



Xilinx Inc.
 2100 Logic Drive
 San Jose, CA 95124
 Phone: +1 408-559-7778
 Fax: +1 408-559-7114
 E-mail: coregen@xilinx.com
 URL: www.xilinx.com




Figure 2: Core Schematic Symbol

Features

- Fully parameterized Numerically Controlled Oscillator
- Input phase increment resolution from 3 to 30 bits
- Output amplitude resolution from 4 to 16 bits
- Frequency resolution is controlled via phase accumulator word length (3 to 30 bits)
- User can control phase noise via programmable phase resolution (3 to 10 bits)
- Sine and cosine outputs available – excellent for high speed I/Q modulation/demodulation
- Drop-in modules for the XC4000E, EX, XL and Spartan families.
- High performance and density guaranteed through Relational Placed Macro (RPM) mapping and

- placement technology
- Available in Xilinx CORE Generator Tool

Functional Description

The Numerically Controlled Oscillator (NCO) module generates a digital “staircase” approximation to a sine (or cosine) wave, the frequency of which is determined by the input phase increment value. The output can either be used directly, for example, by a digital multiplier, or can be passed into a Digital-to-Analog Converter (DAC) for use in the analog domain.

The module comprises an increment register, a phase accumulator, and a sine/cosine ROM look-up table (LUT). The increment register stores the phase increment value,

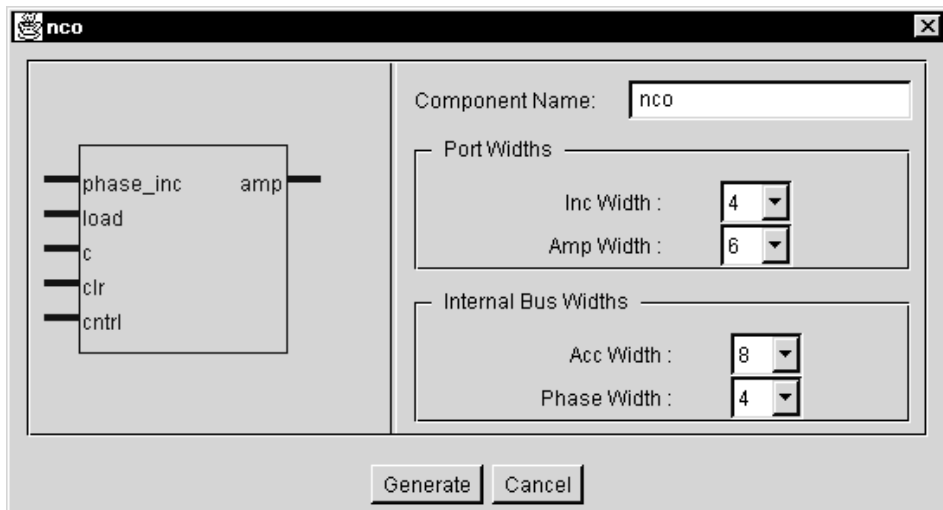


Figure 1: Parameterization Window

which is continually integrated by the phase accumulator. The integrated phase value (or a truncated version of the same) is used to address the sine/cosine LUT, which outputs the amplitude corresponding to the sine (or cosine) of the current phase value.

Theory of Operation

Typically, sinusoids can be expressed as functions of their phase angle, θ :

$$a = \sin\theta, \text{ or } a = \cos\theta$$

However, phase can also be expressed as the time integral of angular frequency, ω , such that a change in phase, $\delta\theta$, is given by:

$$\delta\theta = \omega\delta\tau,$$

where $\delta\tau$ is the time interval over which the phase change took place. Rearranging for ω :

$$\omega = \delta\theta/\delta\tau = 2\pi f,$$

where f is the frequency (in Hertz) of the sinusoid. In the context of a numerically controlled oscillator, the time interval, $\delta\tau$, is determined by the frequency, f_{clk} , at which the phase accumulator updates its output value, by the following equation:

$$\delta\tau = 1/f_{clk}$$

It is therefore possible to express the output frequency of the sinusoid in terms of the phase change and the clocking rate of the phase accumulator:

$$f = \delta\theta \cdot f_{clk} / 2\pi.$$

In this equation, the denominator represents the maximum attainable phase value before the sinusoid “wraps round” and repeats in a periodic manner. In the case of an NCO, the maximum possible phase value is governed by the phase accumulator, which maps the 0 to 2π phase range into a digital word. Thus, for an n -bit accumulator, $2\pi = 2^n$, and the phase change term, $\delta\theta$, when expressed as a digital word, ϕ_{word} , must lie in the range:

$$0 \leq \phi_{word} \leq 2^n - 1.$$

The output frequency of the NCO can therefore be expressed as:

$$f = \phi_{word} \cdot f_{clk} / 2^n.$$

So, for example, with a clock frequency of 50 MHz, a phase word of $0F_{16}$, and an accumulator word length of 8 bits, the oscillator frequency is:

$$f = 0F \times 50 \text{ MHz} / 2^8 = 2.9296875 \text{ MHz}.$$

The implicit assumption here is that the phase word and the accumulator word are both of the same length. Clearly, however, the result of the above example is unaffected if the phase word is chosen to be only 4 bits long, as compared with 8 bits. If, on the other hand, the original phase word had been 10_{16} (giving $f = 3.125 \text{ MHz}$), the result of limiting the precision to 4 bits (i.e. ignoring the four most

significant bits) would have been quite different ($f = 0 \text{ Hz}$). It can be seen therefore that the word length of the phase change term sets an upper limit on the achievable output frequency – the output frequency range is maximized when the word lengths of the phase change term and the accumulator are the same. The accumulator word length effectively sets the frequency resolution of the oscillator.

The ability to set the maximum oscillator frequency, independently of clock rate and frequency resolution, allows the user to guarantee that the Nyquist criterion will not accidentally be violated. This is useful in preventing unwanted aliasing effects.

The other oscillator parameter which is of prime concern is noise. In addition to reducing system dynamic range, noise in an oscillator can generate spurious signals, as well as increasing side-band energy. Both of these latter effects can lead to undesired mixing products (which may be difficult to filter) which can cause increased in-band interference, as well as a generally “dirty” spectrum which may result in interference for other users. In an NCO, noise in the output signal comes from two main sources (other than the master clock): amplitude quantization and phase quantization.

For amplitude quantization, the output signal-to-noise ratio (SNR) as a function of word length, n , is approximately:

$$SNR_{dB} = 6n + 1.8$$

Phase quantization results in the formation of spurious spectral artifacts. The spurious free dynamic range (SFDR) of the NCO (i.e. the ratio of the power of the desired signal to the power of the strongest spur) can be determined from the phase word length, n :

$$SFDR_{dB} = 6n$$

Table 1 tabulates values of SNR and SFDR against different word lengths.

Table 1: Effect of Word Lengths on SNR and SFDR

Number of Bits	SNR _{dB} (amplitude)	SFDR _{dB} (phase)
8	49.8	48
10	61.8	60
12	73.8	72
14	85.8	84
16	97.8	96

The LogiCORE NCO provides parameterized amplitude and phase quantization levels, and so allows the user complete freedom to trade design size and complexity for output SNR and SFDR.

Pinout

Signal names for the schematic symbol are shown in Figure 1 and described in Table 2.

Table 2: Core Signal Pinout

Signal	Signal Direction	Description
phase_inc[p:0]	Input	Phase increment value
load	Input	Active high – allows a new value of phase to be loaded
clr	Input	Asynchronous clear of phase register and accumulator
cntrl	Input	When high, NCO returns sine; when low, NCO returns cosine.
c	Input	Clock – active on rising edge
amp[a:0]	Output	Output data

CORE Generator Parameters

The CORE Generator accepts parameters entered through the dialog box and creates the specific design from the values entered using a parameterized VHDL recipe. VHDL instantiation code and a schematic symbol are created along with the netlist for the design.

The parameters are as follows:

- **Component name:** Enter a name for the component.
- **Phase increment width:** Select an input bit width, P, from the pull-down menu. The valid range is 3 to 30.
- **Accumulation width:** Select a phase accumulation bit width from the pull-down menu. The valid range is 3 to 30, but must be greater than or equal to the phase increment value.
- **Phase width:** Select a phase bit width from the pull-down menu. The valid range is 3 to 10 but must be less than or equal to the accumulation width.
- **Amplitude width:** Select an output bit width, A, from the pull-down menu. The valid range is 4 to 16.

Latency

The module has two pipeline latencies which depend on phase width. The numbers of clock cycles required before a valid output is obtained, after assertion of “load”, or after a toggle of “cntrl”, are shown in Table 3.

Core Resource Utilization

The number of CLBs required for the module is a function of the phase increment width, accumulation width, phase width, and the amplitude width. The following table shows the equations to calculate the maximum number of CLBs

required. In these equations, **x** is the amplitude width, **y** is the accumulation width, and **z** is the phase increment width. (When using these equations, round down to the nearest integer.)

