

University of Southampton Research Repository  
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

**UNIVERSITY OF SOUTHAMPTON  
FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS  
INSTITUTE OF SOUND AND VIBRATION RESEARCH**

**Inverse Filtering for Virtual Acoustic Imaging Systems**

by

**Timoleon Papadopoulos**

Thesis for the degree of Doctor of Philosophy

May 2006

*Στον Πάνο Γεραμάνη  
που με συντρόφεψε όλα αντά τα χρόνια  
αλλά έφυγε πριν προλάβω να τον συναντήσω.*

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

INSTITUTE OF SOUND AND VIBRATION RESEARCH

Doctor of Philosophy

INVERSE FILTERING FOR VIRTUAL ACOUSTIC IMAGING SYSTEMS

by Timoleon Papadopoulos

The research topic of this thesis is the use of inverse filtering for the design and implementation of two-channel virtual acoustic imaging systems that utilise loudspeakers. The basic objective of such systems is to invert the electroacoustic plant between the input to the loudspeakers and the output at the listener's ears and hence make it possible for a pair of binaural signals to be locally reproduced at the position of the listener's ears. As a starting point for the research presented, a previously introduced type of inverse filtering design is considered in which the inverse is implemented with FIR filters. The basic formulation of this design is described and a number of innovative points regarding its implementation are made. An experimental procedure is then formulated for the evaluation of the effectiveness of this inverse filtering design that is based on objective measurements of the inversion process. Unlike previously employed methods that are based on computer simulations or subjective experiments, the introduced experimental procedure is shown to be very efficient in isolating and exactly quantifying the effect on the accuracy of the inversion of a number of errors and approximations typically present in the implementation. A detailed evaluation is thus presented of the inverse filtering design at hand in realistic conditions of implementation. Subsequently, a novel method for the off-line implementation of the inverse filtering is presented that utilises recursive filters of lower order. In this method, the responses of the inverse filters are decomposed into two parts, one realisable in forward time and one in backward time. The effectiveness of this new method for the implementation of the inverse is tested and compared with a small selection of the objective evaluation results described above. Finally, an algorithm for the on-line implementation of the forward-backward inverse filtering is proposed and its computational cost is compared with the currently available frequency-domain block-processing filtering algorithms.

# ***Table of Contents***

Abstract .....	3
Table of Contents .....	4
Acknowledgements .....	7
1 Introduction .....	8
1.1 Overview of Virtual Acoustic Imaging Systems .....	9
1.1.1 The objective of Virtual Acoustic Imaging Systems .....	9
1.1.2 Stereophony .....	10
1.1.3 Global sound-field reconstruction .....	11
1.1.4 Ambisonics .....	12
1.1.5 Multi-channel extensions of stereophony .....	12
1.1.6 The binaural technique .....	13
1.2 The specific model discussed in the thesis – Local sound-field reproduction over loudspeakers .....	15
1.3 The contribution of the thesis and the organisation of the presented material ..	23
2 Inverse filtering in virtual acoustic imaging systems .....	27
2.1 Introduction .....	27
2.2 Review of inverse filtering designs .....	27
2.3 Computation of the inverse in the frequency-domain .....	32
2.3.1 Impulse response of a rational transfer function .....	32
2.3.2 Realisability of the inverse filter matrix .....	36
2.3.3 Regularisation .....	39
2.3.4 Modified method for the determination of the inverse in the frequency-domain .....	41
2.4 Notes on the implementation of the inverse .....	45
2.4.1 Distribution of the zeros of the measured electroacoustic transfer functions .....	45
2.4.2 Filtering through the common denominator .....	50
2.5 Conclusions .....	54
3 Objective evaluation of inverse filtering – Motivation and methodology .....	55
3.1 Introduction .....	55

