ASYNCHRONOUS RESEARCH CENTER

Portland State University

Subject: **Class 1 – Ring Oscillators**
Date: **September 26, 2010**
From: **Ivan Sutherland**
ARC#: **2010-is43**

References:

http://arc.cecs.pdx.edu/teaching
ARC# 2010-is36: Class Blurb, Ivan Sutherland, August 13, 2010
Sproull, R.F., and Sutherland, I.E., "Logical Effort: Designing for Speed on the Back of an Envelope," IEEE Advanced Research in VLSI, C. Sequin, ed., MIT Press, 1991.
Sutherland, I., Sproull, R., Harris, D., Logical Effort, Morgan Kauffman, ISBN 1-55860-557-6, 1999.
Sutherland, I.E., Technology and Courage, Perspectives #1, Sun Microsystems Laboratories, reprinted in 1996 from CMU Computer Science 25th Anniversary Commemorative, 1991. Also on the web.

PURPOSE

This memo offers four ring oscillator circuits for you to simulate using SPICE. Each consists of one or two loops of five inversions. I have chosen the transistor sizes to make each logic gate, usually an inverter, drive a load that is three times its strength. Your simulations will reveal the period of each circuit's oscillations.

INTRODUCTION

Let us connect a number of inverters in a loop. If the number of inverters in the loop is even, the resulting circuit is stable in either of two states. A common form of such a loop has two inversions. The two-inversion loop circuit is useful because it has two stable states. Such a *bi-stable element* forms the basis of flip-flops and latches. Because of its two stable states, it can store one bit of information.

If the number of inverters in the loop is odd, however, the circuit is unstable. A signal passing around a loop with an odd number of inverters experiences a net inversion. Each trip around the loop produces the complement of the previous value and thus the circuit oscillates. The period of the oscillation is the time the signal takes to traverse the loop twice.

Two questions come immediately to mind, one for small rings and one for large rings. First, what if the number of inverters in the loop is exactly one? Will a ring of one inverter oscillate? If you connect the output of an inverter to its input, what does the inverter do? Is the result stable?

The second question involves large rings. Suppose we have 101 inverters in a ring. Is it possible that more than one transition might circulate concurrently? If so, the frequency of oscillation would be a multiple of the fundamental or base frequency. The period of the base frequency is, as we have seen, twice the transit time around the ring. The more ambitious student may wish to simulate a ring of many inversions. Given suitable initial conditions, can you observe harmonic oscillations? Do they persist?

Next week we'll discuss both the ring of one and many inverters.

FIRST ASIGNMENT

Your first assignment, due before class starts next week, is to simulate the circuits of Figures 1 through 4. I'd prefer that you run your simulations in SPICE from netlists you produce using Electric. With this in mind we've posted both Electric and definitions of the circuits on the ARC website. However, if you are more comfortable with other tools, feel free to use them. Copy the circuits carefully, please, numerical accuracy matters. Kash and Nav can provide help in getting started with Electric.

TWO SIMPLE RINGS

Figure 1 is a ring of five inverters of strength 10, each driving the next inverter. In addition to its successor, each inverter in the ring drives a load inverter of strength 20. The wires in the ring are named a[1:5]. Electric interprets a colon inside brackets as meaning all the values between in sequence. Thus a[1:5] means the five wires a[1], a[2], a[3], a[4] and a[5]. I will use Electric's notation in class memos.

Two different forms of inverter appear in Figure 1. One has a bubble at its output and the other has a bubble at its input. Both represent *exactly the same inverter* circuit. I find that the two forms help me keep track of the "odd" and "even" positions in a ring and so I like to alternate them as you see in these figures. DeMorgan's theorem allows you two symbols for every logic gate. You can find the alternative icons in the library of gates provided with Electric. I use the icon that offers best understanding. In Figures 3 and 4 I use the "AND with bubble input" icon to represent a NOR gate because the important action happens when both its inputs go LO.

Near each inverter in Figure 1 you will see the notation X=10 or X=20. The strength, or X value, of an inverter is a measure of the width of its transistors. Bigger X means more output current which requires wider transistors. We use inverters whose P-transistors are twice as wide as their N-transistors. The X value of the inverter

applies a multiplier to the width of the N and P transistors but retains the 2 to 1 ratio of their widths. Pull-up strength is approximately equal to pull-down strength.

