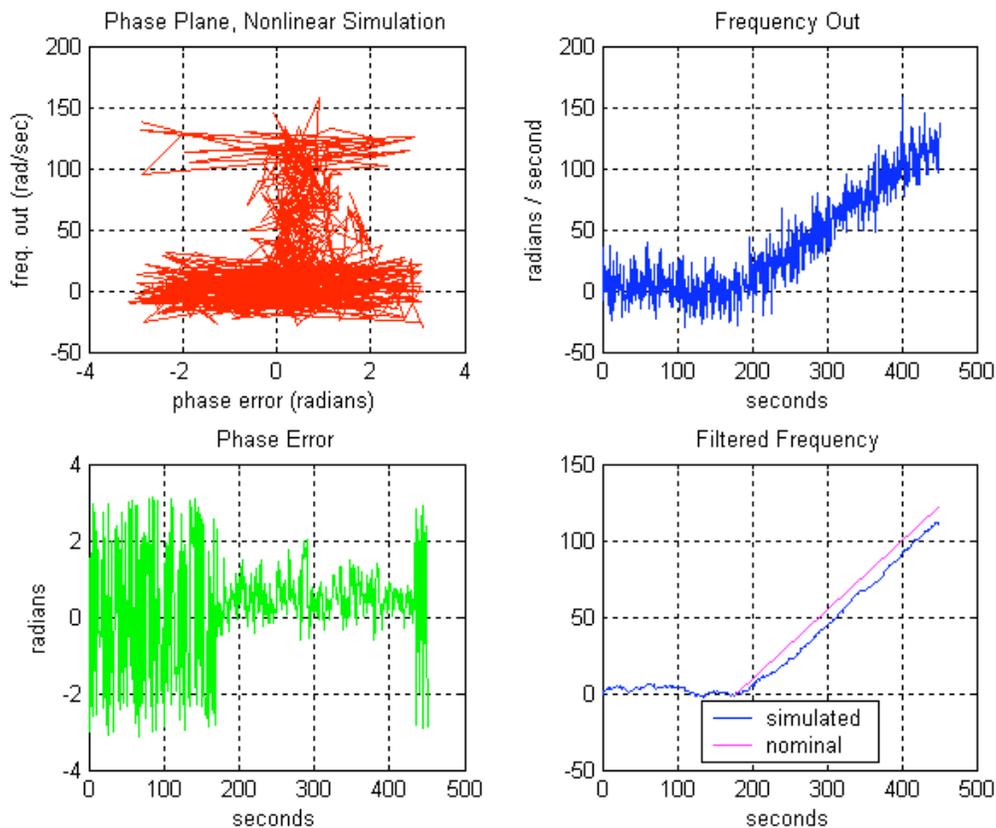


## 19.M APPENDIX: MATLAB SCRIPT, `Swpi`

The MATLAB script `Swpi`, which was used to generate simulations of swept acquisitions for this chapter, is described here. It can be downloaded for a more complete understanding. `Swpi` has several modes of operation.

### 19.M.1 Time Responses

Figure 19.M.1 shows one of the possible time-response outputs of `Swpi`. It is similar to time response outputs previously described except for the new “Filtered Frequency” output. For these plots, the input reference frequency is  $80\omega_n$  below the VCO center frequency at the start of the sweep, the sweep rate is  $0.45\omega_n^2$  and  $\omega_n = 1$ .



