Using LTspice in macOS — an introductory tutorial

In this tutorial, we will learn the details of using LTspice on a Mac to set up and analyze a circuit. In many ways the Mac and Windows versions of LTspice are quite similar, but there are some differences that can be confusing. This tutorial is intended to get us past that initial confusion.

The circuit used in the tutorial is boring and the DC analysis that we will do is also boring. But the focus here is on the setup, not the details of the specific circuit. Once we have become effective at setting up circuits in LTspice, we will be able to look at more interesting circuits and more interesting simulations. After we have finished this tutorial, we should look at other items on the SPICE site — in particular: transient analysis, op amps, and transistors[1]. Also, there are examples that pop up from time to time in the EE 201 (circuits) and EE 230 (electronics) class notes.

The first difference that we note is that the Mac version presents a very spartan user interface. Nearly all interaction occurs through contextual pop-up menus. A two-button mouse with a scroll wheel works well in this case — the pop-up menus are activated by clicking the right mouse button. With a one-button mouse or a trackpad, we would use “control-click” to bring up the menus.

A second difference between the versions is the method for specifying the type of analysis. On Windows, the simulation details are set up using a standard dialog box. On macOS, the simulation is specified using a “directive” statement, which is embedded in the schematic. This is not difficult, but the directive statement must be in a specific format — without either prior knowledge or some guidance, we would be lost.

Finally, the Mac version has some weaknesses in presenting the simulation results. In particular, DC results are a bit awkward on the Mac, as we will see. Also, the tools for manipulating graphs are not quite as complete on the Mac. However, there are work-arounds to the few shortcomings. Certainly, the differences are not serious enough to keep us from using our preferred computing platform.

The circuit

The circuit for this first LTspice tutorial should be familiar — it appeared several times in the EE201 class notes. It is a nice example of a circuit that can be handled using node-voltage, mesh-current, or superposition techniques.

Using LTspice, we would like to find the node voltages and component currents. We will define the bottom node as ground.

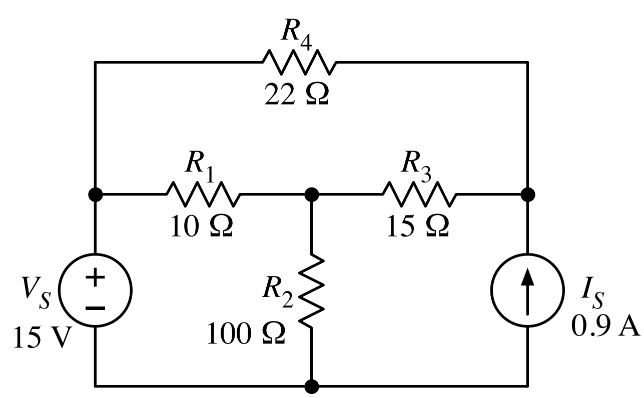

[1] If you are reading this in early 2022, you might note that there are no “other tutorials” on the SPICE site. My apologies for that — the other tutorials are planned and will appear in the coming weeks/months. –GT

Circuit schematic

In building the schematic, we don't need to be overly careful about placing the components in any particular order or arrangement. Names, locations, and orientations can all be changed later, if needed. However, if we know ahead of time how we want the schematic to appear, we can save a tiny bit of time by placing them carefully now.

1. Launch LTspice. A "Welcome to LTspice" dialog appears with some options for opening files. Choose "Start a new, blank schematic". An empty schematic-drawing window opens.

   We should probably give the file a name and save it. If we use the "File→Save" menu option, the file will be saved with a generic name like "Draft1.asc" in a folder named "LTspice" that has been created the documents folder of your user directory. If we use the "File→Save As" option, we can choose the name and location for the file. In the future, we can open existing files using "File→Open" or by double-clicking the file icon in the finder.

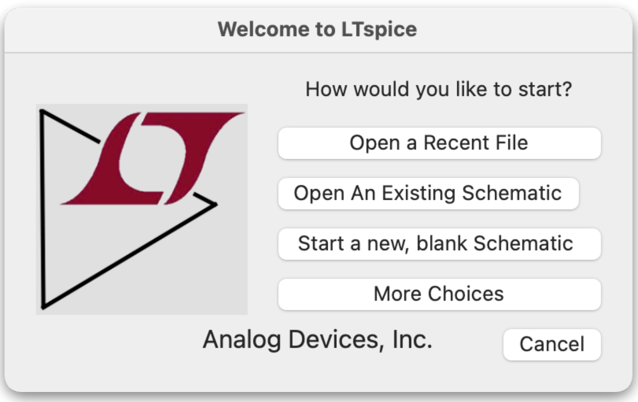

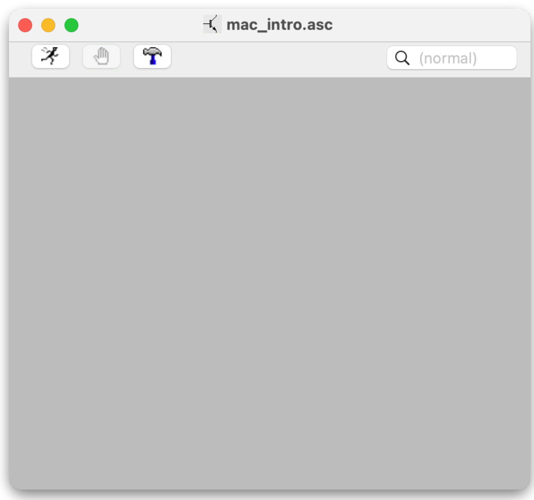

2. Place the voltage source. Right-click somewhere within the window. From the pop-up menu, choose "Draft→Component". The "Select Component Symbol" dialog opens. We could scroll through the long list of items to find the voltage source, but it is easier to use the search function. Start typing "voltage" into the search box, and within a couple of letters, the voltage-source component is selected and displayed at upper left.