Stronger inverters with wider transistors provide more output current, but their wider transistors also present more input load. Both drive and load are proportional to width. Thus we expect a strength 10 inverter driving a total load of 30 to have the same delay as a strength 40 inverter driving a load of 120. The ratio of a logic gate's drive strength to its load is called its **fan-out.** A closely related term, the **step-up** of two successive inverters is just the ratio of their X values. Because fan-out and step-up are both ratios, they carry no units. Both examples in this paragraph have a step-up of 3; 30/10 = 3 and 120/40 = 3.

There is also a linear relationship between fan-out and delay. The delay in a logic gate is very nearly proportional to its fan-out plus a constant delay. This linear behavior is what makes the idea of fan-out useful.

Inverters with X values less than 1 or larger than 100 may present loads or drive strengths that don't exactly follow the linear rule, but that won't concern us here.

Each of the size 10 inverters in the ring of Figure 1 drives a load exactly three times its strength. The three times load is made up of the next strength 10 inverter in the ring plus the load inverter of strength 20. Thus each inverter in the ring has **fan-out** of three, sometimes abbreviated FO3. The abbreviation FO4 means a fan-out of four.

Measure the period of the oscillations you get by simulating Figure 1. If you divide that number by 10 you will get the FO3 gate delay for the technology of your simulation. Enter your result on the answer form attached to this memo. The ambitious will wish to measure the period of a ring of FO4 which you can do by changing the sizes of the load inverters. The truly ambitious may wish to measure the period for other fan-out conditions and make a graph of inverter delay as a function of fan-out. **Figure 5** is such a graph for a typical technology.

Figure 2 is much the same as Figure 1. Figure 2 retains the fan-out of 3 for each stage by making the stages successively stronger. Notice that each strength, 1 3 9 and so forth is 3 times the strength of the previous inverter. Why is the last load inverter of size 242 rather than size 243? Does this small difference matter?

Measure the period of the oscillations you get by simulating Figure 2. Are they the same as you got for Figure 1? Plot a[2] and b[2] on the same graph to see if they are the same. I'll bet they will be very nearly identical.

TWO MORE COMPLICATED RINGS

Figure 3 shows a pair of rings coupled by an AND function. The AND function in Figure 3 has bubbles on its input. You may be more familiar with the NOR

icon for this kind of logic gate; Figure 3 uses the LO-AND form of icon to emphasize that the logic gate produces active HI output only when both inputs are LO.

Because the logical effort of the NOR gate is 5/3, the bottom inverter that drives c[2] has strength 1.66. To accommodate the extra load on c[1] of the strength 1.66 inverter, I reduced the size of big load near the top from 242 to only 240.33. Does the accuracy of the 240.33 strength matter? Are 5 significant digits useful?

There is a second ring in Figure 3 which includes c[6]. Because the NOR gate has a logical effort of 5/3, the NOR gate presents a load of 5 to c[6]. I took this load into account when figuring out how big to make the giant load inverter at the very top. My calculation notices that c[6] already drives a load of 5 in the NOR. So the big load should be (240.33 * 3) – 5. That explains the strength 716. Does the load of 5 on c[6] matter? Are three significant digits necessary?

Measure the period of the oscillations you get by simulating Figure 3. Is this the same as you got for Figures 1 and 2? Very nearly, I'll bet. Notice that each of the paths from c[3], the output of the NOR, back to the inputs of the NOR, c[1] and c[6], have exactly four inverters each with step-up of three. Thus we have a five-inversion ring oscillator, but this one has a split path.

Now plot the two inputs to the NOR, namely c[2] and c[6] on the same waveform display. Do they rise and fall together? Very nearly, I expect. But compare the output of the NOR, c[3], to the output of the inverter in your previous simulations, a[3] or b[3]. You may notice that c[3] rises more quickly than did a[3] or b[3]. That's because there are two P-type transistors in the NOR that act together when driving c[3] towards HI. This little detail turns out to be helpful as we shall see later on.

Figure 4 shows a pair of ring oscillators coupled by a **state wire**. The state wire appears in the figure as a U-shaped wire, e[1], driving a strength 83.3 load inverter. A lonesome P-type transistor of strength 30 at the left of the state wire can only drive the state wire HI. A lonesome N-type transistor of strength 30 at the right of the state wire can only drive the state wire LO. A wire driven differently from its two ends is called a "single track"; it can carry messages alternately in both directions.

If made of the proper sizes, the two lonesome transistors together would be a lot like an inverter. For proper sizing to make such an inverter, the P-transistor should be twice as wide as the N-transistor. In Figure 4, both are shown as size X=30.

Go down inside the icons for the lonesome P- and N-transistors. You will find that the P-transistor is twice as wide as the N-transistor. The X value for any gate, or any transistor, specifies its output drive strength rather than its width. To convert strength into width we have to take into account the fact that N-transistors provide more drive per unit width than do P-transistors. The icons of Figure 4 calculate width from strength internally, taking the different capability of the transistors into account.

