22 A Short Introduction to Bode Plots

v1.4 – June 2021

Prerequisite knowledge required: Frequency Responses

22.1 Introduction

The Bode[1] plot is a method of plotting frequency responses which also provides a graphical technique for estimating the frequency response of a network given the gain at low frequencies and the break frequencies of the poles and zeros. This can be very useful, as it allows a designer to quickly produce a graphical estimate of the frequency response of a system without having to calculate the values for a large number of frequencies.

For the amplitude response (in other words the gain as a function of frequency), the Bode plot is always produced on a graph using dB on the y-axis and a logarithmic scale of frequency on the x-axis. (These logarithmic axes are used for two reasons: they make it possible to show a wide range of gains on the same graph, and they allow accurate approximations to the performance of many simple circuits to be drawn using straight lines.)

For the phase response (in other words the phase difference between the output and the input as a function of frequency), the y-axis is typically degrees, while the x-axis remains a logarithmic function of frequency. An example of a Bode plot is shown in Figure 22.1, for the case of a circuit with two poles with break frequencies at 320 Hz and 360 kHz, and a zero at 10 kHz.

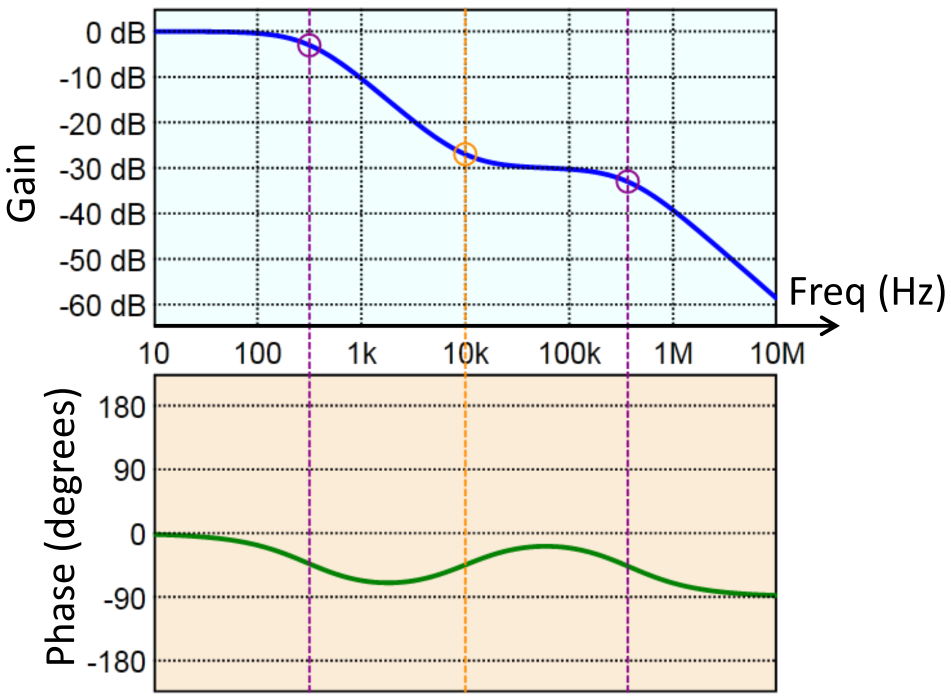

Figure 22.1 Example Bode plot showing amplitude and phase response

[1] Pronounced "Boh-dee", and named after its inventor, Hendrik Wade Bode (1905-82).

The most important (and useful) point about the Bode plot is that frequency responses with a limited number of poles and zeros result in plots that can be conveniently approximated using straight lines. This chapter is about why this happens, and how to produce these estimated frequency responses quickly.

