

Physics 364, Fall 2014, Lab #27 **Name:** _____

(*Field Programmable Gate Arrays: 4*)

Monday, December 8 (section 401); Tuesday, December 9 (section 402)

Course materials and schedule are at <http://positron.hep.upenn.edu/p364>

This file: <http://positron.hep.upenn.edu/wja/p364/2014/files/lab27.pdf>

Useful files: <http://positron.hep.upenn.edu/wja/p364/2014/files/?C=M;O=D>

Part 1

The first state machine we will study today is a kind of player piano. It reads from a memory a sequence of notes and durations, and plays the tune on a speaker. For each note to be played, the memory contains a 16-bit duration (in milliseconds) and a 16-bit half-period (in microseconds). One FPGA output pin will be connected to a speaker. To play a given tone, the FPGA will drive this output pin HIGH (+3.3 V) for one half-period, then LOW (0 V) for one half-period, then HIGH for one half-period, then LOW for one half-period, and so on. It continues to play this tone until the desired duration has elapsed. So to play the *A* note ($f = 440$ Hz, $T = \frac{1}{f} = 2272.7 \mu\text{s}$) above middle *C* for a duration of one second, we would drive the output pin HIGH for $1136 \mu\text{s}$, then LOW for $1136 \mu\text{s}$, and so on, until 1000 ms have elapsed.

A .zip archive of the files for Part 1, including a pre-made ISE project file, are at <http://positron.hep.upenn.edu/wja/p364/2014/files/lab27.zip>

If you want to look at the Verilog file on its own, you can either view it within the ISE software or follow this link:

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part1.v

A key goal of today's lab is for you to see how much more powerful a state machine becomes when it is combined with a memory that contains a list of tasks for it to carry out. This concept may inspire you to go in your own direction. Feel free at any time to modify this program to do something slightly different, or to ask us for help with modifying it to suit your interests. If you're uninspired by today's lab, feel free to ponder project ideas.

ROM

To store the desired information, we use a 256×16 ROM (Read Only Memory), i.e. 256 storage locations, each of which is 16 bits wide. That means that the input to our ROM is 8 address lines, `address[7:0]`, and the output of our ROM is 16 data lines, `dataout[15:0]`. There are several ways to describe a ROM in Verilog. The method used here maps directly onto last week's discussion of a ROM as a special case of a multiplexer.

```
1 // 256x16 ROM, i.e. 256 storage locations, each of which is 16 bits wide;
2 // how exactly you get Verilog to infer a ROM, a RAM, etc. is somewhat
3 // idiomatic and depends on the FPGA vendor's software (e.g. Xilinx);
4 // this is one acceptable way to tell Xilinx that you want a ROM
```

```

5 module rom256x16 ( output [15:0] dataout,
6                   input  [7:0]  address );
7   wire [7:0] A = address; // abbreviation to reduce typing
8   // white piano keys starting from middle C:
9   //   note:      C   D   E   F   G   A   B   C
10  // f   (Hz):    262 294 330 349 392 440 494 523
11  // T/2 (us):   1911 1703 1517 1432 1276 1136 1012 956
12  assign dataout =
13      //   duration      halfperiod
14      A== 0 ? 1000 : A== 1 ? 1517 :
15      A== 2 ? 1000 : A== 3 ? 1703 :
16      A== 4 ? 1000 : A== 5 ? 1911 :
17      A== 6 ? 1000 : A== 7 ? 1703 :
18      A== 8 ? 1000 : A== 9 ? 1517 :
19      A== 10 ? 1000 : A== 11 ? 1517 :
20      A== 12 ? 1000 : A== 13 ? 1517 :
21      A== 14 ? 1000 : A== 15 ? 1703 :
22      A== 16 ? 1000 : A== 17 ? 1703 :
23      A== 18 ? 1000 : A== 19 ? 1703 :
24      A== 20 ? 1000 : A== 21 ? 1517 :
25      A== 22 ? 1000 : A== 23 ? 1276 :
26      A== 24 ? 1500 : A== 25 ? 1276 :
27      A== 26 ? 1000 : A== 27 ? 0 : // rest 1s
28  ...
29      A==250 ? 0 : A==251 ? 0 :
30      A==252 ? 0 : A==253 ? 0 :
31      A==254 ? 0 : A==255 ? 0 : 0 ;
32 endmodule

```

We store **duration** in even-numbered addresses and **halfperiod** in odd-numbered addresses. So we need to read from two successive addresses to play a given note. We use the special case **halfperiod==0** to represent a musical “rest,” i.e. to represent silence, e.g. at $A = 27$.

In the special case **duration==0**, we use the corresponding **halfperiod** value to indicate the address from which the machine should subsequently start reading. This allows us to do a “GOTO” operation, so that at the end of a tune, we can go back and play the same tune again.

State diagram

Shown below is a state diagram for the machine that reads this ROM to play out a tune. We use a Verilog statement called “**localparam**” to assign state names such as **START**, **FETCHDURA**, **FETCHPITCH**, etc., to the integer state numbers 0, 1, 2, etc.

```

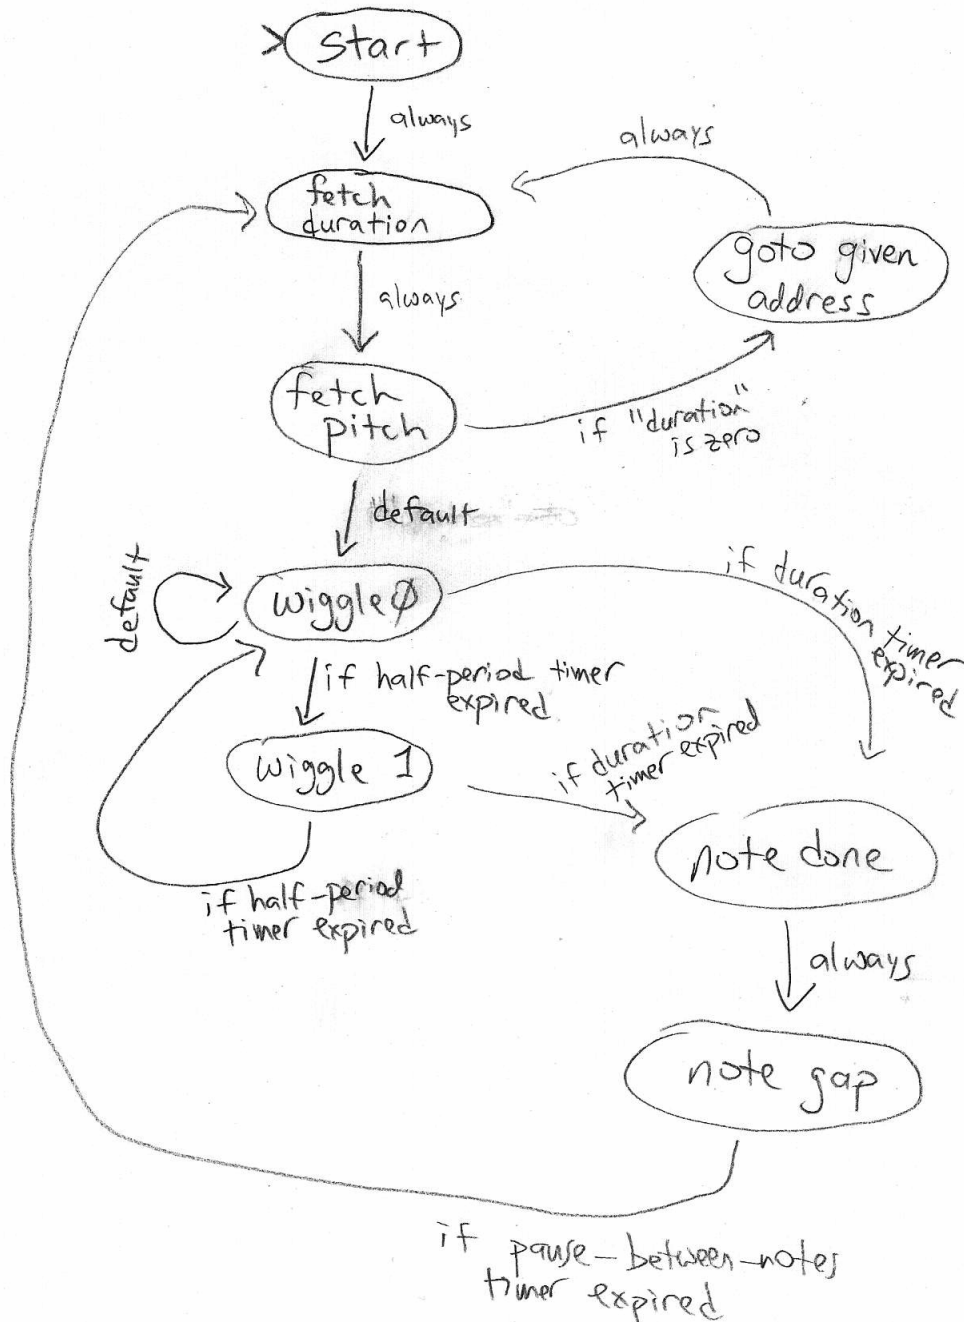
1   // Enumerate the possible states of our state machine
2   localparam
3       START=0,           // initial state: reset goes here
4       FETCHDURA=1,     // fetch next note's duration (unit=millisecond)
5       FETCHPITCH=2,    // fetch next note's half-period (unit=microsecond)
6       GOTO=3,           // special case (duration==0): GOTO new memory address
7       WIGGLEO=4,        // output bit is LOW for one-half period

```

```

8 WIGGLE1=5, // output bit is HIGH for one-half period
9 NOTEDONE=6, // finished playing a note
10 NOTEGAP=7; // pause briefly before moving on to next note

```



State representation in Verilog

Since there are 8 states, we use a 3-bit-wide D-type flip-flop to hold the current state of the machine, called `state`. The logic to decide which state to enter on the next clock cycle manipulates the wires called `nextstate`.

