Simulating Noise in Nonlinear Circuits Using the HP Microwave and RF Design Systems

Product Note 85150-4

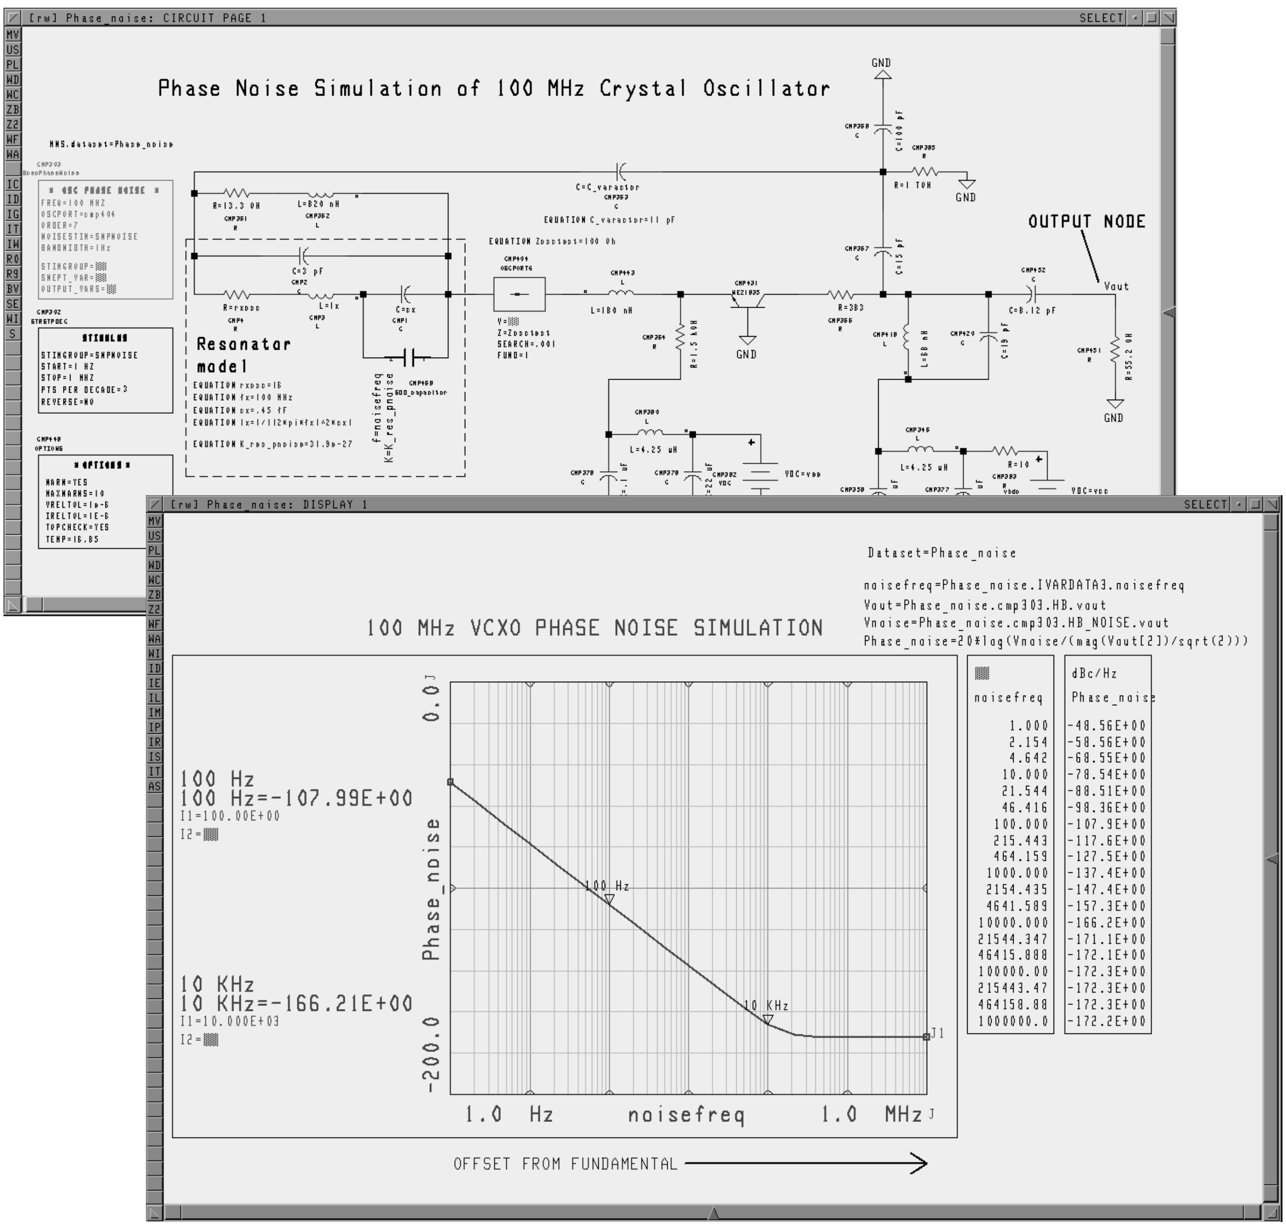

noisefreq	Phase_noise (dBc/Hz)
1.000	-48.56E+00
2.154	-58.56E+00
4.642	-68.55E+00
10.000	-78.54E+00
21.544	-88.51E+00
46.416	-98.36E+00
100.000	-107.9E+00
215.443	-117.6E+00
464.159	-127.5E+00
1000.000	-137.4E+00
2154.435	-147.4E+00
4641.589	-157.3E+00
10000.000	-166.2E+00
21544.347	-171.1E+00
46415.888	-172.1E+00
100000.00	-172.3E+00
215443.47	-172.3E+00
464158.88	-172.3E+00
1000000.0	-172.2E+00

Contents:

Introduction

With the HP Microwave and RF Design Systems (HP MDS and HP RFDS) release B.05 and later revisions, simulation of noise in nonlinear circuits is possible. This simulation capability includes the noise figure of mixers and systems, phase noise of oscillators, and phase noise effects in systems. This product note is divided into several sections, each covering simulation and analysis methods for one of these applications. It covers how HP MDS does the analyses, what the designer needs to include in each simulation, and how to do the simulations.

Oscillator Phase Noise Simulation

<u>Phase noise basics</u>

Ideally, a signal displayed on a spectrum analyzer would appear as a single vertical spike at the oscillation frequency. However, because the signal's frequency varies with time due to phase noise, noise skirts (sidebands) appear around the signal. Phase noise is analyzed theoretically by assuming that the signal is stable and that the noise skirts are due to phase modulation of the signal. The phase noise of an oscillator is usually quantified by the single-sideband (SSB) phase noise, which is defined as the ratio of noise power in a 1 Hz bandwidth at an offset, f_m, to the signal (or carrier) power, as shown in Figure 1

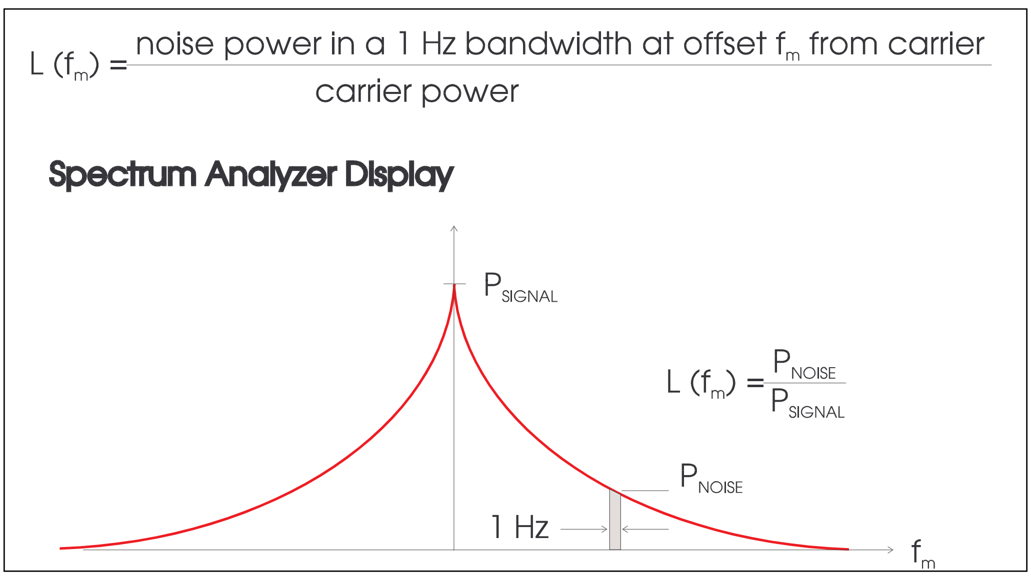

Figure 1: Definition of single-sideband (SSB) phase noise.

A signal at nominal frequency f_0, and unmodulated amplitude, V_s, with phase modulation at a rate, f_m, with peak phase deviation, $\Delta\Phi_{peak}$, can be described by the following equation:

$$v(t)=V_s \cos(2\pi f_0 t + \Delta\Phi_{peak} \sin(2\pi f_m t))$$

This equation describes a signal with the spectrum shown in Figure 2. HP MDS computes the amplitude of the noise sidebands (V_{ssb}, in Figure 2) at each offset frequency, f_m. The phase noise can then be computed as the ratio of this sideband noise power to the signal (or carrier) power, and plotted versus offset frequency, f_m.

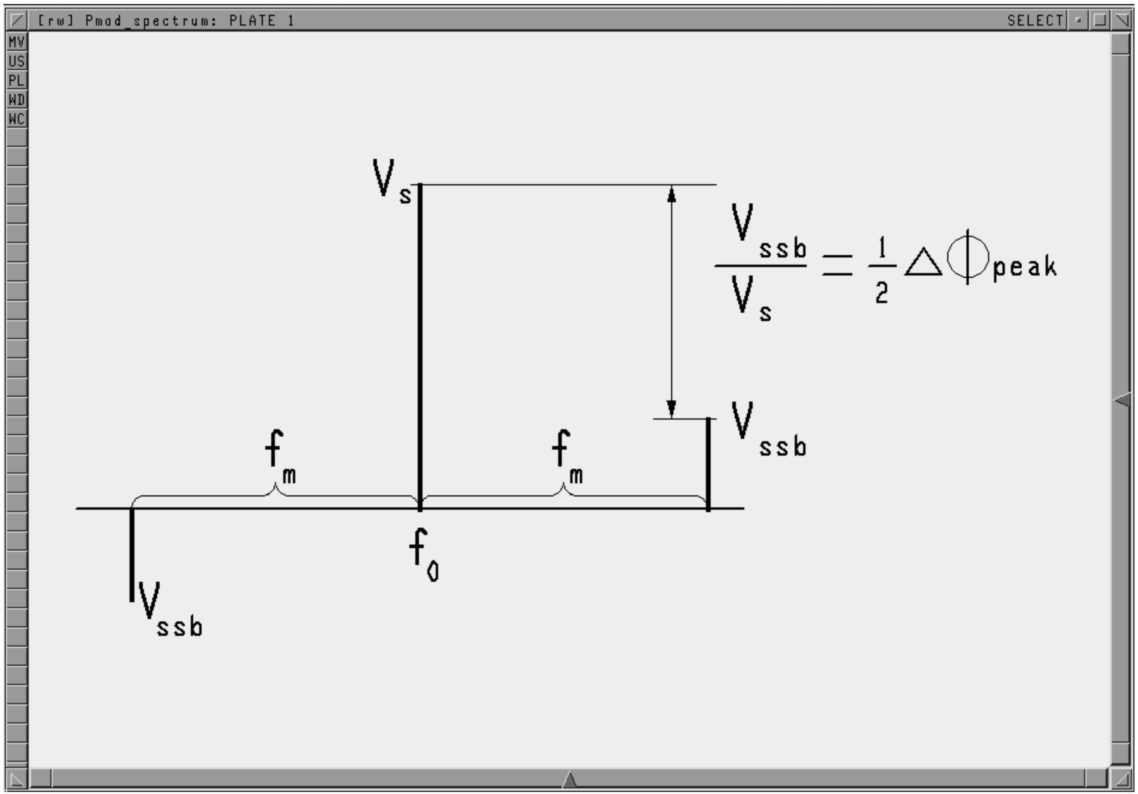

Figure 2: Spectrum of signal with phase modulation.

Please refer to [1], [2], and [3] for a more detailed discussion of phase noise theory and phase noise measurements.

How HP MDS computes oscillator phase noise

HP MDS computes oscillator phase noise by first doing a harmonic balance analysis to solve for the frequency of oscillation and the amplitude of the signal at the circuit output node specified by the user. All phase noise comes from the various noise sources in a circuit, for example, thermal noise from resistors and lossy elements, and shot, 1/f, and burst noise from bipolar junction transistors (BJTs). Any noise source from an active device model, or that the user adds to the circuit will contribute phase noise.

Noise sources contribute to phase noise via two mechanisms. One is direct up-conversion from baseband to the oscillation frequency due to the nonlinear effects in the circuit. This effect is illustrated in Figure 3.

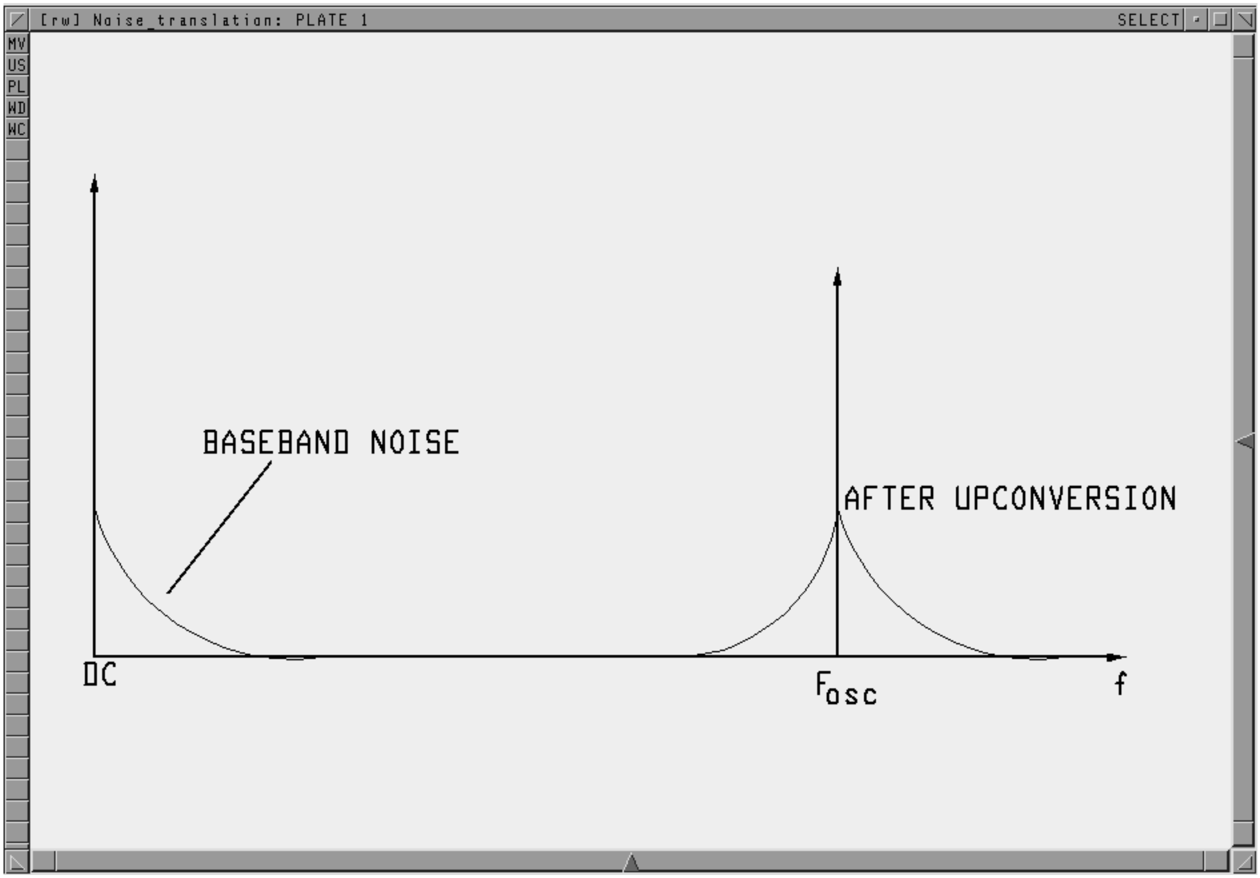

Figure 3: Up-conversion of baseband noise.

Via this mechanism, white noise at DC will contribute white noise around the output signal, and 1/f noise sources will contribute a 1/f noise spectrum around the output signal. The other mechanism is due to the noise sources actually varying the frequency of oscillation. To compute the contribution from this effect, HP MDS computes the derivative of the oscillation frequency with respect to the amplitude of each noise source. From these derivatives, the mean square frequency noise contributed by each noise source is computed. This mean square frequency noise is then converted to phase noise (since phase is the integral of frequency, frequency noise in the frequency domain can be converted to phase noise in the frequency domain by multiplying by $1/f_m^2$, from Fourier transform theory). The equations are shown in Figure 4. Noise sources with white frequency spectra contribute $1/f^2$ phase noise spectra. Noise sources with 1/f frequency spectra contribute $1/f^3$ phase noise spectra. These noise calculations are carried out for each noise source and added independently.

