be 4 wide and the pMOS 16. For each gate, create schematic and layout facets following the same steps you used with the nand2 / and2. Be sure to export the inputs, output, power, and ground. Simulate each to ensure correct operation. Run DRC, ERC, and NCC to check each layout. Be sure to save your work frequently. It is wise to keep a backup of your work from time to time in case your library becomes corrupted.

9. What to Turn In

Please provide a hard copy of each of the following items:

1. Please indicate how many hours you spent on this lab. This will not affect your grade, but will be helpful for calibrating the workload for the future.
2. A printout of your nor2 schematic.
3. A printout of your nor2 layout.
4. A printout of your or2 schematic.
5. A printout of your or2 layout.
6. Simulation waveforms demonstrating correct operation of the or2 layout.
7. What is the verification status of your or2 layout? Does it pass DRC? ERC? NCC?