Here is the 3-bit-wide D-type flip-flop that holds the present state of the machine:

```
1 // We'll use a 3-bit D-type flip-flop to hold the state of the FSM
2 wire [2:0] state, nextstate;
3 dffe_Nbit #(.N(3)) mystatedff
4   (.q(state), // on each rising edge of the clock, the flip-flop
5    .d(nextstate), // copies 'nextstate' (D) to 'state' (Q)
6    .clock(clock), // clocked at 1 MHz
7    .enable(1), // always enabled
8    .reset(reset)); // reset to initial state by pushing button
```

Next-state logic

The `nextstate` logic determines which state we go into on the next clock cycle:

```
1 // Compute next state based on current state and pertinent conditions
2 assign nextstate =
3   (reset || stopthenoise) ? START : // go here if reset
4   (state==START || forcenewaddress) ? FETCHDURA :
5   (state==FETCHDURA) ? FETCHPITCH :
6   (state==FETCHPITCH && zeroduration) ? GOTO : // dur==0 means GOTO
7   (state==FETCHPITCH) ? WIGGLEO :
8   (state==GOTO) ? FETCHDURA :
9   (state==WIGGLEO && durationup) ? NOTEDONE : // end of this note?
10  (state==WIGGLEO && zeroperiod) ? WIGGLEO : // rest vs. tone
11  (state==WIGGLEO && wigglenow) ? WIGGLE1 : // wiggle up & down
12  (state==WIGGLE1 && durationup) ? NOTEDONE :
13  (state==WIGGLE1 && wigglenow) ? WIGGLEO :
14  (state==NOTEDONE) ? NOTEGAP :
15  (state==NOTEGAP && notegapdone) ? FETCHDURA : // pause betw. notes
16 /* default: stay in same state */ state ;
```

If you compare the above code snippet with the state diagram, you should see that the code is expressing the same state transitions as the diagram.

- If we press the **reset** button (`btn[0]`, the right-hand button), or if we slide the **stopthenoise** switch (`sw[0]`, the right-hand switch) upward, then the machine goes to the **START** state. This is also where it starts when we first program the FPGA.
- From the **START** state, the next state is always the **FETCHDURA** state (where we fetch from ROM the duration of the next note)
 - The **FETCHDURA** state goes to the next ROM address (which should be an even-numbered address) and records the 16-bit “duration” (in milliseconds) of the next note to play.
 - We will also go into the **FETCHDURA** state if the **forcenewaddress** button is pressed (which is `btn[1]`, the second button from the right). This button allows the you to directly modify the machine’s address, so that the machine subsequently starts playing music from that address.
- From the **FETCHDURA** state, the next state is always the **FETCHPITCH** state.

- The `FETCHPITCH` state goes to the next ROM address (which should be an odd-numbered address) and records the 16-bit “halfperiod” (in microseconds) of the next note to play.
- From the `FETCHPITCH` state, there are two possible next states:
 - By far the most common transition is to go to the `WIGGLE0` state, which is used along with `WIGGLE1` to make the output bit oscillate (wiggle) back and forth to make the desired musical tone.
 - But in the special case in which the “duration” value is zero, we interpret this to mean that the “halfperiod” value actually represents the next memory address from which we should continue reading musical notes. In that case, the next state will be the `GOTO` state, to handle this change of address.
- If we happen to have wound up in the `GOTO` state, the next state from there is always the `FETCHDURA` state, from which the process begins of fetching from memory the next musical note to play.
 - Again, the point of the `GOTO` state is to make the memory address jump discontinuously to a new location. Normally the memory address just increases by one step at a time, incrementing once to find a new duration and then incrementing again to find a new pitch.
 - We will look at the address-update logic in a moment.
- While a note is playing, we go back and forth between the `WIGGLE0` state and the `WIGGLE1` state.
 - If we’re in `WIGGLE0`, the output bit (to the speaker) is `LOW`; if we’re in `WIGGLE1`, the output bit is `HIGH`.
 - We reset a microsecond counter to zero before entering the `WIGGLE0` or `WIGGLE1` state. When the counter reaches the `halfperiod` value, we know it is time to transition to the other state, so that the output pin goes back and forth between `LOW` and `HIGH` at the correct oscillation period.
 - Meanwhile, in the `FETCHPITCH` state, we reset a millisecond counter to zero. If the number of milliseconds elapsed reaches the `duration` value, we know it is time to transition to the `NOTEDONE` state, to finish up playing this note and move on.
 - There is one more special case: if `halfperiod` is zero, indicating that we should play silence (a rest) rather than a note, then we just stay in the `WIGGLE0` state instead of going back and forth between `WIGGLE0` and `WIGGLE1`.
- From the `NOTEDONE` state, we always go to the `NOTEGAP` state.
 - The purpose of this state is to put a small (currently 25 milliseconds) gap between notes, so that your ear can hear the difference between two consecutive eighth-notes and a single quarter-note at the same pitch.
 - I didn’t realize that I needed this state until I heard how terrible *Mary Had a Little Lamb* sounded without it.

- If not otherwise specified, we just remain in the same state for the next clock cycle.

ROM address & instantiation

Next, the ROM connections. The memory address from which to read is stored in an 8-bit D-type flip-flop. We again use the generalized `dffe_Nbit` with `#.N(8)` to indicate an 8-bit-wide flip-flop.

```

1 // Use an 8-bit D-type flip-flop to hold the memory address
2 // from which the state machine is reading
3 wire [7:0] memory_addr, memory_nextaddr;
4 dffe_Nbit #(.N(8)) memaddrff
5     (.q(memory_addr),
6     .d(memory_nextaddr),
7     .clock(clock),
8     .enable(1), .reset(0));

```

Next, we instantiate the ROM and make connections for address (input) and data (output):

```

1 // We will store the desired tune to be played in a ROM (Read-Only
2 // Memory) with 256 locations, each of which can store a 16-bit number;
3 // declare wires to store the 'dataout' value from the memory, as well
4 // as the 'duration' and 'halfperiod' that will be copied (via flip-flops)
5 // from the 'memory_data' wires
6 wire [15:0] memory_data, duration, halfperiod;
7 rom256x16 myrom1
8     (.dataout(memory_data), // data comes out
9     .address(memory_addr)); // address goes in

```

Here is the logic to decide what ROM address to read from on the next clock cycle:

```

1 // Compute what memory address we will read from on next clock cycle
2 assign memory_nextaddr =
3     (btn[1])      ? sw[7:0]      : // button1: goto switch address!
4     (state==START) ? 0           : // start reading from zero
5     (state==FETCHDURA) ? memory_addr+1 : // after reading duration or
6     (state==FETCHPITCH) ? memory_addr+1 : // pitch, increment address
7     (state==GOTO) ? halfperiod   : // special GOTO command
8     /* default: same */ memory_addr ; // otherwise stay unchanged

```

As a special case, if we hold down `btn[1]`, the address is loaded from the `sw[7:0]` sliding switches. On the `START` state, we start out at address zero. If we are in either the `FETCHDURA` state or the `FETCHPITCH` state, we want to increment the address by one so that the next read from the ROM will happen from the next memory location after this one. If we are in the `GOTO` state, then the `halfperiod` value contains (instead of a musical note) the next address from which we should continue reading.