How hard is it to drive the P-transistor of X=30 compared to how hard it is to drive an inverter of size X=30? The P-transistor represents 2/3 of the total load presented by the inverter, so its easier to drive the P-transistor than the inverter. We say that the lonesome P-transistor has a logical effort of 2/3. Likewise it's a lot easier to drive a N-transistor of size 30 than it is an inverter of size 30, so we say that the lonesome N-transistor has a logical effort of 1/3.

Now you can see how I calculated the sizes of the load inverters in Figure 4. The size 81 transistor whose output is d[1] should drive a total of 3*81 = 243. It drives the load of the lonesome P-transistor which is 2/3 * 30 = 20 plus the top load inverter. So the top load inverter should be size 243 – 20 = 223. Likewise the top load on the right is reduced from 81 to 71 to account for the load of the lonesome N-transistor which is 1/3 * 30 = 10.

I want to put a load on e[1] that is three times the size of each of the two lonesome transistors. The load on e[1] should be a total of 90. The NOR gate presents an input load of 5/3 * 3 = 5 and the inverter presents a load of 1.66, so the extra load on e[1] must be 83.3 as shown. Did I need more than three significant digits?

Measure the period of the oscillations you get by simulating Figure 4. Is this the same as you got for your earlier simulations? How well do the two lonesome transistors drive the state wire, e[1]? Compare waveforms at e[1] and c[6]. What difference do you see?

It would be bad if both of the lonesome transistors turned on at the same time. If both turned on, current would flow directly from the power supply to ground. That kind of current is called "crossover current" and wastes energy. How can we find out if the two lonesome transistors turn on at the same time?

It's easy to discover when the lonesome transistors turn on. Plot d[1] and e[5] in the same waveform viewer. The two waveforms may be remarkably similar. In next week's simulation we'll compare the times at which lonesome transistors that drive a state wire turn on and turn off. We will discover if crossover current is a problem in GasP circuits.

LESSONS

The self-timed circuits we're going to study consist of ring oscillators that are coupled together in various ways. Figures 3 and 4 show two ways to couple adjacent ring oscillators. In Figure 3 a logical AND function couples the two rings. In Figure 4 a shared state wire, e[1], couples the two rings. These two forms of coupling provide a rich variety of ways to string such ring oscillators together into useful circuits.

I have given the name GasP to the family of circuits we will study this term. The GasP initials mean nothing; GasP is not an acronym. Rather, GasP is a name, or,

if you like, GasP is what you are supposed to do when you see how fast these circuits can be. In this course we will see how GasP circuits can control elastic pipelines.

In this course we will study many forms of GasP circuits. We will see how to branch and merge pipelines of GasP circuits. We will see how to make a "wigwag" or alternating branch that sends its inputs alternately to two or more outputs. We will see how to make a demand merge that accepts inputs from multiple pipeline stages on a first-come-first-served basis.

What you should take away from today's lesson is a feeling for fan-out. You should know what FO3 and FO4 mean. You should remember that the delay in a gate depends on the strength of the gate and how much load the gate must drive. Different gates may have different delay.

You should also have a general understanding of ring oscillators. You have simulated four different circuits, all of which oscillate. We have omitted clock inputs to these circuits as unnecessary; these circuits oscillate by themselves.

By doing the homework, you should have perfected your ability to simulate the circuits we provide. You should also have acquired some ability to modify them.

In today's lesson you have had an encounter with Logical Effort. You must remember that the strength of a gate, represented as its X value, is a measure of the drive capability of the gate. Different kinds of gates will need different transistor widths to achieve the specified strength. If you know the strength of a gate and its type, you can calculate how wide its transistors must be and thus how much load the gate puts on its input terminals. I'll provide the logical effort values for different gate types as you need them, but I expect you to remember three particular values: The logical effort of an inverter is 1. The logical effort of a lonesome N-transistor is 1/3. The logical effort of a lonesome P-transistor is 2/3.

Remember that the delay of a gate depends on the ratio of the load it drives to its drive strength. An overloaded logic gate acts slowly.

ANSWER SHEET – due by beginning of class on 5 October 2010 **1**

Name _____

Turned in on Date _____

Simulation in _____ technology – e.g. 180 MOSIS

My simulation shows the period for **Figure 1** to be _____.

The period of a five inverter FO4 ring is _____. (optional)

My graph of delay as a function of load is attached (optional).

My simulation shows the period for **Figure 2** to be _____.

