Analysis Of Ill-Conditioning Of Multi-Channel Deconvolution Problems

Ole Kirkeby, Per Rubak, and Angelo Farina[*]

Department of Communication Techology, Fredrik Bajers Vej 7, Aalborg University, DK-9220 Aalborg Ø, Denmark

[*]**Department of Industrial Engineering, University of Parma, Via delle Scienze, 43100 Parma, Italy**

Deconvolution of single- and multichannel systems is often an ill-conditioned problem whose solution boosts certain frequency bands excessively. A z-domain analysis demonstrates that the degree of ill-conditioning can be assessed by calculating the poles of an ideal solution, and it also shows that frequency-dependent regularisation works by pushing those poles away from the unit circle.

EDICS code: SA 2.3.3

I. INTRODUCTION

Deconvolution in its most basic form can be described as the task of calculating the input to a discrete-time system from its output. It is usually assumed that the system is linear and that its input-output mapping is known with good accuracy. In acoustics and audio, single-channel deconvolution is particularly useful since it can compensate for the response of imperfect transducers such as headphones, loudspeakers, and amplifiers [1]. Multi-channel deconvolution is necessary in the design of cross-talk cancellation systems and virtual source imaging systems [2], [3].

Regularisation is a method that is commonly used when one is faced with an ill-conditioned problem [4, Section 18.4]. The basic idea is to prevent the solution from having some undesirable feature by adding a "smoothness" term to the cost function that we wish to minimise. A suitable choice of the smoothness term can improve the conditioning of the problem substantially but it will inevitably be at the expense of the performance at the troublesome frequencies.

In the following, we show that the effect of frequency-dependent regularisation can be conveniently explained by a pole-zero analysis of a matrix of ideal filters, and that the same pole-zero analysis can be used to indicate the degree of ill-conditioning of multi-channel deconvolution problems. The fast deconvolution method [2], [3] is based on the principles outlined below. This method essentially provides a quick way to solve, in the least squares sense, a linear equation system whose coefficients, right hand side, and unknowns are z-transforms of stable digital filters.

II. SYSTEM DESCRIPTION

The discrete-time multichannel deconvolution problem is shown in block diagram form in Fig. 1. We will use *z*-transforms to denote discrete time filters and signals, and sometimes these *z*-transforms will be referred to as polynomials even though strictly speaking their powers are negative. We define the following column vectors: $\mathbf{u}(z)$ is a vector of *T* observed signals, $\mathbf{v}(z)$ is a vector of *S* source input signals, $\mathbf{w}(z)$ is a vector of *R* reproduced signals, $\mathbf{d}(z)$ is a vector of *R* desired signals, and $\mathbf{e}(z)$ is a vector of *R* performance error signals. The matrices $\mathbf{A}(z)$, $\mathbf{C}(z)$, and $\mathbf{H}(z)$ represent multi-channel filters. $\mathbf{A}(z)$ is an $R{\times}T$ target matrix, $\mathbf{C}(z)$ is an $R{\times}S$ plant matrix, and $\mathbf{H}(z)$ is an $S{\times}T$ matrix of filters whose coefficients we can choose. We assume that the elements of $\mathbf{A}(z)$, $\mathbf{C}(z)$, and $\mathbf{H}(z)$ are finite impulse response (FIR) filters. Note that the vectors read alphabetically $\mathbf{u}$, $\mathbf{v}$, $\mathbf{w}$ along the lower half of the block diagram, and that their dimensions read *T*, *S*, *R* (*R*, *S*, *T*, in reverse); this will make the notation easier to remember. The component z^{-m} implements a so-called modeling delay by shifting all the elements of $\mathbf{u}(z)$ by an integer number of *m* samples [5, Example 7.2.2].

In the single-channel case, *R*, *S*, and *T*, are one, and consequently $A(z)$, $C(z)$, and $H(z)$ are all scalar functions. Perfect equalization of a loudspeaker impulse response, for example, requires that $A(z)$ has flat magnitude response. For multichannel "equalization", $\mathbf{A}(z)$ must be an identity matrix. A perfect cross-talk cancellation network, for example, requires that $\mathbf{A}(z)$ is an identity matrix of order two [3].