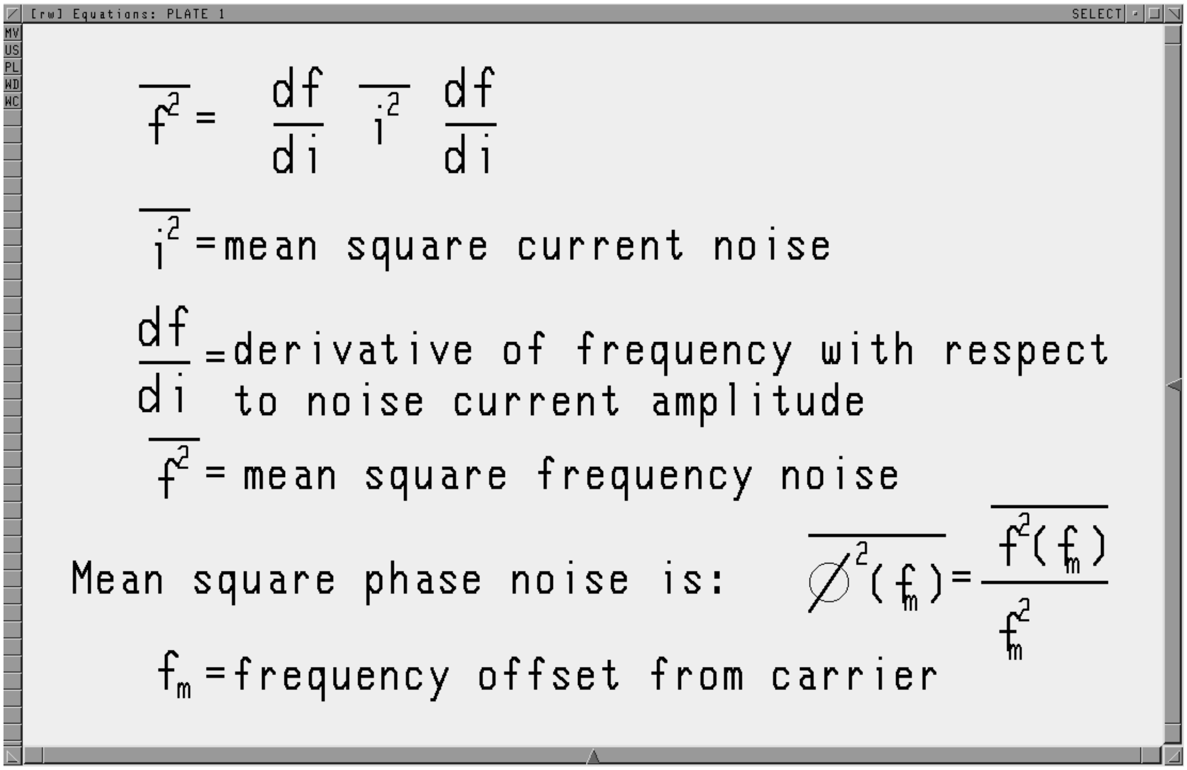

Figure 4: How noise sources cause random frequency fluctuations and phase noise.

Modeling requirements for accurate phase noise simulation:

An accurate phase noise simulation starts with an accurate oscillator simulation. A good nonlinear model for the active device(s) and a model for the resonator (or frequency-selecting element), as well as accurate models for the other passive circuit elements are needed. In addition, it is necessary to model the 1/f noise of the active device. Several considerations for low phase noise design are:

- 1/f noise of the active device(s)
- Resonator 1/f noise or residual phase noise (in some cases)
- Resonator Q
- Circuit design such that the frequency of oscillation is the same as the resonator's center frequency

Resonator simulation:

Generally, the higher the Q of a resonator, the lower the oscillator phase noise. Resonator Q is defined as:

Q=2π*(maximum energy stored)/(energy dissipated per cycle)

Q is also defined as $Q=\frac{\omega_0}{2}\frac{d\omega}{d\omega}$, where ω_0 is the resonant frequency, and

ϕ is the phase of the resonator impedance. For a series RLC resonator, at resonance,

$Q = \omega_0 L/R$, because $\frac{d\phi}{d\omega} = \frac{2L}{R}$ at resonance.

The more frequency-selective the resonator, the larger this derivative, and the better the phase noise. Also, off resonance, this derivative diminishes, causing the Q to decrease.

For simulation purposes, the designer must obtain an equivalent circuit or use measured data for the resonator. Crystal resonator manufacturers usually supply an equivalent circuit model. For oscillators using an LC type resonator, it is important to have good models for the actual circuit elements that will be used. These models can be derived from measurements with a network analyzer, or from libraries such as the HP 85174A RF SMT Library. If axially leaded parts are used at RF frequencies, it is wise to measure their impedance versus frequency to derive an equivalent model, rather than just using their nominal values in a simulation. For dielectric resonators, equations are available for the equivalent circuit elements. These can be used as a first approximation, but again, measurements may be necessary. YIG resonators are commonly used in broadband, tunable microwave oscillators, and equations are also available to describe equivalent circuits based on sphere sizes and coupling geometries. However, it is often necessary to measure these resonators and derive an equivalent circuit. Other resonator elements include varactor diodes and microstrip transmission lines. References [4] and [5] provide information on deriving equivalent circuit models.

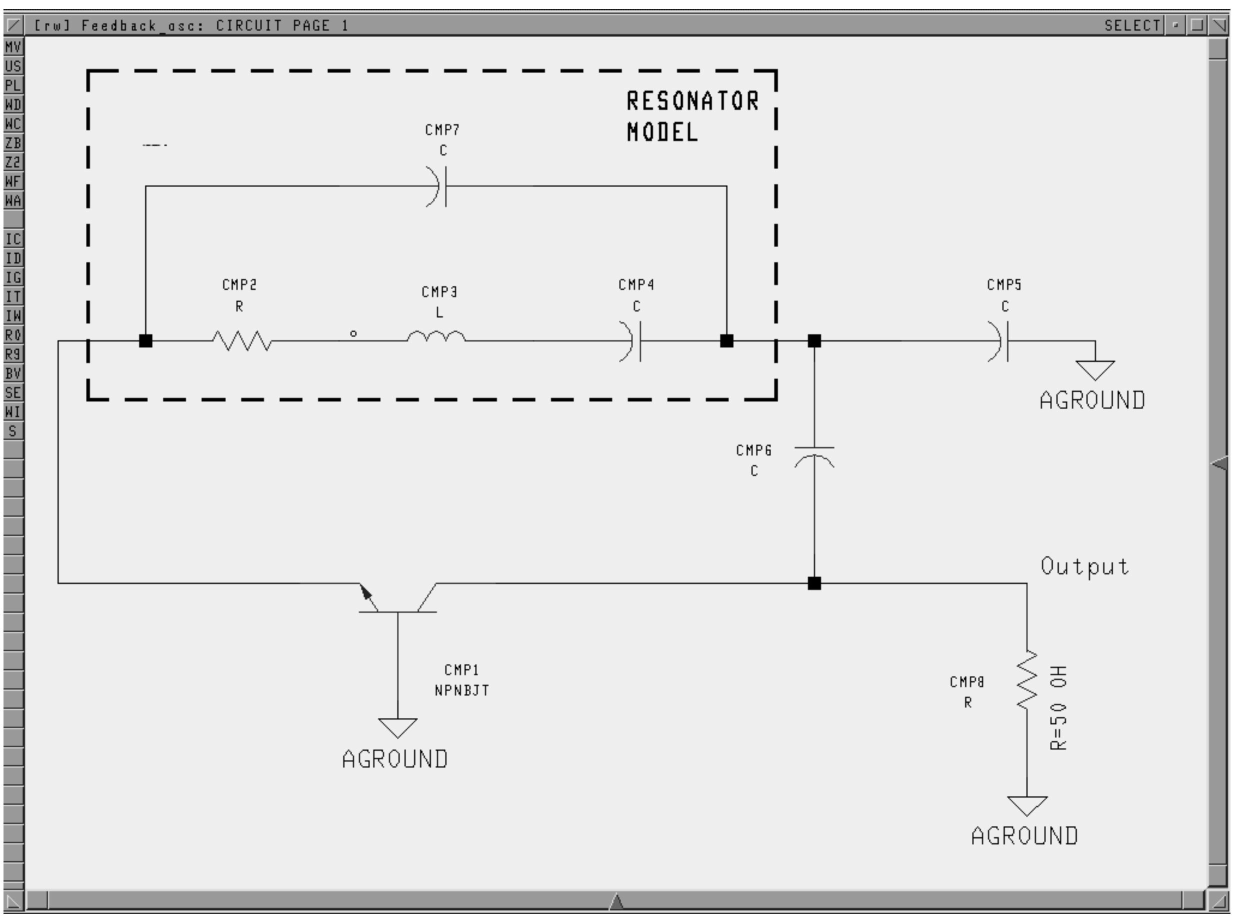

Figure 5: Basic feedback oscillator configuration.

Circuit analysis for optimum phase noise

There are two generic types of oscillators, the feedback type, shown in Figure 5, and the reflection type, shown in Figure 6.

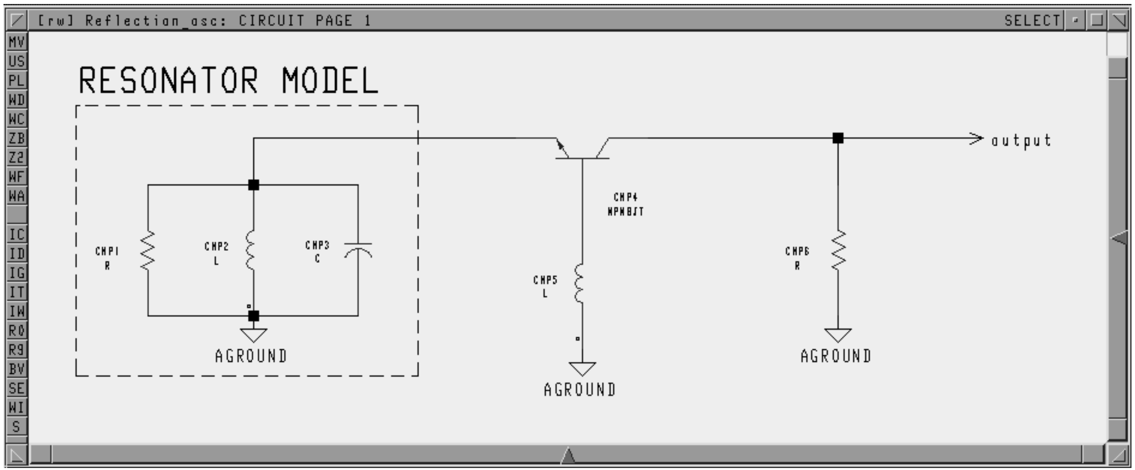

Figure 6: Basic reflection oscillator configuration.

In the feedback type, a signal is amplified by the amplifier, and fed back through the resonator to the input. The circuit will oscillate if the loop gain is greater than one at a frequency where the phase shift around the loop is some multiple of 2π.

Oscillator circuits can be analyzed qualitatively to see if they are designed for optimum phase noise performance. In low phase noise oscillators, a high-Q resonator is used to select the frequency of oscillation. To minimize the oscillator phase noise, the frequency of oscillation should be the same as the resonator's center frequency. The resonator's derivative of phase with respect to frequency will be a maximum at the center frequency, and consequently the resonator will have its highest Q, as described above. A typical resonator phase and magnitude response is shown in Figure 7.

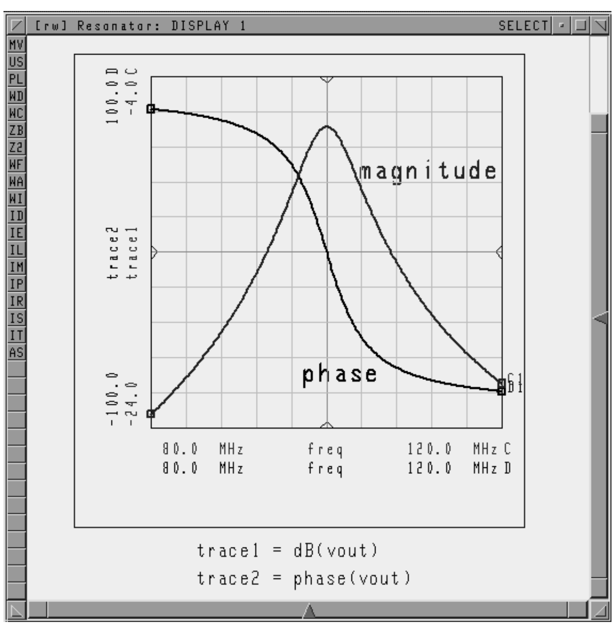

Figure 7: Resonator phase and magnitude response.

This condition will only be achieved if the circuit is designed so that there is no excess phase shift generated by the circuit, excluding the resonator. To accurately model the phase characteristic of the circuit, excluding the resonator, it is necessary to model the components accurately, including transmission line lengths. Generally, the higher the oscillation frequency, the more critical accurate modeling is.

The following describes how the circuit in Figure 5 can be analyzed for optimum phase noise performance. First, the circuit is broken into two parts: the resonator, which determines the frequency of oscillation, and the active circuit that amplifies the oscillating signal. The limiting mechanism of the active circuit determines the amplitude of oscillation. Two simulations are done, 1) a simulation of the large signal impedance of the active circuit seen by the resonator at the resonant frequency, and 2) a simulation of the resonator's impedance versus frequency.

In Figure 8a, the resonator is replaced with a large-signal current source, and the large-signal impedance seen by the resonator is simulated at the resonant frequency, as a function of current source amplitude.

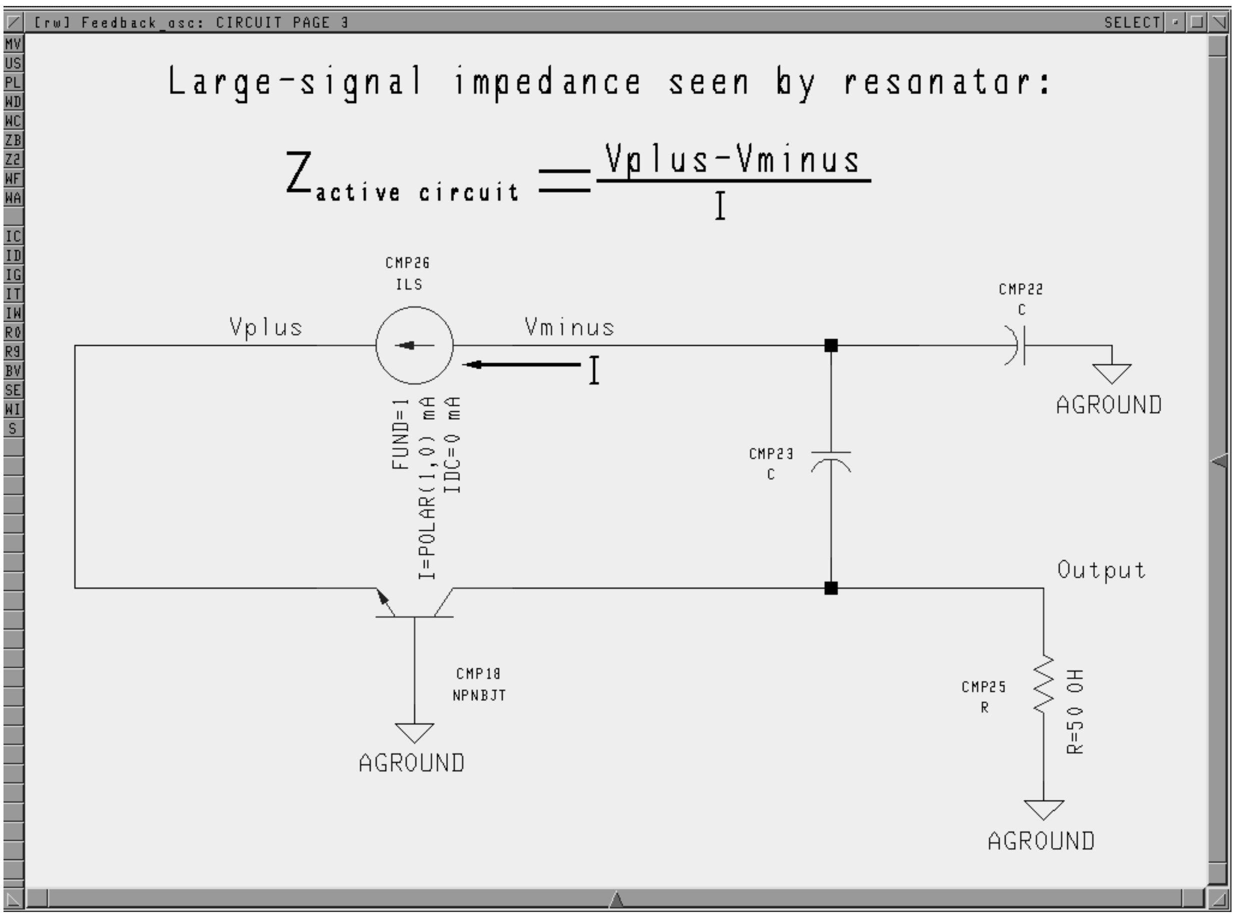

Figure 8a: Resonator replaced with current source, for large-signal impedance analysis.

The impedance of the resonator as a function of frequency is also simulated, as shown in Figure 8b.