22.2 The effect of a single pole

Suppose you have a frequency response which is composed of a single pole, something like:

$$H(j\omega) = \frac{A}{(1 - j\omega / p_1)} \tag{22.1}$$

A Bode plot of this response looks like this (when A = 1):

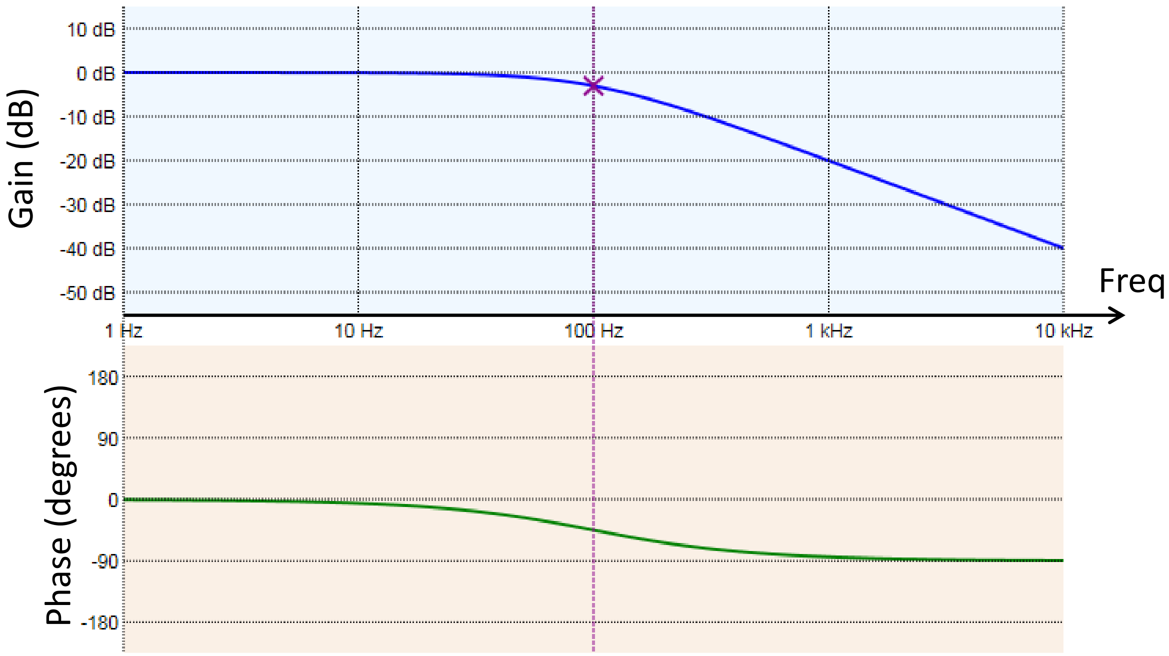

Figure 22.2 Bode plot of a single pole response

First, consider what happens at very low frequencies. Here, the expression $j\omega/p_1$ will be much smaller than one, so we can approximate the frequency response as:

$$H(j\omega) = \frac{A}{(1)} = A \tag{22.2}$$

In other words it's a constant, with value equal to the low-frequency gain of *A*. On the Bode amplitude plot (which plots the gain in decibels), this would be a value of:

$$\text{Gain (dB)} = 20\log_{10}\left(\left|H(j\omega)\right|\right) = 20\log_{10}(A) \tag{22.3}$$

and the phase difference between the output and input would be zero, since the gain is entirely real, the output and input would be in phase. Looking at Figure 22.2 this is exactly what we see.

Now consider high frequencies. At very high frequencies, the expression $j\omega/p_1$ will be much greater than one, so we can approximate the frequency response as:

$$H(j\omega)=\frac{A}{(-j\omega/p_1)}=j\frac{p_1A}{\omega} \tag{22.4}$$

which could be expressed in polar terms as:

$$H(j\omega)=j\frac{p_1A}{\omega}=\frac{p_1A}{\omega}\exp\left(j\frac{\pi}{2}\right) \tag{22.5}$$