III. EXACT LEAST SQUARES DECONVOLUTION

Consider a cost function of the type

$$J = E + \beta V \qquad\qquad\qquad (1)$$

where E is a measure of the performance error **e** and V is a measure of the effort **v**. The positive real number β is a regularization parameter that determines how much weight to assign to the effort term. As β is increased from zero to infinity, the solution changes gradually from minimizing E only to minimizing V only.

It is convenient to consider the regularization to be the product of two components: a gain factor and a shape factor. The gain factor is the conventional regularization parameter β, and the shape factor $B(z)$ is the z-transform of a digital filter that passes through the frequencies that we do not want to see boosted by $\mathbf{H}(z)$. Frequencies that are suppressed by $B(z)$ are not affected by the regularization. Although it is the frequency response, and not the time response, of $B(z)$ that is important, we prefer to design $B(z)$ in the time domain. This prevents its frequency response from changing very abruptly, and it also makes it straightforward to include the z-transform of the filter's impulse response in the analysis. The phase response of $B(z)$ is irrelevant since it does not affect the value of the effort V which is an energy quantity.

The derivation of $\mathbf{H}(z)$ in the general multi-channel case is directly analogous to that presented in [2]. For the purpose of defining $\mathbf{H}(z)$ uniquely, the complex variable z is constrained to be on the unit circle so that the contour $|z| = 1$ is in the region of convergence. This means that the filters contained in $\mathbf{H}(z)$ are constrained to be stable, but not constrained to be either causal or of finite duration. Consequently, they are generally not realisable, and we therefore refer to these filters as *ideal*. Nevertheless, a sufficiently large modeling delay m will ensure that only an insignificant part of the non-causal response cannot be implemented by a realisable digital filter matrix. The ideal filter matrix $\mathbf{H}(z)$ is given by

$$\mathbf{H}(z) = \left[\mathbf{C}^{\mathrm{T}}(z^{-1})\mathbf{C}(z) + \beta \, B(z^{-1})B(z)\mathbf{I}\right]^{-1}\mathbf{C}^{\mathrm{T}}(z^{-1})\mathbf{A}(z) \, z^{-m} \qquad (2)$$

where the superscript T denotes the transpose operator and $\mathbf{I}$ is an identity-matrix of order S. In the single-channel case, this result simplifies to

$$H(z) = \frac{C(z^{-1})A(z)}{C(z^{-1})C(z) + \beta \, B(z^{-1})B(z)} \, z^{-m} \qquad (3)$$

This special case is very important, not only because it is often encountered in practice, but also because it suggests a way to analyze the more complex multichannel result.

IV. POLE-ZERO ANALYSIS

It is seen from Eq. 3 that in the single-channel case $H(z)$ can be naturally expressed in rational form, just like the z-transform of a conventional infinite impulse response (IIR) filter. Consequently, we can learn about the properties of this filter by looking at its poles and zeros in the complex plane. The positions of the poles with respect to the unit circle are particularly important. Poles near the unit circle make the time response of the filter decay away very slowly. The time constant α, in samples, associated with a single pole close to the unit circle is approximately proportional to the reciprocal of the distance r between the two, so

$$\tau = \frac{1}{r} \qquad (4)$$

when $r \ll 1$ [6]. If the pole is just inside the unit circle, the filter's time response will be right-sided and decay away in forward time, whereas if the pole is just outside the unit circle, its response will be left-sided and decay away in backward time [7, Chapter 2].

5

A. General results; the single-channel case

We start by observing that if Eq. 3 is written in the form of a fraction,

$$H(z) = \frac{P(z)}{Q(z)} \tag{5}$$

then the zeros of $H(z)$ are the zeros of the numerator polynomial $P(z)$,

$$P(z) = C(z^{-1})A(z)\,z^{-m} \tag{6}$$

and the poles of $H(z)$ are the zeros of the denominator polynomial $Q(z)$,

$$Q(z) = C(z^{-1})C(z) + \beta\,B(z^{-1})B(z). \tag{7}$$

Note that $Q(z)$ has the form of an auto-correlation function since it is symmetric in z and z^{-1}.