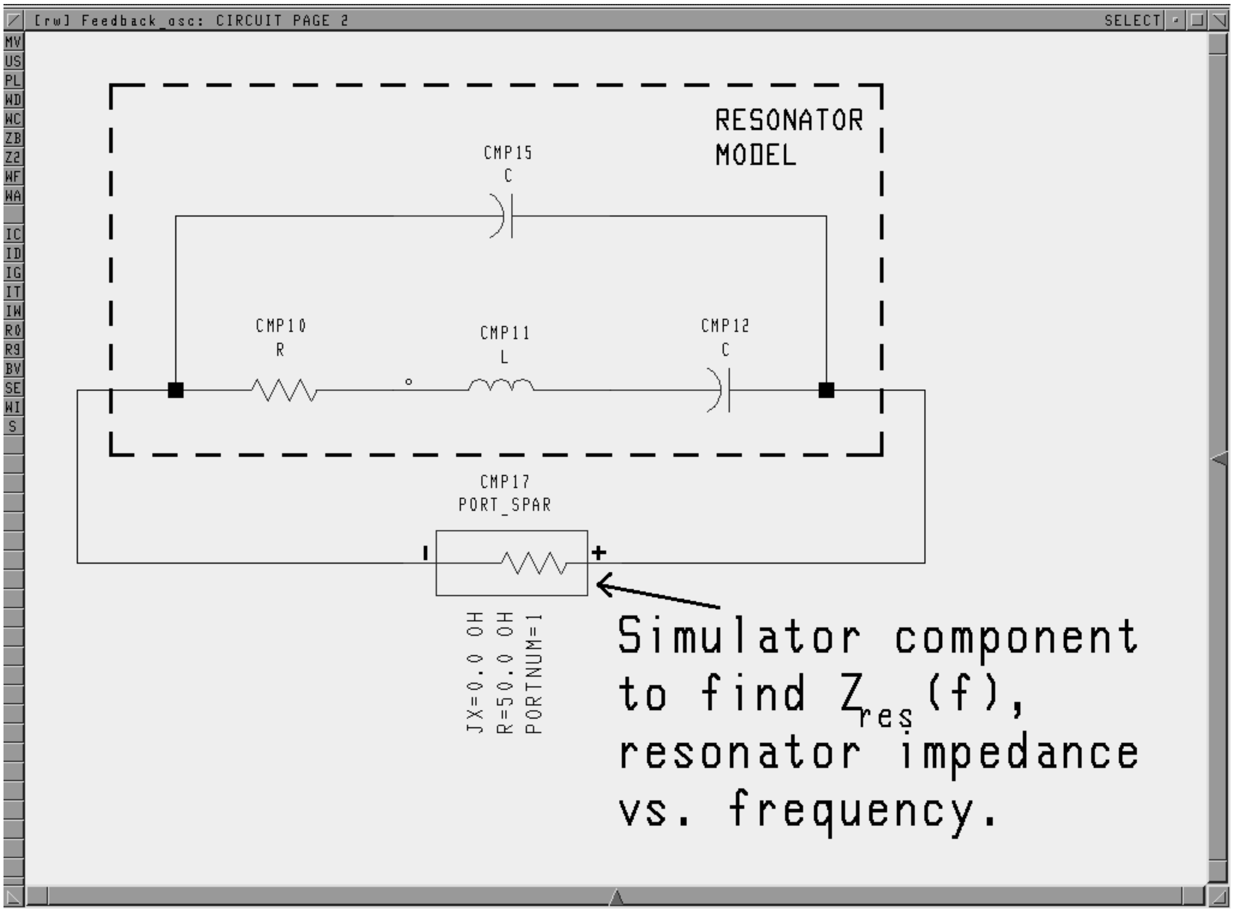

Figure 8b: Simulating resonator impedance as a function of frequency.

At the frequency and current amplitude where the sum of the two impedances is 0, the steady-state oscillation conditions are satisfied:

$$Z_{\text{res}}(\text{f})+Z_{\text{active circuit}}(\text{amplitude})=0$$

If this occurs at the center frequency of the resonator, as shown in Figure 9, and the impedance trajectories intersect at a right angle, then the oscillator's phase noise should be optimum (for the particular device and resonator being used). A right angle intersection is optimum because in this condition, AM noise that might alter the large-signal impedance seen by the resonator (corresponding to a movement along the large-signal impedance trajectory) does not lead to a shift in oscillation frequency.

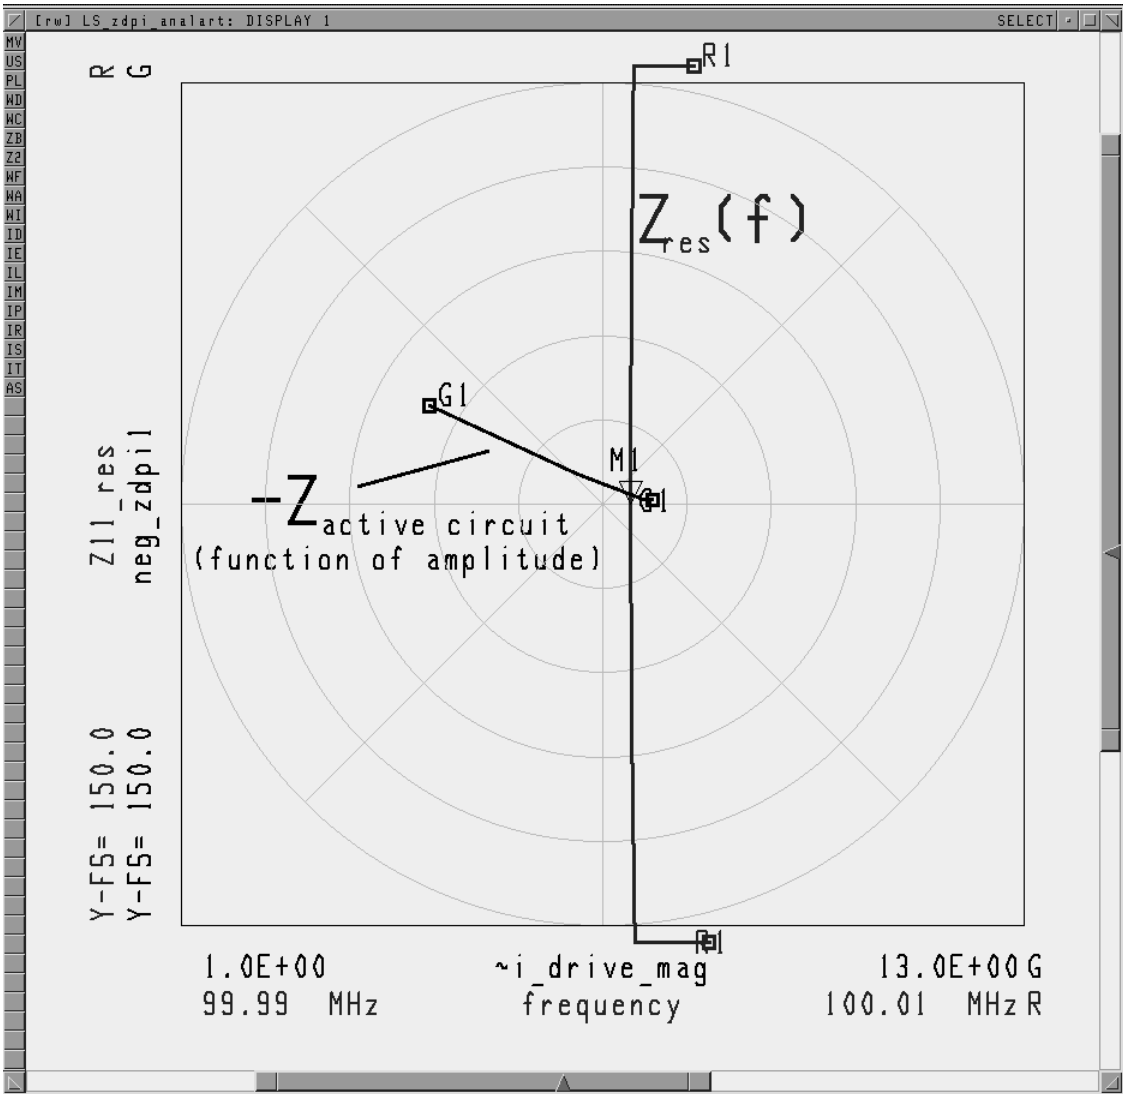

Figure 9: Plot of resonator impedance versus frequency, and negative of active circuit impedance versus drive level.

If the two impedance trajectories do not intersect at the resonator's resonant frequency, then circuit element values can be adjusted to improve the design.

For the reflection oscillator, for optimum phase noise the resonator should be on resonance at the frequency where the oscillation conditions are satisfied. This can usually be achieved by carefully choosing the length of transmission line between the resonator and active circuit. For broadband, tunable oscillators, this condition of optimum phase noise can only be achieved at a single frequency, and the further away from this frequency the oscillator is tuned, the worse the phase noise gets.

Device low-frequency noise modeling

Perhaps the most critical thing to model when simulating oscillator phase noise is the 1/f noise of the active device(s). Unfortunately, device manufacturers do not supply 1/f noise parameters or models for their devices. The only way for a designer to model 1/f noise is to measure it, and extract the coefficients for the noise models that are part of the device models included with HP MDS.

Bipolar junction transistors

Most low-frequency or RF oscillators use BJTs as the active device. A commonly used BJT noise model is shown in Figure 10 (reference [6]). This noise model is implemented in HP MDS.

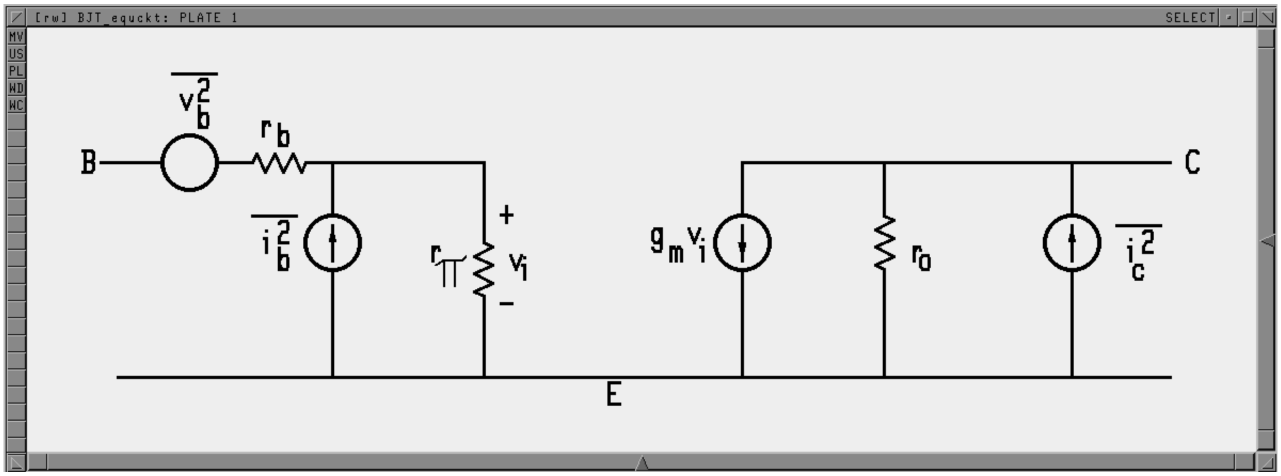

Figure 10: BJT noise model.

In Figure 10, $\overline{v_b^2} = 4kTr_b\Delta f$ is the base resistance thermal noise,

$\overline{i_c^2} = 2qI_c\Delta f$ is the collector current shot noise, and

$\overline{i_b^2} = 2qI_B\Delta f + K_f\frac{I_B^{A_f}}{f}\Delta f + K_B\frac{I_B^{A_B}}{1+\left(\frac{f}{f_B}\right)^2}\Delta f$ is the base current shot, flicker, and burst noise.

In most cases the base current shot and 1/f (flicker) noise sources dominate the noise at the output. The shot noise depends only on the DC currents flowing in the device, and is computed automatically. Thermal noise due to various device resistances (and all lossy circuit elements) is also computed automatically. To include flicker and burst noise in a simulation, it is necessary for the designer to input values for the flicker and burst noise coefficients:

K_f, A_f, K_B, A_B, and f_B.

These coefficients can be derived by making measurements using the block diagram shown in Figure 11a.

Figure 11a shows a block diagram for making 1/f noise measurements on a BJT device, manually, using a spectrum analyzer. This measurement can be automated by using HP IC-CAP software. IC-CAP's instrument driver for noise measurements (HP 85199G Noise Measurement Drivers),supports the HP 35670A Dynamic Signal Analyzer. The instrument driver allows the user to specify the start frequency, the stop frequency, the number of averages, the noise measurement units (such as $(\text{volts})^2/\text{Hz}$, or $(\text{volts})/\sqrt{\text{Hz}}$, etc.), and the number of measurement points.

Extracting the values of K_f and A_f involves measuring the output noise spectrum of the device at two different base current bias levels. From the two different base currents and noise corner frequencies (frequency at which the 1/f noise and the broadband white noise are equal), K_f and A_f can be computed using the equations on page 14.

If either the HP 4141 DC Source/Monitor or HP 4142 Modular DC Source/Monitor product is used instead of batteries or other power supplies, then IC-CAP's HP 85199D DC Measurement Drivers may be used to enable IC-CAP to set and monitor the supply voltages and currents.

IC-CAP does not have the equations built-in to extract A_f and K_f directly from measurements. However, IC-CAP does allow the user to enter equations (via transforms and functions) for parameter extraction from measured data. Therefore, if the user writes an equation to calculate the noise corner frequency, then A_f and K_f can be extracted.

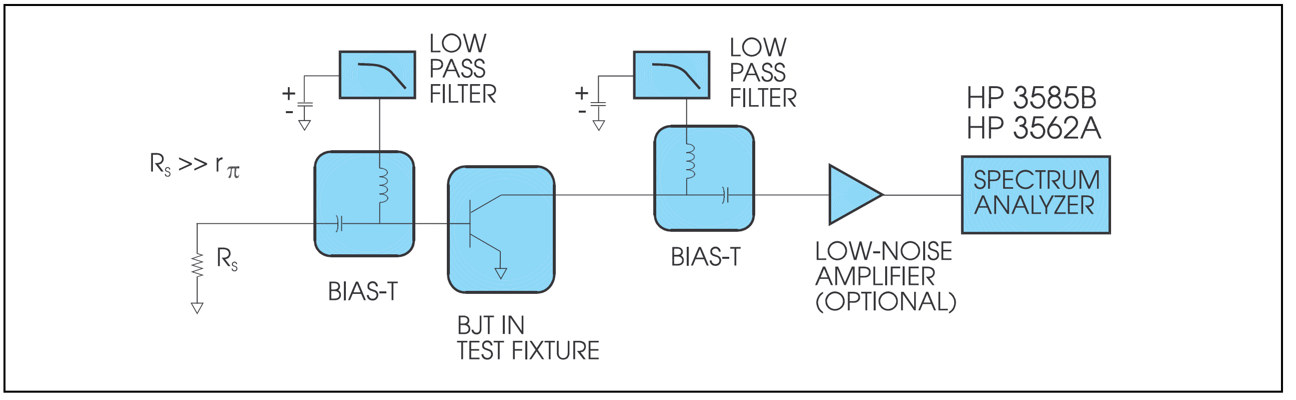

Figure 11a: Block diagram for measuring low-frequency device noise.

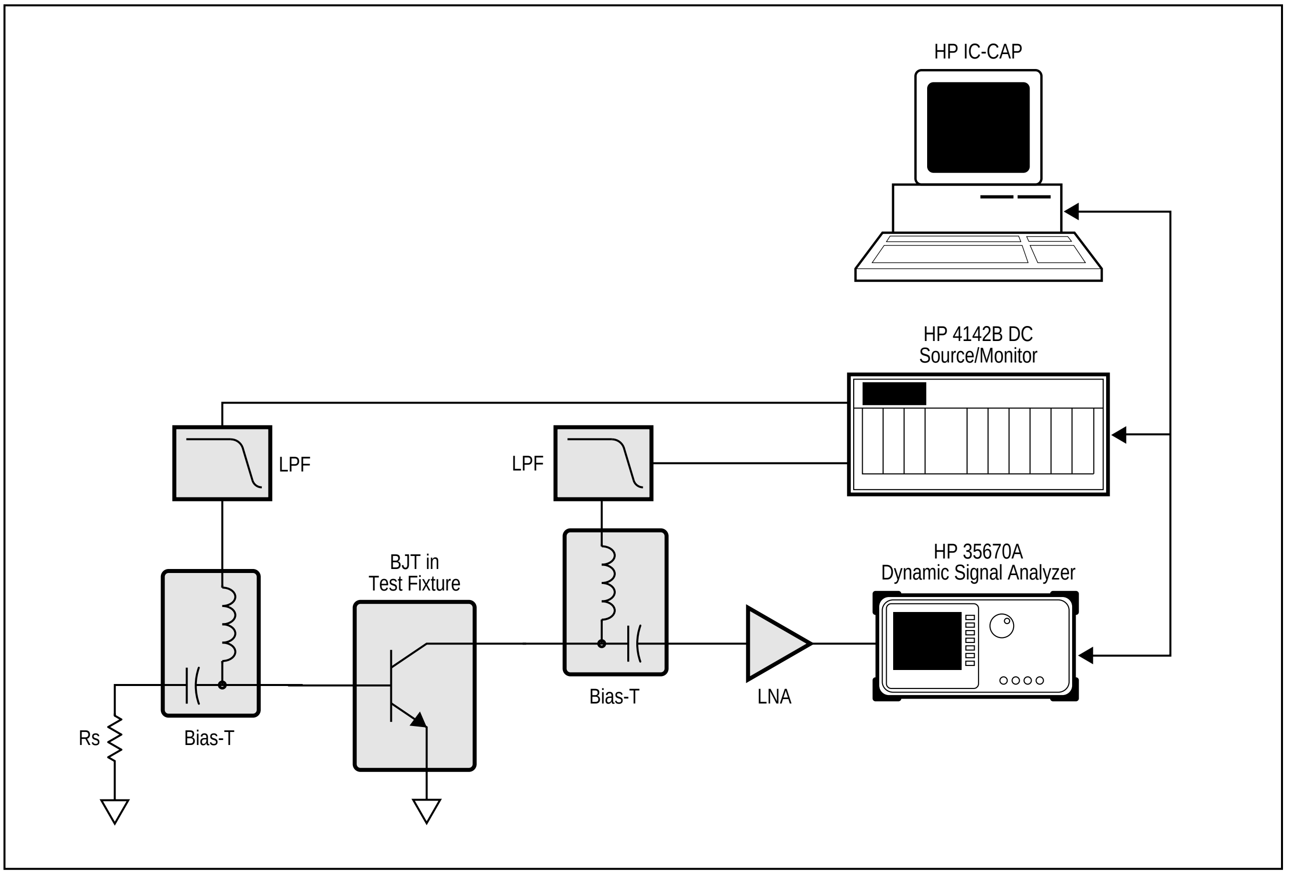

Figure 11b: 1/f noise measurement system.

IC-CAP now allows noise simulations, using HP MDS. These simulations are useful for extracting device model parameters via optimization, and for verifying that a model agrees with measured data.

Figure 11b shows a 1/f noise measurement system, utilizing an HP 35670A Dynamic Signal Analyzer and an HP 4142 Modular DC Source/Monitor.