   Click the "OK" button to close the dialog and return to the schematic-drawing window. A voltage-source symbol is attached to the cursor.

   Move the cursor to a point somewhere in the left-center portion of the window and left-click to place the voltage source. The source has a default name of "V1" and no value — it's just "V". We will change these later.

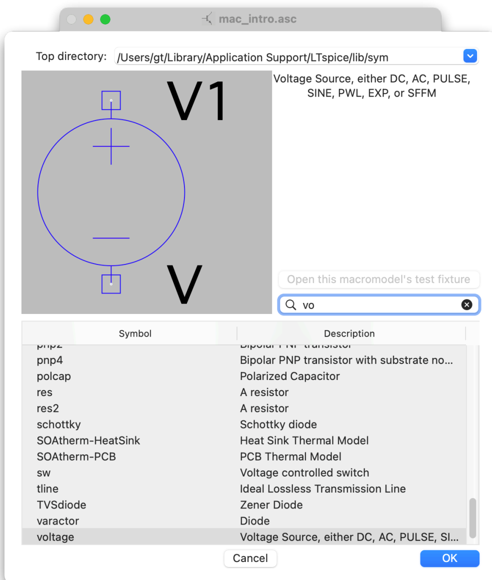

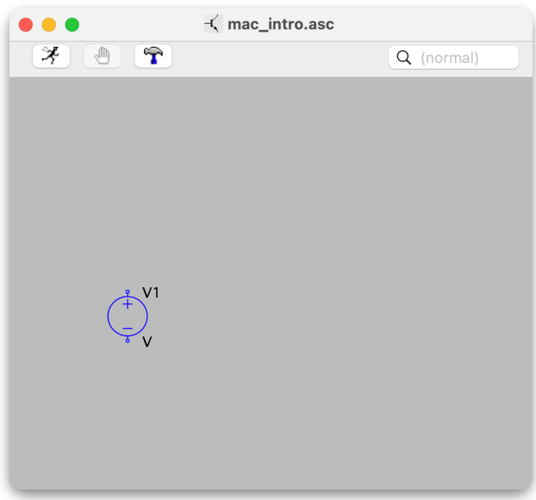

   After placing the first voltage source, the cursor still has a source icon attached to it. We could place more sources if we wanted, but since we need only one, hit the escape key. This takes us out of "voltage-source placement mode" and the source icon disappears.

   While a component is still attached to the cursor, there are a couple of actions that we could use to change the orientation.

   - Option-R (or command-R) rotates the component clockwise by 90°.
   - Option-E (or command-E) flips the component left-right (i.e. around a vertical axis).
   - Use control-Z to undo.

   All these actions can be done later during editing.

3. Place the current source. Repeat the process just used for the voltage source. Right-click and choose “Draft→Component”. Start typing “current” into the search box. Once the current source symbol appears, click “OK”. The current-source symbol is attached to the cursor. Move it somewhere to the right of the voltage source and left-click to place it. Hit the “escape” key to end placement mode.

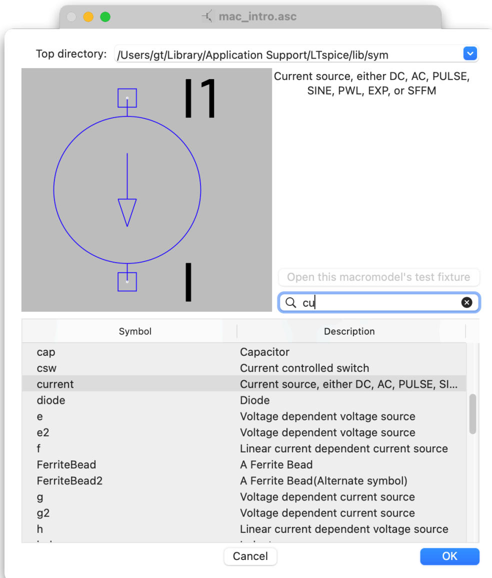

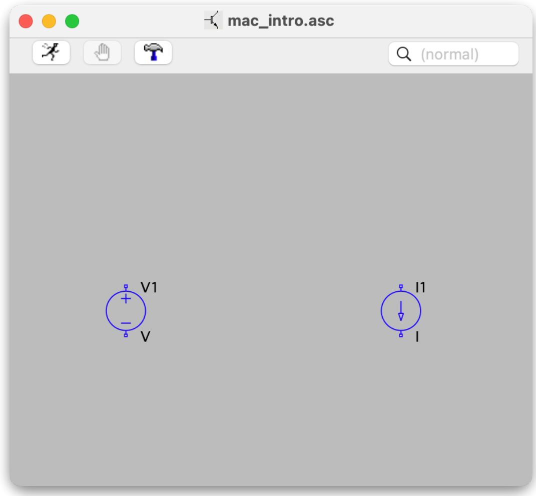

4. Using the same component-selection process, place the four resistors. Right-click and choose “Draft→Component” from the pop-up menu. Start typing “resistor” into the search box. Once the resistor source symbol shows up, click “OK”. The resistor symbol is attached to the cursor. Rotate it (“option-R”), move it somewhere to the right and above the voltage source, and left-click to place it. Place three more resistors, rotating as needed. The resistors will be numbered sequentially as you place them. Once all four are placed, hit the “escape” key to end the mode.

