Phase Locking for the SNS Event Link

Larry Doolittle and Jozsef Ludvig

October 9, 2002

1 Introduction

One function needed throughout SNS is an Event Link receiver. The Event Link signal is broadcast site-wide over fibers and twisted pair, and provides the master synchronization to the storage ring revolution frequency, and therefore the chopping time structure on the linac beam current.

LBNL instrumentation and control has several applications for the reconstructed clock signal, some of which require a low jitter clock, which runs at a large multiple (64) of the ring revolution rate. This note covers the analysis and design of the phase lock circuitry now under construction for a general purpose Event Link receiver in the form of an Industry Pack Module.

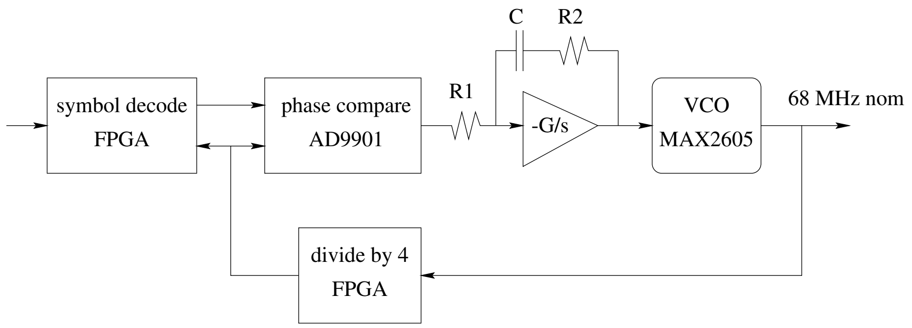

First we need to know what frequency lock range might be needed in operation. Beam energy in the ring can vary from 840 MeV to 1300 MeV. This energy range is bounded from below by the lowest politically acceptable output neutron flux, and at the high end by the size of magnets in the transfer lines and ring. The beta for protons (m_0=938.232 MeV) at these two extremes of energy are 0.8495 and 0.9079. The ring is 248 m in nominal circumference, and actual path length can vary from that by $\pm$0.05 %. That places the range of revolution frequencies from 1.021 to 1.102 MHz. The Event Link transmits bits at 16$\times$ this frequency, and the symbol decode logic needs to sample the input stream at 4$\times$ that rate, yielding a final VCO span of 65.34 to 70.53 MHz. The Maxim MAX2605 VCO discussed in the next section appears to meet this requirement.

2 VCO Evaluation

The Maxim MAX2605 was hooked up to a Tek 'scope for jitter measurement. Analysis of a 160 μs waveform sampled at 1.25 GHz showed an average frequency of 68.5795 MHz. Interpolating the zero crossings and computing their difference from that average rate gives us the following picture.

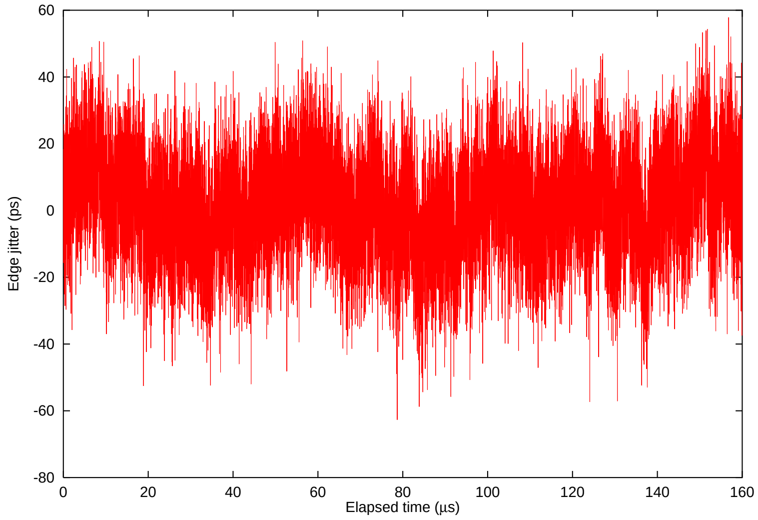

Investigation of the frequency spectrum shows the overall 17.1 ps rms jitter is made up of about 5 ps of sub-1 MHz $1/f$ noise, 11.5 ps of a 15.58 MHz line (that probably originates from the unshielded test environment), and 11.3 ps of broad-band noise that continues all the way to the 34 MHz Nyquist frequency. The relative contribution of the VCO chip and the 'scope to this jitter measurement is unknown.