B. Two single-channel examples

In order to illustrate how the regularisation modifies the pole-zero structure of $H(z)$, we consider the simple system

$$C(z) = 1 - z^{-1} \tag{8}$$

$C(z)$ has a single zero on the unit circle at $z = 1$, and so its magnitude response is zero at DC. If we set $A(z) = 1$ (flat target response with zero phase), $m = 0$ (no modeling delay), and $B(z) = 1$ (frequency-independent regularisation), then from Eq. 3 we find

$$H(z) = \frac{1 - z^{-1}}{(1 - z)(1 - z^{-1}) + \beta} \tag{9}$$

$H(z)$ has a single zero at $z = 1$, and two poles, also on the real axis, at

$$z = 1 \pm \frac{\sqrt{\beta^2 + 4\beta}}{2} + \frac{\beta}{2} \qquad (10)$$

When $\beta << 1$, the power series expansion of this expression is given by

$$z = 1 \pm \sqrt{\beta} + O(\beta) \qquad (11)$$

which shows that for small values of β, the distance from the two poles of $H(z)$ to the unit circle is proportional to the square root of β. For example, if $\beta = 0.0001$, then the two poles of $H(z)$ are on the real axis at 1.01 and 0.99. According to Eq. 4, this corresponds to a time constant α of approximately 100 samples since both poles are a distance of approximately 0.01 away from the unit circle.

It is seen that the zero of $H(z)$ is at the position of the singularity, $z = 1$. This illustrates the general principle of regularisation: if the system is very ill-conditioned at a particular frequency, $H(z)$ wont attempt to deconvolve the system at this frequency because the effort penalty far outweighs the performance error. We now consider how frequency-dependent regularization modifies the pole-zero structure of a slightly more complex system.

Fig. 2 shows the properties of the sequence $c(n) = \{1,0,0,0,0.96\}$. This filter has been constructed from its four zeros at 0.99, $\pm 0.99i$, and -0.99. Note that the zeros are evenly spaced around the unit circle at a distance of 0.01 away from it. Fig. 2a shows the modulii of the zeros of $C(z)$ plotted against their arguments in radians (note the scaling of the y-axis), and Fig. 2b shows the magnitude response $|C|$ of $C(z)$.

Fig. 3 shows the pole-zero map of the filter $H(z)$ calculated from Eq. 3 when $\beta = 0.003$ and

$A(z) = 1$. In Fig. 3a, the shape factor is constant as a function of frequency which means that $B(z) = 1$. In Fig. 3b, the shape factor is a first order high-pass FIR filter, $B(z) = 1 - z^{-1}$ (see Eq. 8), whose single zero is on the unit circle at 1. The circles are the zeros, given by the zeros of $P(z)$ (see Eq. 6), and the crosses are the poles, given by the zeros of $Q(z)$ (see Eq. 7). If β was zero, such a plot would show half the poles being cancelled out exactly by zeros. The positions of the surviving poles would correspond exactly to the zeros of $P(z)$ (see Fig. 2a). When β is increased, however, the poles move away from the unit circle in the regions where $B(z)$ contains energy. In Fig. 3a, $B(z)$ is an all-pass filter, and so all the poles have moved. In Fig. 3b, $B(z)$ is a high-pass filter, and so the poles bend away from the unit circle at high frequencies (at arguments near $\pm\pi$) whereas the pole just outside the unit circle at zero radians is cancelled by a zero of $P(z)$ because the regularisation does not have any effect at low frequencies.

Fig. 4 shows the magnitude response $|H|$ of $H(z)$ calculated with frequency-dependent regularisation (solid line) and with no regularisation (dashed line). Thus, the solid line in Fig. 4a corresponds to the frequency response of the signal whose pole-zero map is plotted in Fig. 3a whereas the solid line in Fig. 4b corresponds to the frequency response of the signal whose pole-zero map is shown in Fig. 3b. It is seen that the frequency-dependent regularization has succeeded in attenuating high frequencies without affecting low frequencies.