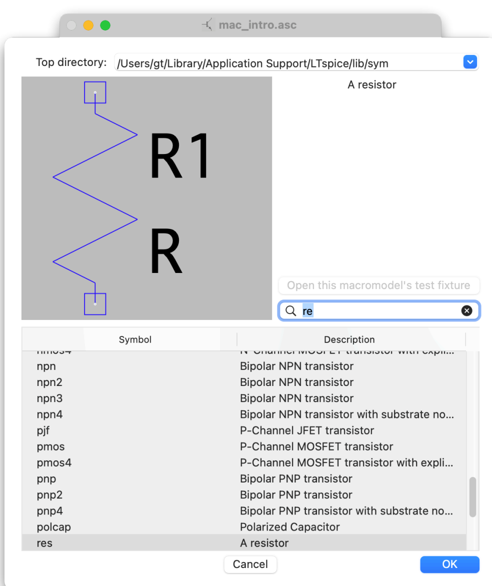

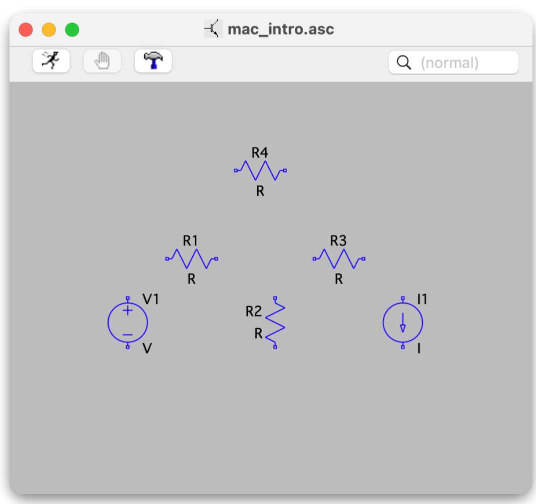

5. Wire the components together using the wire tool. Right-click and choose “Draft→Wires” from the pop-up menu. The cursor turns into a giant pair of cross-hairs that extend horizontally and vertically across the entire window. These are helpful in aligning the placement of the wires.

   Move the cross-hairs to the top of the voltage source and left-click on the terminal to start the wire. Move the cursor up to draw a vertical section of wire, left-click again to introduce a “corner”, and then move right towards the left end of the resistor. Left-click on the left-hand terminal of the resistor to complete that section of wire.

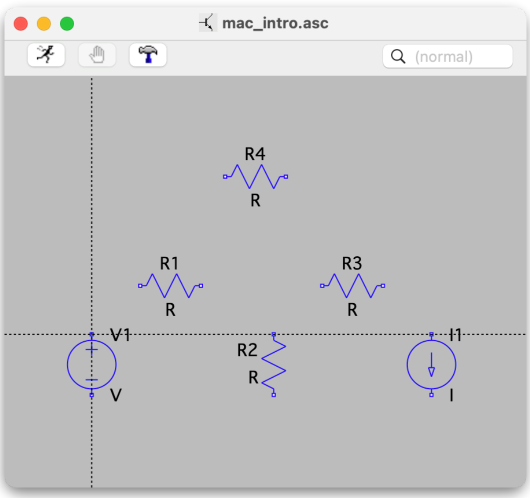

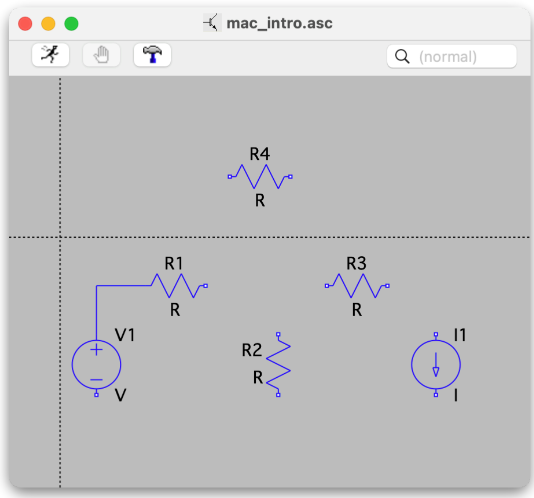

Continue wiring to connect all the other components according to the original circuit diagram. A wire can be started or stopped on an existing wire to make a T-connection, which is indicated by a small square dot. (Similar to the circular dots used to show connections on circuits in class.) To get out of wiring mode, hit the “escape” key.

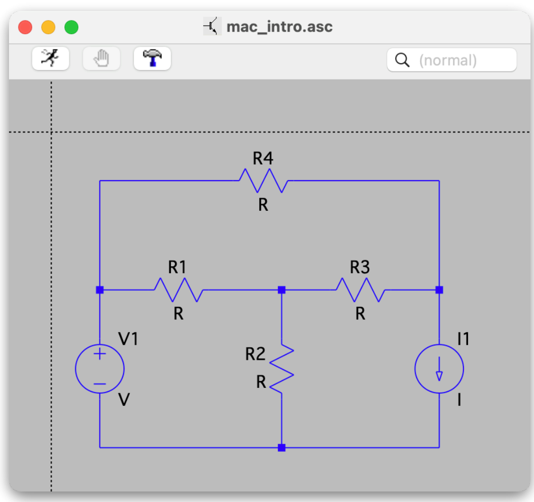

Pro tip regarding wiring: It is not necessary to wire from terminal-to-terminal, stopping and clicking at each end point. LTspice has a “speed wiring” mode, in which wires can be drawn right through components. It looks wrong as we do it, but once we stop to click at a corner or intersection, the program automatically makes connections for the components that were “drawn through”. The figure at right shows a “speed wire” in mid-process as it is drawn through resistors R1 and R3. Using this trick can speed things up a bit. In fact, in LTspice we can draw the wires first and then add the components on top of the drawn wires. If a circuit is laid out in nice rectangular meshes — as our 201 circuits often are — drawing the wires first might be faster.