Duration and halfperiod flip-flops

We store the duration and the half-period for the current note in a pair of 16-bit-wide

flip-flops. The `duration` wires and the `halfperiod` wires connect the outputs of these flip-flops, respectively. The `duration` flip-flop is only enabled in the `FETCHDURA` state; and the `halfperiod` flip-flop is only enabled in the `FETCHPITCH` state. In both cases, the **D** (data) inputs of the flip-flops are from the `memory_data` outputs of the ROM. We're just writing down the value that we read from the ROM in either the `FETCHDURA` or the `FETCHPITCH` state. The fact that `memory_data` (the ROM output) is connected to the **D** inputs of these two flip-flops is a key point — if you don't see why it is connected this way, please ask.

```

1 // Use a 16-bit D-type flip-flop to hold the desired duration (in ms)
2 // the note we are playing (or are about to play)
3 dffe_Nbit #(.N(16)) durationdff
4   (.q(duration), // output goes to 'duration' wires
5    .d(memory_data), // input data come from memory output data
6    .enable(state==FETCHDURA), // enabled only in the FETCHDURA state
7    .clock(clock), .reset(0));
8 assign zeroduration = (duration==0); // indicates special 'GOTO' command
9
10 // Use a 16-bit D-type flip-flop to hold the half-period (in us)
11 // of the note we are playing (or are about to play)
12 dffe_Nbit #(.N(16)) halfperioddff
13   (.q(halfperiod), // output goes to 'halfperiod' wires
14    .d(memory_data), // input data come from memory output data
15    .enable(state==FETCHPITCH), // enabled only in the FETCHPITCH state
16    .clock(clock), .reset(0));
17 assign zeroperiod = (halfperiod==0); // indicates rest (silence) vs. note

```

Counting microseconds and milliseconds.

Now we have two 16-bit counters. One of them increments once per microsecond. (This is easy, because we set up our master clock to run at 1 MHz.)

The second counter increments once per millisecond. But we still clock it with the same 1 MHz clock. We use a once-per-millisecond pulse called `pulse_1kHz` to enable the millisecond counter. We do this because we want all of the logic in our state machine to be synchronous to a single clock.

```

1 // Use a 16-bit counter to count off microseconds until the next
2 // time the wire driving the speaker needs to wiggle up or down
3 wire [15:0] count_usec;
4 wire reset_usec;
5 counter_Nbit #(.N(16)) myuseccounter
6   (.q(count_usec),
7    .clock(clock), // clocked by 1 MHz clock
8    .enable(1), // always enabled
9    .reset(reset_usec)); // reset to zero when starting new half-period
10 // We start a new half-period immediately after FETCHPITCH or once
11 // the number of microseconds exceeds the desired half-period
12 assign reset_usec = (state==FETCHPITCH || wigglenow);
13 assign wigglenow = (count_usec==halfperiod);
14
15 // Use a 16-bit counter to count off milliseconds until the end
16 // of the note that we are playing (or are about to play)

```



```

17  wire [15:0] count_msec;
18  wire      reset_msec = (state==FETCHPITCH || state==NOTEDONE);
19  counter_Nbit #(N(16)) mymseccounter
20    (.q(count_msec),      // elapsed millisecs (i.e. counter value)
21     .clock(clock),      // counter operates from 1 MHz clock
22     .enable(pulse_1kHz), // enable only once per millisecond
23     .reset(reset_msec)); // reset when starting or ending a note
24  assign durationup = (count_msec==duration);
25  assign notegapdone = (count_msec==25); // 25 msec gap between notes

```

Other connections

Finally, notice that the left-hand pin on the **JC** connector is the wire that wiggles between LOW and HIGH when the state machine goes back and forth between the **WIGGLE0** and **WIGGLE1** states. You will want to connect this pin (the left-hand pin of **JC**) to one side of a small speaker. Connect the other side of the speaker to the GND wire (the 2nd pin from the right on the JC connector).

Once you compile and load the FPGA program, notice that the LEDs display in binary what note we are currently playing (i.e. its half-period).

If you hold down `btn[1]`, then the 7-segment LEDs display the 16 bits of `memory_data`, so that you can see what is stored inside the ROM. Also, when you hold down `btn[1]`, the 8 bits of `memory_address` are updated to contain whatever value is on the sliding switches. These two features together allow you to inspect the ROM contents at a given address.

If you are not holding down `btn[1]` (normally you won't be), then the digits do this:

- Digit 0 (on the right) indicates how much time is left until the current note is finished playing.
- Digit 1 indicates what state the machine is in. It looks as if it is displaying “9” all the time, but really it is just going back and forth between “4” and “5” (**WIGGLE0** and **WIGGLE1**).
- Digits 3 and 2 display the 8-bit `memory_address` from which the machine is currently playing.

```

1  // Connect useful values to the inputs of 'myd4d1' (displaydigits module)
2  wire [15:0] timeleft = duration-count_msec;
3  assign digit0 = btn[1] ? memory_data[3:0]   : timeleft[14:7];
4  assign digit1 = btn[1] ? memory_data[7:4]   : state;
5  assign digit2 = btn[1] ? memory_data[11:8]  : memory_addr[3:0];
6  assign digit3 = btn[1] ? memory_data[15:12] : memory_addr[7:4];
7  assign dots[2:0] = btn[3:0];
8  assign led = halfperiod[7:0];
9
10 // Use 4 JC pins and 4 JD pins as outputs to drive LEDs, etc.
11 assign jc[1] = (state==WIGGLE1);

```



```

12     assign jc[4:2] = 0;
13     assign jd[4:1] = {clock,pulse_1kHz,count_msec[0],count_msec[15]};

```

Spend some time checking out the features described above, looking through the ROM contents by moving the switches, and annoying your neighbors with my horrible rendition of *Mary Had a Little Lamb*. Go through the first several lines of the ROM contents and see how I managed to code in E, D, C, D, E, E, E. D, D, D. E, G, G, and so on.

Other tunes: the pips

If you slide the switches to contain the decimal value 64 (binary 01000000) and momentarily hold down `btn[1]`, then you will hear over and over again, every ten seconds, the pips¹ that are played e.g. by the BBC World Service on the hour. See if you can understand how this is coded into the ROM:

```

1     A== 64 ? 1000 : A== 65 ? 0 : // BBC news GMT "pips"
2     A== 66 ? 100 : A== 67 ? 500 : // 1 kHz for 100 ms at :59:55
3     A== 68 ? 900 : A== 69 ? 0 :
4     A== 70 ? 100 : A== 71 ? 500 : // 1 kHz for 100 ms at :59:56
5     A== 72 ? 900 : A== 73 ? 0 :
6     A== 74 ? 100 : A== 75 ? 500 : // 1 kHz for 100 ms at :59:57
7     A== 76 ? 900 : A== 77 ? 0 :
8     A== 78 ? 100 : A== 79 ? 500 : // 1 kHz for 100 ms at :59:58
9     A== 80 ? 900 : A== 81 ? 0 :
10    A== 82 ? 100 : A== 83 ? 500 : // 1 kHz for 100 ms at :59:59
11    A== 84 ? 900 : A== 85 ? 0 :
12    A== 86 ? 500 : A== 87 ? 500 : // 1 kHz for 500 ms at :00:00
13    A== 88 ? 4500 : A== 89 ? 0 :
14    A== 90 ? 0 : A== 91 ? 64 : // goto address 64

```

Do you see now why we need a `GOTO` instruction so that after reading from addresses 90 and 91 the machine can be told to loop back to address 64?

Other tunes: invention 13

Here (in the figure below) is a much more interesting tune. I actually transcribed the whole thing a few years ago (and made a couple of mistakes), but there was a limit to what would fit into the BASYS2 board's available RAM.

It starts at ROM address 96 (decimal), or 01100000 (binary). If you key this into the `sw[7:0]` and hold down `btn[1]` for a moment, it should start to play.