3.2 Motivation for the undertaking of the chosen objective evaluation experimental procedure.....	55
3.3 Experimental arrangement.....	60
3.4 The measured model of the plant matrix .....	67
3.5 On the accuracy of the computer-simulation of the inversion process.....	70
3.6 Conclusions.....	74
4 Objective evaluation of inverse filtering – Inversion efficacy under ideal conditions .....	76
4.1 Introduction.....	76
4.2 Length of the model of the plant matrix .....	77
4.3 Regularisation .....	92
4.4 Length of the inverse when regularisation is applied .....	102
4.5 Comments on inversion results under ideal conditions .....	108
4.6 Conclusions.....	125
5 Objective evaluation of inverse filtering – Inversion efficacy under non-ideal conditions.....	128
5.1 Introduction.....	128
5.2 Presence of reflections in the plant .....	128
5.3 Listener position misalignment.....	143
5.4 Non-individualised plant model.....	160
5.5 Conclusions.....	170
6 Recursive inverse models – Forward-backward time filtering .....	173
6.1 Introduction.....	173
6.2 Convolution in backward time.....	174
6.3 Decomposition of the inverse into minimum and maximum phase parts.....	177
6.3.1 Parallel connection.....	177
6.3.2 Cascade connection.....	180
6.3.3 Computation of the decomposition determining the poles and residues of the inverse transfer function.....	181
6.3.4 Two-channel case.....	184
6.4 Numerical method for the decomposition of the inverse transfer function .....	188
6.4.1 Determination of $C_{\min}(z)$ and $C_{\max}(z)$ .....	189
6.4.2 Determination of $B_{\min}(z)$ and $B_{\max}(z)$ .....	194

6.5 Issues regarding the practical implementation of the method – Reduced order inverse models .....	199
6.6 Inversion results .....	202
6.7 Conclusions .....	215
7 On-line implementation of the forward-backward time inverse filtering .....	218
7.1 Introduction .....	218
7.2 On-line implementation of the inverse filtering algorithm – Filtering of the input in segments .....	218
7.3 Computational cost and input-output latency of the on-line algorithm .....	224
7.4 Conclusions .....	231
8 Conclusions – Proposed further research .....	233
Appendix .....	239
A1. Convolution and filtering in backward time .....	239
A2. Proof of the matrix identity of §6.2 .....	245
References .....	247

## ***Acknowledgements***

First and foremost, I would like to thank my supervisor, Prof. P. A. Nelson, for his invaluable guidance towards the completion of this thesis. I have benefited immensely not only from his unique knowledge and deep understanding of the subject matter of this thesis but equally from the truly generous and sincere manner in which he has always approached my aspirations or concerns.

My parents have always been a constant source of unassuming encouragement and support. Over the years I have come to realise that they have also been a true source of inspiration to me and I really thank them for that.

A big part of this thesis was written during one year I spent in Cyprus. I am really grateful to my parents-in-law, Zacharias and Dona, and to grandma Eleni for all the help, but most importantly for making me feel at home.

During my study at ISVR there is probably not a single member of staff, technician or student that I have not turned to at one time or another for help or advice and I would like to thank them all. I would particularly like to thank Antony Wood, Rob Stansbridge and Keith Sims for the time and effort they have spent on my various experimental projects, Maureen Strickland and Sue Brindle for all the solutions they have provided and Yiannis Andrianakis for the breaks and the discussions.

Elena and Orestis, thank you for making it all happen!

# ***1 Introduction***

Virtual Acoustic Imaging Systems are sound reproduction systems that allow us to control arbitrarily the auditory perception of one or more listeners located in a given space. The design of such systems has been a research topic of increasing interest over the last few decades. With a diverse range of applications including the fields of home entertainment, film production, personal computer multimedia, car audio, telecommunications, even aircraft cockpit human interface (see (Begault 1994) and references therein), the relevant research currently focuses on a number of different implementation strategies. These vary with respect to the underlying theoretical models of the acoustics and psychoacoustics, the type and number of acoustic sources and audio channels used, and the audio equipment deployed.

A specific type of implementation is considered in this thesis. This is based on a two-channel reproduction system using two full audio frequency range loudspeakers. As is described in the following, under certain assumptions, the creation of an arbitrarily chosen auditory impression in a single static listener using such a system, essentially amounts to the inversion of a  $2 \times 2$  matrix that contains the transfer functions relating the input to the sources to the output at the positions of the listener's ears. The inverse of such a transfer function matrix, albeit inherently recursive, is typically implemented in currently available realisations with non-recursive filters. Such an implementation incurs an excessively high computational cost requirement. The aim of this thesis is: (i) to provide a closer investigation of the feasibility and the effectiveness of this inversion and (ii) to propose an alternative realisation that makes use of the recursive character of the inversion, thus achieving a reduction in the required computational cost.