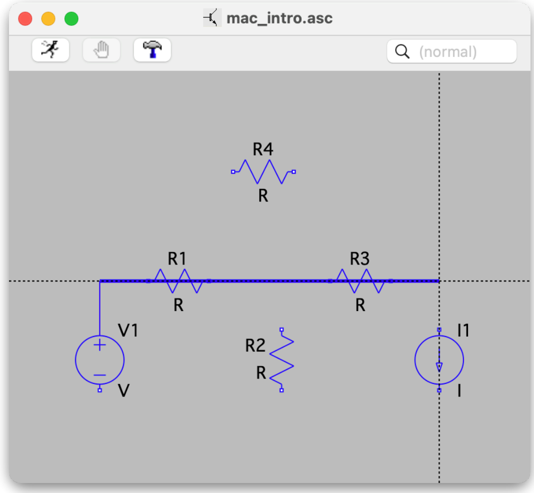

6. Add ground. Every SPICE circuit is required to have a ground connection, which SPICE labels as node 0. (A common SPICE user error is forgetting to include a ground. The result will be an error message when the simulation is started.) Right-click and choose “Draft→Net Name” from the pop-up menu — the “Enter Net Name dialog” appears. Select “GND (global node 0)” and click OK. Back in the drawing window, a ground symbol is attached to the cursor — move it to the desired location and left-click to place it. Then use the wire tool (from step 5 above) to wire it to the rest of the circuit.

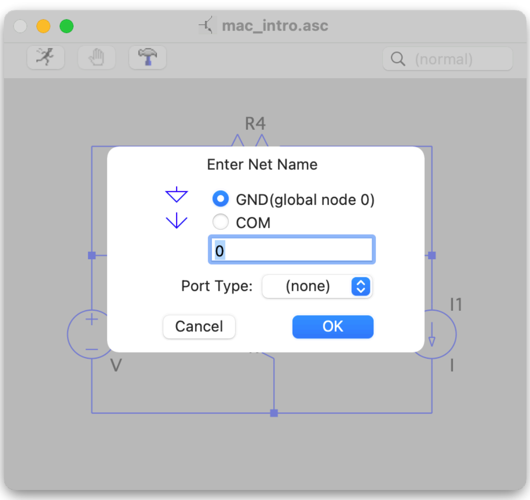

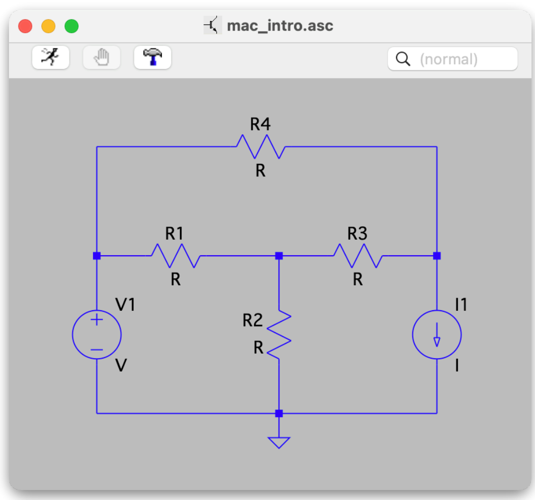

There is only one node 0 in a circuit, but we can use multiple ground symbols in the schematic. Every ground symbol will be treated as being connected to node 0. There does not need to be an explicit wire connecting all the ground points. For example, in our circuit we could have used three ground symbols — one at the negative terminal of the voltage source, another at the bottom terminal of R2, and a third at the bottom terminal of the current source — and left off the bottom wire. Using multiple ground symbols can simplify a complicated schematic.

After all the ground connections have been placed, press the “escape” key to end “ground mode”.

Add values and edit

Next, we need to enter specific values for the components. If desired, we can change the reference names. If needed, we can make changes to the circuit layout.

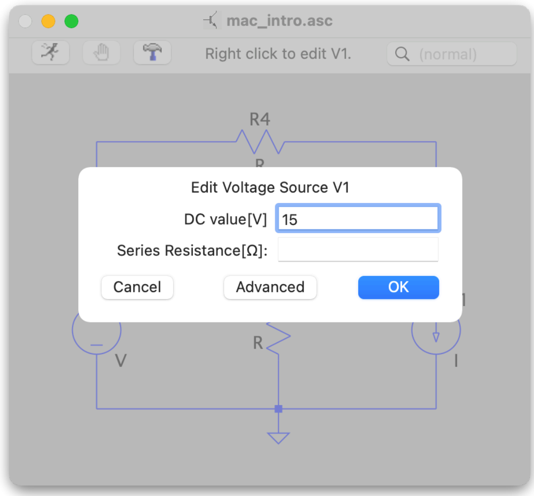

To change component values, hover the mouse over the component— the cursor changes to a "flying finger" icon — and then right-click to bring up a dialog where the value can be changed. Enter the desired value and click "OK". As opposed to what is preached in class, SPICE treats units as optional. SPICE knows that a voltage source value needs to be in volts, a resistor value needs to be in ohms, etc. (It is OK to include the units with the value if it makes us feel better, but it is not necessary.) Unit prefixes can be used. For example, a 1000-Ω resistor could be listed in SPICE as "1k".