On the Bode plot for the amplitude response (which plots the amplitude gain in decibels), this would be a value of:

$$\text{Gain (dB)}=20\log_{10}\left(\left|\frac{p_1A}{\omega}\right|\right)=20\log_{10}(|p_1A|)-20\log_{10}(\omega) \tag{22.6}$$

Plotted against frequency on a logarithmic axis, this is a straight line, with a gradient of -20 dB per decade. (If that's not obvious, consider two frequencies a factor of 10 apart: the gain at the lower frequency ω_1 would be:

$$\text{G}_1\text{ (dB)}=20\log_{10}(|p_1A|)-20\log_{10}(\omega_1) \tag{22.7}$$

and at the higher frequency $10\omega_1$ it would be:

$$\begin{aligned}\text{G}_2\text{ (dB)}&=20\log_{10}(|p_1A|)-20\log_{10}(10\omega_1)\\&=20\log_{10}(|p_1A|)-20\log_{10}(10)-20\log_{10}(\omega_1)\\&=\text{G}_1\text{ (dB)}-20\log_{10}(10)\\&=\text{G}_1\text{ (dB)}-20\end{aligned} \tag{22.8}$$

which is 20 dB lower than the gain at ω_1. This happens for all high frequencies, so the graph at high frequencies must be a straight line that drops 20 dB for each increase in frequency of a factor of ten.)

In terms of phase, at high frequencies the phase difference between the output and input is given by the phase term in the polar response, which is:

$$\arg(H(j\omega))=\arg\left(\frac{-p_1A}{\omega}\exp\left(j\frac{\pi}{2}\right)\right)=-j\frac{\pi}{2} \tag{22.9}$$

and this is not a function of frequency. Again, Figure 22.2 shows that this approximation is accurate at high frequencies.

It's only within a factor of ten of the pole's break frequency (100 Hz in the example) that neither the high-frequency or low-frequency approximations are particularly accurate, however even here an estimate based on the low- and high-frequency approximations does give a reasonable estimate of the actual response.

All you have to do is extrapolate the low- and high-frequency approximations to the amplitude response until they meet, and estimate the phase by assuming that the phase is zero up to 1/10th of the break frequency, and 90 degrees after 10 times the break frequency, and changes linearly between these two frequencies.

These Bode approximations are shown on the figure below. As you can see, the approximations are quite accurate (for the case of a single pole, the maximum error in the amplitude response is 3 dB, and occurs at the break frequency; see later for an explanation of why this is the case).