The objective is to bias the device without adding any noise, and to provide an input impedance large enough to force the base current noise to go through r_π, which causes it to appear at the output, where it is measured by the spectrum analyzer. The noise at the output might typically be measured from 10 Hz to 1 MHz or until the device's noise floor is reached. Because of the presence of large spurious signals in this frequency range (especially 60 Hz and its harmonics), shielding the test system may be necessary. A battery can also be used to avoid 60 Hz noise that might appear on the signal from AC power supplies.

The noise at the output will be dominated by the base current shot and flicker noise, assuming the burst noise is neglected [15]. The flicker noise coefficients can be derived as follows. The noise spectrum of the BJT, will have a corner frequency, f_c, at which the base shot and flicker noise currents are equal:

$$2qI_B = K_f \frac{I_B^{A_f}}{f_c} \quad \text{or} \quad f_c = \frac{K_f}{2_q} I_B^{(A_f - 1)}$$

By taking the log of the equation for f_c, a linear equation can be derived relating the base current and the corner frequency:

$$\log(f_c) = \log\left(\frac{K_f}{2_q}\right) + (A_f - 1)\log(I_B)$$

Measuring the corner frequency at two different base currents
(since the corner frequency will vary with base current) gives the data to solve the linear equation for the two unknowns, K_f and A_f. This derivation is from [15].

JFETs and MOSFETs

Junction field effect transistors and MOSFETs are modeled by a shot current noise generator between the gate and source, and a thermal and 1/f noise current generator between the drain and source, as shown in Figure 12 (reference [6]).

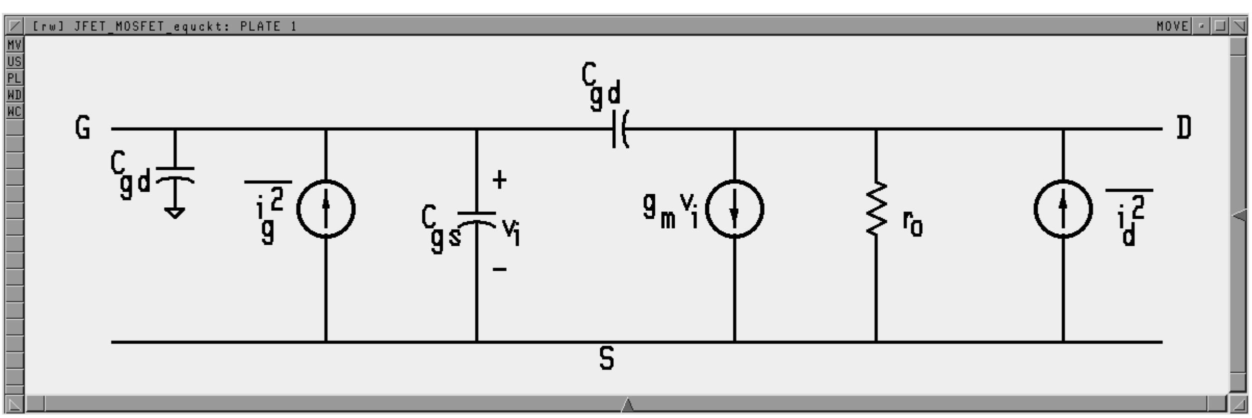

Figure 12: Small-signal equivalent circuit of JFET and MOSFET, with noise sources.

$\overline{i_g^2} = 2qI_G\Delta f$ is the gate current shot noise, and

$\overline{i_d^2} = 4kT\left(\frac{2}{3}g_m\right)\Delta f + K\frac{I_B^{A_f}}{f}\Delta f$ is the drain thermal and flicker noise.

I_D is the drain bias current,
I_G is the gate leakage current,
K_f is the flicker noise coefficient,
A_f is the flicker noise exponent (between 0.5 and 2),
g_m is the device transconductance at the operating point.

The shot and thermal noise are modeled automatically, but modeling flicker noise requires inputting flicker noise coefficients, K_f and A_f, similar to those for a BJT. These coefficients can be derived using a measurement system and procedure similar to that described above for BJTs. These noise sources are described in [6].

MESFETs

Two types of MESFET noise are modeled in HP MDS. The broadband noise is important in nonlinear circuits such as mixers, and has an effect on oscillator phase noise at high offset frequencies from the signal. It is modeled using an equivalent circuit that is described in [7]. The equivalent circuit includes an intrinsic noiseless transistor, gate and drain noise current sources, and thermal noise sources for the gate and source resistances, as shown in Figure 13.

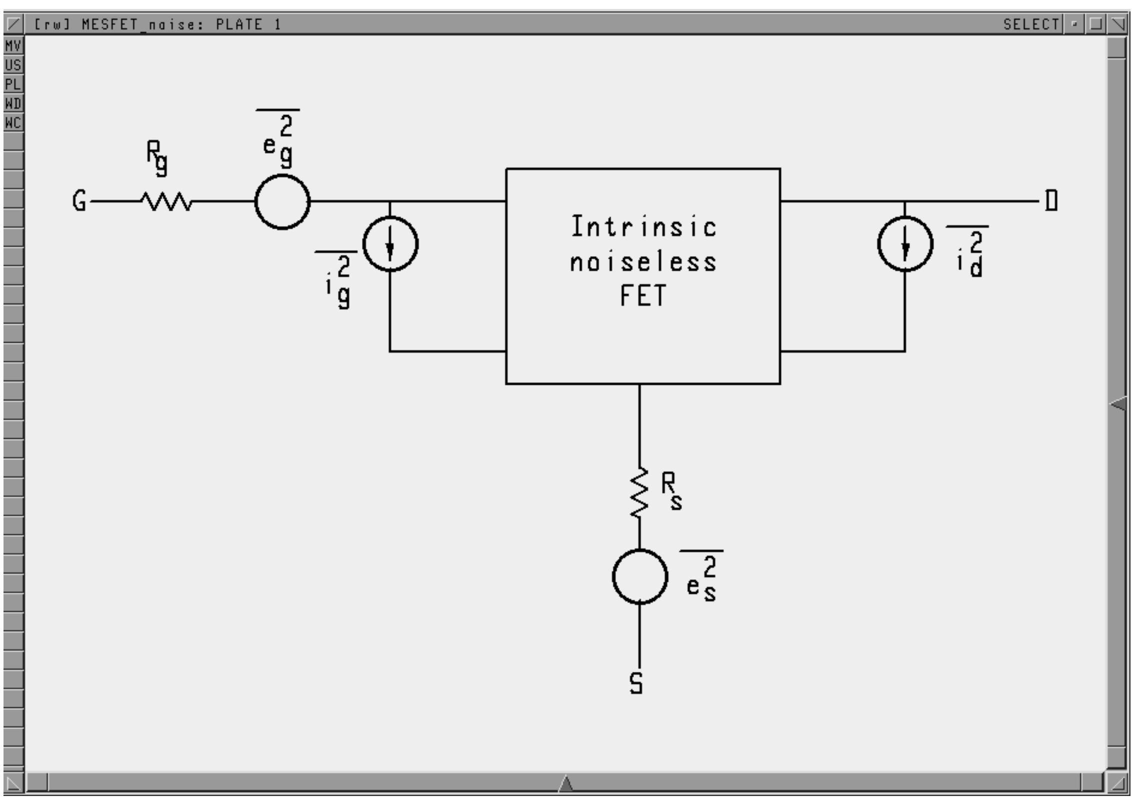

Figure 13: MESFET noise equivalent circuit model.

$\overline{i_d^2} = 4kTg_m P\Delta f$ is the drain noise current, excluding correlation effects, and

$\overline{i_g^2} = 4kT\Delta f C_{gs}^2 \omega^2 R/g_m$ is the gate noise current, excluding correlation effects.

Three parameters that are bias and process dependent are used to describe the gate and drain current noise sources, and the correlation between them:

R is the gate noise coefficient (dimensionless, close to 1-3, depending on the device technology and bias conditions).

P is the drain noise coefficient (dimensionless, and depends on bias conditions and device parameters).

C is the correlation coefficient between the gate and drain current noise, defined as: $C = \dfrac{\overline{i_g i_d^*}}{\sqrt{\overline{i_g^2} * \overline{i_d^2}}}$

These parameters are described in more detail in [7]. The above noise coefficients can be obtained using the procedure described in [8]. The procedure involves measuring the admittance matrix for a FET and the four spot noise parameters, F_{min}, G_0, B_0, and R_N, which are related:

$F = F_{min} + \frac{R_N}{G_S} \left| Y_S - Y_0 \right|^2$, where F_{min} is the minimum noise figure, R_N is the equivalent noise resistance, $Y_0 = G_0 + jB_0$ is the optimum source admittance, and $Y_S = G_S + jB_S$ is the actual source admittance.

Using a de-embedding procedure described in [8], it is possible to obtain the R, P, and C parameters. However, because these parameters vary with bias, it is necessary to make measurements and do an extraction for each set of bias conditions. These noise parameters affect the phase noise floor, at high offsets from the signal.

Close to the signal, the phase noise is affected by 1/f noise, which HP MDS models using a parameter, FNC, which is the corner frequency at which the 1/f noise equals the broadband noise. For some MESFET devices under certain bias conditions, this assumption that the noise simply rises at a 1/f slope below the noise corner frequency is inaccurate. Reference [9] shows that a MESFET operated in the ohmic region has a 1/f noise current spectral density. However, when operated in the saturation region, depending on the bias current and the frequency, the slope of the noise current spectral density changes. Assuming that the spectral density has a 1/f characteristic is probably reasonable as a first order approximation only.

MESFET 1/f noise can be measured using a set-up similar to that described above for BJTs, except that setting the input resistance to the device equal to 50 ohms (or a value it will see in actual operation) is better. Note that this set-up is not suitable for extracting the R, C, and P parameters.

MODFETs, etc.

If a designer is using a device for which no model exists, or is modeling a device using the HP Root Model, the 1/f noise can be measured using a set-up similar to that used for BJTs described above, and noise current and voltage generators can be used at the input or output to model the measured output noise spectrum. Figure 14a shows the measured equivalent input noise voltage spectrum for a MODFET device developed by HP. Figure 14b shows the equations that were used to generate the equivalent input noise voltage source. These equations were generated via curve-fitting.

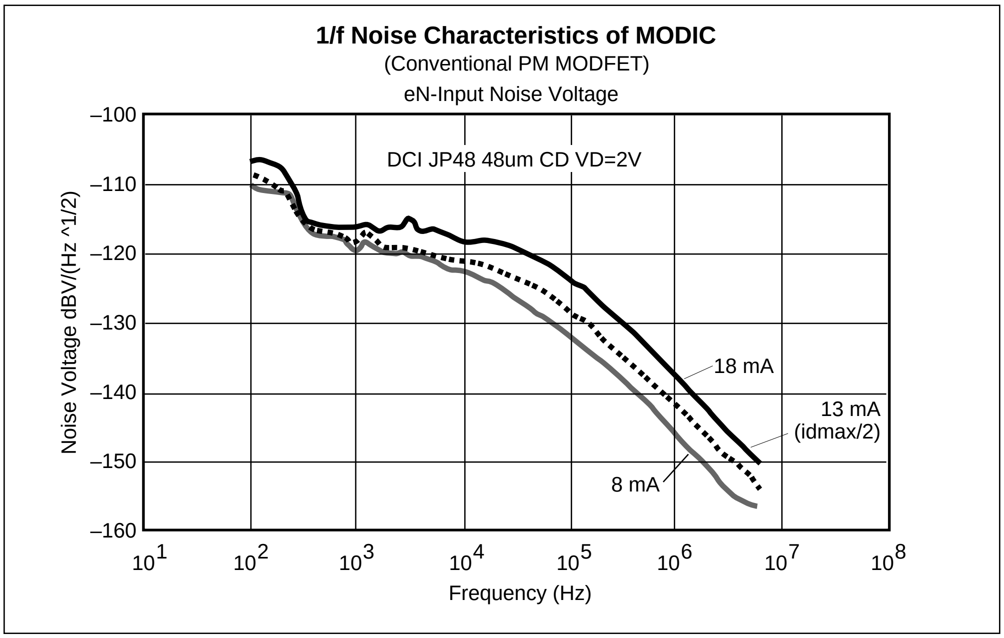

Figure 14a: Measured low-frequency noise spectra of a MODFET.

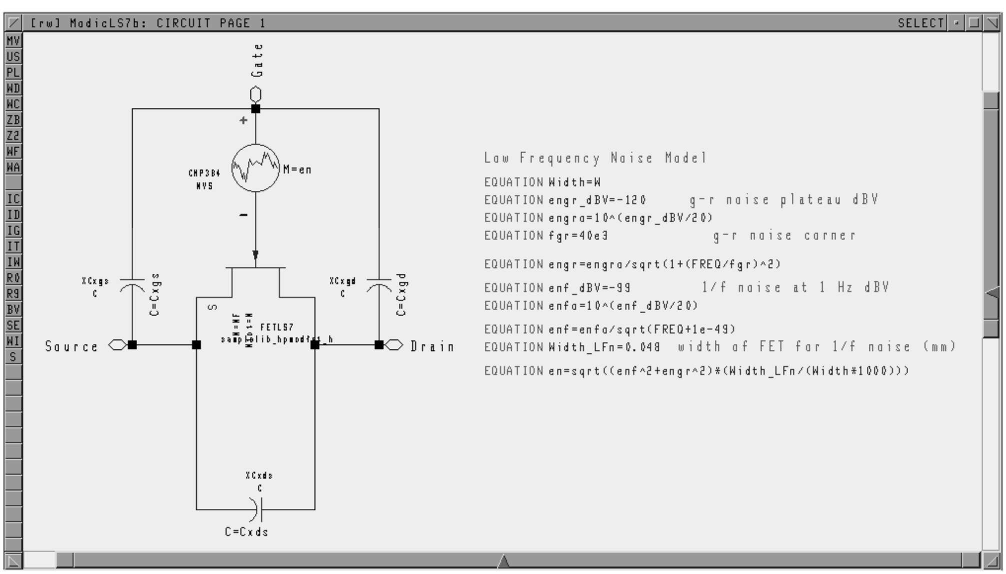

Figure 14b: Equations that define the equivalent input noise source voltage.

Modeling Resonator Phase Noise

Measurement techniques

Some resonators (SAW and crystal) add phase noise to signals that pass through them. The added phase noise can be measured using residual phase noise measurement techniques described in [2] and [3]. Good correlation between the phase noise of measured SAW resonators and the oscillators into which they are assembled is described in [10].

The measurement set-up described in [10] is shown in Figure 15. A signal source (typically a HP 8662A or a different source, depending on the frequency) is split, with one signal passing through the resonator under test, and the other signal passing through a variable delay line. The output from the variable delay line becomes the L input to the HP 3048A phase noise measurement system, and the output from the resonator under test becomes the R input. The variable delay line is adjusted until the signals at the input to the HP 3048A are 90 degrees out of phase. In this condition, the HP 3048A will detect phase differences

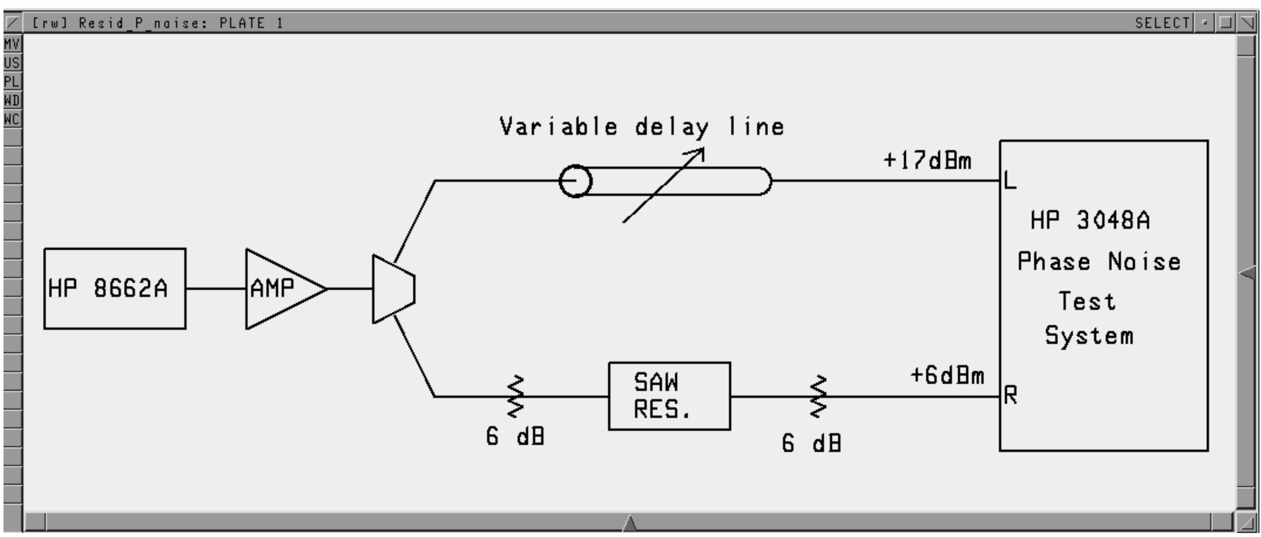

Figure 15: Set-up for measuring resonator phase noise.

(phase noise) between the two signals as a function of offset frequency from the carrier. Crystal resonators can be measured using a similar technique, as described in [2]. These measurements require care to prevent mechanically induced phase fluctuations, and the components may require RFI (radio frequency interference) shielding.

Modeling the noise

Using the symbolically-defined device (SDD, reference [11]) in HP MDS, it is possible to construct a component that adds phase noise to a resonator. There are at least two ways of modeling resonator phase noise in an oscillator. One is to add a component that adds phase shift to a signal, with the amount of phase shift dependent on a control voltage, with a 1 rad./volt transfer function. The SDD to implement this is shown in Figure 16a.