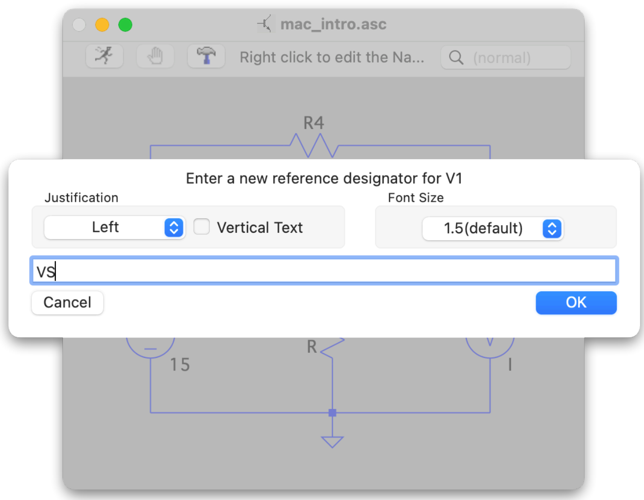

Component reference names can be changed, if desired. To change a name, hover the mouse over the text. The cursor turns into and "I-beam" icon. Right-click to bring up a dialog where the new name can be entered. Also, details about font size and orientation can be changed.[2]

This same method can be used to change the component values. (This is in addition to the method for changing values described above.)

[2] In the standard version of SPICE, reference names should have the first letter correspond to the type of component, followed by letters and numbers to distinguish individual parts. For example, Vxxx for a voltage source, Iyyy for a current source, Rzzz for a resistor, etc. However, in LTspice the standard convention is not mandatory — if you would like to name the voltage source "Zelensky", it will work. LTspice will add a "V" in front of "Zelensky" when the netlist is built. Which seems appropriate.

Use tools in the "Edit" item of the pop-up menu to move components, move text (names and values), reconfigure the wiring, or delete anything. A few comments/hints:

- There are some mouse actions that can be used to change how the circuit is situated within the drawing window.

  - Resizing the window will scale the circuit in a corresponding fashion.

  - Using the scroll wheel will scale the circuit without changing the window size.

  - Holding down the left mouse button and dragging will move the entire circuit within the window without any scaling.

- To edit a particular part in the circuit, first move the cursor away from any components to an open area of the drawing window. Then right-click and select the desired editing tool. (If the cursor is over a part when right-clicking, a dialog will open to change values, as described above.)

- We can move or drag components or text to shift their location within the circuit. "Move" is intended for labels and not-yet-connected components. "Drag" is intended for components that are already connected by wires. To invoke either, right-click and choose "Edit→Move" or "Edit→Drag". The cursor will turn into a little "hand" icon. Left click on the particular item to attach it the cursor and shift it to the new location. Left-click again to lock it into place. There is virtually no difference between "Move" and "Drag" for text items. The main difference shows up when attempting to move components that are already wired into the circuit. Using "Move" in that case *disconnects* the component from any wires before it is moved. (The dangling wires must then be deleted and the moved part rewired back into the circuit.) Using "Drag" will move the component along with the wires to the new location. The wires are rerouted during the move in order to maintain the connections. (However, the results are often look rather janky, and we may end deleting and replacing wires anyway.)

- If we want to shift a group of components, we can select either "Move" or "Drag" to get the hand icon. Then, while holding down the left mouse button, drag the hand across the components to select the group. The selected group can be shifted as a unit.

- The functions "Delete" and "Duplicate" work as expected. Right-click and choose either from the "Edit" pop-up, and then left-click on the component, text, or wire to be deleted or duplicated. The duplicated item can then be moved to the desired location.

- Don't forget about "Undo" (command-z). It can be your best friend when editing.

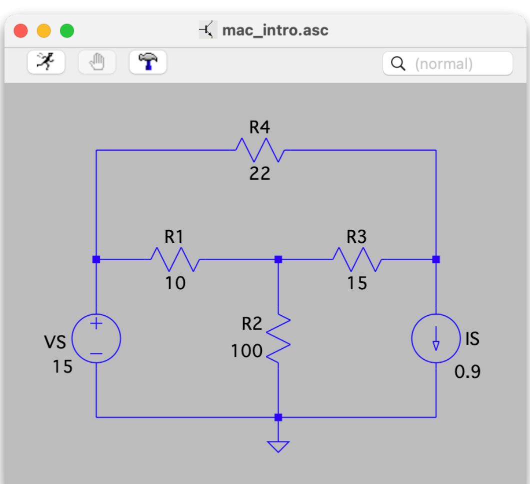

The circuit after editing.

Analysis directive

Next we tell LTspice what type of analysis we would like to do. On the Mac, LTspice is a bit unique in that it uses the text commands from the original text-file-input version of SPICE and embeds them right into the schematic. (The Window version of LTspice, PSPICE, and most other flavors of SPICE use dialog boxes to enter information about analysis type.) The Mac LTspice approach is cute, but it does require that we know the text descriptions of the analysis options.