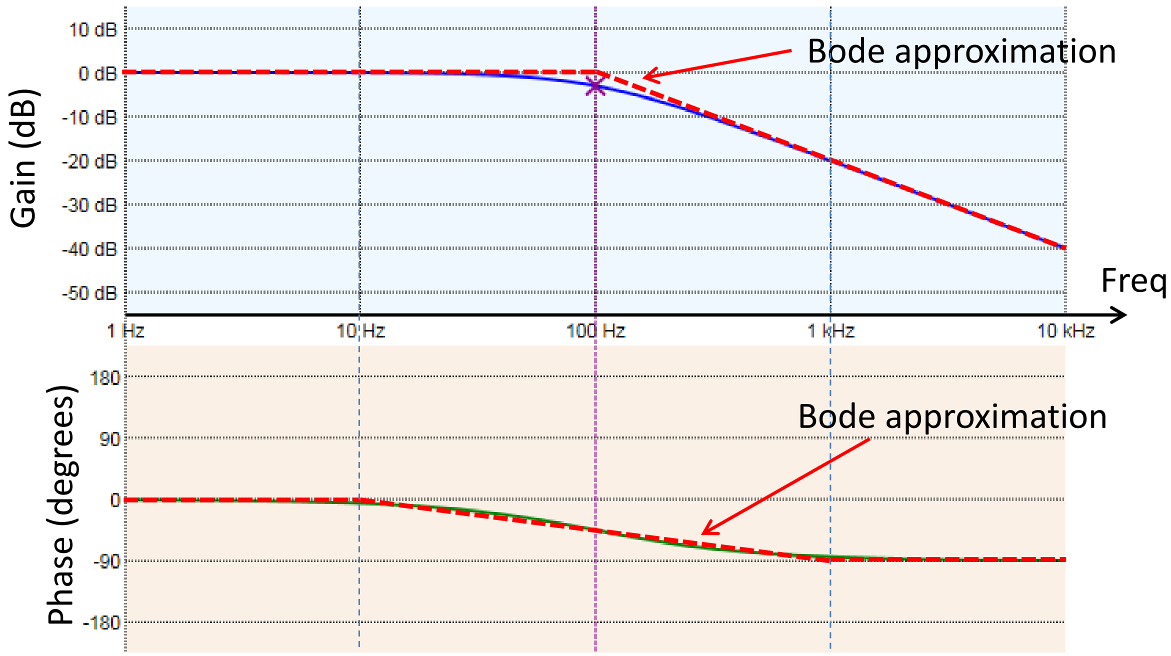

Figure 22.3 Bode plot for single-pole filter showing the Bode approximations

22.3 The effect of a single zero

Similar results can be obtained for zeros. A circuit with a single zero would have a frequency response of:

$$H(j\omega) = A(1 - j\omega / z_1) \tag{22.10}$$

which when plotted looks like:

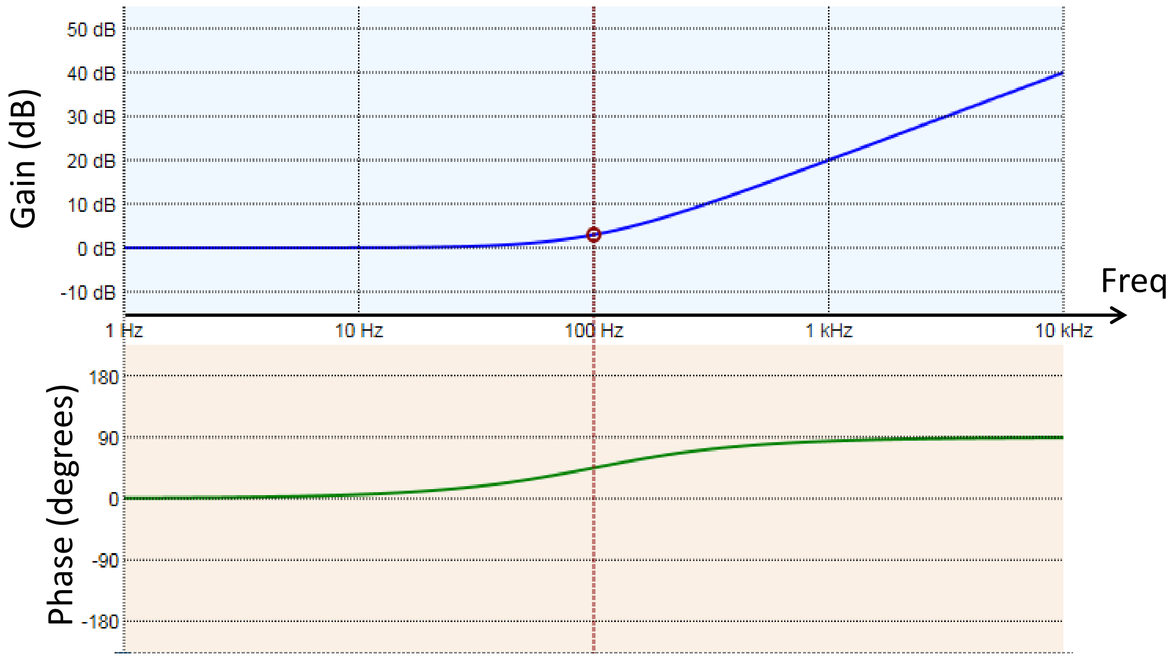

Figure 22.4 Bode plot of a single-zero response

With a single zero in the response, the low-frequency approximation is:

$$H(j\omega) = A \tag{22.11}$$

which is a constant gain of *A* with no phase difference between the input and output, and the high-frequency approximation is (in polar form):

$$H(j\omega) = -A\frac{\omega}{z_1}\exp\left(j\frac{\pi}{2}\right) \tag{22.12}$$

so the amplitude response at high-frequencies in decibels is:

$$\text{Gain (dB)} = 20\log_{10}\left(\left|\frac{A\omega}{z_1}\right|\right) = 20\log_{10}\left(\left|\frac{A}{z_1}\right|\right) + 20\log_{10}(\omega) \tag{22.13}$$

which is a straight line on a Bode plot, increasing by 20 dB per decade, and the phase response at high-frequencies is:

$$\arg(H(j\omega)) = \arg\left(-A\frac{\omega}{z_1}\exp\left(j\frac{\pi}{2}\right)\right) = \frac{\pi}{2} \tag{22.14}$$

which is constant, and equal to 90 degrees. Both these are consistent with the graph of the single-zero response, and again, an approximation can be made to the frequency response by extrapolating the straight lines of the amplitude response, and producing a piecewise-linear approximation to the phase response from 1/10th of the zero's break frequency to 10 times the zero's break frequency:

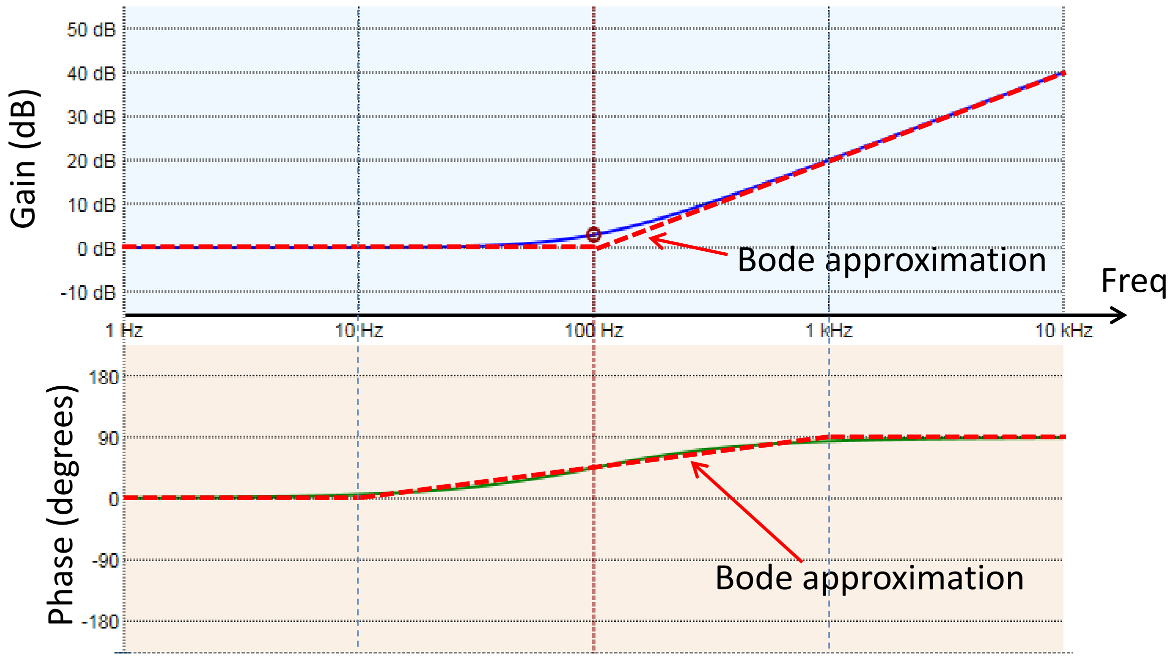

Figure 22.5 Bode plot for single-zero filter showing the Bode approximations

Once again, the approximations are quite good over a wide range of frequencies.