Table 3: NCO Latency

Phase Width	Load	Cntrl
3 to 4 bits	3	1
5 bits	4	2
6 to 10 bits	5	3

Table 4: NCO Resource Use

Phase Width	CLB Count
3	$(x+1)/2 + (y+1)/2 + 1 + (z+1)/2$
4	$x + (y+1)/2 + 1 + (z+1)/2$
5	$x + (x+1)/2 + (y+1)/2 + 1 + (z+1)/2$
6	$x + 4 + (y+1)/2 + 1 + (z+1)/2$
7	$2x + 4 + (y+1)/2 + 1 + (z+1)/2$
8	$2x + 6 + (x+1)/2 + (y+1)/2 + (z+1)/2$
9	$4x + 7 + (x+1)/2 + (y+1)/2 + (z+1)/2$
10	$8x + 8 + (x+1)/2 + (y+1)/2 + (z+1)/2$

Performance Characteristics

The combination of the number of parameters and their ranges results in a large design space for this core. When coupled with the need to run oscillators for a significant time in order to perform an accurate frequency domain analysis, the difficulties inherent in fully characterizing such a design become all too apparent (especially when different combinations of Xilinx parts and speed grades are also considered). For this reason, no attempt is made here to characterize fully the NCO core. Rather, what follows is an example of the sort of performance which can be expected from one particular point in the design space.

The following spectral plots were obtained by performing a MatLab FFT (with a Blackman window) on raw (i.e. an interpolation filter was not employed) data generated by a ModelSim simulation of an NCO core. From plots such as these, estimates of spurious free dynamic range (SFDR), total harmonic distortion (THD), and the like can be made, as application requirements dictate.

Applications Information

In the previous section, the example spectral plots were obtained from raw simulation data. It can be clearly seen that some of the plots (e.g. Figures 3, 4, and 5) exhibit strong spurious signals at odd harmonic multiples of the NCO output frequency. It should be noted also that higher order harmonics have been aliased back into the base band (the Nyquist frequency being 25 MHz).




Figure 3: Power Spectral Density, 2.5 MHz Tone




Figure 6: Power Spectral Density, 10 MHz Tone




Figure 4: Power Spectral Density, 5 MHz Tone




Figure 7: Power Spectral Density, 12.5 MHz Tone




Figure 5: Power Spectral Density, 7.5 MHz Tone




Figure 8: Power Spectral Density, 15 MHz Tone

The strong harmonic content is due entirely to the fact that the oscillator output has not been filtered. As a consequence, the Fourier analysis is performed on a “staircase” approximation to a sine wave, which remains harmonically rich because there is no band-limiting in the simulation environment. It is recommended therefore that, in applications requiring high spectral purity (and especially where mixing is involved), the NCO output should be passed through an interpolation filter, which “smooths” the steps in the staircase, and suppresses the harmonics. (The interpolation filter is essentially a low pass filter which will pass the fundamental frequency and block the harmonics.)

Some applications may require that the NCO drive a DAC, in order to produce an analog output. In such cases, careful consideration should be given to the specification of the DAC. In particular, DAC non-linearities will introduce harmonics, and mixing products of harmonics. Higher order harmonics will fold back in-band, and cannot be filtered. Furthermore, DAC glitch energy will introduce spurs which can also degrade performance. It is important therefore to choose a DAC with high linearity (to suppress harmonics and mixing as much as possible), and large spurious free dynamic range (to minimize spurs). The DAC must also have adequate resolution and bandwidth for the intended application.

In addition to these considerations, it is advisable to use an interpolation filter on the DAC output, to further attenuate harmonics and mixing products. The filter must pass the fundamental, but reject the first spurious harmonic (which may have been aliased), and must of course be analog. As the output frequency of the NCO approaches the Nyquist rate, the degree of separation between the fundamental and any aliased components is greatly reduced, placing ever more stringent requirements on the interpolation filter design. In practice, these filtering constraints make it difficult to achieve an oscillator output frequency greater than $f_{clk}/3$, even though the Nyquist criterion sets an upper bound of $f_{clk}/2$.

Typical uses of the NCO core would be in frequency synthesis and digital modulator applications.

Ordering Information

This macro comes free with the Xilinx CORE Generator. For additional information, contact your local Xilinx sales representative or e-mail requests to coregen@xilinx.com.

Parameter File Information

Parameter	Type	Notes
COMPONENT_NAME	String	
INC_WIDTH	Integer	3-30
ACC_WIDTH	Integer	3-30
PHASE_WIDTH	Integer	3-10
AMP_WIDTH	Integer	4-16