For this circuit, we are keeping it simple and want only to find the DC values of the voltages and currents. In SPICE, a simple DC analysis is known as an “operating point" simulation, and the relevant command is “.OP”. To invoke it, right-click out away from the circuit, and in the pop-up menu, choose “Draft→SPICE Directive”. In the dialog box that opens, type “.OP”.

Click “OK” to return to the drawing window. The text “.OP” is attached to the cursor. Move it off to the side somewhere and click to place it. When the simulation starts, LTspice will look for the text to determine what type of analysis to do. We will look at other types of analysis (transient, DC sweep, AC sweep) in future examples, later in EE 201 or in EE 230.

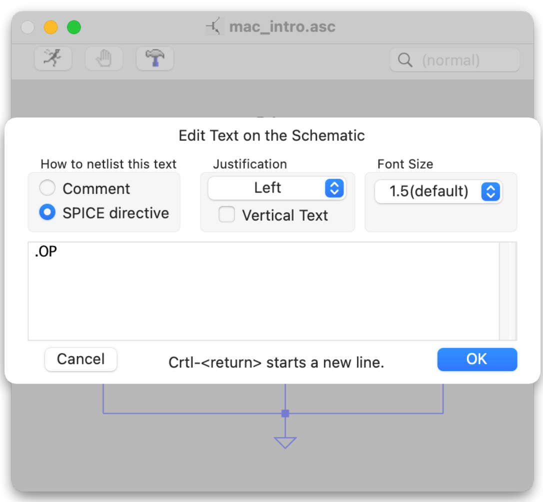

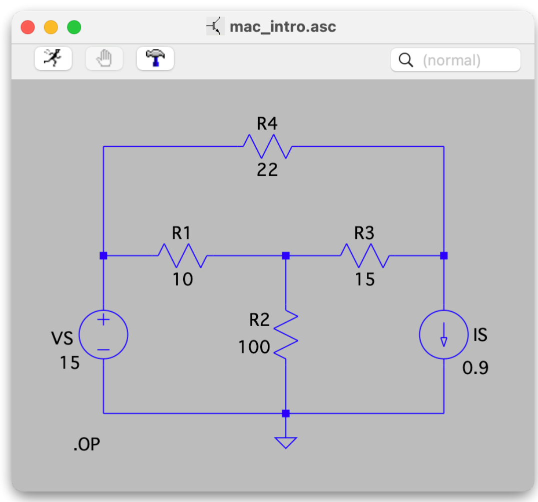

Run the simulation — view the results

Finally, we are ready to simulate and see the results. There are three ways to see DC voltages and currents, and each will be described below.

To begin the simulation, right-click in an open area of the drawing window and choose the "Run" menu item from the pop-up. Or click the "Run" icon at the top of the window. The simulation will be done in a second or two.

Assuming that there were no run-time errors in the simulation, a blank window will open. The window is for plotting the simulation results. Since making a plot for DC values is a bit weird, we will just close the window for now. (We will return to it shortly.)

In the macOS version of LTspice, the simplest way to see all of the DC results is to examine the log file that is created during the simulation. Right-click the mouse and from the pop-menu choose "View→SPICE Error Log". The window that opens has the DC node voltages and component currents.

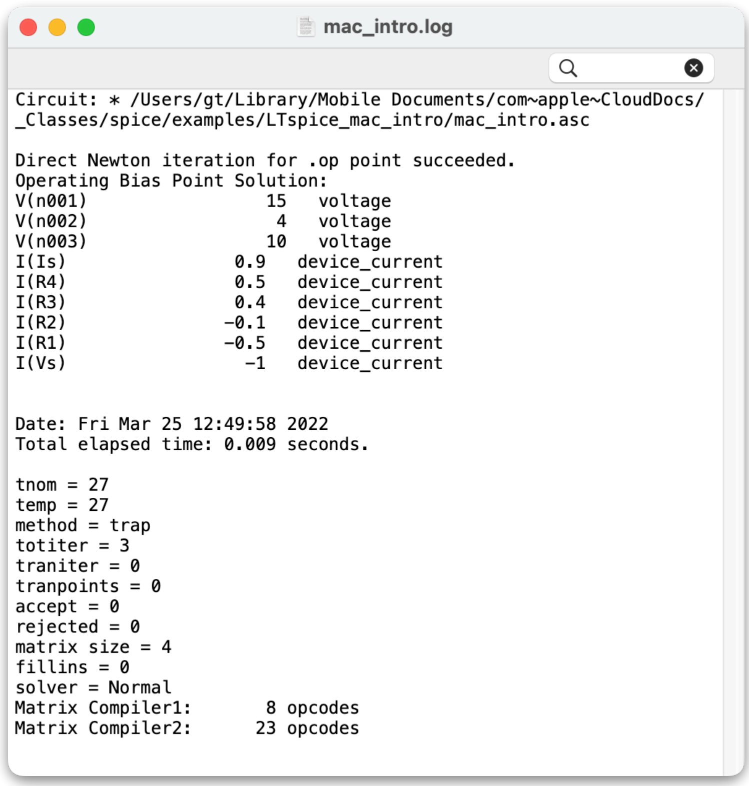
mac_intro.log

```
Circuit: * /Users/gt/Library/Mobile Documents/com~apple~CloudDocs/
_Classes/spice/examples/LTspice_mac_intro/mac_intro.asc

Direct Newton iteration for .op point succeeded.
Operating Bias Point Solution:
V(n001)                  15    voltage
V(n002)                   4    voltage
V(n003)                  10    voltage
I(Is)                   0.9   device_current
I(R4)                   0.5   device_current
I(R3)                   0.4   device_current
I(R2)                  -0.1   device_current
I(R1)                  -0.5   device_current
I(Vs)                    -1   device_current


Date: Fri Mar 25 12:49:58 2022
Total elapsed time: 0.009 seconds.

tnom = 27
temp = 27
method = trap
totiter = 3
traniter = 0
tranpoints = 0
accept = 0
rejected = 0
matrix size = 4
fillins = 0
solver = Normal
Matrix Compiler1:        8 opcodes
Matrix Compiler2:       23 opcodes
```

