6.004 Metstability Mathematics Gill Pratt March 4, 1997

We can model the arbiter simply as a saturating linear amplifier with positive feedback:

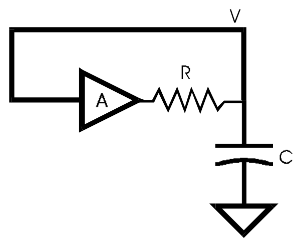

This turns out to be a pretty good model. To make the following analysis simple, I'll make the amplifier's saturation points be $V_{oh} = +1$ volt and $V_{ol} = -1$ volt.

If the amplifier is not saturated, the voltage *V* evolves according to:

$$V(t) = V_0 \; e^{\frac{(A-1)}{RC}t}$$

Where V_0 is the initial voltage at time $t = 0$. We can save ourselves some mess by compressing the following constants into one time constant variable:

$$t = \frac{RC}{A-1}$$

so we can re-write the equation for $V(t)$ as:

$$V(t) = V_0 \; e^{\frac{t}{t}}$$

This is an increasing exponential that curves up. At what point does the amplifier saturate? This occurs when the amplifiers **output** saturates. So, if we assume the amplifier output has just become +1 volt at time $t = t_{sat}$, then the amplifier's input (which is hooked to the system output *V)* will have just become *1/A* volts. As *V* continues to climb above *1/A* volts, the amplifier's output acts just like a wire connected to +1 volts, and the system output *V* evolves as:

$$V(t) = 1 - \left(\frac{A-1}{A}\right) e^{-\frac{t - t_{sat}}{t}}$$

Note that this is an increasing exponential that curves down. Here’s a plot of V(t):

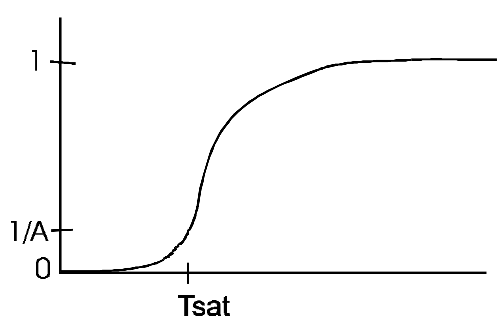

Because the linear gain *A* of the amplifier used in the demo is very large, *1/A* is very small and the upward curving part of the exponential is all but hidden from view.

Luckily, the downward curving part of the exponential after Tsat doesn't affect the probability analysis. Once *V* has reached a value of *1/A*, the convergence to a valid output is really over, and we just have to wait a few time constants for *V* to approach its final value of +1 or -1. Note that this takes a fixed amount of time to happen following *Tsat*, regardless of how long we had to wait until *Tsat* occurred.

To figure the probability of **synchronization failure** we assumed that the arbiter is started at a random V_0 between 0 and +1 Volts (the same argument works from 0 to - 1 Volts). If V_0 happens to be large, the system will converge quickly. If V_0 happens to be small, it will take longer.

Let's say we want the system to converge before a time *Tpd* has elapsed. What is the minimum value of V_0 for which this will occur? Let's declare the system to be converged when the amplifier's output first saturates (i.e. when its input is *1/A*). In this case, we can find the V_0 that will converge in exactly *Tpd* seconds by solving:

$$\frac{1}{A} = V_0 \; e^{\frac{T_{pd}}{t}}$$

This gives us:

$$V_0 = \frac{1}{A} \; e^{-\frac{T_{pd}}{t}}$$

We know that any V_0 less than this will result in a convergence time longer than *Tpd*.

If we assume the random choice of V_0 has a uniform probability distribution between 0 and 1, then the value of V_0 which takes exactly *Tpd* time to converge is equal to the probability that the convergence will take longer than *Tpd*. To see this, imagine a large number of experimental trials being conducted. Those experiments begun with V_0 below the value given above will result in failure, those above this value will result in success. Since V_0 is randomly distributed between 0 and 1 with uniform probability, the probability of synchronization failure given time *Tpd* to obtain a result is:

$$p_{failure} = \frac{1}{A} \; e^{-\frac{T_{pd}}{t}}$$

In actuality, V_0 is probably a digital waveform, which spends most of its time near +1 or -1 volts. Thus, to obtain a more accurate probability, we multiply the above formula by the probablity of the input voltage actually being in transition (i.e. rising or falling):

$$p_{failure} = \frac{p_{transition}}{A} \; e^{-\frac{T_{pd}}{t}}$$

Given a desired failure probability, we calculate the required delay by solving for *Tpd*:

$$T_{pd} = t \; \ln\left(\frac{p_{transition}}{A\, p_{failure}}\right)$$

Let's take a look at some real numbers to get a feel for this equation. Lets say we have a 100 MHz data signal that spends 10% of its time in transition$(p_{transition} = 0.1)$

We wish to synchronize this signal to an independent 100 MHz clock. Lets say we use a flip-flop with a gain *A* of 10 and an *RC* value of 10^{-8} (10 nanoseconds). This yield a system time constant of

$$t \approx 10^{-9}$$

How long must we wait if we want an average synchronization failure rate of 1 failure every year?
1 Year has about 31536000 seconds. Since we are running at 100 MHz, we want a failure probability of:

$$p_{failure} = \frac{1}{10^8 \times 31536000} \approx 3.17 \times 10^{-16}$$

This yield a *Tpd* of:

$$T_{pd} = t \ \ln\left(\frac{p_{transition}}{A\,p_{failure}}\right) \approx 10^{-9} \ln\left(\frac{0.1}{10 \times 3.17 \times 10^{-16}}\right) \approx 31 \times 10^{-9}$$

or 31 nanoseconds. How much longer must we wait if we want to have 1 failure every 10 years? This divides $p_{failure}$ by a factor of 10, which multiplies the argument to the natural log by a factor of 10, yielding an increase of only $\ln(10) \approx 2.3$ nanoseconds! In TTL circuits, we typically allow 100 nanoseconds for convergence. This yields a failure probability of:

$$p_{failure} = \frac{p_{transition}}{A} \ e^{-\frac{T_{pd}}{t}} \approx \frac{0.1}{10} e^{-100} \approx 10^{-2} \times 10^{\frac{-100}{\ln(10)}} \approx 10^{-45.4}$$

which at a clock rate of 100 MHz is equal to about 1 failure every

10^{30} YEARS!

(The oldest hominid fossil is about 5 x 10^6 years old. The earth is about 5 x 10^9 years old)