C. General results; the multichannel case

Just as in the single-channel case, we start by writing the z-transform of the ideal filter matrix $\mathbf{H}(z)$ (see Eq. 2) in a form which is equivalent to a fraction,

$$\mathbf{H}(z) = \mathbf{Q}^{-1}(z)\mathbf{P}(z). \tag{12}$$

Here, the numerator *P*(*z*) is an *S*×*T* matrix of *z*-transforms given by

$$\mathbf{P}(z) = \mathbf{C}^{\mathrm{T}}(z^{-1})\mathbf{A}(z)\, z^{-m}, \tag{14}$$

and the denominator *Q*(*z*) is a square *S*×*S* matrix of *z*-transforms given by

$$\mathbf{Q}(z) = \mathbf{C}^{\mathrm{T}}(z^{-1})\mathbf{C}(z) + \beta\, B(z^{-1})B(z)\mathbf{I}. \tag{13}$$

We can invert $\mathbf{Q}(z)$ formally by dividing its adjoint adj[$\mathbf{Q}(z)$], which is a matrix whose dimensions are the same as those of $\mathbf{Q}(z)$, by its determinant *Q*(*z*) [8, Section 0.8.2],

$$\mathbf{H}(z) = \frac{\text{adj}\big[\mathbf{Q}(z)\big]}{Q(z)}\mathbf{P}(z) \tag{15}$$

It is seen that the determinant of $\mathbf{Q}(z)$ is a common denominator of all the elements of $\mathbf{H}(z)$. This implies that the elements of $\mathbf{H}(z)$ share a common set of poles given by the zeros of the polynomial *Q*(*z*), and in addition that those poles are not related in a simple way to the zeros of the elements of $\mathbf{Q}(z)$. Consider, for example, the two-by-two system

$$\mathbf{C}(z) = \begin{bmatrix} 1 & 0.5z^{-1} \\ 0.5z^{-1} & 1 \end{bmatrix}. \tag{16}$$

If β is zero and $\mathbf{A}(z)$ is an identity-matrix of order 2, then

$$\mathbf{Q}(z) = \begin{bmatrix} 1.25 & 0.5(z+z^{-1}) \\ 0.5(z+z^{-1}) & 1.25 \end{bmatrix}. \tag{17}$$

The two off-diagonal elements of $\mathbf{Q}(z)$ each have two zeros on the imaginary axis at +i and -i

whereas the diagonal elements do not have any finite zeros at all. Nevertheless, $Q(z)$ has two zeros on the real axis at ± 0.5 and another two at ± 2, and none of these coincides with the zeros of any of the individual elements of $\mathbf{Q}(z)$, or $\mathbf{C}(z)$. Unfortunately, this means that even if the individual elements of $\mathbf{C}(z)$ are well-conditioned there is no guarantee that $Q(z)$ will be well-conditioned. Cross-talk cancellation systems, for example, are generally ill-conditioned at low frequencies because the transfer functions contained in $\mathbf{C}(z)$ are very similar when the acoustic wavelength is long, and so one ends up essentially having to invert a 2-by-2 matrix that contains all ones [3].

D. The roots of the denominator Q(z)

It is no trivial matter to find the roots of a polynomial of high order, and there is an overwhelming amount of literature available on the subject (see [9] for a very comprehensive bibliography). As a rule of thumb, a polynomial whose degree is less than 500 can easily be factored on a fast PC. There are few general results that are useful to us but the following three rules give some idea about what to expect in practice. The rules concern symmetry, clustering, and attractors.

First, the roots of $Q(z)$ always appear in groups of two or four. Since $Q(z)$ is equal to $Q(z^{-1})$ it follows that if z_0 is a zero of $Q(z)$ then so is $1/z_0$. Furthermore, zeros off the real axis must appear in complex conjugate pairs since the coefficients of $Q(z)$ are real. Note that the symmetry in z and z^{-1} means that for each zero inside the unit circle, there is a corresponding zero outside, and this effectively spoils our chances of implementing $\mathbf{H}(z)$ as a stable causal filter matrix (apart from in a few special cases).