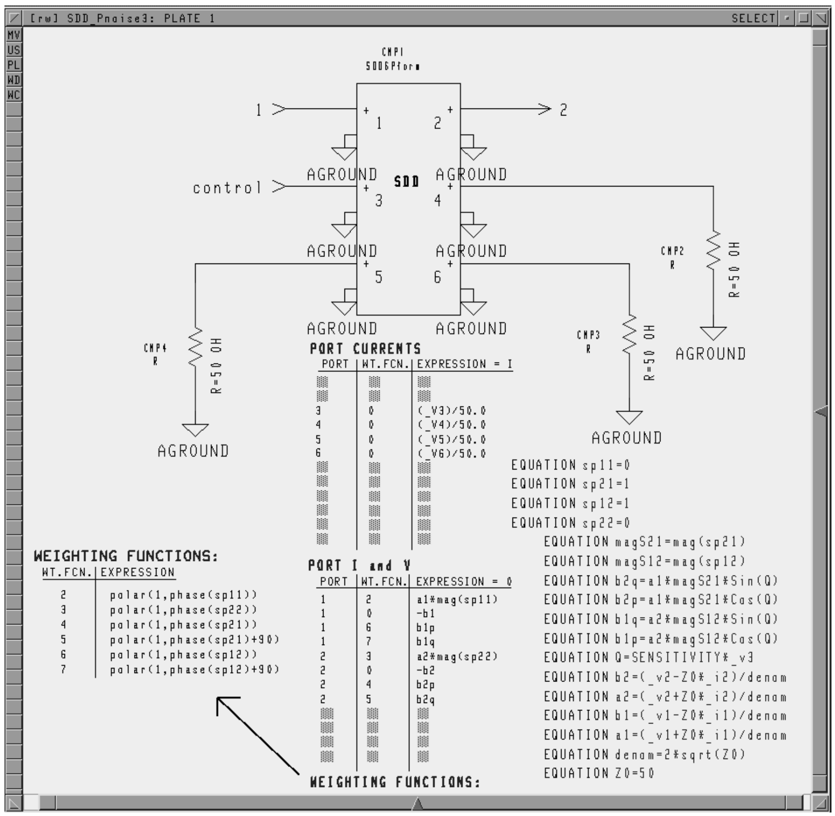

Figure 16a: SDD to implement resonator phase noise.

Figure 16b shows the SDD's symbol, and how it would appear on a circuit page, with the necessary equations and noise voltage source. The SDD would be inserted next to the resonator. Pnoise_at_fm is the residual phase noise of the resonator (in dBc/Hz), at offset frequency fm.

While this SDD looks complicated, all it does is shift the phase of the signal at port 2 relative to the signal at port 1 by Q radians, where

Q = SENSITIVITY (1 radian/volt) * voltage at port 3

Because the voltage at port 3 is

rms noise with a $K/\sqrt{f_{offset}}$ spectrum, the peak phase deviation at offset

frequency, f_{offset}, will be $K/\sqrt{f_{offset}}$.

This adds phase noise to the signal, given by:

$$\text{Phase noise} = 20 * \log (0.5 * K/\sqrt{f_{offset}})$$

The value for K is computed from the residual phase noise of the resonator (Pnoise_at_fm in Figure 16b) at a particular offset frequency (fm in Figure 16b):

K=2*sqrt(fm)*10^(Pnoise_at_fm/20)

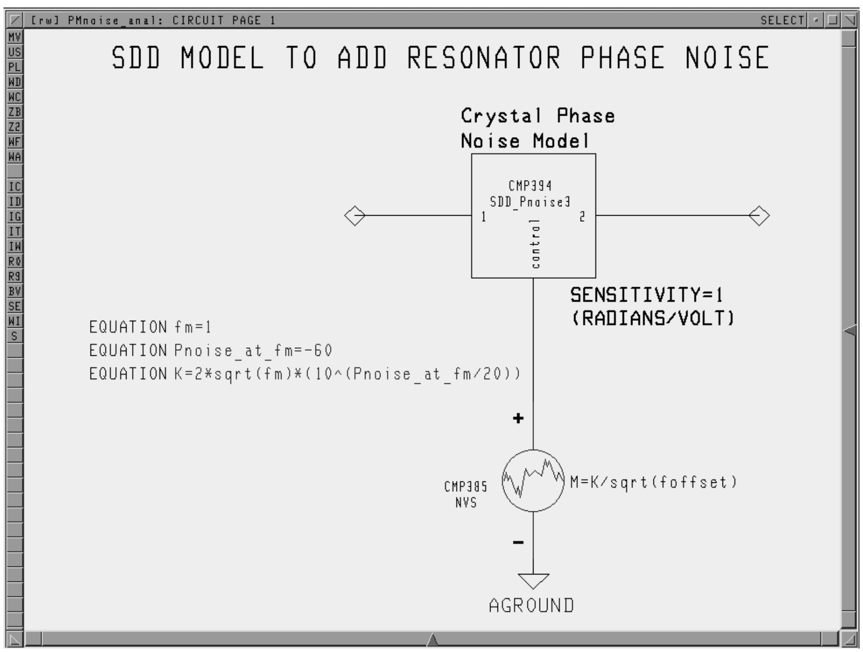

Figure 16b: SDD symbol, equations, and noise source.

The above SDD implements 1/f phase noise. Another approach is to use an SDD to implement1/f frequency noise. If a resonator is modeled with a series RLC circuit, the resonant frequency can be shifted slightly by adding, for example, a small capacitor in parallel with the resonator's series resonant capacitor, as shown in Figure 17. If the frequency of oscillation is the same as the resonator's resonant frequency, then changes in ΔC also change the frequency of oscillation, which models frequency noise.

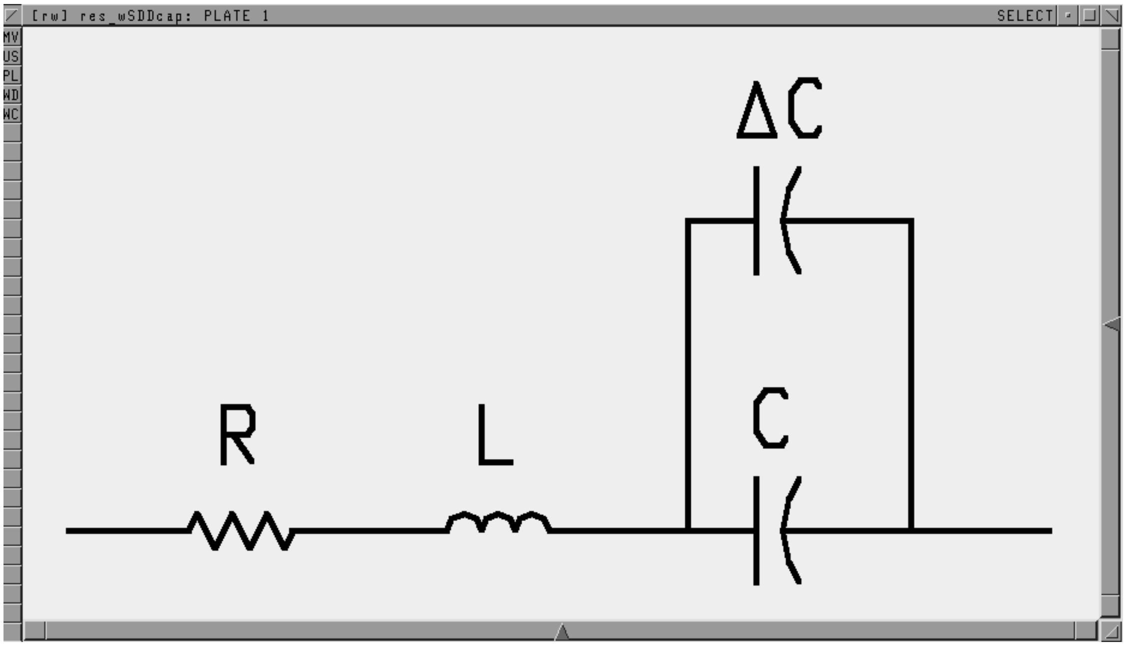

Figure 17: Resonator equivalent circuit, with differential capacitor, ΔC, to model 1/f frequency noise.

If the phase noise is modeled via a signal with sinusoidal frequency modulation, then the instantaneous frequency is:

$f = f_0 + \Delta f_{peak} \cos(2\pi f_m t)$ and the phase is:

$\Phi = \int 2\pi f \, dt = 2\pi f_0 t + \frac{\Delta f_{peak}}{f_m} \sin(2\pi f_m t)$. The peak phase deviation is:

$\Delta\Phi_{peak} = \frac{\Delta f_{peak}}{f_m}$. For small phase deviations, phase noise $= \frac{1}{4}\Delta\Phi^2_{peak}$

(reference [1]).

Δf_{peak} is found by computing the derivative of the oscillation frequency w.r.t. the total series resonant capacitance, and then

multiplying by ΔC: $\Delta f_{peak} = \frac{df_0}{dC_{tot}} \Delta C$. For the resonator in Figure 17,

$$\Delta f_{peak} = -\frac{1}{4\pi} \frac{1}{C\sqrt{LC}} \Delta C.$$

If the small capacitor is made to be voltage-dependent, and the controlling voltage is a noise source with a $1/\sqrt{f_{offset}}$ rms noise voltage spectrum ($f_{offset} = f_m$), as shown in Figure 18, then 1/f frequency noise can be modeled.

The value of ΔC is $K/\sqrt{f_{offset}}$. K can be found by computing the peak phase deviation (from Δf_{peak} above), and the equation relating the peak phase deviation and oscillator phase noise:

$$\Delta\Phi_{peak} = \frac{\Delta f_{peak}}{f_{offset}} = -\frac{1}{4\pi}\,\frac{1}{C\sqrt{LC}}\,\frac{K}{\sqrt{f_{offset}}}\,\frac{1}{f_{offset}}.$$

Since Phase noise (dB) = 20*log($\Delta\Phi_{peak}$)-6, K is computed:

$$K = 10^{\frac{\text{phase noise}+6}{20}}\left(-4\pi C\sqrt{LC}\,f_{offset}^{3/2}\right).$$

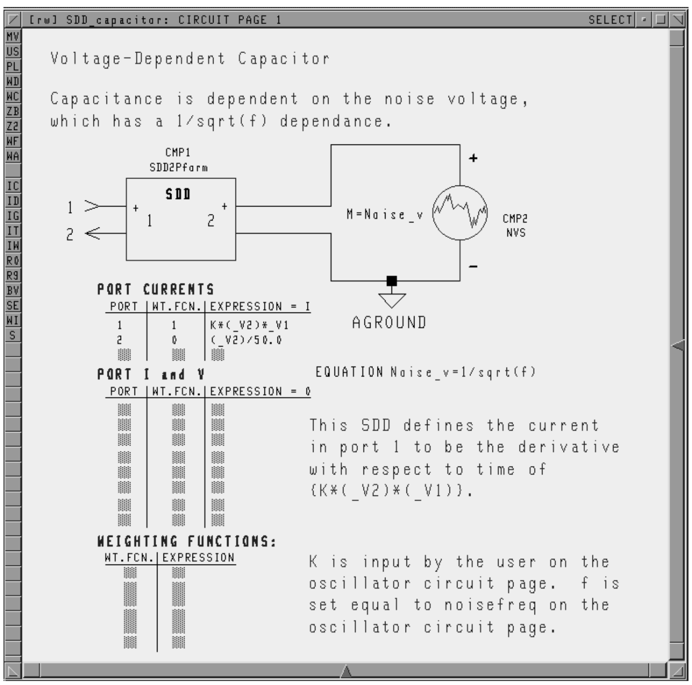

Figure 18: SDD to implement voltage-dependent capacitor.

In this case, the "phase noise" term in the exponent is the phase noise of an oscillator that has the resonator. f_{offset} is the offset frequency from the oscillator's fundamental.

In some cases, crystal suppliers include in the data sheet a crystal's phase noise, which is actually the phase noise of an oscillator with the crystal. This phase noise and its corresponding offset frequency, f_{offset}, along with the series inductance and capacitance for the resonator are plugged into the above equation to determine K. The advantage of this method of simulating resonator phase noise is that it is not necessary to make a residual phase noise measurement. However, it is necessary to have measured phase noise on an oscillator with the resonator, and it does depend on the assumption that at low offset frequencies the resonator's frequency noise is the dominant contributor to the oscillator's phase noise.

Oscillator Phase Noise Simulation Example

Figure 19 shows the schematic of a 100 MHz crystal oscillator whose phase noise will be analyzed.

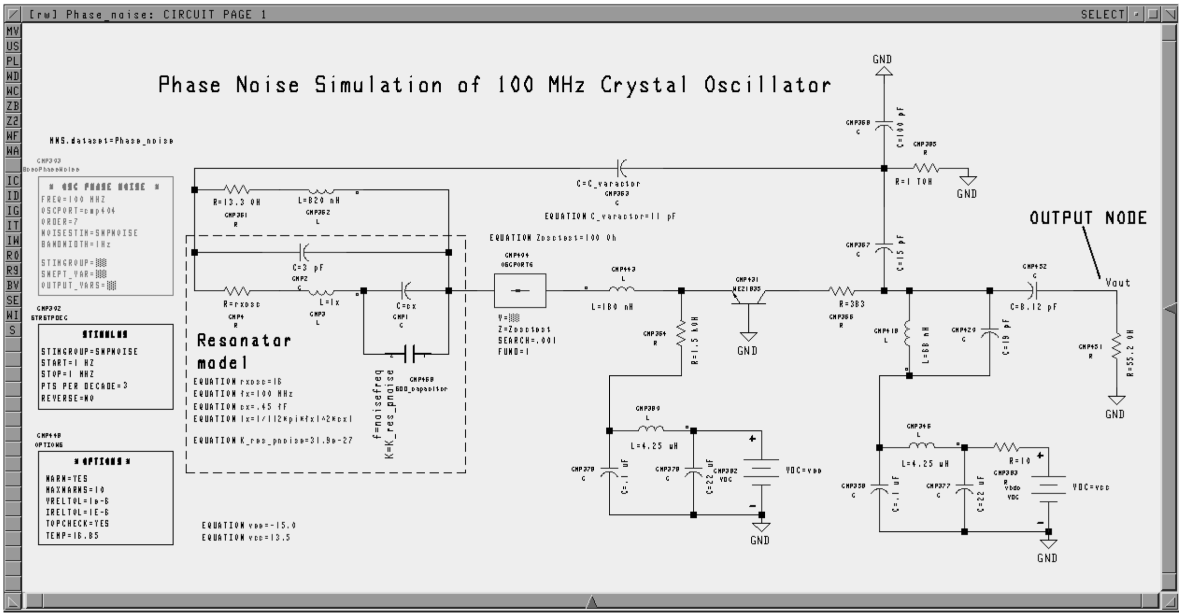

Figure 19: 100 MHz crystal oscillator.

The circuit consists of a common-base BJT as the amplifier, a crystal resonator in the feedback path to choose the frequency of oscillation, a tunable varactor diode in series with the crystal, to provide frequency tuning, and a resonant circuit at the collector to turn the collector voltage waveform into a sinusoid. After choosing a suitable oscillator configuration (a step omitted here), the circuit can be analyzed to see whether the oscillation frequency is the same as the resonator's center frequency or not, using the procedure described earlier, in the "Circuit analysis for optimum phase noise" section.

Once the oscillator's oscillation conditions have been analyzed, the OSCPORT component can be used to do a harmonic balance analysis, which will find the frequency of oscillation and harmonic levels. The OSCPORT component is just to the right of the resonator in Figure 19. A phase noise analysis can be done, automatically, instead of a separate harmonic balance analysis. The phase noise analysis will be described shortly. The OSCPORT component is required because, prior to the analysis, the actual oscillation frequency is not known exactly because oscillators are not "forced" circuits (i.e., there is no external source driving the circuit) and the signal levels are also unknown.

Usually the designer knows the approximate frequency of oscillation, and it is better if this frequency is entered into the oscillator harmonic balance analysis component, especially if the oscillator has a very high Q resonator. Once the OSCPORT has found a potential oscillation frequency (via the OSCTEST analysis described below), it does harmonic balance analyses, gradually increasing the amplitude of the fundamental signal until the loop gain is 1 and the amplitudes and phases of the harmonics are equal on both sides of the OSCPORT component.

There are two ways to find the potential oscillation frequency. One was described earlier, where the resonator and active circuit were separated and their impedance trajectories were plotted. The other is to use the OSCTEST component, which performs a small-signal loop gain analysis on the oscillator. If the impedance trajectory analysis described earlier is done, then it should not be necessary to use the OSCTEST component to find the frequency of oscillation. The OSCTEST component does the initial analysis that is performed by the OSCPORT component, and it acts like a signal source, load, and ideal circulator, and is used to determine the small-signal closed loop gain and phase shift, as a function of frequency. The impedance of the source and load are the same, and are set via the Z parameter on the OSCTEST component. For some oscillators, the loop gain and phase shift vary with this Z parameter, so it is sometimes a good idea to check the loop gain and phase shift versus frequency for different values of the Z parameter. An example simulation set-up for doing this is shown in Figure 20. In this case, the impedance of the OSCTEST component is swept from 10 to 100 ohms, and the closed loop gain and phase responses are simulated over a narrow frequency range around 100 MHz, for each value of OSCTEST impedance.

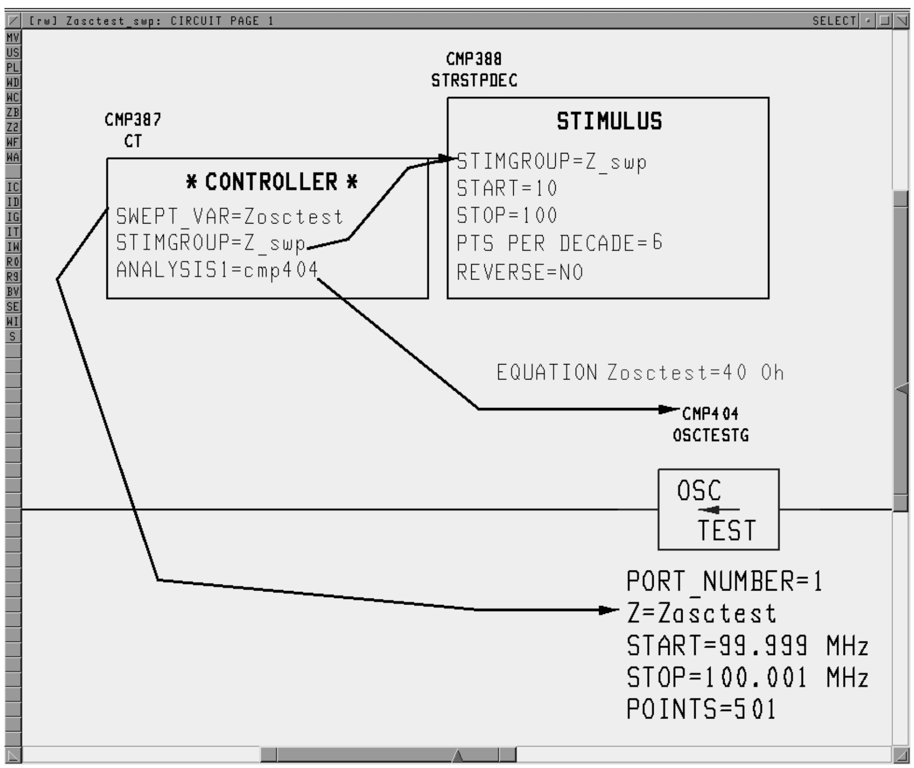

Figure 20: Set-up for simulating closed-loop gain and phase as a function of OSCTEST component's impedance.

