Phase Detectors*

Christopher Schmitz

Abstract

This module presents an overview of Phase Detectors for use in Phase-Locked Loops.

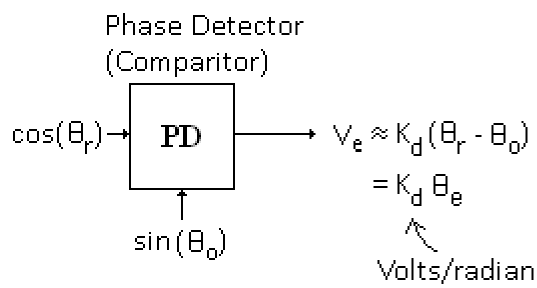

Figure 1: The Phase Detector of a common PLL.

An ideal phase detector (PD) compares a reference phase θ_r to the output phase θ_o of the closed-loop-PLL voltage-controlled oscillator (VCO). Ideally, the PD outputs an error voltage proportional to the phase error $\theta_e = \theta_r - \theta_o$, at least for some range of θ_e. If the proportionality constant is given by K_d (volts/rad), then the error voltage is given by $V_e \approx K_d \theta_e$. The value of K_d can be found from the slope of V_e vs. θ_e.

*Version 1.7: Feb 23, 2010 5:16 pm -0600

Relationship of Kd to Ve and θe

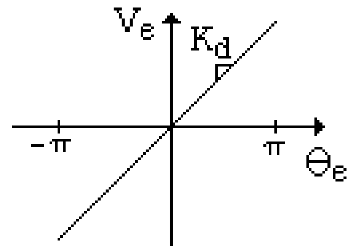

Figure 2: Ideal Phase Detector characteristic.

Variation in phase detector designs result due to the tradeoffs in hardware complexity with accuracy of the phase error estimate in analog implementations. When θ_e can be measured directly, like in software implementations, the phase detector is trivially derived from the known values of θ_r and θ_o. We see this in the QAM receiver discussed below.

1 Calibration

Calibration to determine the gain constant of a phase detector is made possible by sending the same sinusoid into both the reference and VCO ports of the phase detector (so that both ports have the same reference frequency) while controlling the phase of one input through a variable delay. Alternately, one could control the phase by using a FIXED delay (SAW device) and making slight changes in the reference frequency.

2 Best Case: Piecewise Linear Phase Detectors

Sawtooth PD

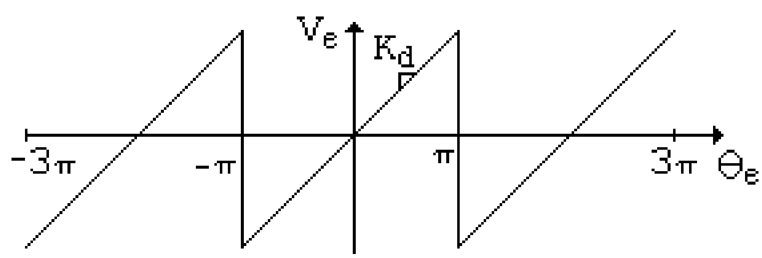

Figure 3: A sawtooth Phase Detector.

A PD would only be a true-linear PD only if the phases are all "unwrapped" with respect to some absolute location in time, say, by referencing both the input reference and the PLL output to a common clock. Unfortunately, in a system for which the reference signal and the PLL are operated on different clocks, there is no way to define and track such an absolute time reference. Yet, outside of perhaps a non-critical phase-ambiguity and some modeling details, an absolute time reference is not needed. Knowing θ_r and θ_o modulo-2π is both practical and sufficient. For data-modulated signals, say BPSK, the main effect of piecewise (rather than absolute) linearity is an ambiguity between between the +1s and the -1s which must be cleared up by the use of a preamble or through differential encoding. We will also find later, that piecewise linearity results in depressing the Phase Detector gain constant when the reference signal is noisy. The piecewise-linear sawtooth PD outputs the result of ($\theta_r - \theta_o$)mod-2π.

Consider the case where the VCO "mixes" the reference signal to complex-valued baseband (also called zero-IF) signal I+jQ where the real-valued I is called the in-phase channel and the real-valued Q is called the quadrature channel. This is known as the QAM (Quadrature Amplitude Modulation) Receiver. In the figure below, $\theta_r = 2\pi$f_c t and $\theta_o = 2$ π f_c_hat t + theta_hat. The received reference signal is rotated by $-\theta_o$ to produce the new complex-valued signal $re^{j\theta_e}$ which can written in rectangular coordinates as the complex-valued point (I Q) where $I + jQ = re^{j\theta_e}$. If the PLL has perfectly tracked the input in frequency and phase, then θ_e (and V_e) will be equal to 0. This also implies I equals +/- r and Q is zero. In other words, the sawtooth PD outputs exactly $\theta_e = \tan^{-1}$ (q/i).