## 1.1 Overview of Virtual Acoustic Imaging Systems

### 1.1.1 The objective of Virtual Acoustic Imaging Systems

From everyday experience we know that, when in the presence of an existing sound-field, humans are able to extract a great amount of information regarding the sources that create this field. We are, in other words, based solely in the sense of hearing, able to create an acoustic image<sup>1</sup> of the physical phenomena that are responsible for the creation of the given sound-field. This acoustic image contains information relating to:

- The number, the type and the individual timbre of the sources that create the sound-field. When among a group of people in conversation, a listener can identify each particular speaker. Similarly, a trained listener can identify the contribution of each instrument when listening to an orchestra performing a musical piece.
- The type of acoustic signal emitted by the sources and the information that it carries, be it for example speech, music or a specifically identifiable kind of noise.
- The position, i.e. the direction and the distance, of the sources relative to the listener. When crossing the street, the noise from a car is usually enough to tell us where the car is and whether it is approaching or moving away.
- The characteristics of the space in which the sound-field exists. Within a reasonable margin of error, a listener can identify whether the sound that they hear exists in an open space, in a big cathedral, in a small, furnished room or in a concert hall.

The objective of a virtual acoustic imaging system is then to recreate convincingly the perception of this identical acoustic image, termed a virtual acoustic image, in a listener that is located in a space that does not contain the original physical

---

<sup>1</sup> In the relevant literature the term *acoustic image* is often used to denote the perception related to a single source of sound whereas in this instance we use it to refer to the total impression created to the listener by the soundfield that surrounds him/her (this impression that comprises a collection of individual acoustic images is sometimes referred to as the *soundstage*). The actual meaning of the term is self-evident from the context and here we use the term *acoustic image* loosely to describe both notions.

phenomena responsible for the creation of the acoustic image. Using the terminology introduced by Blauert (1996), the goal is to recreate the “auditory event” which in a natural listening situation would correspond to a given “sound event” without the presence of the “sound event” itself.

To take an often used example of such an application, we can imagine a listener who uses such a system in his/her living room in order to enjoy the full sensation of a musical performance given by a famous orchestra in a music hall of exceptional acoustics. Assuming that the stimuli of the other senses are perceptually irrelevant or somehow blocked, the design must be able to convey to the listener the musical text up to the finest detail supported by the orchestra’s performance and the room’s acoustics, as well as to create convincingly in him/her the image of the stage in all its dimensions with the positions of the instruments clearly defined (at least as well defined as they would be in the natural listening situation) and also to recreate the perception of the acoustics of the hall. The natural extension of the above principle is the ability to control arbitrarily the auditory perception of the listener, thus being able to create acoustic images that do not necessarily originate from an existing listening situation. An obvious area of application of such a design would be that of multimedia applications and video gaming.

### *1.1.2 Stereophony*

The first sound reproduction systems that aspired to achieve the aforementioned goal had their roots in A. D. Blumlein’s theoretical work<sup>2</sup> in the early 1930’s and first appeared as commercial implementations in the 1950’s under the name “stereophonic systems”. As the name suggests, the new aspect introduced by two-channel stereophony compared to the then known monophonic systems was the goal of creating a three-dimensional auditory impression with acoustic images positioned at any intended point in the space that surrounds the listener. As has been shown in its practically universal use by sound engineers and its widespread application in home-audio hi-fi systems, with proper use of the “stereo-panning” technique, stereophony does indeed succeed in controlling the positioning of images inside the angle spanned by the two loudspeakers. On the other hand, it shows little promise as regards the

---

<sup>2</sup> For Blumlein’s original British Patent document and for other works describing the early days of stereophony see (AES 1986).

creation of a full three-dimensional sound-stage. Images at locations outside the loudspeaker span, or at elevations other than the position of the loudspeakers appear erratically, if at all. The depth of the sound-stage (i.e. the source distance information) is similarly poorly controlled.

It should be noted that despite the existence of a comprehensive theoretical foundation for the stereophonic technique, the methods used in its everyday implementation are in many cases based on the practical experience and personal taste of the individual recording and producing engineer<sup>3</sup>. Similarly, at the reproduction end of stereophonic systems' implementation, even though a standard exists specifying that the two loudspeakers must be placed at the corners of an equilateral triangle facing a listener positioned at the third corner of the triangle, such practice is not always followed. Instead, it is usually left to the user to find the placement that gives the best results for the specific system and room. Hence it is difficult to undertake an exact overall evaluation of the effectiveness of stereophonic reproduction in controlling the spatial aspect of the auditory experience of the listener.