After doing the simulations with the OSCTEST component and finding the potential frequency of oscillation, the OSCTEST component can be replaced with an OSCPORT component and oscillator harmonic balance simulations can be run. If the oscillator has a high Q resonator, then the oscillation frequency (found by the OSCTEST simulations) should be specified accurately on the harmonic balance oscillator analysis component. It may also be necessary to narrow the number of octaves over which the simulator searches for a potential oscillation frequency, by reducing the SEARCH parameter on the

OSCPORT component. It may also be necessary to increase the number of steps to be taken per octave during the search. This value can be changed by editing the OSCPORT component (PERFORM/EDIT COMPONENT), and increasing the STEPS_PER_OCTAVE parameter.

When oscillator phase noise is computed, the simulator first does a harmonic balance oscillator analysis to determine the frequency of oscillation and the harmonic levels, as described above. Then it computes the noise voltage at all labeled nodes, due to the phase noise mechanisms described earlier. The noise voltages are computed at offset frequencies from the fundamental, as specified by the designer. The simulator control elements of Figure 19 are repeated here in Figure 21. A harmonic balance analysis is run on the oscillator, and then its phase noise is computed at offset frequencies from 1 Hz to 1 MHz, as specified in the stimulus block. The options block is used to set the temperature to 16.85° C, which is 290° K, the standard temperature for noise analysis.

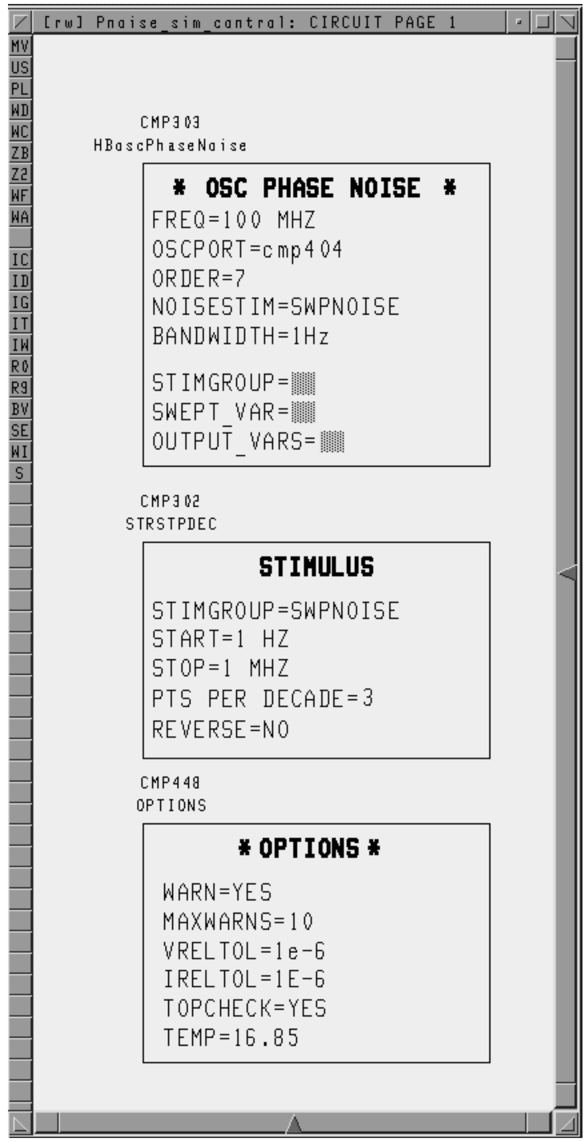

Figure 21: Example simulator control blocks for oscillator phase noise analysis.

An oscillator harmonic balance simulation can be changed to a phase noise simulation just by editing the name of the harmonic balance oscillator control block and changing it from HBosc to HBoscPhaseNoise. A stimulus block (normally start, stop, and points per decade) to define the offset frequencies (at which the phase noise will be computed) from the fundamental, as shown in Figure 21, must also be added.

The simulated phase noise must be computed on the PRESENTATIONS page, as shown in Figure 22. Phase noise is computed by taking the log of the ratio of noise power (at each offset frequency) to the signal power at the fundamental. However, because the simulator outputs the rms noise voltage and the peak fundamental signal voltage, the signal must be converted to rms by dividing by $\sqrt{2}$, before taking the logarithm

Phase_noise = 20*log(Vnoise/(mag(Vout[2])/sqrt(2))),

where mag (Vout[2]) is the voltage amplitude at the fundamental. The computed phase noise is then plotted versus the noise frequency, as shown in Figure 22. The signal amplitude and harmonic levels, as well as the time waveform may also be plotted.

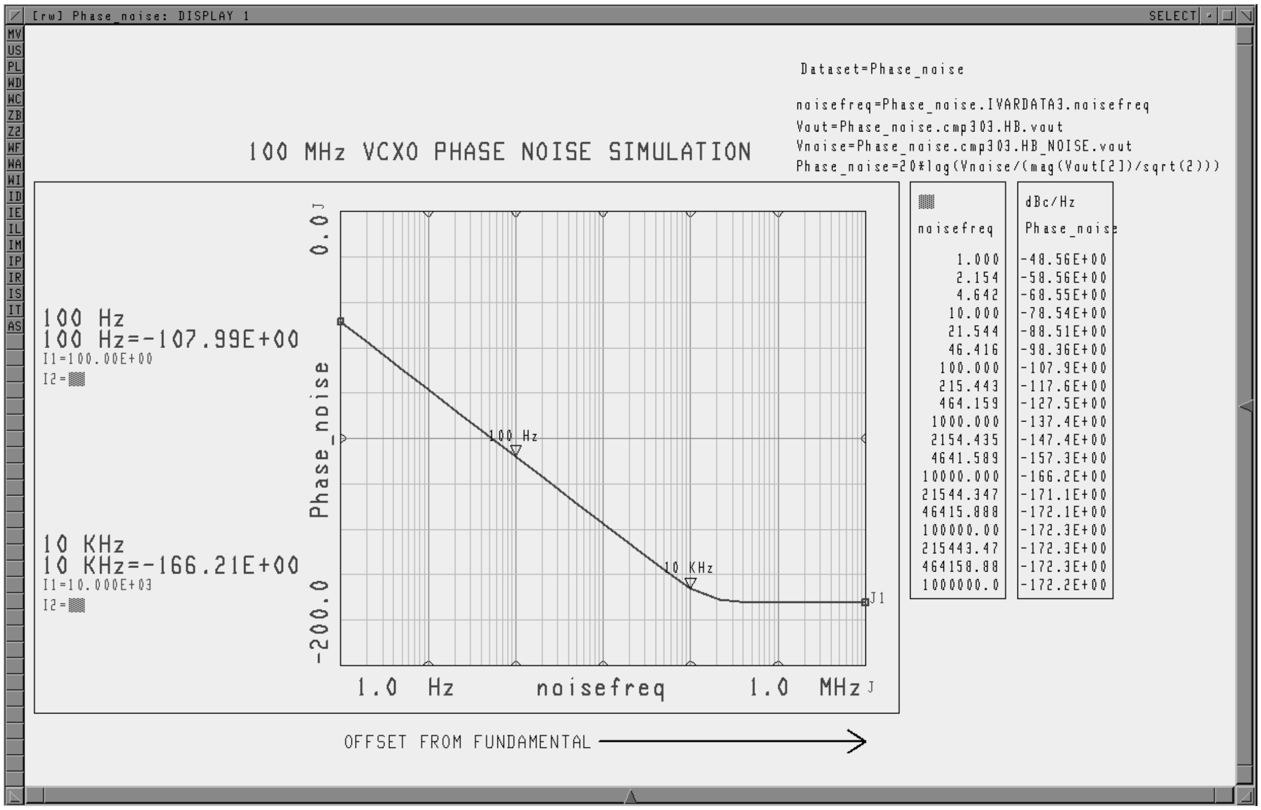

Figure 22: Phase noise versus offset frequency from fundamental.

More complex simulations can be done, for example, phase noise versus varactor tuning capacitance, versus bias voltage, or versus some other parameter.

Mixer Noise Figure Simulation

Why simulate mixer noise figure?

Mixers are often used in receivers to down-convert signals from an RF to baseband. The noise figure of a receiver determines the lowest signal level it can receive, and mixer noise figure is important in determining a receiver's noise figure.

How HP MDS Simulates Mixer Noise Figure

Mixer noise figure is defined as the total noise at the output port per Hz of bandwidth divided by the noise at the output port per Hz of bandwidth due to thermal noise at the input port and at the input frequency specified by the designer:

$$NF = \frac{\text{Total noise at output port /Hz}}{\text{Noise at output port due to noise at input port /Hz @ a particular input frequency}}$$

The noise power contributed at the input port is the maximum thermal noise available in a bandwidth, B Hz, from a resistor at temperature, T degrees Kelvin, or kTB, where k is Boltzmann's constant (1.38×10^{-23} joules/K) . At T = 290 K, kTB is 4×10^{-21}Watts, or -174 dBm in a 1 Hz bandwidth.

Mixers translate signals from one frequency to another, as intended, due to their nonlinear behavior. They also translate noise, via a process known as noise mixing. In HP MDS, noise is assumed to be a small-signal perturbation on the quasi-periodic harmonic balance solution. Noise is mixed (changes frequency) due to circuit nonlinearities and the presence of large-signal harmonic balance sources. Since noise is small-signal, it cannot exercise circuit nonlinearities. This means that even if a circuit has nonlinearities, there will be no frequency translation of noise without the presence of large-signal sources. Figure 23 shows noise mixing, conceptually.

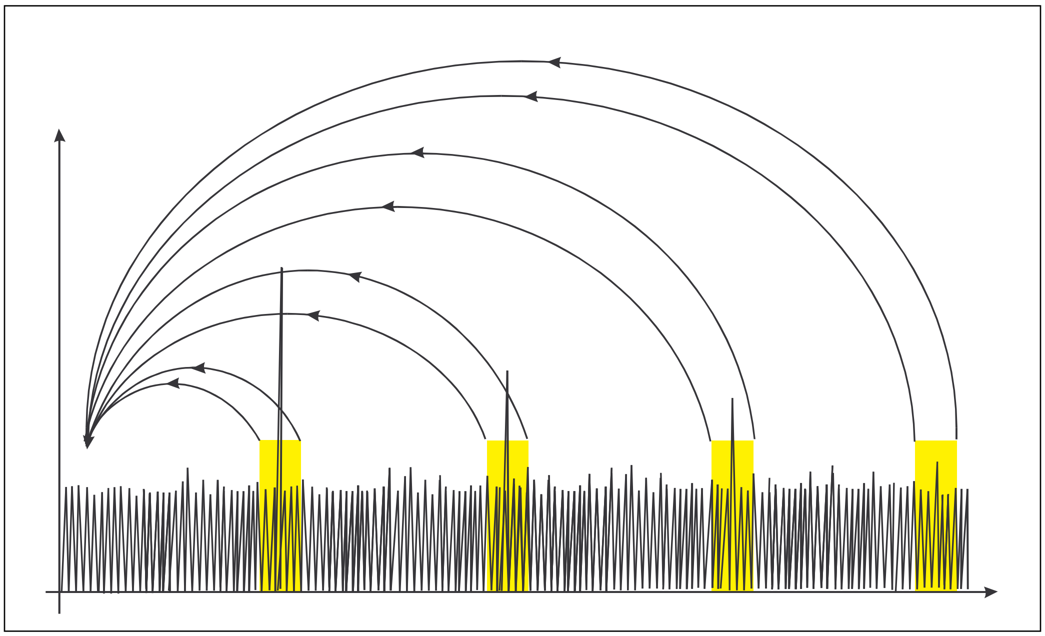

Figure 23: Frequency translation of noise due to noise mixing.

Assume this is the frequency spectrum inside a mixer, down-converting an RF signal to an IF. The large spikes are the LO and its harmonics, inside the mixer. The RF signal is one IF above the LO, i.e. RF=LO+IF. The LO and its harmonics mix noise from the frequencies LO+/-IF, 2*LO+/-IF, 3*LO+/-IF, etc. down to the IF. So HP MDS computes mixer noise figure by calculating the total noise mixed down to the IF, divided by the noise at the IF that was down-converted from the RF signal band, only.

In order to do the calculations described above, HP MDS computes a frequency translation transimpedance matrix (reference [14]). The elements of this matrix determine how anoise current at some frequency (in the example above, the frequencies would be LO +/-IF, etc.) produces noise at some other frequency (usually the IF or desired output frequency). To do a mixer noise figure simulation, the designer must specify a “noisefreq”, which is the output frequency where the noise will be calculated. In the above example, this would be the IF frequency. The designer must also specify an “inputfreq”, which is the frequency of the signal to be down-converted. Usually this is the RF frequency. HP MDS does a single-sideband noise figure calculation, which means that the simulator assumes the mixer converts only one of the LO’s sidebands to the desired frequency.

Figure 24 shows the harmonic balance noise analysis control block required for a mixer noise figure calculation. This is the most simple case, as HP MDS will compute the noise figure for fixed LO and IF frequencies. Note that the order of the LO tone is set to 7, and the order of the RF tone is set to 1. This assumes that the LO signal is much larger than the RF, and that the RF signal is very small.

For some mixer circuits, it may be necessary to increase the order of the LO and the MAXORDER term, but simulation time will also increase. (MAXORDER is the maximum combined order for all combined frequency components. For example, if a circuit has 2 tones, an LO and an RF, and the ORDER of each one is 4, but MAXORDER is 5, then the frequency component 3*LO+2*RF would be included, but 3*LO+3*RF would be neglected.) Some variable could be swept if desired, by setting SWEPT_VAR equal to the variable name, and defining the range using a stimulus block.

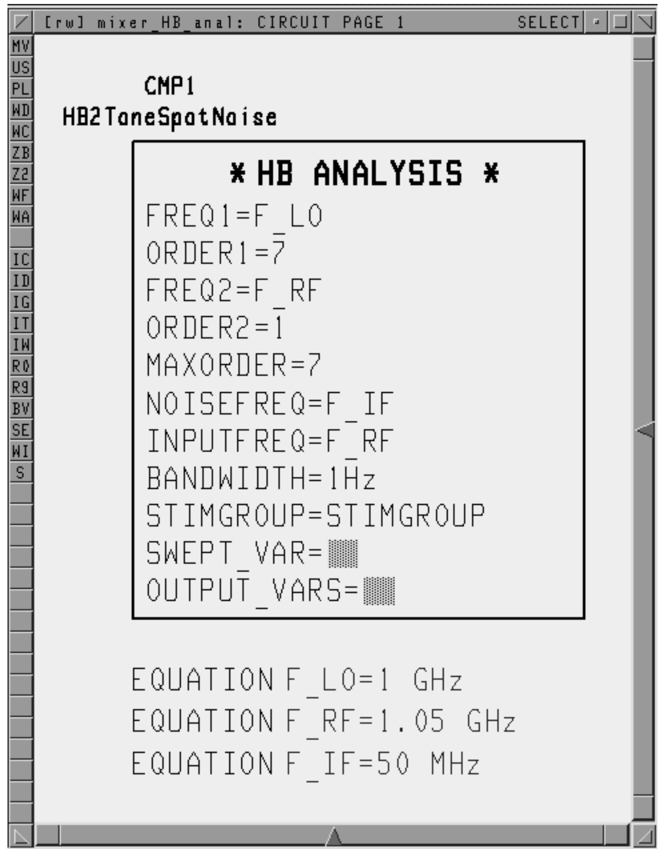

Figure 24: Harmonic balance noise analysis control block for mixer noise figure computation.

If the RF signal is small compared to the LO, then the noise mixed by the LO and its harmonics will dominate, and simulations can be run much more quickly by just including the LO and its harmonics in the harmonic balance noise analysis. A comparison was done on a simple diode quad double-balanced mixer, with the LO power at 10 dBm, and the RF at -10 dBm. Doing a 1-tone harmonic balance simulation (the simulation set-up is shown in Figure 25), including only the LO, HP MDS required 7 seconds to compute the noise figure at 11 frequencies in a 20 MHz bandwidth centered at the IF. Doing a 2-tone harmonic balance simulation (as shown in Figure 26), including the LO and the RF, required 372 seconds for the same noise figure computation. The difference in the computed noise figure was less than 0.1 dB, as shown in Figure 27. The 2-tone simulation was repeated after increasing the RF power to 0 dBm. The simulation time was again 372 seconds, but the noise figure now differs by about 0.6 dB from that computed when including only the LO. Note that for these simulations, only 5 harmonics of the LO were used. Increasing the number of harmonics (via the ORDER parameter in the HB NOISE ANALYSIS block) will improve the accuracy of the computed noise figure, but the simulation time will increase.