Secondly, the roots of a polynomial of high order are not scattered all over the complex plane,

but rather "...the roots of a random polynomial tend to be evenly distributed in angle and tightly clustered near the unit circle as the degree of the polynomial increases" (quote from [10]). Since most signals we come across in practice, such as measured impulse responses, have a certain degree of randomness built into them, this asymptotic result accurately describes what is most often observed with experimental data.

Thirdly, the roots of $B(z^{-1})B(z)$ act as a kind of attractor set for the roots of $Q(z)$ for large values of β. It is easily verified that when β is very small, the roots of $Q(z)$ are those of the determinant of $\mathbf{C}(z^{-1})\mathbf{C}(z)$ whereas when β is very large, the roots of $Q(z)$ are those of $B(z^{-1})B(z)$. Consequently, for some choices of $B(z)$ it is possible that excessive use of regularisation can cause some of the poles to be pulled back towards the unit circle. Note that when frequency-independent regularisation is used, $B(z)$ is a constant and consequently it has no roots. The implication of this is that for large β the attractors are zero (the origin) and complex infinity (points very far away from the origin), and so in this case the regularisation will push all the poles away from the unit circle.

V. CONCLUSIONS

The degree of ill-conditioning of a single- or multichannel deconvolution problem can be assessed by mapping out the zeros of the determinant of a matrix of z-transforms. Those zeros become the poles of a matrix of ideal filters, and any poles close to the unit circle will result in sharp peaks in the magnitude responses of the ideal filters. The sharpness of those peaks can be reduced selectively by using frequency-dependent regularisation which works by pushing some of the poles away from the unit circle. However, any regularisation modifies the pole-zero structure of the ideal filters in such a way that strictly speaking they cannot be

both causal and stable, and so a modeling delay must be used in the design of a realisable filter matrix.

It is straightforward to relax the constraints on the target matrix $\mathbf{A}(z)$, the shape factor $B(z)$, and the plant transfer functions $\mathbf{C}(z)$ so that they can have infinite impulse responses. For example, if $B(z)$ is written $B_{\text{FIR}}(z)/B_{\text{IIR}}(z)$, fomal manipulations of the z-transforms still allow a pole-zero analysis of $\mathbf{H}(z)$.

REFERENCES

[1]   M. Karjalainen, E. Piirilä, and J. Huopaniemi, "Comparison of loudspeaker equalization methods based on DSP techniques", *J. Audio Eng. Soc.*, Vol. 47, No. 1/2, pp. 14-31, 1999

[2]   O. Kirkeby, P.A. Nelson, H. Hamada, and F. Orduna-Bustamante, "Fast deconvolution of multichannel systems using regularization", *IEEE Trans. Speech Audio Processing*, Vol. 6, pp. 189-194, 1998

[3]   O. Kirkeby, P.A. Nelson, P. Rubak, and A. Farina, "Design of cross-talk cancellation networks by using fast deconvolution", presented at the 106th *Audio Eng. Soc. Convention* in Munich, May 8-11, 1999. AES preprint 4916 (J1). Submitted to *J. Audio Eng. Soc.*

[4]   W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992

[5]   R.A. Roberts and C.T. Mullis, *Digital Signal Processing*, Addison-Wesley, 1987

[6]   M. Bellanger, *Digital Processing of Signals*, 2nd ed. John Wiley & Sons, 1989

[7]   A.V. Oppenheim and R.W. Schafer, *Digital signal processing*. Englewood Cliffs, NJ: Prentice-Hall, 1975

[8]   R.A. Horn and C.R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985

[9]   J.M. McNamee, "A bibliography on roots of polynomials", *J. Comput. Appl. Math.*, vol. 47, pp. 391-394, 1993. New version available on the World Wide Web at http://www.elsevier.nl/locate/cam

[10]   K. Steiglitz and B. Dickinson, "Phase unwrapping by factorization", *IEEE Trans. Acoustics Speech Signal Processing*, Vol. ASSP-30, pp. 984-991, December 1982