By the way, if you get tired of the music at some point, you can slide `sw[0]` up and your machine should sit quietly in the `START` state.

```

1     A== 96 ? 125 : A== 97 ? 0 : // start of RHS of invention 13
2     A== 98 ? 125 : A== 99 ? 1516 :
3     A==100 ? 125 : A==101 ? 1136 :
4     A==102 ? 125 : A==103 ? 955 :

```

¹http://en.wikipedia.org/wiki/Greenwich_Time_Signal

5 A==104 ? 125 : A==105 ? 1012 :
 6 A==106 ? 125 : A==107 ? 1516 :
 7 A==108 ? 125 : A==109 ? 1012 :

The image shows two pages of handwritten musical notation. The left page is titled 'Inventio 13' and features a complex piece of music with multiple staves. The right page is numbered '16' and also contains musical notation with staves. Both pages are heavily annotated with handwritten letters and symbols, likely representing chords or specific musical instructions, such as 'EACB EBD C C#A', 'A A EACB EBD C A G#', 'EACB EBD C A C A D#A D#A C', 'F#B G#D F', 'EACFD D#B#E C A F A', 'GCE#D#E G B G', 'E G C E', 'B B C E D D F', 'G G F G C G C F', 'C G C E G D F E C G E', 'C A B A C E A C D F', 'B G D G B G B C E G B', 'E C G E G C E G B A C E G F#B#A D#A G B D F#E G C E C E G', 'A#F#B D#F# G G#E A F#F#', 'D#B G E ECAC F#G#F#F#D#E', 'F#A B D# EC#ACEG F#D#G#B#F#E#C#A#F#A C B C A B B', 'E#B#G#B#G#B', 'G#B G#E G#E#A', 'F#A#F#D#F#D#B#X', 'E#G#E#C#E#C#D#C#A#', 'D#F#D#D#B#B#B#E', 'E E#A B#D#C# G#E D D#F#A#B FED C C#E#F#A E#F#C# B B D F G# D C B', 'EAC B E D C A# E ACEC AC#F#AC#A#D#C#A C#B#D#B#D#F#C# F#D#F#E#D CE#E#C#E#C#D#C#A#F#C#E#A', 'C#A#E#E#A#C#B#E#D#C#E#C#E#C#A#F#A#F#E#G#B#G#E#B#C#E#A C#E#C#A#C#E#A', 'G#B#C#E#E#A#C#B#E#D#C#A#C#E#D#B#F#E#C#E#G#F#E#D#C#B#D#F#G#B#A#F#D#D#D#B#A#C#B#A#E#A#G#A#E#A', 'B#G#E#B#D#C#E#G#E#A#F#B#G#C#A#D#B#G#F#D#B#C#A#D#E#F#E#E#A'

Part 2

I hope that by now you get the idea that a state machine is what allows a digital logic circuit to step through the kind of sequence of operations that you might normally expect a computer to carry out. The state machine from Part 1 was much more complicated than the vending-machine or traffic-signal state machines from last week, though it does not really have such a huge number of distinct states. The main thing that makes this machine capable of doing something so complicated is that it is essentially reading a program from a memory and carrying out the instructions given in that memory.

Here is the Verilog file for Part 2:
http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part2.v
 If you prefer to avoid compiling you own copy, you can use this pre-compiled .bit file:
http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part2.bit

The main point of this part is to remind you that a Verilog program is not really a program at all, but just a representation of a schematic diagram. I've taken the state machine from Part 1 and encapsulated it into a module called `playerpiano`, which you can look through in the Verilog source code, either in your web browser or in the ISE software.

```

1 // put the state machine logic from Part 1 into a module, so that we can
2 // instantiate TWO music machines to play at the same time!
3 module playerpiano
4   ( output          wigglewire,      // output wire wiggles to play music
5     input          clock,           // 1 MHz clock
6     input          reset,           // reset button (start play from addr 0)
7     input          gotobutton,      // button to start playing from 'startaddr'
8     input          stopthenoise,   // switch to shut off the music
9     input          pulse_1kHz,     // pulse (1 us duration) once per ms
10    output [7:0]   memaddr,        // address at which to read from ROM

```

```
11     input  [15:0] memdata,      // the data value returned by the ROM
12     input  [7:0]  startaddr ); // addr to jump to if 'gotobutton' is pressed
```

Then on the top level of the schematic diagram, I've instantiated two copies of `playerpiano`, reading from two separate ROMs and sending their output music to two separate FPGA pins: `JC[1]` (the left-hand pin on connector JC) and `JD[1]` (the left-hand pin on connector JD).

If you load the program and plug your speaker from `JC[1]` to `GND`, you will hear the same tune as at the end of Part 1. If you move your speaker wire from `JC[1]` to `JD[1]`, then you will hear the corresponding left-hand part of Invention 13.

Notice that the two tunes are playing at the same time! The 7-segment display is showing the left-hand memory address with digits 3,2 and the right-hand memory address with digits 1,0.

If you have time (if you reach this point by 3:30pm or so), then **try the following**:

Using one opamp and several resistors, wire up a circuit that allows you to combine the audio signals from the two FPGA output pins and send the combined (summed) signal to a single speaker. As an option, you might also contemplate a way to mix the two inputs with a ratio other than 1:1. Maybe you want the right-hand tune to be twice as loud as the left-hand tune, for instance. **Please try very hard not to expose the FPGA output pins to any voltage lower than 0 V or higher than +3.3 V, to avoid cooking the FPGA.**

Fallback option: If you want to try something quicker, just borrow a second speaker and use linear superposition of sound waves to sum the two signals.

If you want to see this circuit playing through a big speaker (borrowed from Bill Berner) via a transistor push-pull buffer, here's a video from Fall 2012:

<https://www.youtube.com/watch?v=tHcDawYmWtE> .

Part 3

The main idea for this part is go to a step beyond the music machine, to see that the microprocessor at the heart of your Arduino, iPhone, notebook computer, etc. is really just a fancy state machine connected to a large memory.

There's no Verilog programming for you to do here, but I'd like you to study and tinker with this circuit enough that you can follow how it works. In Part 4 you will modify the FPGA's memory contents to alter the program's behavior.

If you're feeling ambitious, in Part 4 you can load the FPGA's memory with an entire instruction sequence of your own design, but that's probably only feasible if you've encountered assembly-language programming somewhere in the past.

Feel free to modify my Verilog program to do something different, if you like. Don't feel obligated to follow the script. Ask us for help if you have an idea you want to try.

Here is the Verilog file for Part 3:

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part3.v

If you decide to compile this for yourself, then you will also need to download

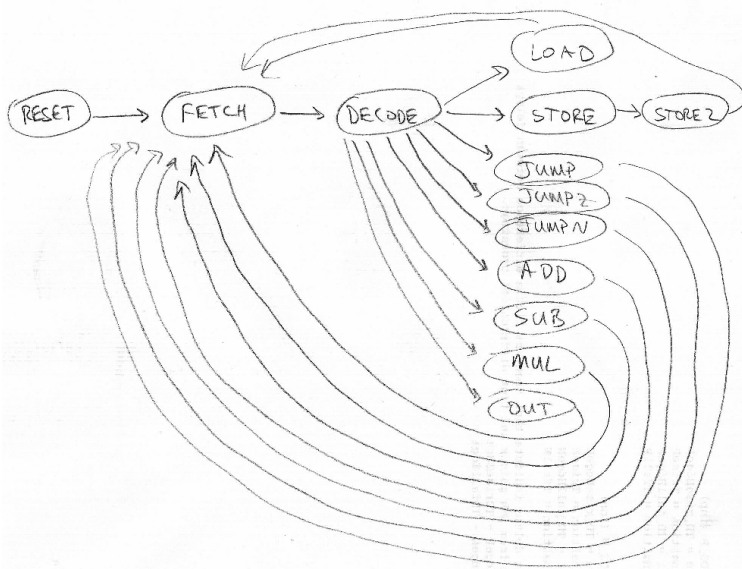
<http://positron.hep.upenn.edu/wja/p364/2014/files/asm.hex>

and put it into the same ISE project directory that contains `lab27_part3.v`. Or you can avoid compiling your own copy by using this pre-compiled `.bit` file:

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part3.bit

If you use ADEPT to load this `lab27_part3.bit` file into your BASYS2 board, you will see your board display a sequence of prime numbers. (Make sure the sliding switches are all in the DOWN positions.)

This state diagram for our simple computer should look familiar by now:



Here is what this list of states (i.e. the association from word-like names to integer state numbers) looks like in Verilog. I just added a JUMPNZ state (“jump if nonzero”), which hasn’t yet made it into the state diagram or the reading.

```

1 // Enumerate the possible states of CPU's state machine
2 localparam
3     RESET = 0, // initial state: reset goes here
4     FETCH = 1, // fetch next instruction from memory
5     DECODE = 2, // decode instruction: what are my orders?!
6     LOAD = 3, // execute LOAD: AC := memory[argument]
7     STORE = 4, // execute STORE: memory[argument] := AC
8     STORE2 = 5, // store gets an extra clock cycle for write to finish
9     JUMP = 6, // execute JUMP: PC := argument
10    JUMPZ = 7, // execute JUMPZ: if (AC==0) PC := argument
11    JUMPN = 8, // execute JUMPN: if (AC<0) PC := argument
12    JUMPNZ = 13, // execute JUMPNZ: if (AC!=0) PC := argument
13    ADD = 9, // execute ADD: AC := AC + memory[argument]
14    SUB = 10, // execute SUB: AC := AC - memory[argument]
15    MUL = 11, // execute MUL: AC := AC * memory[argument]
16    OUT = 12; // execute OUT: display AC on 7-segment LEDs

```

Since there are 14 states (more than 8 but fewer than 16), we use a 4-bit-wide D-type flip-flop to hold the current value of the CPU state:

```

1 // Use a 4-bit D-type flip-flop to hold the state of the CPU's FSM
2 dffe_Nbit #(N(4)) state_ff (.q(state), .d(state_next),
3                             .clock(clock), .enable(run), .reset(reset));