3 PLL Stability Analysis

Output of the MAX2605 Voltage Controlled Oscillator is a nonlinear function of the input voltage, as shown on the data sheet. The derivative of this curve is called the tuning gain, and is shown alongside.

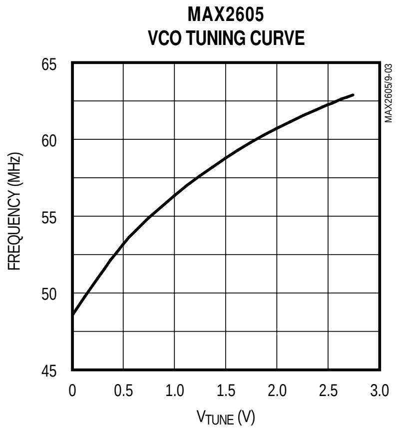

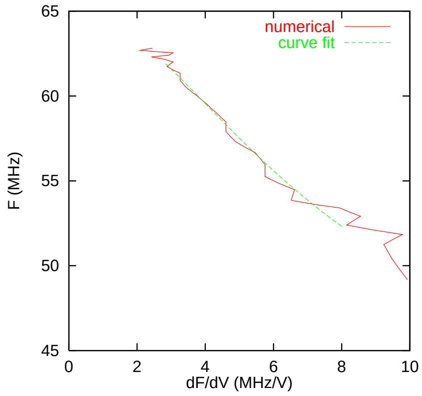

The small signal analysis of stability will refer to a particular operating point, around which changes in the output frequency ω_1 are linearly related to changes in the control voltage v_c according to

$$\omega_1 = K_f v_c \ .$$

Where K_f varies from 1.7×10^7 to 5.0×10^7 $\mathrm{V^{-1}s^{-1}}$ over the recommended span of control voltage (0.4 to 2.4 V).

The AD9901 phase comparator is linear over a broad range of phases, and we will only consider stability over a narrow range near zero anyway. Its output voltage v_p is related to the input phase difference ϕ according to

$$v_p = K_p\phi, \ \text{where } \phi = \frac{\omega_1}{ns}$$

and n is the divide ratio, 4 in this case. In its normal configuration, with $\mathrm{R_{set}}=50\,\Omega$, K_p=0.29 Volt/radian.

The filter circuit is based on an op-amp, which is modeled as an integrator with gain-bandwidth product G. This model breaks down at both low and high frequencies, but will be valid over the "interesting" frequency range. The active filter satisfies

$$-v_c\frac{s}{G} = v_p\frac{R_2+\frac{1}{sC}}{R_1+R_2+\frac{1}{sC}} + v_c\frac{R_1}{R_1+R_2+\frac{1}{sC}}$$

which can be rearranged to show the voltage gain

$$\frac{v_c}{v_p} = -\frac{A_0(1+\tau_1 s)}{s(1+\tau_2 s)}$$

$$A_0 = \frac{G}{1+GR_1C}\ , \quad \tau_1 = R_2C\ , \quad \tau_2 = \frac{(R_1+R_2)C}{1+GR_1C}\ .$$

Thus the overall loop gain for the PLL is

$$-\frac{A_1(1+\tau_1 s)}{s^2(1+\tau_2 s)}, \text{ where } A_1 = \frac{K_f K_p A_0}{n} .$$

A conventional PLL design strategy is to place $1/\tau_1$ and $1/\tau_2$ each a factor of α away from the zero-dB crossing frequency. Larger values of α give greater stability, typical choices are 3 to 10. Since we have a rather large gain variation (a factor of $\beta = 3$), due to the nonlinear response of the MAX2605 control input, we need to explicitly account for variability in the zero-dB crossing frequency. Doing so gives $\tau_1/\tau_2 = \alpha^2\beta \approx 100$. Even at high gain, we want to suppress the $f_0 = 16\,\text{MHz}$ component of the AD9901 output by a large amount, say $\gamma = -60\,\text{dB}$. So $|A_1\tau_1/s^2\tau_2| = \gamma = 10^{-3}$ at $s = 2\pi j f$, and $A_1 \approx 10^{11}\,\text{second}^{-2}$. At the low gain end, A_1 is β times smaller.