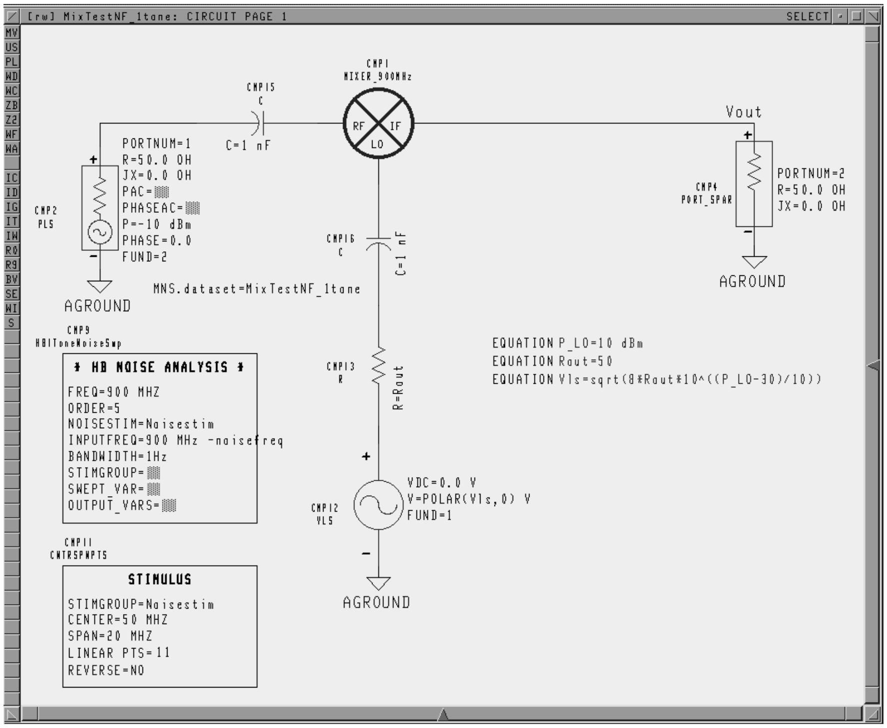

Figure 25: Noise figure simulation with only one harmonic balance tone, the LO.

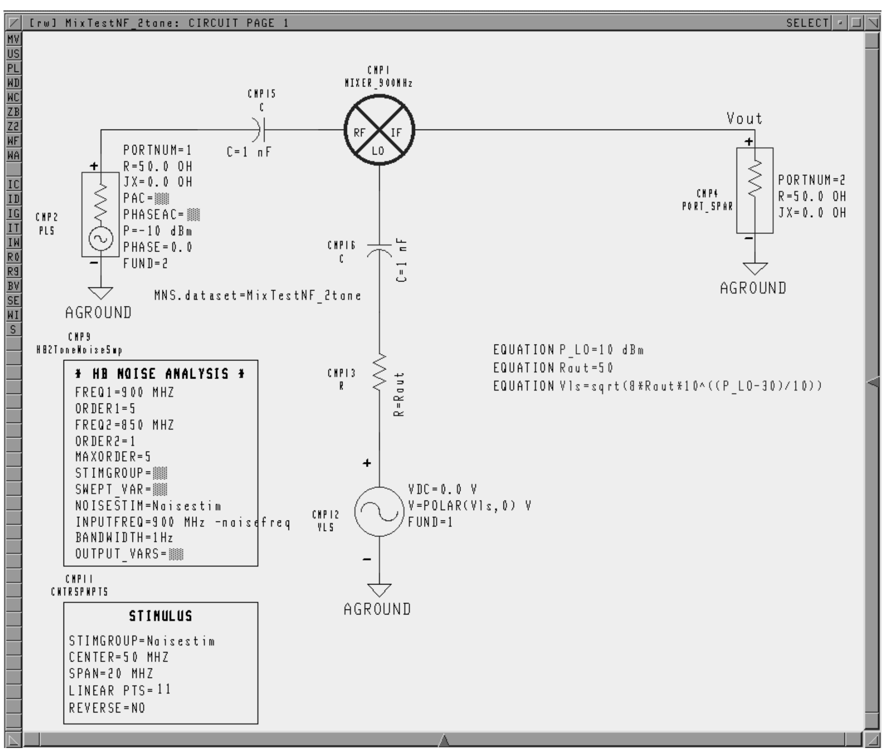

Figure 26: Noise figure simulation with two harmonic balance tones, the LO and RF.

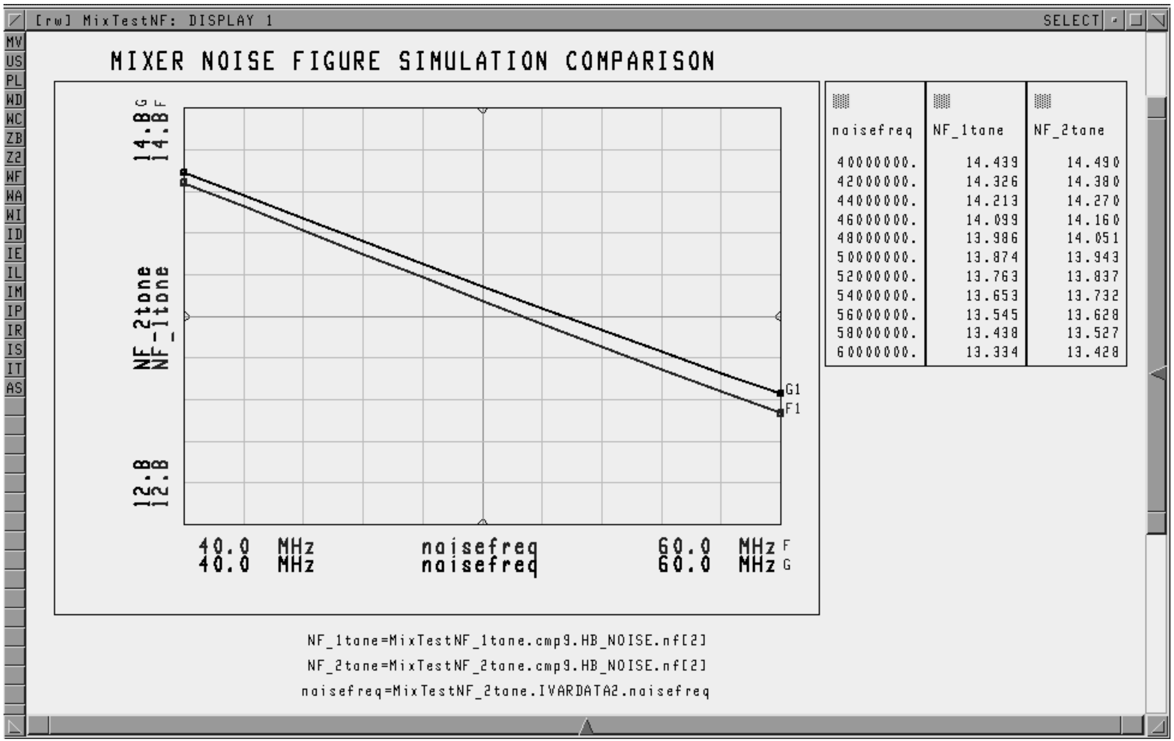

Figure 27: The computed noise figure is virtually the same for 1-tone or 2-tone harmonic balance analysis.

In these simulations, NOISESTIM defines the noise frequencies (the variable, noisefreq) at the output, and the INPUTFREQ is the difference between the LO frequency and noisefreq. Basically, the RF signal frequency and the output or IF frequency are swept together so the noise figure is calculated properly. The above simulation would be suitable for evaluating the noise figure of a mixer with a fixed LO and a variable RF and IF. Since the LO and its harmonics are the only large signals present, and they are not changing, it is valid to just sweep the input frequency and the output frequency at which the noise is calculated.

Modeling requirements for mixer noise figure simulation

For a mixer noise figure simulation, perhaps the most important factor is a good nonlinear model of the device being used, whether it is a diode, a FET, or a BJT. For diodes and BJTs no noise parameters need to be included, as shot and thermal noise are computed automatically by HP MDS. 1/f noise is usually not important in mixer noise analysis, because the RF is offset far enough from the LO that the broadband noise dominates the 1/f noise, after translation to the output port at the IF. For example, if a designer is using diodes that have 1/f noise out to 10 KHz, but the IF is at 5 MHz, the 1/f noise will be far below the shot and thermal noise at the IF. For FETs, it is necessary to include the R, C, and P parameters described earlier in the device low-frequency noise modeling section.

Mixer noise figure can be degraded by LO noise. Usually AM noise is most important and can be included by adding a broadband noise source as shown in Figure 28.

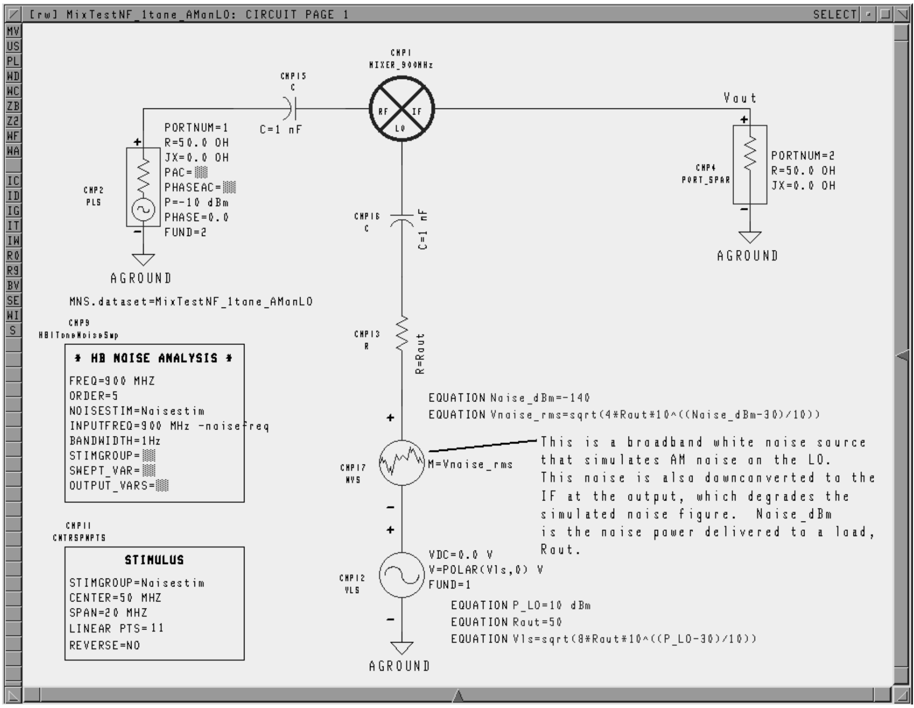

Figure 28: Simulating the effect of broadband noise from the LO on mixer noise figure.

There is a new (HP MDS B.05.01) model parameter, LOTOIFCONV, that has been added to the system model library [16] mixer components. This parameter is the LO-to-IF conversion gain factor, given in dB. The default value is infinity (i.e., no conversion gain). This parameter only influences noise analysis, and models the degradation in mixer noise figure due to the LO mixing its own white noise to the IF port. Note that this is not the same as LO-to-IF isolation, which is the leakage of the LO signal that appears at the IF port (but still at the LO frequency). Unfortunately, LO-to-IF conversion gain is not normally specified by mixer manufacturers, but it can be measured, by terminating the RF input to a mixer, and supplying at the LO port, an LO and a small signal at LO +/- IF. The difference between the signal amplitude at the IF port at the IF frequency and the small signal added at the LO port is the LO-to-IF conversion gain.

Simulating Phase Noise in Systems

In some applications, system designers must be concerned with the effects of oscillator phase noise. One is the down-conversion of RF signals in a receiver. A local oscillator's phase noise gets translated directly to the IF signal. So if two adjacent RF signals, one much larger than the other, are to be down-converted, the phase noise on the larger IF signal can mask the smaller IF signal. This is described in [12].

<u>How HP MDS Simulates Phase Noise in Systems</u>

HP MDS can simulate phase noise in systems, using either data taken from an oscillator data sheet, or data taken directly from an oscillator phase noise simulation. Because of the way HP MDS computes oscillator phase noise, if a designer wants to simulate anything more complex than an oscillator/amplifier combination, it is usually better to simulate the oscillator's phase noise separately, then use the computed phase noise data in a system phase noise simulation, as described below.

Signal sources with phase noise are modeled using a large-signal voltage source and a noise voltage source. Due to the noise mixing process described earlier, the phase noise gets translated down (or up) to appear around the output signals. The symbol that is used on a circuit page is shown in Figure 29. This symbol and circuit, as well as example files are in the demo file, “Phase_noise,” that is included in HP MDS B.05.01 and HP RFDS B.05.01 and later revisions.

The symbol allows the designer to specify the output resistance, Rout, the power delivered to a matched load, Pout, and the temperature of the resistor, Tnoise. The fundamental number, Fund, also must be specified.

On the circuit page are equations that compute the required rms noise voltage amplitude, vnoise, versus offset frequency from the fundamental. There is an equation for vls, the source voltage amplitude required to deliver power, Pout, to a matched load, Rout. The equation for vnoise computes the rms noise voltage required to produce a phase noise, Pnoise (dBc/Hz), given a source voltage amplitude, vls. Pnoise is phase noise data, and in this case it is from an oscillator data sheet. The phase noise data, as a function of offset frequency, was entered into a CITIfile (a UNIX file for transferring data [13]), as shown in Figure 31. In this CITIfile, FREQOFFSET is the independent variable, and phase_noise is the dependent variable.   The CITIfile data was read into a HP MDS DATASET, using the command PERFORM/READ/CITIFILE, and the DATASET VARIABLE allows the DATASET data to be used in MDS simulations. In simulations, for each value of FREQOFFSET, Pnoise takes a different value, as originally defined in the CITIfile.

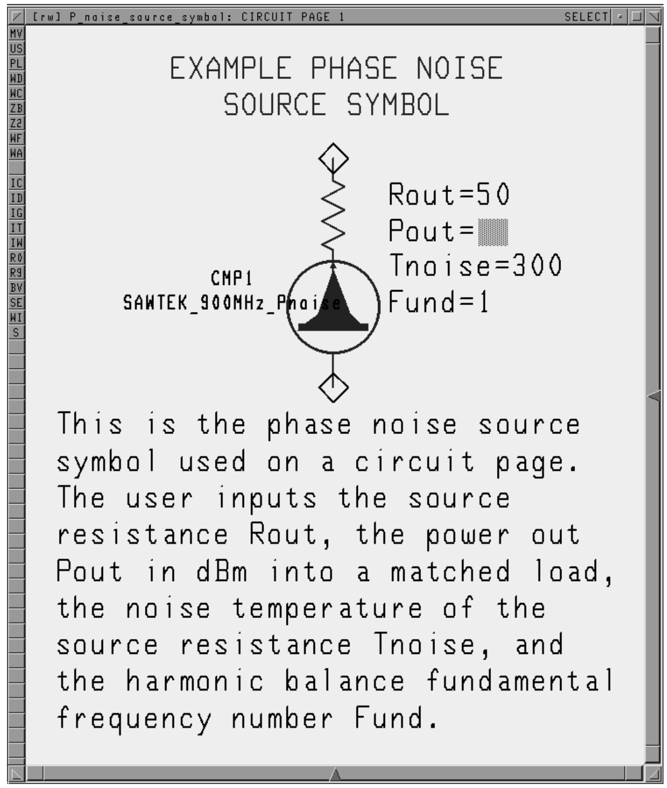

Figure 29: Symbol of signal source with phase noise.

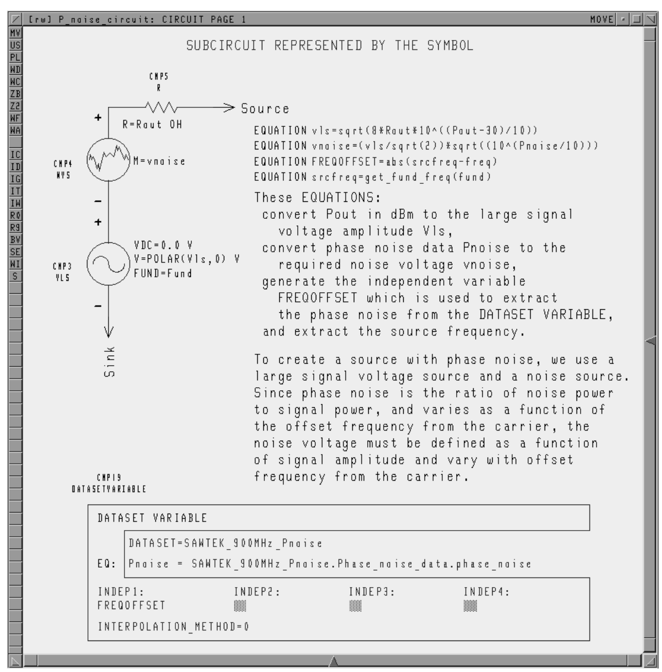

Figure 30: Circuit represented by symbol in Figure 29.

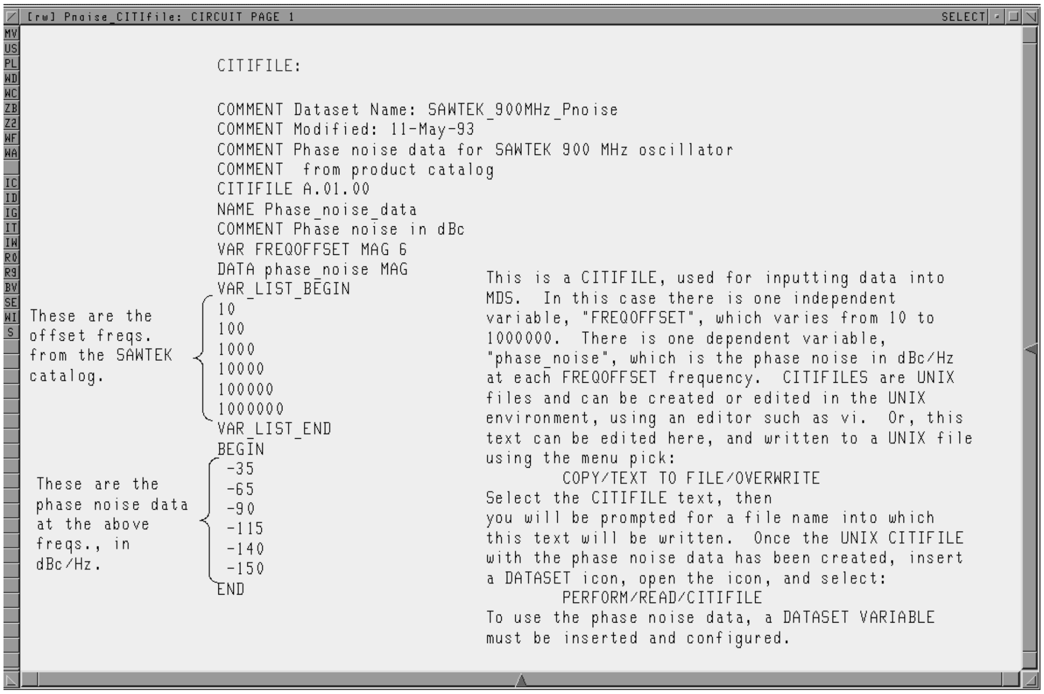

Figure 31: CITIfile example, for entering oscillator phase noise.