```

and here is the next-state logic to set up the transitions from state to state:

```

1 // Compute next state based on current state and pertinent conditions
2 assign state_next =
3     reset          ? RESET          : // reset line => RESET
4     state==FETCH  ? DECODE          : // FETCH => DECODE
5     state==DECODE ? (IR[15:8]==0 ? LOAD : // DECODE => execute decoded
6                     IR[15:8]==1 ? STORE : // instruction (LOAD, STORE,
7                     IR[15:8]==2 ? JUMP  : // JUMP, etc.)
8                     IR[15:8]==3 ? JUMPZ :
9                     IR[15:8]==4 ? JUMPN :
10                    IR[15:8]==9 ? JUMPNZ :
11                    IR[15:8]==5 ? ADD   :
12                    IR[15:8]==6 ? SUB   :
13                    IR[15:8]==7 ? MUL   :
14                    IR[15:8]==8 ? OUT   :
15                    FETCH) : // unknown => FETCH
16     state==STORE ? STORE2          : // STORE => STORE2
17     /* default */ FETCH             ; // default => FETCH

```

As the reading described, the CPU points its “Program Counter” (PC) at the next address in memory, FETCHes the instruction from memory[PC], and then DECODEs it to figure out what to do next. The upper 8 bits of the instruction contain the “opcode,” i.e. they determine whether the instruction is a LOAD operation, a STORE operation, an ADD operation, etc. The

lower 8 bits of the instruction are the “argument” for the operation, and generally refer to a memory address.

Random Access Memory

This CPU uses a read/write memory, a.k.a. a RAM, both to store instructions that it will execute and to store the intermediate results of its calculations. The module that defines the RAM begins like this:

```
1 // 256x16 RAM, i.e. 256 storage locations, each of which is 16 bits wide
2 module ram256x16 (
3     input        clock,          // clock (pertinent for writes only)
4     input        writeenable,    // write-enable
5     input  [7:0]  address,        // address at which to read/write
6     input  [15:0] datain,        // data to store next clock (if write-enabled)
7     output [15:0] dataout       // current memory contents at address A
8 );
```

The other details of this `ram256x16` module are not worth studying, except to note that the initial (“power-up”) contents of the memory, when the FPGA is first configured by the ADEPT software, are read from a file called `asm.hex`.² (This file is read when the FPGA program is compiled to produce the `.bit` file, not when the FPGA is loaded.) Here’s the rest of the `ram256x16` module:

```
1     reg [15:0] memory [255:0];
2     always @ (posedge clock) begin
3         if (writeenable) memory[address] <= datain;
4     end
5     assign dataout = memory[address];
6     // power-up memory contents come from text file 'asm.hex'
7     initial $readmemh("asm.hex", memory);
8 endmodule
```

We connect the 8-bit `memory_addr` lines, the 16-bit `memory_datain` lines, the single `memory_write` line (which is HIGH only when we want to update the value stored in the memory at the current address), and the 16-bit `memory_dataout` lines. The first three of these are outputs of the CPU state machine (they are inputs to the memory), so they are connected like this:

```
1     // This CPU uses a RAM consisting of 256 16-bit words. In the
2     // FETCH state, the memory address is the Program Counter, so that
3     // we can fetch the next instruction. Otherwise, the memory
4     // address is the "argument" of the decoded instruction, i.e. the
5     // low 8 bits of the IR.
6     assign memory_addr = (state==FETCH) ? PC : IR[7:0];
7
8     // The memory is only written in the STORE state; the data written
9     // to the memory always come from the accumulator (AC).
10    assign memory_write = (state==STORE);
11    assign memory_datain = AC;
```

²<http://positron.hep.upenn.edu/wja/p364/2014/files/asm.hex>

The `memory_dataout` lines are an input to the CPU (they are an output of the memory). The logic for `memory_addr` is worth thinking about for a moment. If the CPU is `FETCHING` the next instruction to be run, then the relevant memory address is contained in the Program Counter (PC). Otherwise, the relevant memory address is contained in the low 8 bits of the Instruction Register `IR[7:0]`, because each instruction consists of an opcode (upper 8 bits) and an address (lower 8 bits). The opcode is “what operation to perform,” and the address indicates “what data to operate on” (in addition to the accumulator).

Accumulator

The **accumulator** is the register that does nearly all of the CPU’s work. It is stored in a 16-bit-wide D-type flip-flop. Here is the Verilog code that instantiates the flip-flop and connects it:

```

1      // Accumulator (AC) is this CPU's primary register; all math
2      // instructions operate on the accumulator.
3      //
4      // Accumulator next-value logic:
5      //   ADD   => AC := AC + memory
6      //   SUB   => AC := AC - memory
7      //   MUL   => AC := AC * memory
8      //   LOAD  => AC := memory
9      //   RESET => AC := 0
10     //
11     // Note that the multiply happens in a single clock cycle, so it
12     // will compile to an entirely combinational multiplier -- the
13     // one you would write down using an adder and a multiplexer for
14     // each bit of the multiplicand.
15     wire [15:0] product =
16         (AC*memory_dataout > 'hfff) ? 'hfff : AC*memory_dataout;
17     dffe_Nbit #(.N(16)) AC_ff (.q(AC), .d(AC_next), .clock(clock),
18         .enable(AC_enable), .reset(reset));
19     assign AC_next = (state==ADD) ? AC + memory_dataout :
20         (state==SUB) ? AC - memory_dataout :
21         (state==MUL) ? product :
22         (state==LOAD) ? memory_dataout : 0;
23     assign AC_enable = run &&
24         (state==ADD ||
25         state==SUB ||
26         state==MUL ||
27         state==LOAD ||
28         state==RESET );

```

It is actually really neat (I think) to see how the contents of the accumulator are transformed when the CPU is in the various states like `LOAD`, `ADD`, `SUB`, `MUL`. The expression `AC + memory_dataout` represents a 16-bit-wide adder, whose two inputs are the present contents of the accumulator and the present contents of `memory_dataout`. The present contents of `memory_dataout` are whatever happens to be stored at the memory address referred to by `memory_addr` on the previous page. As described above, `memory_addr` is normally the low 8 bits of the Instruction Register, i.e. the “operand” or “argument” of whatever instruction is currently being carried out by the CPU. If the current state is `ADD`, then on

the next clock cycle, the contents of the accumulator will be replaced by the output of the adder. The net result of this will be $AC \leftarrow AC + \text{memory}[\text{IR}[7:0]]$, i.e. whatever is stored in the memory address pointed to by the low 8 bits of the Instruction Register will be added to the accumulator. If you don't fully understand this, then now is the time to stop one of us and demand a better explanation!

Similarly, the expression `AC - memory_dataout` represents a 16-bit subtraction. To subtract two numbers, you take the first number and add to it the negative of the second number. To compute the negative of a number, in two's complement representation, you first invert all of the bits (all 1's become 0's and vice-versa) and then add 1 to the result. So the logic that Verilog needs to generate for subtraction is similar to that for addition.

Accumulator's multiplication logic Multiplication of two binary numbers is less complicated than you might think. Suppose you want to compute $9 \times 5 = 45$, which in binary would be $1001_2 \times 101_2 = 101101_2$. You can do this using the binary version of the "long multiplication" technique that you learned (in decimal!) as a kid:

```

      1001
x   0101
-----
      1001
     0000
    1001
+ 0000
-----
   0101101

```

So multiplying together two 16-bit numbers requires 16 multiplexers (the thing to be added in each row is either zero or a shifted copy of the first multiplicand) and 15 adders. In general, the result of multiplying two 16-bit integers together is a 32-bit integer, so in general the answer won't fit back into the accumulator. A real computer needs a mechanism to detect **overflow** conditions in a way that the programmer can handle correctly. To keep things simple, my `product` logic uses this crude workaround: if the product is too big to fit into 16 bits, then I simply report 65535 (the largest 16-bit number) as the answer. That kludge is implemented by multiplexing the output of the multiplier with FFFF_{16} :

```

1   wire [15:0] product =
2       (AC*memory_dataout > 'hffff) ? 'hffff : AC*memory_dataout;

```

Output register

The **output register** is how this little CPU communicates its results to the outside world. When the `OUT` instruction is executed, the current contents of the accumulator are copied to the `OUT` register, whose contents are always displayed on the 4-digit 7-segment LED display.