The nodes are labeled n001, n002, and n003. (The fourth node is ground, which is always "0".) The voltages at each node are given. The currents through all six components are also listed.

There are two minor annoyances in looking at the results. First, since LTspice assigned the nodes when converting the schematic to the SPICE input text file, we may not know which node is which. Fortunately, it is easy to find out. Go back to the schematic window and hover the mouse over a node. In the toolbar at the top of the window, there will be a message: "This is node Nxyx". Easy enough.

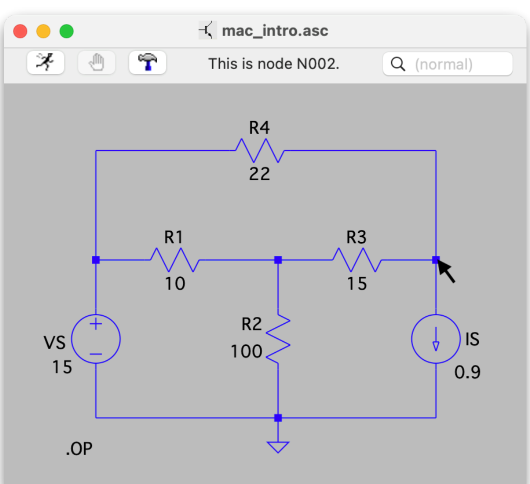

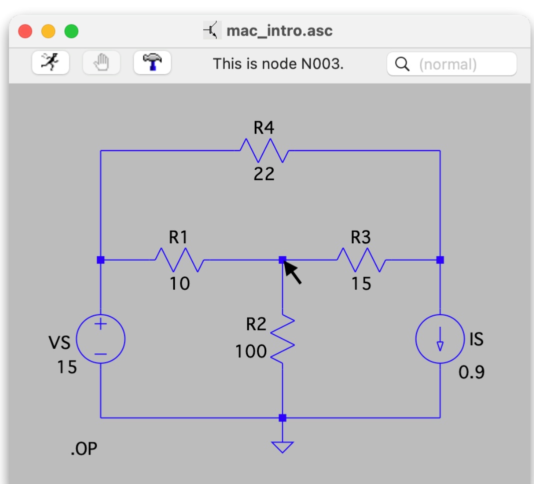

The second annoyance relates to the currents — it is unclear what the signs mean. Again, this is a result of LTspice assigning the node connections on its own. If the currents seem confusing, we can sort things out by looking at the net list created for the simulation. To view it, right-click and from the pop-up menu choose “View→SPICE Netlist”. A window opens, showing the text file that defines the circuit for SPICE. Each line corresponds to one component, with a name, the node connections, and the value. SPICE always defines current as positive flowing from the first node to the second. For example, resistor R1 is defined between N003 and N001. So for R1, the current is considered positive flowing from N003 to N001. (In other words, the current arrow for R1 points to the left.) In the Error Log list, the value for the R1 current is given as –0.5 A, meaning that 0.5 A is actually flowing to the right.

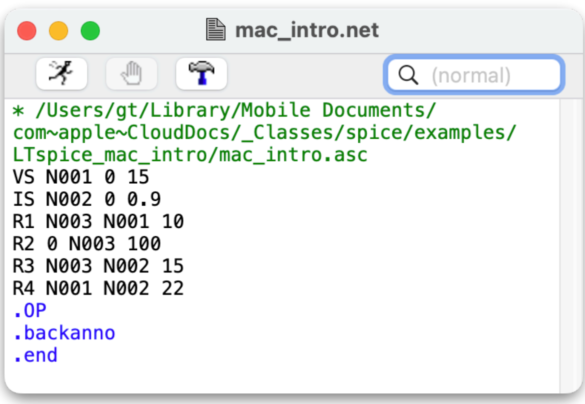

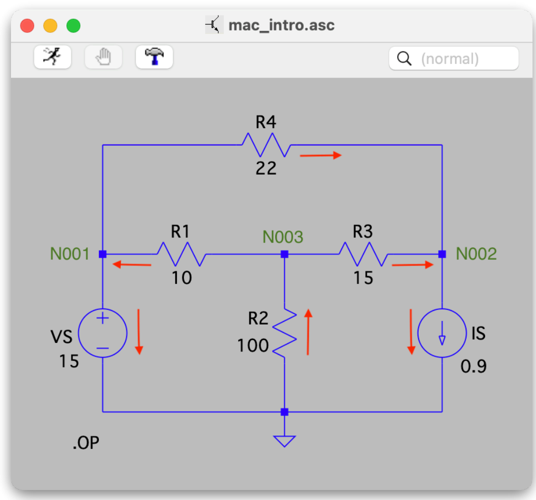

The second method for viewing the DC results works well if we are primarily interested in the node voltages. LTspice will print the node voltage right on the circuit schematic, similar to what we might do if analyzing the circuit with pencil and paper.