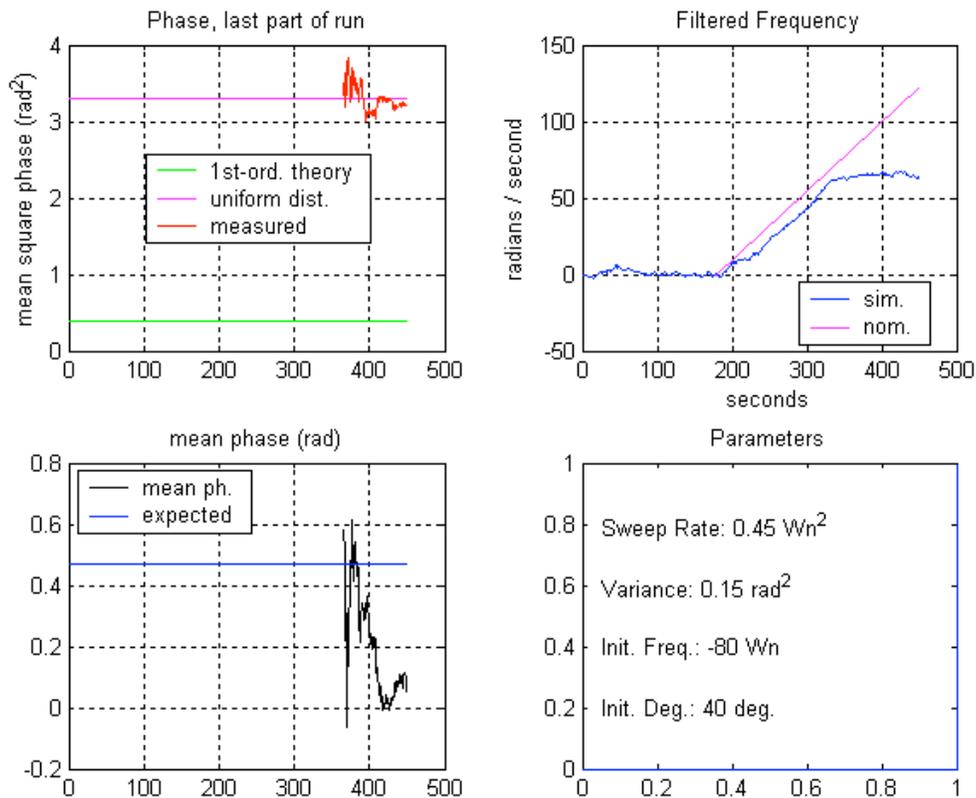
**Fig. 19.M.1 Time responses from `Swpi`.**

Frequency Out shows the VCO frequency following the input beginning around  $(80/0.45 \approx) 178$  seconds, indicating that the loop has locked. However, the Phase Error noise, which diminishes at the same time, becomes noisy again near the end of the simulation, leading us to suspect that the lock has broken. This is not apparent from Frequency Out. Filtered Frequency is Frequency Out after passing through a simple first-order filter to reduce the noise so we can better see the average frequency. The filter also

causes a delay in the response, which makes it lag the nominal value and helps to mask a loss of lock near the end of the simulation. This illustrates how the declaration of a successful lock depends on the length of the simulation.

### 19.M.2 Running Statistics for Individual Sweeps

The two left plots in Fig. 19.M.1 can be replaced with statistical information as in Fig. 19.M.2.



**Fig. 19.M.2 Display with statistical information.**

The lower right box shows information about the simulation. These are parameters that can be iterated when `Swpi` is operating in that mode, so it is helpful to be able to read them when we observe the responses as they are displayed.

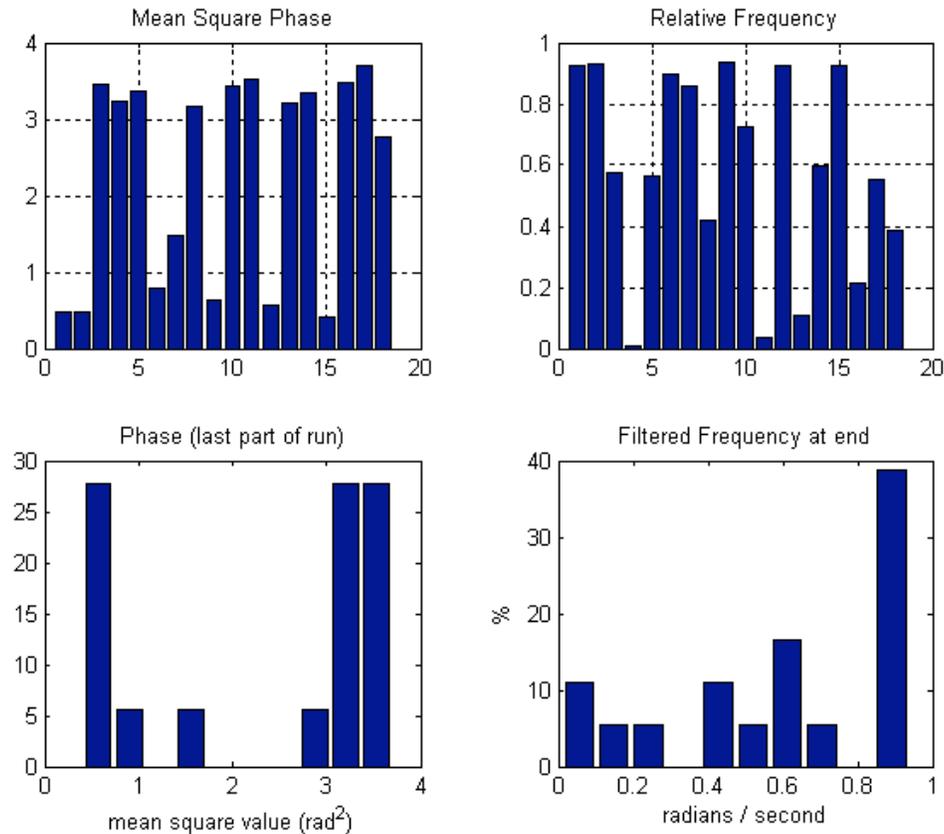
Above that is the Filtered Frequency. In this run, the lock broke in mid sweep. On the upper left is the mean square phase, which has been set to begin averaging at 80% of the sweep duration. This is done to emphasize conditions near the end of the sweep. We can see that the mean square phase is near a value corresponding to a uniform distribution of phase, an indication that the loop is not locked.