22.4 Multiple poles and zeros

Where the Bode approximations really start to be useful is when there are multiple poles and zeros in the response. For example, consider a network with a series of poles and zeros. The general form of the frequency response is:

$$H(j\omega)=\frac{G\,(j\omega)^n(1-j\omega/z_1)(1-j\omega/z_2)(1-j\omega/z_3)\ldots}{(1-j\omega/p_1)(1-j\omega/p_2)(1-j\omega/p_3)\ldots} \tag{22.15}$$

and to work out the amplitude response, we can multiply through by the complex conjugate to get the square of the gain in amplitude, and since:

$$(1+j\omega/z_1)(1-j\omega/z_1)=(1+\omega^2/z_1{}^2) \tag{22.16}$$

this gives:

$$\begin{aligned}|H(j\omega)|^2&=\frac{G^2\omega^{2n}(1+\omega^2/z_1{}^2)(1+\omega^2/z_2{}^2)(1+\omega^2/z_3{}^2)\ldots}{(1+\omega^2/p_1{}^2)(1+\omega^2/p_2{}^2)(1+\omega^2/p_3{}^2)\ldots}\\|H(j\omega)|^2&=G\omega^n\sqrt{\frac{(1+\omega^2/z_1{}^2)(1+\omega^2/z_2{}^2)(1+\omega^2/z_3{}^2)\ldots}{(1+\omega^2/p_1{}^2)(1+\omega^2/p_2{}^2)(1+\omega^2/p_3{}^2)\ldots}}\end{aligned} \tag{22.17}$$

If expressed in terms of decibels, this gives:

$$
\begin{aligned}
Gain(dB) &= 10\log_{10}\left(\left|H(j\omega)\right|^2\right) \\
&= 20\log_{10} A - 10\log_{10}\left(1+\omega^2/p_1^{\,2}\right) - 10\log_{10}\left(1+\omega^2/p_2^{\,2}\right) - \ldots \\
&\qquad + 10\log_{10}\left(1+\omega^2/z_1^{\,2}\right) + 10\log_{10}\left(1+\omega^2/z_2^{\,2}\right) - ..
\end{aligned}
\tag{22.18}
$$

and as you can see the effects of all the poles and zeros can be separated out, and the total amplitude response expressed as the sum of the effects of each pole and each zero.

The Bode approximation assumes that poles and zeros have no effect on the amplitude response below the break frequency, and produce a slope of 20 dB/decade above, so this suggests a simple way of deriving the shape of the frequency response given the low-frequency gain, and the pole and zero break frequencies. The simplest form of the method (which works provided there are no zeros at 0 Hz in the frequency response) goes like this:

1) Start with a gain of 20 $\log_{10}$(A) at low frequencies, and continue with a constant gain until you reach the first pole or zero break frequency
2) If the break frequency is due to a pole, bend the line down by 20 dB/decade
3) If the break frequency is due to a zero, bend the line up by 20 dB/decade
4) Continue on to the next breakpoint, and go to step 2).

However, for responses with zeros at zero this approach doesn't work. With one or more zeros at zero the Bode approximation (and the actual response) will start increasing at 20N dB/decade at low frequencies (where N is the number of zeros at 0 Hz). Sometimes in these cases you can easily work out the amplitude response at very high (infinite) frequencies and work backwards, however in general it's necessary to work out the gain at a low frequency (a long way below the lowest break frequency of any pole or zero) and start from there, with the amplitude response moving up at 20 dB/decade.