Finally, $\tau_2 = 1/(2\pi f_0 \gamma \sqrt{\alpha}) = 0.14\mu\text{s}$, and $\tau_1 = \alpha^2\beta\tau_2 = 14\mu\text{s}$. This analysis is enough to produce a plot of loop gain L, and the normalized response to input and output noise, $1/(1-L)$ and $L/(1-L)$ respectively.

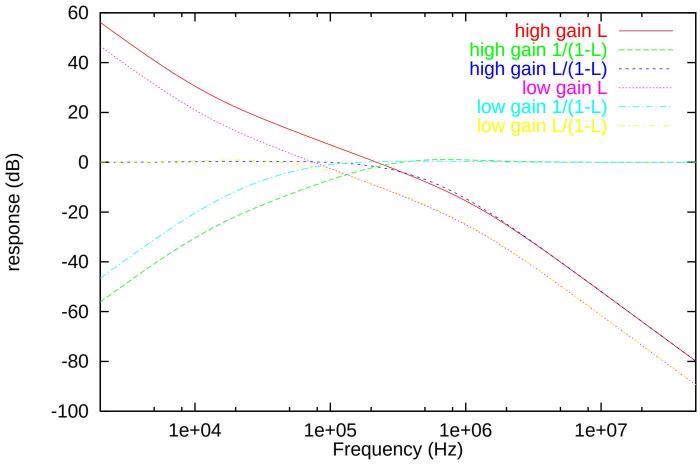

Now we want actual component values. Assume $GR_1C >> 1$, and take R_2 as the free scaling parameter, to compute:

A_0	nA_1/K_fK_p	$4\times10^4\ \text{s}^{-1}$
R_2		$330\ \Omega$
R_1	$R_2/A_0\tau_1$	$600\ \Omega$
C	τ_1/R_2	$0.04\ \mu\text{F}$
G	$(1+A_0\tau_1)/\tau_2$	$2\pi \cdot 1.8\,\text{MHz}$

A larger value of G will slightly increase stability, but diminish the circuit's rejection of 16 MHz. Since $GR_1C = 280$, the assumption above is confirmed valid.

4 Symbol Decode and Clock Reconstruction

So far, we have only assumed that we can extract a useful rising edge from every symbol of the input data stream. Verilog gives us a compact representation of six flip-flops and a combinatorial output that accomplishes this, at least once the frequency is locked. Refer to the timing diagram.

```
reg phase1, phase2, ilatch, iprev, data, error;
always @(posedge clk) begin
   phase1 <= ~phase2;
   phase2 <= phase1;
   if (~phase1 &  phase2) ilatch <= inp;
   if (~phase1 & ~phase2) iprev  <= ilatch;
   if (~phase1 & ~phase2) data   <= ilatch ^ iprev;
   if ( phase1 & ~phase2) error  <= error | ~( inp ^ ilatch );
end
assign phcomp = ilatch ^ inp;
```