Quadrature Channels

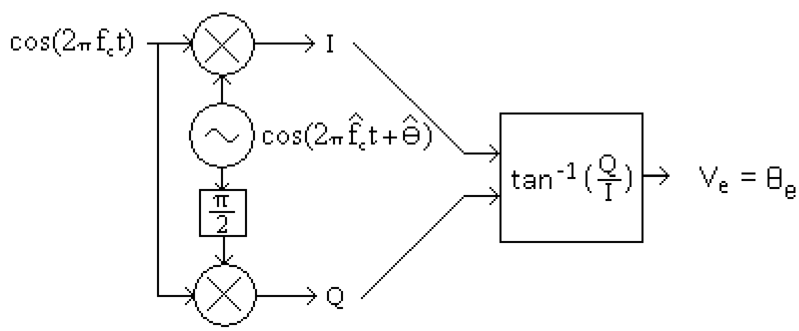

Figure 4: Obtaining quadrature components.

Relationship of θe to I,Q

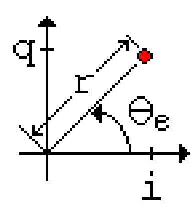

Figure 5: Relationship of θe to quadrature components I and Q. Any received complex-valued data point at time t, let's call it (i,q), is directly related to the phase error in a cartesian-to-polar transformation.

Of course, such a device must be easily implemented in hardware for an analog PD or must be compuatationally efficient for a software (or FPGA-based) PD. As $\tan^{-1}(x)$ generally is not, a suitable approximation is often sought. The first solution might be to utilize a lookup table for the argument x mapping it to $\tan^{-1}$. As accuracy is improved via a larger lookup table, more memory is required and lookup time increased.

3 Traditional Small-Angle Approximation: Ve = q

Another solution involves replacing $\tan^{-1}(q/i)$ with the mathematically equivalent $\sin^{-1}\left(\frac{q}{\sqrt{i^2+q^2}}\right)$. Assuming the PLL is locked and the phase error θ_e reasonably small, $\sin^{-1}\left(\frac{q}{\sqrt{i^2+q^2}}\right)$ can be approximated by the argument $\frac{q}{\sqrt{i^2+q^2}}$. Since $\sqrt{i^2+q^2}$ is simply the sqrt of the input power, it can be removed using Automatic Gain Control (AGC) for a flat-fading channel or considered part of the fixed value K_d for a non-fading channel. In this case the error voltage becomes $V_e = q \approx K_d\theta_e$. This corresponds to the bottom "leg" of the QAM receiver.

Traditional PD

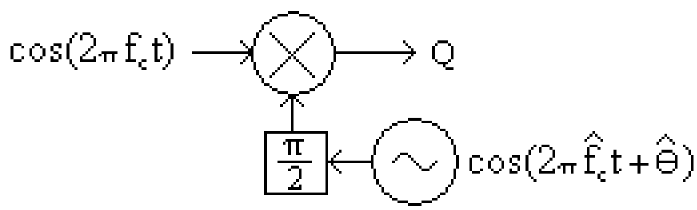

Figure 6: The "traditional" mixer-type phase detector.

Of course, cosine shifted by $\pi/2$ is just a sinusoid and the mixer produces both DC and double-frequency terms. Therefore, this figure can be more accurately represented as

Traditional PD (Common Representation)

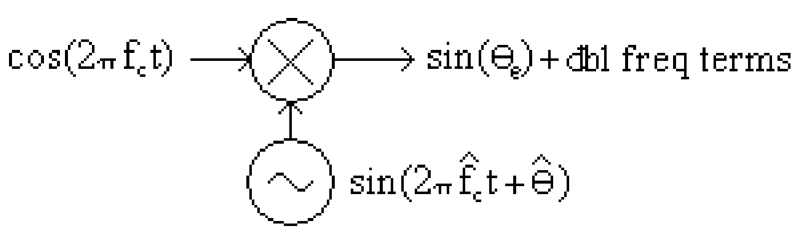

Figure 7: The more common representation of the traditional phase detector.

In a PLL, the double-frequency terms will likely be removed by the loop filter and are of little concern.