To display a node voltage, hover the mouse over the node and right click to choose “Place .op Data Label” from the pop-up menu. Repeat at each node of interest.

(DC data labels can also be placed by choosing the “Draft→.op Data Label” item from the usual pop-up menu.)

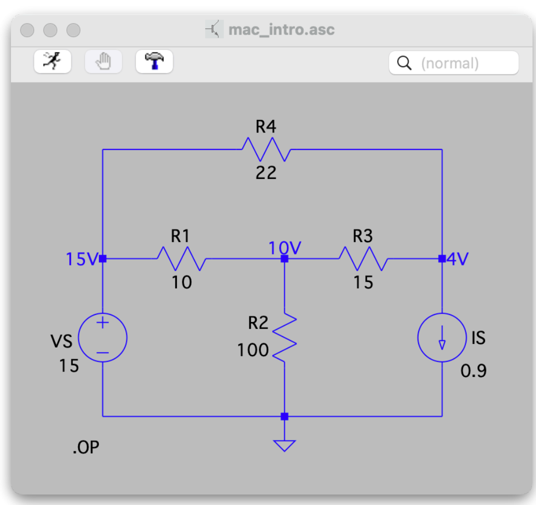

Unfortunately, the macOS version seems to be lacking a similar method for displaying currents.[3] If needed, we can use the node voltages and KCL to calculate a branch current.

[3] The Windows version can display DC currents. It would be nice if the folks minding the macOS version could add this feature.

The third method for displaying the DC values is to make a graph. This is not very satisfying. First, plotting DC quantities doesn't make a lot of sense.[4] Secondly, it can be difficult to read the values from the graph without doing some rescaling or using the cursor. It's not impossible to use the plots — it is just a bit strange.

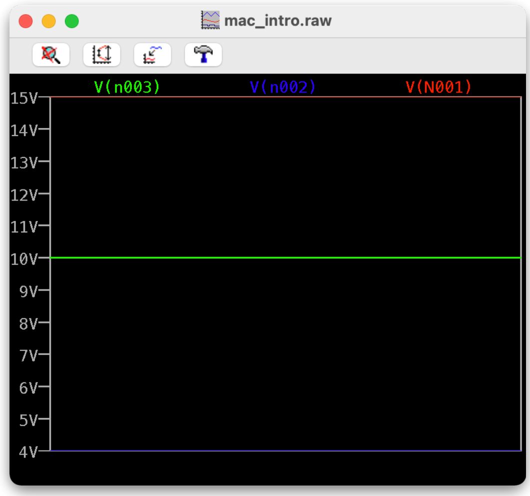

To plot the DC values, we must run the simulation again. (Click the run icon or right-click and choose "Run" from the pop-up menu.) The blank graph will re-appear. To plot a voltage on the graph, move back to the schematic window and click in an open area to make it active. Hover the cursor over a node — the cursor turns into a little "voltage probe" icon. Click on the node with the "probe" and a trace is added to the plot corresponding to the node voltage. Click on the other nodes to make traces for those. The figure shows the DC traces for the three node voltages — the red trace is for the node N001 (top of VS), green for N002 (top of IS), and blue for N003 (top of R2). As seen previously, the values are 15 V, 4 V, and 10 V, respectively. The red and blue traces are hard to see because they are at the very top and bottom of the graph.

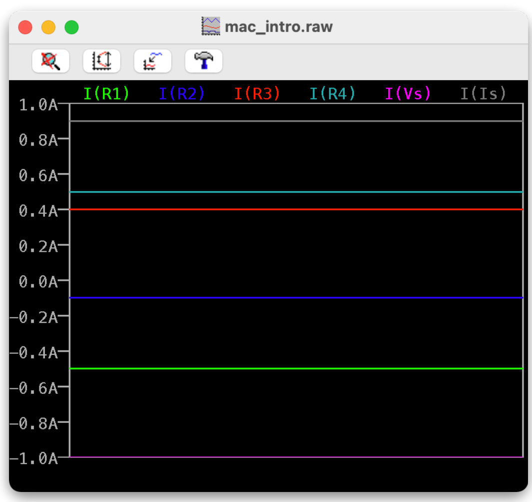

We can plot the DC currents. Close the voltage trace window and re-run the simulation to create another blank graph window. (We could simply add the currents to voltage graph. However, the addition of even more horizontal lines will just make a bad situation worse.) In the schematic, hover the cursor directly over a component (source or resistor) and the cursor turns into an arrow with a little "current-sensing loop" circling it. The arrow points in the direction that is defined as positive for that component. (Thus, we have an alternative method for determining the defined current directions.) Click on the component with the "probe" and the value is plotted as a horizontal line on the graph. Repeating for the other components gives traces for each current. The resulting graph has many horizontal lines with values that can be hard to read.

We could fiddle with the graphs if we wanted — changing scales, using the cursor to read values to several significant digits, etc. But, as we have stated, it is probably better to use one of the other methods to see the output of a .OP analysis. We will make better use of plots when we do other types of simulations.

Done!

That's it. We have used LTspice on a Mac to set up and run a simple DC simulation. We will look at circuits with a bigger variety of components and other simulation modes in other tutorials, and we will need the skills learned here for setting up those circuits.

[4] What does the horizontal axis represent? Perhaps, the x-axis is time, in which case these graphs might look something like DC voltages shown on an oscilloscope.