My simulation shows the period for **Figure 3** to be _____.

My simulations made waveforms c[2] and c[6] match don't match.

If don't match - my mismatch between c[2] and c[6] is _____ psec.

My simulation shows the period for Figure 4 to be _____.

My waveform for e[1] does does not look much like the others.

If different please describe difference: _____.

The logical effort of an inverter is _____.

The units of logical effort are _____.

The logical effort of a lonesome N-transistor is _____.

The logical effort of a lonesome P-transistor is _____.

The number of inversions in a ring oscillator must be _____.

My age at Birth _____ (answer may be culturally dependent).

ringOsc Experiment

nsk 1 October 2010

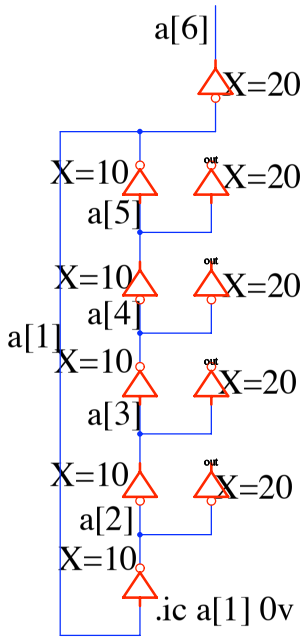

Fig 1

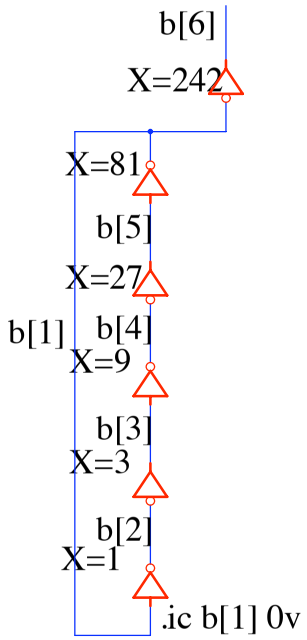

Fig 2

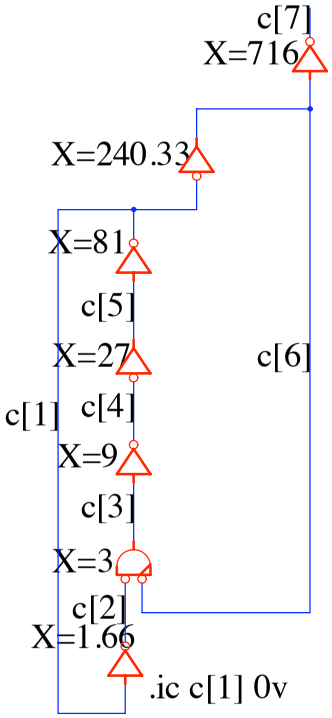

Fig 3

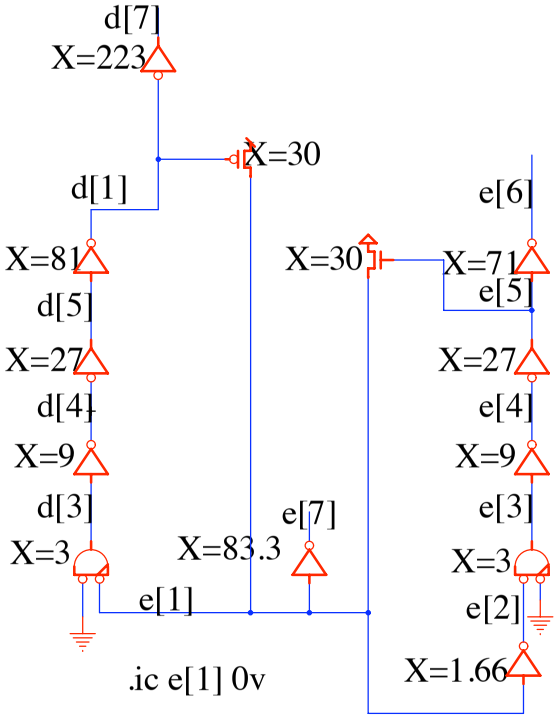

Fig 4

	A	B	C	D	E	F
1	**Spread sheet of delay vs load**			26-Sep-10		
2						
3	inverter strength	8.0				
4						
5	Load in X units	Load ratio	Delay (psec)			
6	8.0	1.0	325			
7	16.0	2.0	525			
8	24.0	3.0	725			
9	28.0	3.5	825			
10	32.0	4.0	925			
11	40.0	5.0	1125			

Figure 5: Inverter delay as a function of load
(for a fake technology)