On a real computer, a mechanism similar to this would be used to interface the computer to a digital-to-analog converter, or to external devices like printers, network interfaces, disk drives, etc. A real computer would also have an opcode like IN to receive input from external devices like keyboards, mice, analog-to-digital converters, disk drives, etc. Anyway, here is the Verilog code for the output register:

```
1 // The output register is this CPU's way to report its results to
2 // the outside world. The only path to the 'out' register is from
3 // the accumulator. The 'out' register is only enabled while the
4 // OUT instruction is executing.
5 dffe_Nbit #(16) out_ff (.q(out), .d(AC), .clock(clock),
6                       .enable(out_enable), .reset(reset));
7 assign out_enable = (state==OUT) && run;
```

Instruction Register

The **Instruction Register** (IR) holds the 16-bit instruction that is currently being executed, as described in the textbook chapter. Here is the Verilog code for the IR. Notice that the IR is a flip-flop that is only enabled during the FETCH state. Fetching the next instruction from the memory into the IR is analogous to the music machine's fetching the next *duration* and *halfperiod* from the memory into the corresponding registers (a.k.a. flip-flops). We need the IR to hold a copy of the instruction we are currently executing, because the process of carrying out the current instruction will in general require us to read or write other memory locations.

```
1 // The Instruction Register (IR) holds the instruction that is
2 // currently being executed. The only path into the IR is from
3 // the memory; the IR is only enabled in the FETCH state, i.e.
4 // while fetching the next instruction from memory.
5 dffe_Nbit #(16) IR_ff (.q(IR), .d(memory_dataout), .clock(clock),
6                      .enable(IR_enable), .reset(reset));
7 assign IR_enable = (state==FETCH) && run;
```

Program Counter

The **Program Counter** (PC) holds the memory address from which the next instruction will be read, the next time the CPU enters the FETCH state. Normally the Program Counter just increases by one on each subsequent instruction, corresponding to running a program that has no GOTOs, no loops, etc. But the JUMP, JUMPZ, JUMPN, and JUMPNZ opcodes can overwrite the contents of the Program Counter, thus changing the flow of the program.

If you have ever traced through the sequence of operations carried out by a computer program, either by running your pencil down a program listing line-by-line as your mind mimics the computer's actions, or by using a "symbolic debugger" go watch your program run line-by-line, then Program Counter corresponds to "what line of the program is running now." Since each instruction of an assembly-language program is stored at its own memory address, the Program Counter holds the address in memory corresponding to the line of the program that is currently running. (More precisely, the PC actually holds the address of the *next* line of the program to run, since the PC is incremented at the end of the FETCH state.)

The **conditional** jump instructions are the most interesting ones, as they are what permit the computer to make decisions: it can e.g. (JUMPZ instruction) go to a different address if the accumulator currently equals zero, or else continue along its current path if the accumulator contents are non-zero. Similarly, we can check whether the accumulator is negative (i.e. the highest bit is set), and jump or not jump accordingly. **Important point:** Without these conditional jump instructions, you couldn't write programs that use **IF** or **WHILE** or **FOR** to make decisions or to loop until some condition occurs.

Here is the Verilog code for the Program Counter:

```

1 // The Program Counter (PC) holds the address from which the next
2 // instruction will be fetched. Here is the program counter
3 // update logic:
4 // RESET => PC := 0
5 // FETCH => PC := PC+1 (after fetching from PC, point to PC+1)
6 // JUMP => PC := low byte of IR
7 // JUMPZ => PC := low byte of IR if AC == 0, else unchanged
8 // JUMPN => PC := low byte of IR if AC < 0, else unchanged
9 // JUMPNZ => PC := low byte of IR if AC != 0, else unchanged
10 dffe_nbit #(N(8)) PC_ff (.q(PC), .d(PC_next), .clock(clock),
11 .enable(PC_enable && run), .reset(reset));
12 assign PC_next = (state==RESET) ? 0 :
13 (state==FETCH) ? PC+1 : IR[7:0] ;
14 assign PC_enable = run &&
15 ((state==RESET) ||
16 (state==FETCH) ||
17 (state==JUMP) ||
18 (state==JUMPZ && AC==0) ||
19 (state==JUMPNZ && AC!=0) ||
20 (state==JUMPN && AC[15]));

```

Top-level connections

That's basically all it takes to make a simplified computer. You can see that it's just a state machine connected to a memory, not so different from the music machine we studied earlier today. To make it easy for you to see all of the computer's inputs and outputs, I put the guts of the computer (a.k.a. CPU, Central Processing Unit) into a module called `simple_cpu`. Here are the connections needed from the top-level module:

```

1 // These wires connect to the inputs/outputs of the CPU module
2 wire [15:0] memory_dataout, memory_datain, IR, AC, out;
3 wire [7:0] PC, memory_addr;
4 wire [3:0] state;
5 wire memory_write;
6
7 // Determine when the CPU will run (do its normal thing) and when
8 // it will pause to wait for the user.
9 wire run = !sw[0];
10
11 // Button 1 will reset the CPU to its initial state.
12 // function
13 wire reset = btn[1];

```

```

14
15 // Instantiate the CPU and its memory
16 simple_cpu cpu (.clock(clock), .reset(reset), .run(run),
17               .PC(PC), .AC(AC), .IR(IR), .state(state),
18               .memory_dataout(memory_dataout), .memory_addr(memory_addr),
19               .memory_datain(memory_datain), .memory_write(memory_write),
20               .out(out));
21 ram256x16 ram (.clock(clock),
22               .writeenable(memory_write),
23               .address(memory_addr),
24               .datain(memory_datain),
25               .dataout(memory_dataout));
26
27 // The green LEDs will display the Program Counter
28 assign led = PC;
29
30 // The 7-segment display will show the OUT register, i.e. the most
31 // recent prime number.
32 wire [15:0] leddat = out;
33 ....
34 assign {digit3,digit2,digit1,digit0} = leddat;

```

If you look at the complete Verilog program at this URL:
http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part3.v
you will see that there is really not that much to it.

Now what can this computer do?

To prove to you that this computer is capable of doing a real computation, I coded the stupidest imaginable algorithm for calculating all of the prime numbers from 2 to 9973. Here is how the algorithm would look if I were to write it in the C programming language.

```

1 #include <stdio.h>
2
3 int main(void)
4 {
5     int i, j, k, product;
6     // loop 'i' over candidate prime numbers, from 2 to 9999
7     for (i = 2; i<10000; i = i+1) {
8         // loop 'j' over possible first factors, from 2 to i-1
9         for (j = 2; j<i; j = j+1) {
10            // loop 'k' over possible second factors, from j to i-1
11            for (k = j; k<i; k = k+1) {
12                product = j*k;
13                // if j*k equals i, then i must not be prime: jump to 'iloop'
14                if (product==i) goto iloop;
15                // if j*k > i, then skip the rest of the k loop
16                if (product>i) break;
17            }
18        }
19        // if we reach this point, then i is prime: print it out
20        printf("%d\n", i);
21        iloop:; // this label 'iloop' allows the 'goto' to jump here

```

```

22     }
23     return 0;
24 }

```

On my Mac, the program outputs 2, 3, 5, 7, 11, 13, 17, 19, 23, 29, ..., 9887, 9901, 9907, 9923, 9929, 9931, 9941, 9949, 9967, 9973. Since we have no C compiler for our home-made computer, we have to write our program directly in the computer's **assembly language**, i.e. using the opcodes LOAD, STORE, ADD, JUMP, etc.

```

1
2     #
3     # prime.sasm
4     # coded 2010-11-11 by Bill Ashmanskas, ashmansk@hep.upenn.edu
5     #
6     # The purpose of this program is to demonstrate that the CPU
7     # implemented in simple_cpu.v is capable of carrying out a
8     # non-trivial computation.
9     #
10    # This is probably the dumbest imaginable algorithm to compute
11    # prime numbers. Its execution time scales as the third power
12    # of the number of candidates to evaluate. For every candidate i,
13    # loop over possible factors j and k, testing whether i==j*k. If
14    # no such j and k are found, then display i on the LEDs.
15    #
16    # Note that with storage for a mere 5000 boolean values (which
17    # I could have easily cooked up), one can use a much more efficient
18    # algorithm, the Sieve of Eratosthenes. It scales (with number of
19    # candidate integers N) as N*log(N)*log(log(N)), while my algorithm
20    # scales as N**3. I point this out only so that you don't think
21    # that I think the N**3 algorithm is a good way to compute primes.
22    #
23    start: load    istory #
24           store   i      # i := istory (nominally 1)
25    iloop: load    i      # loop i from istory+1 to 9999
26           add     one    #
27           store   i      # i := i+1
28           load    d9999 #
29           sub     i      #
30           jumpn   done   # if (i>9999) goto done
31           load    one    #
32           store   j      # j := 1
33    jloop: load    j      # loop j from 2 to i-1
34           add     one    #
35           store   j      # j := j+1
36           load    i      #
37           sub     j      #
38           jumpz   jdone  # if (j==i) goto jdone
39           load    j      #
40           sub     one    #
41           store   k      # k := j-1
42    kloop: load    k      # loop k from j to i-1
43           add     one    #
44           store   k      # k := k+1