4 Costas BPSK: Ve = i times q

For the pure sinusoidal input considered up to this point, zero phase error is synomonous with a single point in the I-Q plane, located on the positive I-axis. For a Binary Phase-Shifted Signal (BPSK), the signal with zero phase error actually exists at two points within the I-Q plane, as the data stream varies between a positive and negative I-axis value. This results in θ_e being properly defined as zero in two locations on the I-Q plane. The approximation $V_e = q$ indeed gives these two zeros, but for the left-half plane has a negative slope! We must modify the approximation by taking into account the sign of I. In a first attempt, we may multiply by the value i directly: $V_e =$ iq.

Costas PD Response

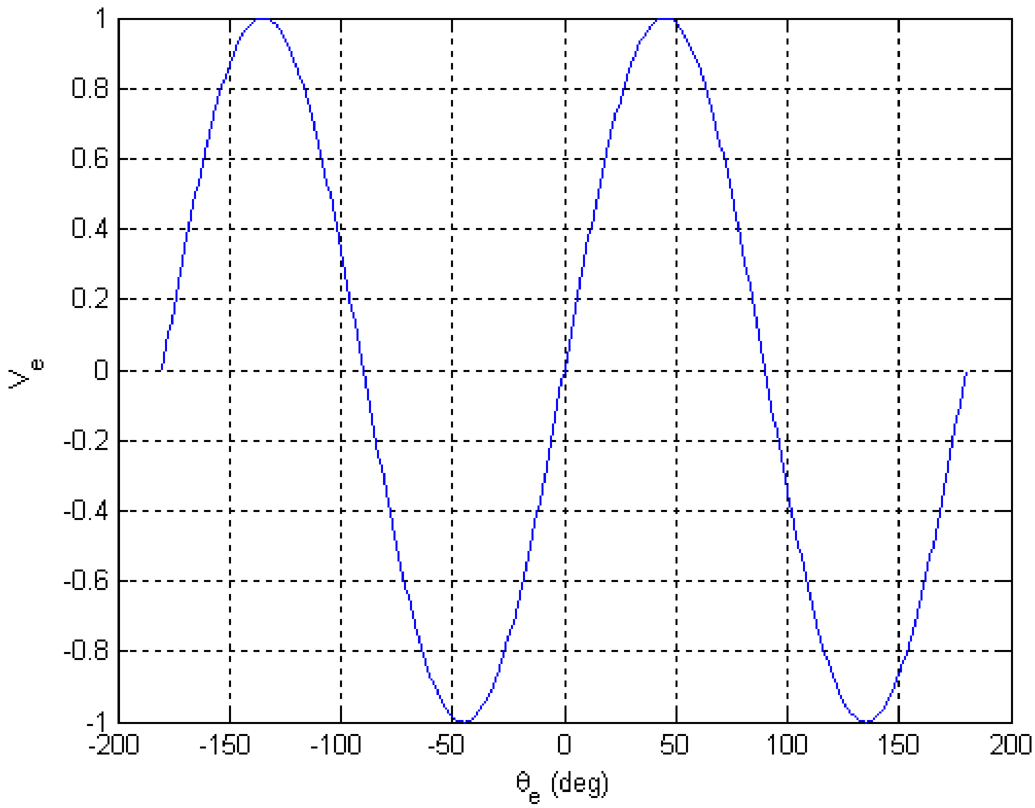

Figure 8: The "S-curve" of a Costas Phase Detector.

5 Modified Costas BPSK: Ve = Sign[i]q

The simplicity of the Costas solution has resulted in widespread use. However, allowing the magnitude of i to affect the approximation of Ve has limited the linear range of the phase error estimate to about (-45,+45) degrees. With the added complexity of a hard-limiter, the I-channel data can be made to affect only the sign, resulting in a much improved phase error estimate: $V_e = \text{Sign}\,[\text{i}]\,\text{q}$. The linear range of the phase error estimate now approaches (-90,+90) degrees as seen in Figure 10.