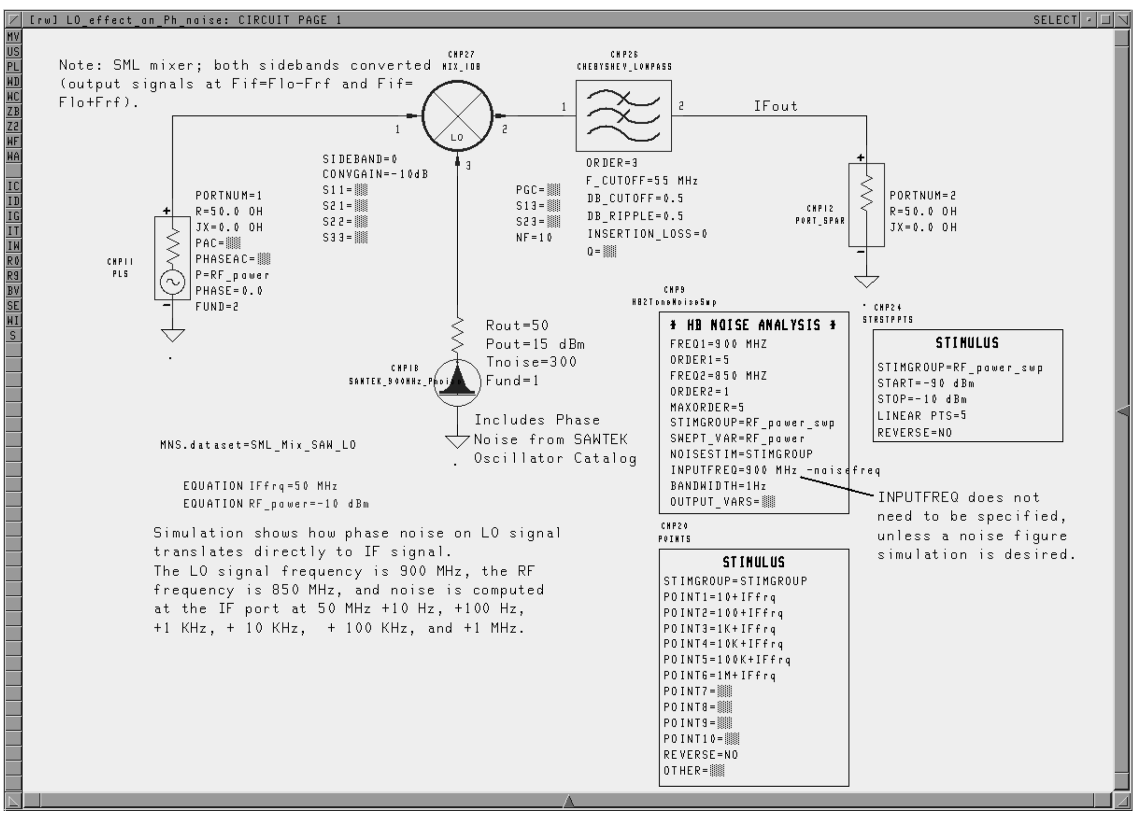

Figure 32: Translation of an LO's phase noise through a mixer to the IF output.

The equation for FREQOFFSET computes the difference between the phase noise signal source's fundamental frequency and "freq", a variable internal to the simulator. The simulator automatically, sequentially sets "freq" to any frequency that would mix (translate) to the output noise frequency specified by the designer. Refer to Figure 32 for an example showing how to simulate LO noise translated from a mixer's LO to its IF output frequency. In this case, the noise at the output will be computed at the frequencies: IFfrq + 10 Hz, IFfrq +100 Hz, etc., where IFfrq = 50 MHz. When the output frequency is IFfrq + 10 Hz, the simulator will compute how much noise is translated to this frequency due to mixing from large-signal tones (i.e., mixing from the RF, LO, RF+LO, 2RF-LO, 2LO-RF, etc.). The "freq" variable described above is set to RF +10 Hz, RF+LO +10 Hz, etc., and the noise translated from these frequencies to the output frequency, IFfrq +10 Hz, is computed.

In this simulation, the RF is 850 MHz, the LO is 900 MHz, and the RF power is swept from -90 dBm to -10 dBm. The simulation results, showing how the noise at the IF port (plotted as an offset from the IF frequency) depends on the RF signal power, are shown in Figure 33.

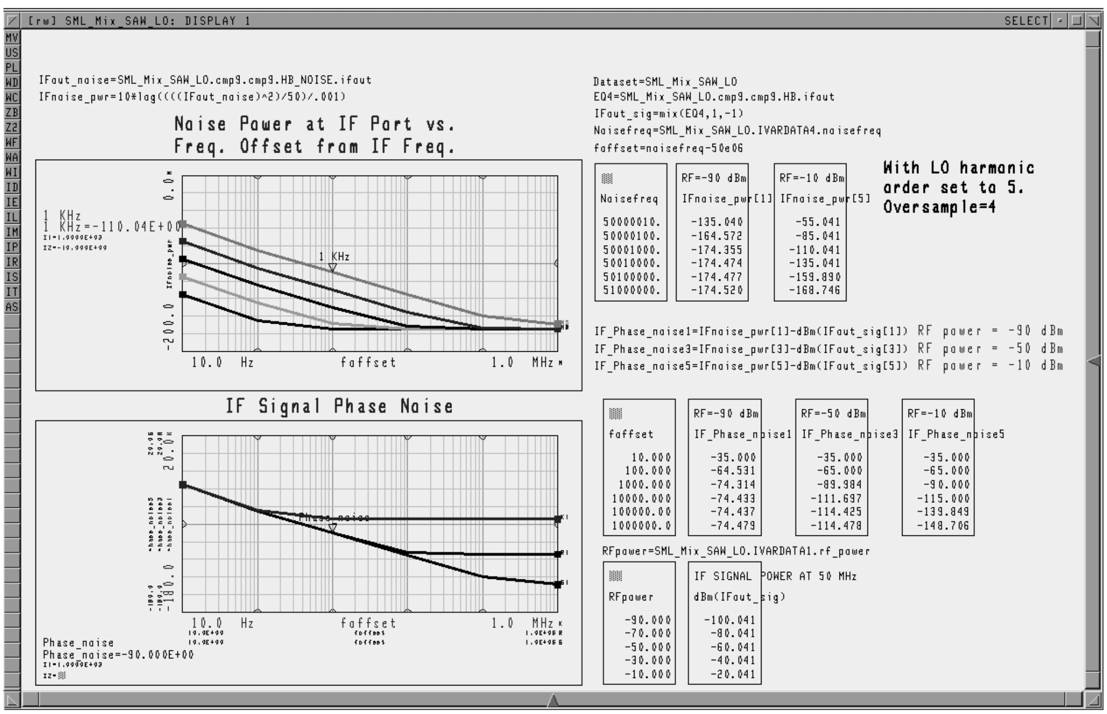

Noisefreq	RF=-90 dBm IFnoise_pwr[1]	RF=-10 dBm IFnoise_pwr[5]
50000010.	-135.040	-55.041
50000100.	-164.572	-85.041
50001000.	-174.355	-110.041
50010000.	-174.474	-135.041
50100000.	-174.477	-159.890
51000000.	-174.520	-168.746

foffset	RF=-90 dBm IF_Phase_noise1	RF=-50 dBm IF_Phase_noise3	RF=-10 dBm IF_Phase_noise5
10.000	-35.000	-35.000	-35.000
100.000	-64.531	-65.000	-65.000
1000.000	-74.314	-89.984	-90.000
10000.000	-74.433	-111.697	-115.000
100000.00	-74.437	-114.425	-139.849
1000000.0	-74.479	-114.478	-148.706

RFpower	dBm(IFout_sig)
-90.000	-100.041
-70.000	-80.041
-50.000	-60.041
-30.000	-40.041
-10.000	-20.041

Figure 33: Noise power and phase noise at the IF port, plotted versus offset frequency from the IF, and for different RF input powers.

There are more system phase noise examples in the demo file, "Phase_noise," described above.

Using results from a different phase noise simulation

With HP MDS B.05.01 and later releases, the results of one noise simulation can be used in another simulation more easily, using the data-based large signal voltage source. This could be used, for example, when simulating an oscillator/mixer combination. If the oscillator is being designed, and its phase noise is unknown, the oscillator could be simulated by itself, driving a load impedance equal to the mixer LO's input impedance (assuming the oscillator will be the mixer's LO). (The designer can assume that the mixer's LO input impedance is 50 ohms, for simplicity, but some error may be introduced.) The oscillator's simulation results (its harmonics and phase noise voltage versus offset from the fundamental) are stored in a dataset. To use this data in a subsequent simulation, for example, as a mixer's LO, insert a data-based large signal voltage source using the menu pick: INSERT/MDS SOURCES/INDEPENDENT VOLTAGE/DATA-BASED LARGE SIGNAL. Enter the previously simulated oscillator's dataset name on the data-based large signal voltage source component. Also enter the variable name using the menu pick: PERFORM/DATASET VARIABLE/EDIT, and select the data-based large signal voltage source. A list of variables from the dataset specified will appear. Select the desired output variable. Usually this is ...HB.vout, where "vout" is the output node label on the oscillator simulation circuit page.

Figure 34 shows an example of how a data-based large signal voltage source is used in a system simulation. The simulated noise data could be plotted similar to that shown in Figure 33.

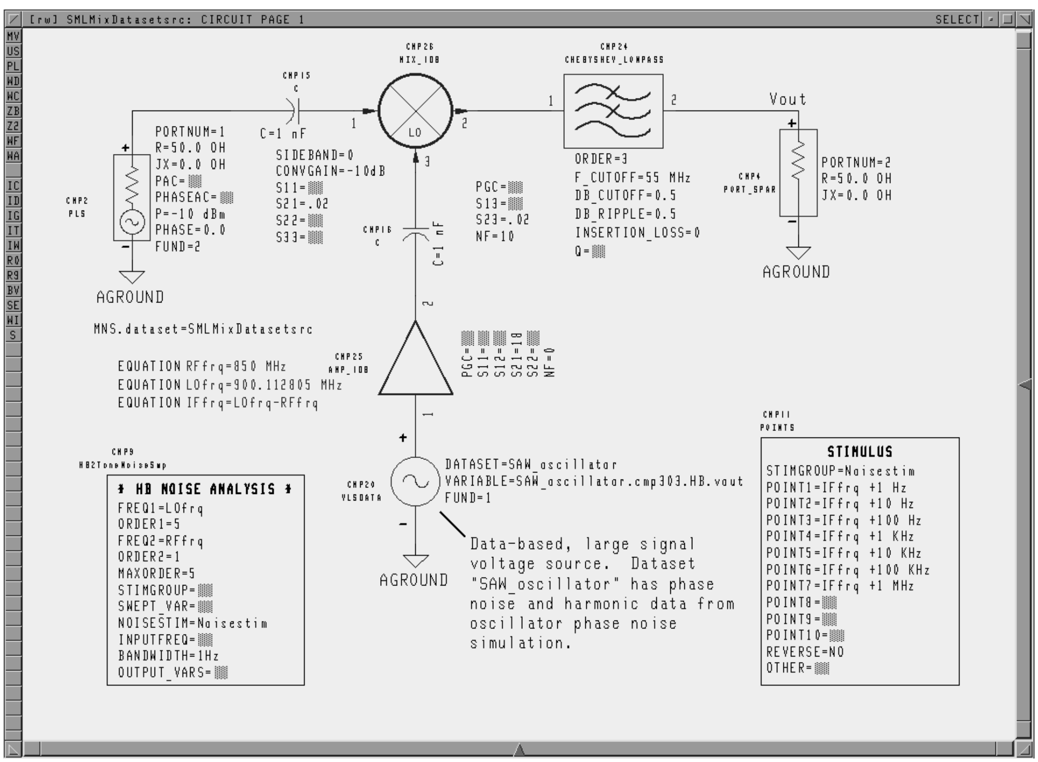

Figure 34: System simulation using previously simulated oscillator phase noise data.

Additional Tips on Nonlinear Noise Simulation:

Nonlinear noise simulation takes more memory than harmonic balance simulation because it accounts for all sidebands and because of the frequency conversion transimpedance matrix that must be computed. Noise computation time can be reduced by only labeling circuit nodes at which the noise must be computed. The fewer the labeled nodes, the smaller the transimpedance matrix will be, and the faster the simulation.

If memory usage or speed prevents simulation, use PERFORM/EDIT COMPONENT to change these parameters on the simulation component:

set the OTHER parameter to OTHER = ALLSSFREQS=NO

set SSTHRESH simulation parameter to a larger non-zero value, such as 1E-4.

The default value for SSTHRESH is zero (all mixing terms are accounted for). As this value is increased, more mixing terms are neglected, and the simulation speeds up. However, as mixing terms are neglected, accuracy will degrade.

When simulating noise figure of nonlinear circuits with multiple tones, be aware of the warning, “Warning detected by MNS during HB analysis ‘hp’. More than one mixing term has landed on frequency XX MHz.” For example, this warning would occur if a simulation were run with an LO at 1000 MHz, and an RF at 800 MHz, if MAXORDER were 7 or higher. This is because there are two ways the LO and RF tones can mix together to generate the 200 MHz IF frequency. One is LO-RF, which is desired, and the other is 4*RF-3*LO. The simulator computes different conversion gains, and consequently, different noise figures for each of the two possibilities. However, there is no way to be sure that the noise figure associated with the desired conversion will be the computed result. To avoid this problem,change one of the frequencies slightly. For example, set the LO to 1000.01 MHz, and specify the NOISEFREQ at 200.01 MHz.

There are additional notes on nonlinear noise simulation in [14].

References:

1. Dieter Scherer, “Design Principles and Test Methods for Low Phase Noise RF and Microwave Sources,” Hewlett-Packard RF& Microwave Measurement Symposium and Exhibition, 1978.
2. Grant Moulton, “Analysis and Prediction of Phase Noise in Resonators and Oscillators,” Hewlett-Packard RF & Microwave Measurement Symposium and Exhibition, 1985.
3. T.R. Faulkner and R.E. Temple, “Residual Phase Noise and AM Noise Measurements and Techniques,” Hewlett-Packard RF & Microwave Measurement Symposium and Exhibition, 1987.
4. Bahl and Bhartia, “Microwave Solid State Circuit Design,” John Wiley and Sons, Inc., 1988.
5. Les Besser, “RF Component Modeling for CAD/CAE,” RF Design, July 1993.
6. Paul R. Gray and Robert G. Meyer, “Analysis and Design of Analog Integrated Circuits,” John Wiley and Sons, Inc., 1984.
7. Alain Cappy, “Noise Modeling Measurement Techniques,” IEEE Transactions on Microwave Theory and Techniques, Vol. 36, No. 1, January 1988.
8. Rizzoli, Mastri, and Cecchetti, “Computer-Aided Noise Analysis of MESFET and HEMT Mixers,” IEEE Transactions on Microwave Theory and Techniques, Vol. 37, No. 9, September 1989.
9. Rucker and Hellums, “Low-Frequency Noise Characteristics of Gallium Arsenide MESFETs,” Solid-State Electronics Vol. 27, No. 11, pp. 947-948, 1984.
10. Elliott and Bray, “Direct Phase Noise Measurements of SAW Resonators,” IEEE 1984 Ultrasonics Symposium, pp. 180-184.
11. Thomas Parker, “Nonlinear Modeling With User-Defined Devices,” available from HP-EEsof.
12. Hewlett-Packard Co., Application Note 283-3, “Low Phase Noise Applications of the HP 8662A and 8663A Synthesized Signal Generators.”
13. HP 85150B Microwave and RF Design Systems, Designer’s Task Reference, Volume 2.
14. HP 85150B Microwave and RF Design Systems, Designer’s Task Reference, Volume 4, Chapter 5.
15. Jim Jaffee and Julio Costa, “Extracting 1/f Noise Coefficients for BJTs”, IEEE Journal of Electron Devices, submitted for publication, October, 1993.
16. HP 85170A System Model Library.

For more information on Hewlett-Packard Test & Measurement products, applications or services please call your local Hewlett-Packard sales offices. A current listing is available via Web through AccessHP at http://www.hp.com. If you do not have access to the internet please contact one of the HP centers listed below and they will direct you to your nearest HP representative.

United States:
Hewlett-Packard Company
Test and Measurement Organization
5301 Stevens Creek Blvd.
Bldg. 51L-SC
Santa Clara, CA 95052-8059
1 800 452 4844

Canada:
Hewlett-Packard Canada Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
(905) 206 4725

Europe:
Hewlett-Packard
European Marketing Centre
P.O. Box 999
1180 AZ Amstelveen
The Netherlands

Japan:
Yokogawa-Hewlett-Packard Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192, Japan
(81) 426 48 3860

Latin America:
Hewlett-Packard
Latin American Region Headquarters
5200 Blue Lagoon Drive, 9th Floor
Miami, Florida 33126, U.S.A.
(305) 267 4245/4220

Australia/New Zealand:
Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130, Australia
131 347 Ext. 2902

Asia Pacific:
Hewlett-Packard Asia Pacific Ltd.
17-21/F Shell Tower, Time Square,
1 Matheson Street, Causeway Bay,
Hong Kong
(852) 2599 7070

Data Subject to Change

Printed in U.S.A. 10/95
5091-9582E