```

```

45     sub    i      #
46     jumpz  jloop  #      if (k==i) goto jloop
47     load   j      #
48     mul    k      #
49     store  prod   #      product := j*k
50     sub    i      #      // if j*k==i then i is not prime
51     jumpz  iloop  #      if (product==i) goto iloop // skip to next i
52     jumpn  kloop  #      if (product<i) goto kloop // keep looping k
53     #      #      // k exceeds i/j, so skip to next j
54     jump   jloop  #      goto jloop
55     #
56     # If we reach here, then i is a prime number.  Display it.
57     #
58 jdone:
59 ... (uninteresting stuff suppressed) ...
60     jump   iloop  # go back up to try next candidate i
61 done:  jump   start # go back and start counting again from i==2
62 #
63 # This is where we define all of the constants and variables that
64 # our program will use when it runs.
65 #
66 zero:  .data    0      # store the constant '0'
67 one:   .data    1      # store the constant '1'
68 i:     .data    0      # store the loop variable 'i' (prime number cand.)
69 j:     .data    0      # store the loop variable 'j'
70 k:     .data    0      # store the loop variable 'k'
71 prod:  .data    0      # store the product 'prod' = j*k
72 outnum: .data    0      # compute/store binary-coded-decimal conversion of i
73 remain: .data    0      # store remainder used in BCD computation
74 hdigit: .data    0      # store hex value used to display one decimal digit
75 h1000: .data    1000   # store hexadecimal constant 0x1000
76 h100:  .data    100    # store hexadecimal constant 0x100
77 h10:   .data    10     # store hexadecimal constant 0x10
78 d10000: .data    2710  # store decimal constant 10000
79 d1000:  .data    3e8   # store decimal constant 1000
80 d100:   .data    64    # store decimal constant 100
81 d10:    .data    a     # store decimal constant 10
82 d9999:  .data    270f  # store decimal constant 9999 (= 270f in hexadecimal)
83 istart: .data    1     # starting value for i (i.e. first prime to check)
84 Jdelay: .data    1000  # delay factor (in hexadecimal)
85 Kdelay: .data    300   # additional delay factor (in hexadecimal)

```

The part of the code that is shown above does the prime number calculation. The whole program, including the parts that I omitted above, is at

<http://positron.hep.upenn.edu/wja/p364/2014/files/prime.sasm>

There are two parts that I didn't show:

First, the conversion of the prime number from a 16-bit hexadecimal integer into four decimal digits (thousands, hundreds, tens, ones), so that the primeness of the prime numbers looks more convincing to a human observer.

Second, the brief delay before displaying each new prime number, so that the numbers do not overwrite each other too quickly for you to see. The delay is implemented as a “nested loop:” the outer loop repeats $1000_{16} = 4096$ times, and the inner loop repeats $300_{16} = 768$ times for each repetition of the outer loop. So it’s just wasting time by counting up to about 3 million. The nested loop is needed because our tiny computer’s integers are only 16 bits wide.

Machine-readable format

The program above is still in human-readable form. We need to convert the instructions into hexadecimal memory contents. Here is the output of that process, i.e. an annotated file that is identical in content to `asm.hex`.³ The annotated version is at `prime_assembled.txt`.⁴ The file looks like this (with boring parts suppressed):

```

1      mem[‘h00] = ‘h006e; // start:  load  istart
2      mem[‘h01] = ‘h015f; //          store  i
3      mem[‘h02] = ‘h005f; // iloop:  load  i
4      mem[‘h03] = ‘h055e; //          add   one
5      mem[‘h04] = ‘h015f; //          store  i
6      mem[‘h05] = ‘h006d; //          load  d9999
7      mem[‘h06] = ‘h065f; //          sub   i
8      mem[‘h07] = ‘h045c; //          jumpn done
9      mem[‘h08] = ‘h005e; //          load  one
10     mem[‘h09] = ‘h0160; //          store  j
11     mem[‘h0a] = ‘h0060; // jloop: load  j
12     mem[‘h0b] = ‘h055e; //          add   one
13     mem[‘h0c] = ‘h0160; //          store  j
14     mem[‘h0d] = ‘h005f; //          load  i
15     mem[‘h0e] = ‘h0660; //          sub   j
16     mem[‘h0f] = ‘h031f; //          jumpz jdone
17     mem[‘h10] = ‘h0060; //          load  j
18     mem[‘h11] = ‘h065e; //          sub   one
19     mem[‘h12] = ‘h0161; //          store  k
20     mem[‘h13] = ‘h0061; // kloop: load  k
21     mem[‘h14] = ‘h055e; //          add   one
22     mem[‘h15] = ‘h0161; //          store  k
23     mem[‘h16] = ‘h065f; //          sub   i
24     mem[‘h17] = ‘h030a; //          jumpz jloop
25     mem[‘h18] = ‘h0060; //          load  j
26     mem[‘h19] = ‘h0761; //          mul   k
27     mem[‘h1a] = ‘h0162; //          store  prod
28     mem[‘h1b] = ‘h065f; //          sub   i
29     mem[‘h1c] = ‘h0302; //          jumpz iloop
30     mem[‘h1d] = ‘h0413; //          jumpn kloop
31     mem[‘h1e] = ‘h020a; //          jump  jloop
32     mem[‘h1f] = ‘h005d; // jdone: load  zero
33     ...
34     mem[‘h5c] = ‘h0200; // done:  jump  start
35     mem[‘h5d] = ‘h0000; // zero:  .data 0
36     mem[‘h5e] = ‘h0001; // one:   .data 1

```

³<http://positron.hep.upenn.edu/wja/p364/2014/files/asm.hex>

⁴http://positron.hep.upenn.edu/wja/p364/2014/files/prime_assembled.txt


```

37     mem['h5f'] = 'h0000; // i:      .data 0
38     mem['h60'] = 'h0000; // j:      .data 0
39     mem['h61'] = 'h0000; // k:      .data 0
40     mem['h62'] = 'h0000; // prod:   .data 0
41     mem['h63'] = 'h0000; // outnum: .data 0
42     mem['h64'] = 'h0000; // remain: .data 0
43     mem['h65'] = 'h0000; // hdigit: .data 0
44     mem['h66'] = 'h1000; // h1000: .data 1000
45     mem['h67'] = 'h0100; // h100:  .data 100
46     mem['h68'] = 'h0010; // h10:   .data 10
47     mem['h69'] = 'h2710; // d10000: .data 2710
48     mem['h6a'] = 'h03e8; // d1000: .data 3e8
49     mem['h6b'] = 'h0064; // d100:  .data 64
50     mem['h6c'] = 'h000a; // d10:   .data a
51     mem['h6d'] = 'h270f; // d9999: .data 270f
52     mem['h6e'] = 'h0001; // irstart: .data 1
53     mem['h6f'] = 'h1000; // Jdelay: .data 1000
54     mem['h70'] = 'h0300; // Kdelay: .data 300

```

So the memory contents that cause the computer to calculate this big sequence of prime numbers look like this: 006e 015f 005f 055e 015f 006d 065f 045c 005e 0160 0060 055e 0160 005f 0660 031f 0060 065e 0161 0061 055e 0161 065f 030a 0060 0761 0162 065f 0302 0413 020a 005d 0163 005f 0164 0669 0426 025c 005d 0165 0064 066a 0430 0164 0065 0566 0165 0228 0065 0563 0163 005d 0165 0064 066b 043d 0164 0065 0567 0165 0235 0065 0563 0163 005d 0165 0064 066c 044a 0164 0065 0568 0165 0242 0065 0563 0564 0163 0800 006f 0160 0070 0161 0061 065e 0161 0953 0060 065e 0160 0951 0202 0200 0000 0001 0000 0000 0000 0000 0000 0000 0000 1000 0100 0010 2710 03e8 0064 000a 270f 0001 1000 0300 0000 0000 ... (more zeros). There's not much to it! To convert the first ("assembly") format into the second ("machine") format, I wrote this Python program:

<http://positron.hep.upenn.edu/wja/p364/2014/files/sasm.py.txt>

which is called an **assembler**, because it converts human-readable **assembly language** into hexadecimal (or binary) **machine code**.

You can run the assembler yourself from your web browser, at this link:

<http://positron.hep.upenn.edu/wja/p364/2014/files/assembler.html>

Again, for Part 3, all you really need to do is to load this pre-compiled **.bit** file:

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part3.bit

into your BASYS2 board with ADEPT and to see your board display a sequence of prime numbers. (Make sure the sliding switches are all in the DOWN positions.)