Modified Costas PD Response

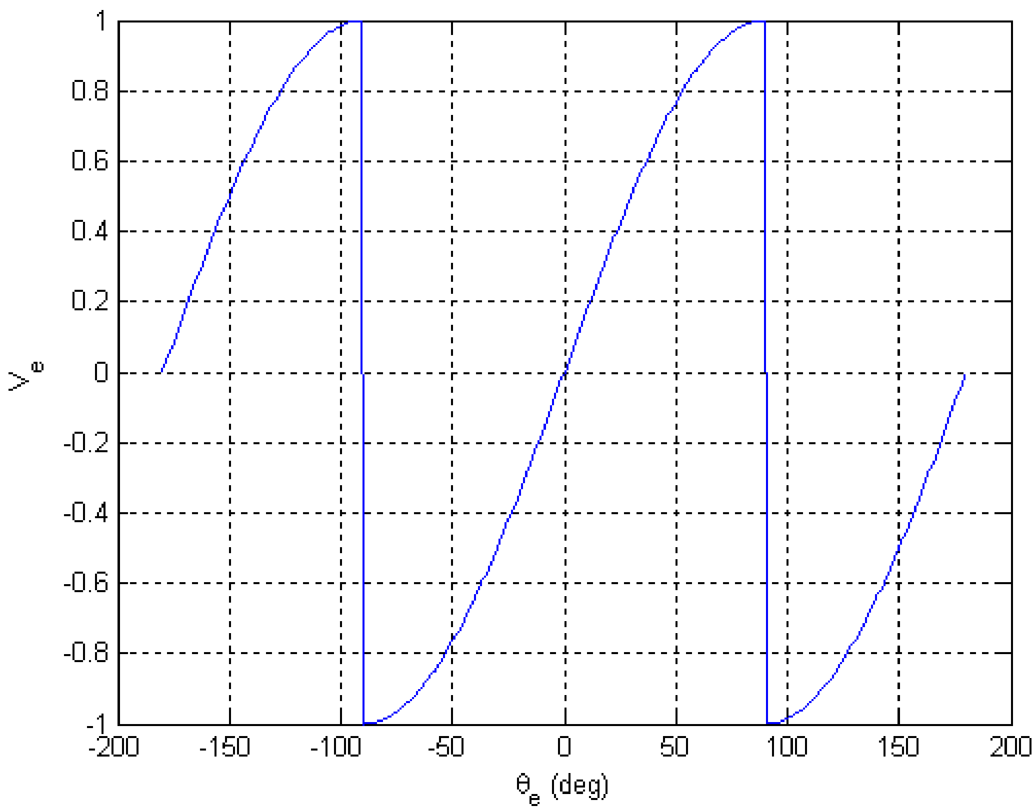

Figure 9: The "S-curve" of a Modified Costas Phase Detector.

The resulting S-curve is nearly equivalent to the ideal sawtooth. In fact, even in low noise, the performance of the modified Costas phase detector and the ideal phase detector is indistinguishable.

6 Linearity and the Effect of Noise

When the input is noisy, the effect is for the PD to occassionally map the true phase error across the discontinuity of the defining S-curve. Large-magnitude phase errors (which sit closer to the discontinuity) are affected more so than small ones. Therefore, as the SNR falls, the first effect is a softening of the sawtooth's "teeth." The sawtooth becomes more sinusoidal. As the SNR continues to fall, eventually even small-magnitude phase errors are affected and the apparent slope (defining the PD gain constant K_d near θ_d=0) decreases in value.

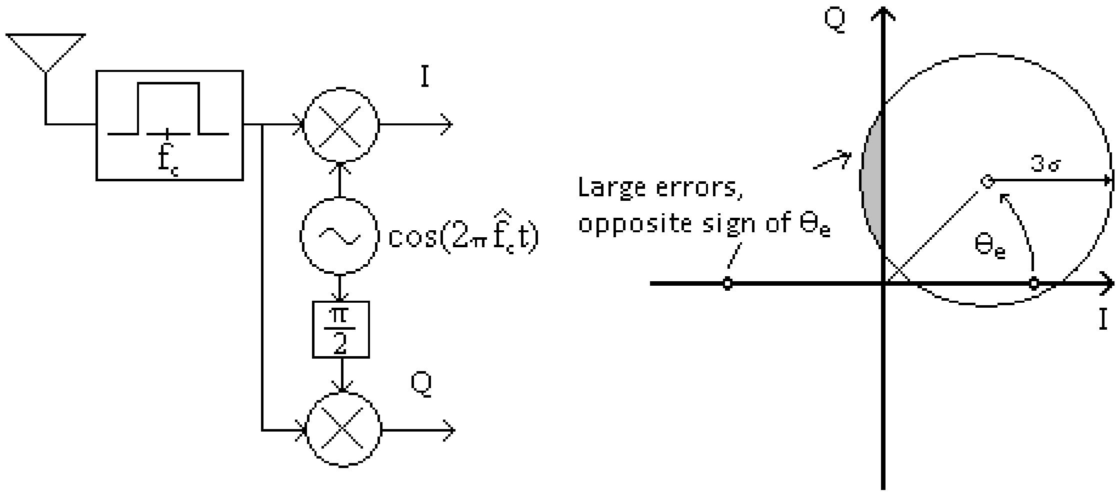

Figure 10: This is why noise effects the S-Curve by rounding off the peaks.