The Mean Phase, also averaged over the last 80% of the run, jumps toward the value expected during the sweep, but then rapidly moves toward zero as more points are averaged, again suggesting that the loop is out of lock.

The purpose of these displays is to help us judge whether lock is attained, but many sweeps must be run to gather statistical information and `Swpi` provides aids in judging the probability of lock taken over a large number of initial phases.

### 19.M.3 Obtaining Probability of Lock

`Swpi` is set up for iteration but, as usual, each iterated variable can be set to a single value. Normally we use a set of equally spaced phases (do not end at  $360^\circ$  if you start at  $0^\circ$  because that will sample the  $0^\circ$  condition twice), sometimes spaced as closely as  $0.1^\circ$ . If more than two phases are used in iteration, the script will switch to a display as shown in Fig. 19.M.3 at the conclusion of the phase iteration. For clarity of illustration, this example uses only eighteen increments (of  $20^\circ$ ).



**Fig. 19.M.3 Display after iterated run.**

The upper boxes show a history, over the phase iterations, of the last values of the Mean Square Phase (in  $\text{rad}^2$ , as in the corresponding box in Fig. 19.M.2) and of the Relative Frequency. The lower boxes are histograms, showing the percent of these two variables that fell into each of ten equally spaced bins. We expect lock to correspond to a high

value of Relative Frequency and a low value of Mean Square Phase. Note that 39% of the final Relative Frequencies were in the highest bin. The frequencies in this bin correspond to the seven highest bars in the display above. The lower-left histogram indicates that, in order to reach 39% of mean square phases, the lower four bins (including the empty one just above 1 rad<sup>2</sup>) must be summed. The corresponding seven lowest bars shown above include one representing a mean square phase that is three times the lowest values. Since these data represent only eighteen simulations, we can see that both the highest of the seven low mean square phases and the lowest of the seven high relative frequencies occurred on the same run. Did all seven runs represent locks? This illustrates some of the difficulties in judging the results, even with so few points.

When many more initial phases are included, bars in the upper part of the display can become too dense to differentiate. It is easier to interpret the text output, as shown in Fig. 19.M.4 for the same eighteen points.

```
(rad^2: random 3.29, expect 0.3817)
rel. freq.    %      integ.    rad^2      %      integ.
 0.0556      11.11   100.00    3.5341    27.78   100.00
 0.1484       5.56    88.89     3.2050    27.78    72.22
 0.2412       5.56    83.33     2.8760     5.56    44.44
 0.3339       0.00    77.78     2.5470     0.00    38.89
 0.4267      11.11    77.78     2.2180     0.00    38.89
 0.5195       5.56    66.67     1.8889     0.00    38.89
 0.6123      16.67    61.11     1.5599     5.56    38.89
 0.7050       5.56    44.44     1.2309     0.00    33.33
 0.7978       0.00    38.89     0.9018     5.56    33.33
 0.8906      38.89    38.89     0.5728    27.78    27.78
Init. Deg.:   0.0   20.0  40.0  60.0  80.0 100.0 120.0 140.0
Rel. Freq.:  0.925 0.928 0.577 0.009 0.566 0.899 0.857 0.419
M.S. Phase:  0.488 0.467 3.446 3.240 3.360 0.800 1.476 3.160
              1      2      3      4      5      6      7      8
```

Fig. 19.M.4 Text display after iterated run (modified to fit).

The display has two halves, divided vertically, one for relative frequency and one for mean square phase. Percentages are shown for each of the ten bins and cumulative percentages are shown beside them. We can more easily see that the isolated bin for maximum relative frequency contains 38.89% of the runs and the same percentage requires the four lowest bins of mean square phase. The individual values for each run are shown below and extend beyond the figure. The value of mean square phase that represents random phase and the value expected for the locked and sweeping, but noiseless, loop are shown at the upper right.

The text output also includes a detailed description of the simulation parameters. Output is also written to a file, a very important safeguard here because of the likelihood of long runs. These often involve iteration of noise level, or some loop parameter, and sweep rate, as well as initial phase.