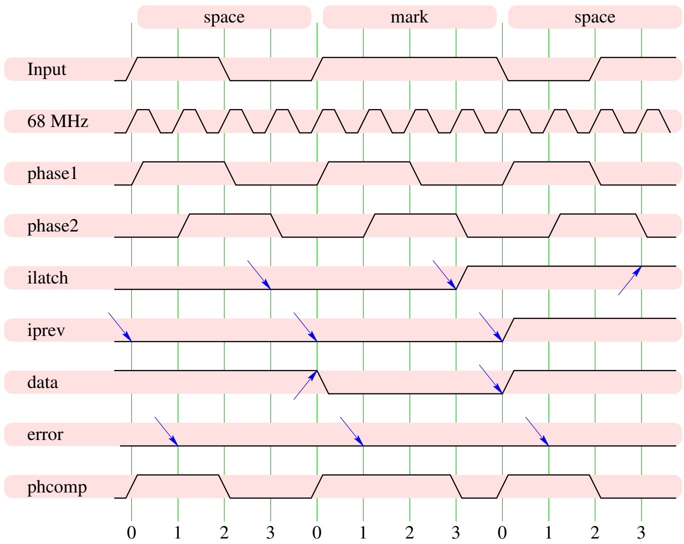

phase1 and phase2 form a simple divide by four. phase1 and phcomp are the two inputs to the AD9901 phase detector. phcomp is guaranteed to have a rising edge at the start of each symbol (when the loop is locked). The falling edge of phcomp happens at different times for mark and space symbols, but the AD9901 logic ignores that edge. The error output is one of many possible consistency checks for proper loop lock and symbol reception.

The next section includes modifications to the above logic that does not get confused by repeated, large phase excursions, as are experienced during initial acquisition of the frequency.

5 Frequency Locking

A slightly more complicated set of logic, based on the one shown in the previous section, was simulated along with the AD9901, analog filter, and MAX2605 VCO. When given a steady stream of “marks”, the system locks in less than 80 μs.

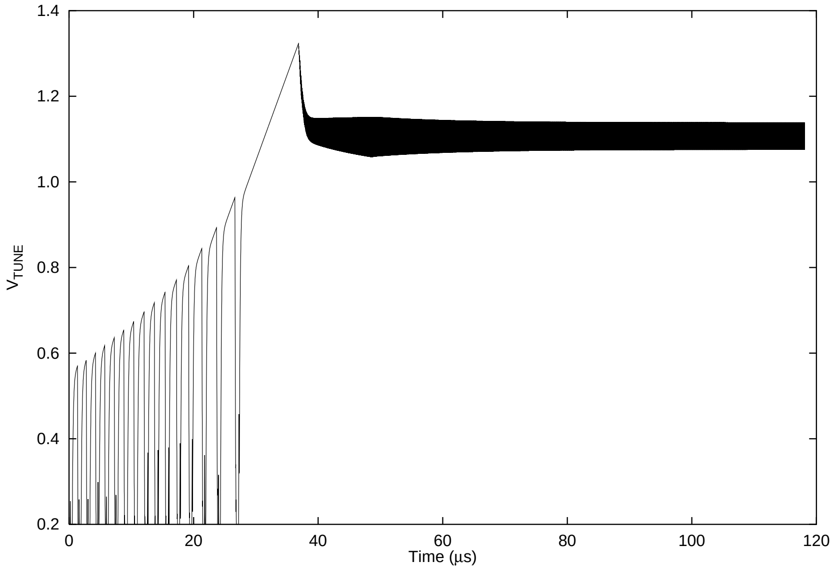

The condensed Verilog code that has to be implemented in the FPGA to achieve this lock is as follows:

```
reg phase1, phase2, d1, d2, d3, retl, advl, prev, data, error;
assign adv = d1==d3 & d1 != d2 & phase1 & phase2;
assign ret = d1==d3 & d1 == d2 & phase1 & phase2 & ~retl;
always @(posedge clk) begin
   phase1 <= ~phase2 ret;
   phase2 <= phase1 & ~adv;
   advl <= adv;
   if (  phase1 & ~phase2       ) d1 <= inp;
   if (  phase1 &  phase2 & ~ret) d2 <= inp;
   if ((~phase1 &  phase2) | adv) d3 <= inp;
   if (  phase1 &  phase2) retl  <= ret;
   if ( ~phase1 & ~phase2) prev  <= d1;
   if ( ~phase1 & ~phase2) data  <= d1 ^ prev;
   if (  phase1 &  phase2) error <= ~( d3 ^ d1 );
end
assign phcomp = d3 ^ inp ^ adv;
```