Part 4

For this final part, the idea is for you to watch the tiny computer step through some of its operations in detail. You can also write your own programs in the tiny computer's assembly language, if you like. To make the operation of the computer less opaque, I've coded in some features for watching and changing what the computer is doing. To get all of these debug features, you'll want to load

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part4.bit

into your BASYS2 board. The Verilog source code is at

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part4.v

And if you want to make changes, you will probably find it easiest to start from this complete Xilinx ISE project in .zip form:

http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part4.zip

Reset button. If you push down `btn[1]`, you will reset the CPU to its initial state, and it will restart whatever program it was running. (If it's the prime-number program, you'll see it restart counting out prime numbers from 0002.)

Pause switch. If you slide `sw[0]` up, you will pause the CPU's activity.

Program Counter display. The 8 green LEDs display the current value of the Program Counter.

New functions for 7-segment LEDs:

- If you slide `sw[7]` up, then the left two digits will display the Program Counter, and the right two digits will display the `memory_addr`, i.e. the address that is currently sent into the RAM.
- If you slide `sw[6]` up (but not `sw[7]`), then the four digits will display the Instruction Register.
- If you slide `sw[5]` up (but not `sw[7:6]`), then the four digits will display `memory_dataout`, i.e. the data that are currently read out of the RAM.
- If you slide `sw[4]` up (but not `sw[7:5]`), then the four digits will display the Accumulator contents.
- If you slide `sw[3]` up (but not `sw[7:4]`), then the right-hand digit will display the state number of the CPU's finite-state machine, i.e. 0=RESET, 1=FETCH, 2=DECODE, etc.
- If you press `btn[3]`, it will rotate between **four modes** of operation for the CPU.
 - The decimal-point dots will display `---*`, `--*-`, `-***`, `*---` respectively for modes 0,1,2,3.
 - **In mode 0**, the CPU runs freely as it did in Part 3. You can use `sw[0]` to pause the CPU and `btn[1]` to reset it.

- **In mode 1**, the CPU will pause at each `OUT` instruction until you press `btn[0]` to continue. This results in each prime number being displayed until you push `btn[0]`. While the CPU is stopped, you can use `sw[7]` through `sw[3]` to explore the contents of 7=`(PC,memaddr)`, 6=`IR`, 5=`memdataout`, 4=`AC`, 3=`state`.
- **In mode 2**, the CPU will pause at each `DECODE` state until you press `btn[0]`. You can single-step through individual instructions and examine the CPU's registers using `sw[7]` through `sw[3]` as above.
- Also in mode 2, if `sw[0]` is up, then the CPU will pause on every clock cycle, no matter what state it is in. You can watch the CPU's finite state machine go from `FETCH` to `DECODE` to (`LOAD`, `STORE`, `ADD`, etc.).
- In mode 0, 1, or 2, pressing `btn[2]` will copy the current 8-bit value represented by the sliding switches into a register called `useraddr`. When you are in mode 1 (not mode 2, this is mode 1 again), if the Program Counter equals the 8-bit value in `useraddr`, then the CPU will pause until you press `btn[0]`. This allows you to explore what the CPU is doing at a particular section of the program code. (If you are a programmer, it is like setting a breakpoint in the debugger.)
- So you can use the `useraddr` feature in mode 1 to make the program stop next time it reaches e.g. the instruction at address 18 (hexadecimal), then switch to mode 2 to single-step through the subsequent instructions.
- **In mode 3**, the CPU does nothing, and you can read/write the memory.
 - * If all of the sliding switches are down, then `btn[0]` will step forward through the memory one address at a time.
 - * **In memory mode (mode 3)**, the 8 green LEDs show the address, and the digits show the memory data stored at that address.
 - * While you are holding down `btn[0]`, the digits will momentarily show the address also.
 - * If the switches are all in the down (zero) position, then pressing `btn[0]` will add 1 to the address. This makes it easy for you to step sequentially through the entire memory to check its contents.
 - * If the switches are set to a non-zero position, then pressing `btn[0]` will use the switches as the address into the memory. This address is remembered once the button is released, so you can then move the switches without changing the address.
 - * If you push `btn[1]`, the 8-bit value from the switches will replace the lower 8 bits of the memory contents at the current memory address.
 - * If you push `btn[2]`, the 8-bit value from the switches will replace the upper 8 bits of the memory contents at the current memory address.
 - * The “current memory address” in memory mode (mode 3) is stored in the same `userdata` register mentioned above.
 - * This allows you to look through the memory contents, compare them with the assembler output, and even modify the program (or more likely the `.data` words).

If you want to write your own program, you will probably find it convenient to run it in

mode 1 so that you can display a sequence of values on the 7-segment LED displays using the OUT instruction, then push `btn[0]` to let the program run until the next OUT instruction.

A few annoying sources of potential confusion:

- I recently modified `lab27_part3.v` so that the CPU is clocked at 50 MHz, but I did not have time to make the same modification to `lab27_part4.v` (because making that modification would break other features in `lab27_part4.v`). So the `Kdelay` constant at the bottom of `prime.sasm` needs to be much smaller (e.g. `Kdelay: .data 10`) for `lab27_part4.v` than in `lab27_part3.v` (where the value is 300). So if you use the `asm.hex` from Part 3 and compile it into the Part 4 Verilog program, the delay between prime numbers will be nearly one minute, unless you reduce this `Kdelay` value.
- The file `asm.hex` must contain exactly 256 hexadecimal values, one value per line. The Xilinx compiler is not very smart (or very forgiving) about the `.hex` file format. If your file is the wrong length, the compiler will not even warn you; it will just initialize the memory contents with all zeros.
- If you change the `asm.hex` file and want to generate a new `lab27_part4.bit` file to reflect this change, you need to right-click “Generate Programming File” and then choose “ReRun All.” The Xilinx compiler isn’t smart enough to know what steps need to be re-run when the `asm.hex` file is changed, as it does not consider it to be a source-code file.
- For these last two reasons, I think that in the future I will avoid using the `.hex` memory file. I only recently learned just how limited the Xilinx `$readmemh` support is. Sorry!

Your Challenge (if you have the time and energy):

Here is a very short program that just counts up, starting from 1, and OUTputs each number to the 7-segment display. About half of the program is there just to delay about a second after displaying each number. You can leave those lines out if you prefer to use the “mode 1” feature described above.

```
1 start:  load   zero      # we'll do a loop over i
2         store  i        # i:= 0
3 iloop:  load   i
4         add    one
5         store  i        # i := i+1
6         out    # send accumulator to the 7-segment display
7
8 # these next six lines are just here to pause after showing each number
9         load   Kdelay   # loop k from Kdelay downto 1
10        store  k        # k := Kdelay
11 kdelay: load   k        # loop k from Kdelay downto 1
12        sub    one      #
13        store  k        # k := k-1
14        jumpnz kdelay   # if (k!=0) keep looping over k
15
16        jump   iloop    # go back up to loop over the next value of i
```

```

17
18 zero:  .data  0      # store the constant '0'
19 one:   .data  1      # store the constant '1'
20 i:     .data  0      # store the loop variable 'i'
21 j:     .data  0      # store the loop variable 'j'
22 k:     .data  0      # store the loop variable 'k'
23 Kdelay: .data  f000  # delay factor (in hexadecimal)

```

Copy this source code (there is an easier-to-copy version at http://positron.hep.upenn.edu/wja/p364/2014/files/lab27_part4_asm.txt) into the **Assembler Code** window at <http://positron.hep.upenn.edu/wja/p364/2014/files/assembler.html> and then click **Assemble it**. Then copy and paste the 256 lines of “.hex format” values to replace the current contents of `asm.hex`, which you can edit using the Xilinx ISE source code editor. Then do “Generate Programming File” → “ReRun All” (you **must** select “**ReRun All**” to make Xilinx ISE use the updated `asm.hex` contents) and load the resulting `.bit` file into your board. You should see the board counting (in hexadecimal).

Now modify the assembler source code in some way that interests you. You could make the counter count up in steps of 3 instead of steps of 1. Or you could display the square of each integer (in hexadecimal) instead of the integer itself. Or you could take one more stab at displaying the Fibonacci sequence! The idea is just to give yourself a chance to write a very short assembly-language program and to load it into the tiny computer that we programmed into the BASYS2 board. Ask us for help if you have an idea that you’re not sure how to implement.

If you plan to spend more than a few minutes editing the assembler source code, then you should take precautions to **avoid losing your work**. Use some sort of text editor (e.g. the ISE source code editor, or even NotePad) to make your changes, and then copy/paste them into the web browser window. (The web browser window is not smart enough to save your file anywhere. Your text is only there until you close or reload the web page.)

If you make only a tiny change to the `asm.hex` contents (e.g. only affecting one or two memory locations), then you might find it quicker to use the “mode 3” feature described above, to avoid the time (about 1 minute) needed to recompile your Xilinx project.