FIGURE CAPTIONS

Fig. 1   The discrete-time deconvolution problem in block diagram form

Fig. 2   The properties of a sequence whose z-transform is $C(z)=1-0.96z^{-4}$. a) The zeros of $C(z)$ in the complex plane, and b) its magnitude response $|C|$. Fig. 2a is a close-up of a thin strip that covers the unit circle

Fig. 3   The positions of the poles (crosses) and zeros (circles) in the complex plane of the ideal inverse $H(z)$ of $C(z)$ whose properties are shown in Fig. 2. $H(z)$ is calculated with a) frequency-independent regularisation, and b) frequency-dependent regularisation

Fig. 4   The magnitude response $|H|$ corresponding to the pole-zero maps plotted in Fig. 3. $H(z)$ is calculated with a) frequency-independent regularisation, and b) frequency-dependent regularisation. The dashed lines show $|H|$ calculated with no regularisation
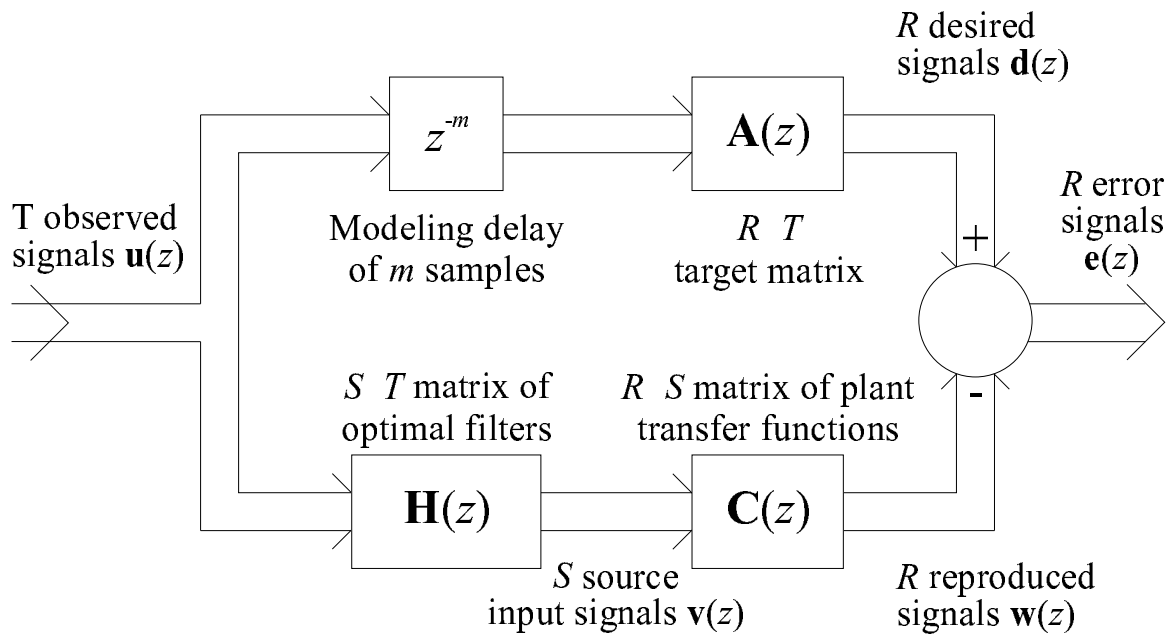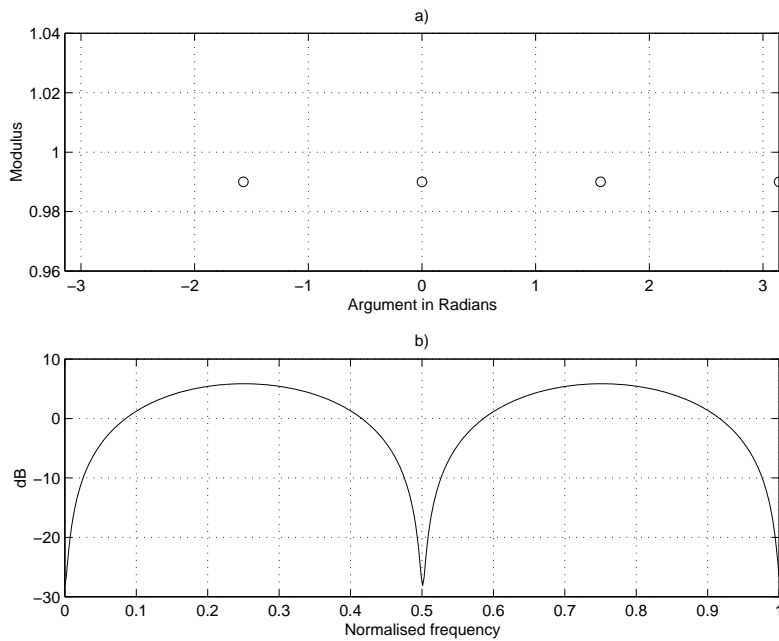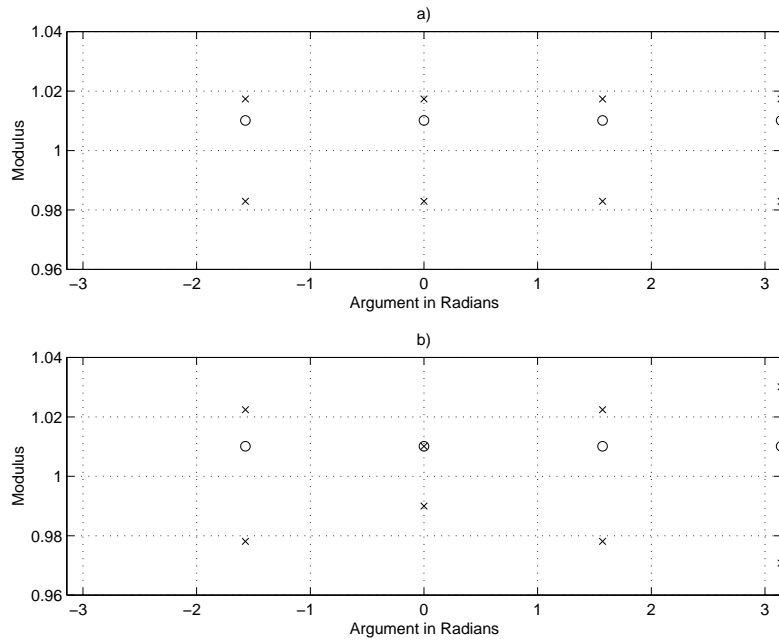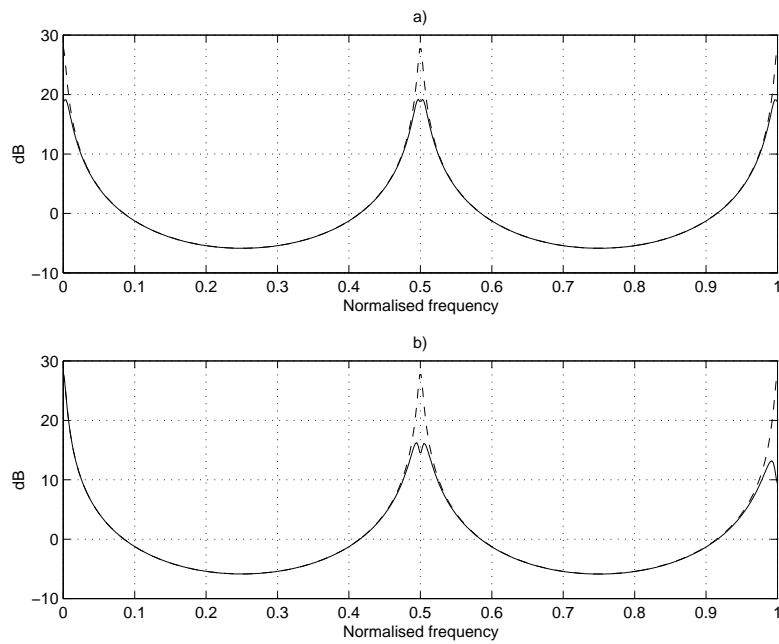
Fig. 1

Fig. 2

Fig. 3

Fig. 4