22.4.1 Approximating the phase response

Bode plots can also be used to estimate the phase shift of a circuit. For the phase, we can note that

for a single pole:

$$
1 - j\omega/p_1 = \left|1 - j\omega/p_1\right| \exp\left(j\tan^{-1}\left(\frac{-\omega}{p_1}\right)\right) \tag{22.19}
$$

and for a single zero:

$$
\left(1 - j\omega/z_1\right) = \left|1 - j\omega/z_1\right| \exp\left(j\tan^{-1}\left(\frac{-\omega}{z_1}\right)\right) \tag{22.20}
$$

and therefore expressing every pole and zero in the general expression for a frequency response in polar form, gives:

$$\frac{(1-j\omega/z_1)(1-j\omega/z_2)\ldots}{(1-j\omega/p_1)(1-j\omega/p_2)\ldots}$$

$$=\frac{|1-j\omega/z_1|\exp\left(j\tan^{-1}\left(\frac{-\omega}{z_1}\right)\right)|1-j\omega/z_2|\exp\left(j\tan^{-1}\left(\frac{-\omega}{z_2}\right)\right)\ldots}{|1-j\omega/p_1|\exp\left(j\tan^{-1}\left(\frac{-\omega}{p_1}\right)\right)|1-j\omega/p_2|\exp\left(j\tan^{-1}\left(\frac{-\omega}{p_2}\right)\right)\ldots}$$

$$=\frac{|1-j\omega/z_1||1-j\omega/z_2|\ldots}{|1-j\omega/p_1||1-j\omega/p_2|\ldots}\exp\left(j\left(\tan^{-1}\left(\frac{-\omega}{z_1}\right)+\tan^{-1}\left(\frac{-\omega}{z_2}\right)-\tan^{-1}\left(\frac{-\omega}{p_1}\right)-\tan^{-1}\left(\frac{-\omega}{p_2}\right)+\ldots\right)\right)$$

(22.21)

from which it's clear that total phase change resulting from the poles and zeros is just the sum of the individual phase changes caused by each pole and zero. Therefore we can just consider the effects of each pole and zero on the phase, and then add them all up.

Step-by-step, the method is:

1) For all poles, determine one-tenth of the pole's break frequency, and ten times the pole's break frequency.
2) Draw a line that starts at zero at very low frequencies, starts to bend down at one-tenth of the pole's break frequency, goes through -45 degrees at the break frequency, reaches -90 degrees at ten times the pole's break frequency, then continues horizontally. This is the phase contribution of this pole.
3) For all zeros not at 0 Hz, determine one-tenth of the zero's break frequency, and ten times the zero's break frequency.
4) Draw a line that starts at zero at very low frequencies, starts to bend up at one-tenth of the zero's break frequency, goes through +45 degrees at the break frequency, reaches +90 degrees at ten times the zero's break frequency, then continues horizontally. This is the phase contribution of this zero.
5) Determine how many zeros there are at 0 Hz, and note that their contribution is to add 90 degrees to the phase at all frequencies (since all interesting frequencies are more than ten times the frequency of these zeros).
6) Add up the phase contributions of all the poles and all the zeros in the system. This gives the total phase response.

For simple systems where there is just one pole, this indicates that the input and the output would be expected to be in phase at low frequencies (a long way below the break frequency) and 90 degrees out of phase at high frequencies (a long way above the break frequency). To get more than 90 degrees of phase shift, you need at least two poles (or two zeros).

22.5 An example of a multiple-pole Bode approximation

Consider the example we started with: a circuit with two poles with break frequencies at 320 Hz and 360 kHz, and a zero with a break frequency at 10 kHz. Plotting the Bode approximation of the amplitude response of this circuit would give:

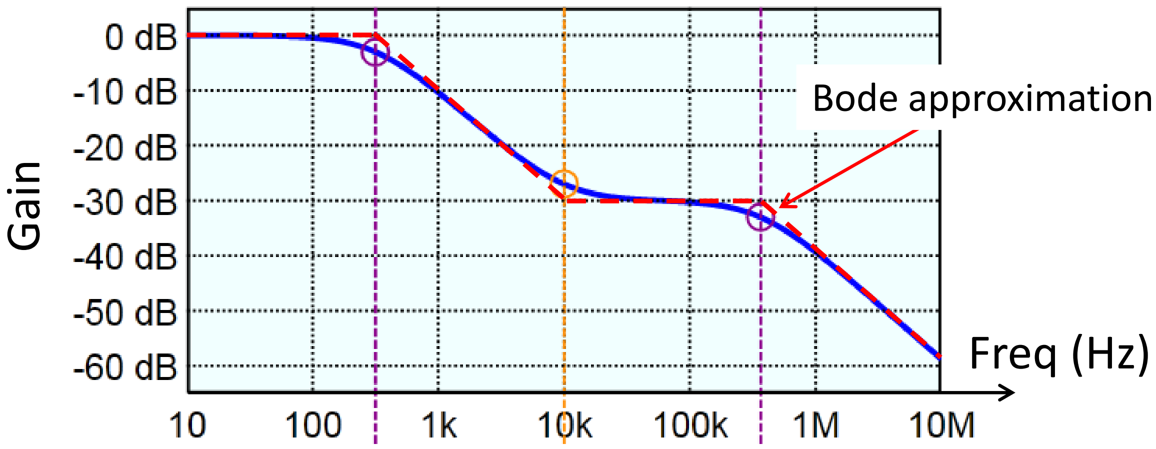

Figure 22.6 Bode approximation of the amplitude response of a two-pole one-zero circuit

Note that at the zero's break frequency at 10 kHz, the act of bending the line up by 20 dB/decade results in the line continuing horizontally again.

22.6 How accurate are the Bode approximations?

This is an interesting and important question.

Looking at Figure 22.3, it appears that the greatest difference between the Bode approximation and the actual amplitude response occurs at the break frequency of the pole. This is true in this case, and here the difference is the difference between the low-frequency gain and the gain of the circuit at the break frequency of the pole, which is:

$$H\left(|p_1|\right)=\frac{A}{\left(1-j|p_1|/p_1\right)}=\frac{A}{\left(1+j\right)}=\frac{A}{\sqrt{2}}\exp\left(-j\frac{\pi}{4}\right) \tag{22.22}$$

In terms of the loss in decibels, this is:

$$E\left(dB\right)=20\log_{10}\left(\frac{A/\sqrt{2}}{A}\right)=20\log_{10}\left(\frac{1}{\sqrt{2}}\right)=10\log_{10}\left(\frac{1}{2}\right)\approx-3 \tag{22.23}$$

In other words, the greatest inaccuracy of this Bode approximation is 3 dB, and occurs at the pole's break frequency. For a greater number of real poles, the maximum possible error is just 3 dB times the number of poles, but this maximum error only occurs when all poles have the same break frequency. In most practical cases 3 dB is a good estimate of the maximum error.

However, if there are any complex poles or zeros (poles and zeros at complex frequencies), the maximum error can be much greater around the resonant frequency of the pairs of complex poles or zeros. While the Bode approximation is not very good around the resonant frequencies of complex poles or zeros, it is still a good approximation for other frequencies.

(See the chapter on "Second-order frequency responses" for more details about this effect.)

22.7 Summary: the most important things to know

- The frequency response of a network $H(j\omega)$ is the ratio of the phasor representing the output to the phasor representing the input.
- The Bode approximation is a piecewise-linear approximation to the amplitude (or phase) response of a network. It can be quickly produced with knowledge of the gain at any frequency, and the poles and zeros in the response.
- The amplitude approximation assumes that poles and zeros have no effect below their break frequency, but produce a 20 dB/decade drop (for poles) or rise (for zeros) after their break frequency.
- Provided the poles and zeros are real, the maximum error in the Bode approximation occurs at the break frequencies, and is around 3 dB per pole or zero.
- The phase approximation assumes that the poles and zeros have no effect below one-tenth of the break frequency, introduce a phase change of 90 degrees (positive for a pole, negative for a zero) above ten times the break frequency, and that the phase changes linearly between 0.1 and 10 times the break frequency.