During this “era of stereophony” of the last 50 years, numerous theoretical formulations and experimental designs have been proposed in order to address the inadequacy of stereo systems in fully reproducing the spatial characteristics of the recorded sound-field. Despite the fact that none of them has yet succeeded in replacing stereophony as the mainstream audio reproduction design<sup>4</sup>, some of them have shown real promise in their application. A description of those designs follows which includes the specific design investigated in this thesis.

### *1.1.3 Global sound-field reconstruction*

The most ambitious among the related theoretical formulations is unquestionably that of global sound-field capture and reconstruction, which is based on the Kirchhoff-Helmholtz integral equation. In simple terms this states that the sound-field inside the whole volume of a given enclosure that does not contain sound sources is uniquely described by the values of the pressure and the pressure gradient on a surface that

---

<sup>3</sup> For a review of stereo recording techniques and of the use of the “stereo panning” method see (Borwick 1990).

<sup>4</sup> This is with the possible exception of the recently introduced multi-channel designs described in §1.1.5.

encloses the volume. The sound-field inside the volume can thus be controlled by a continuous distribution of monopole and dipole sources on the surface. The practical aspects of a sound recording and reproduction system that could realise this principle have been investigated by Berkhout (1988), Berkhout et al. (1993) and Nelson (1994). Although it is recognised that the required scale for the implementation of the full audio range, three-dimensional formulation is excessively large more recent theoretical work has been published which shows that the design can be conveniently downscaled to two dimensions and thus become applicable to problems such as those of direct sound reinforcement and hall acoustics auralisation (Hulsebos et al. 2002), (Start et al. 1999) and (de Vries 1996).

#### *1.1.4 Ambisonics*

Another approach, equally attractive in its theoretical formulation but much less demanding in its practical implementation was presented by Gerzon (1973) in the early 70's under the name “ambisonics”. In the core of this sound recording and reproduction technique lies a specific microphone design (Gerzon 1975), commonly referred to as the “soundfield” microphone after its commercial realisation by Calrec Audio. This microphone allows the recording of four signals corresponding to the sound pressure and the pressure gradients on the three Cartesian axes at one point in the original soundfield (recording in B-Format). This information can then be decoded to drive an arbitrary number of loudspeaker channels (typically four or more) so as to reconstruct a soundfield with the same directional characteristics as that of the original soundfield. Even though the various commercial implementations of this principle never succeeded in establishing a critical mass of installations, mainly due to the severe restraints in multi-channel signal storing and transmission capabilities of the time, the system seems to enjoy a wide appreciation by users and researchers alike and it is widely believed to represent a serious possibility for future commercial implementations.

#### *1.1.5 Multi-channel extensions of stereophony*

Somewhat in contrast to the formulations described above, there also exists a category of implementations that are much simpler in their theoretical justification. These essentially amount to the expansion of the two-channel stereo and the associated

stereo-panning method to more channels and sources. The first system to appear in this category in the 1960's was the “quadraphonic” system, a 4-channel implementation which, despite the original enthusiasm it created, was soon abandoned. This was partly due to the inability of the analogue storage and distribution media of the time to efficiently handle 4-channel signals (the same problem as with ambisonics) but more importantly due to the limited capability of the system to create a genuinely surround sound perception that soon became apparent. Later on, a series of film-sound reproduction systems, that utilise similar technology, were developed by Dolby Laboratories first for cinema installations (Dolby Stereo) and then for home installations (Dolby Surround).

With the passing of commercial sound reproduction systems from the analogue to the digital domain in the 1980's and 1990's, that practically lifted all restraints in the storage and distribution of multi-channel sound signals, this “augmented stereo” principle gave rise to a number of commercially successful multi-channel systems both for film-sound reproduction (5.1 Dolby Digital and DTS, 6.1 or 7.1 Dolby Digital EX and DTS-ES etc.) and for music reproduction (DVD-Audio and SACD formats)<sup>5</sup>. With the inclusion of more channels and sources (5 or 6 full audio range channels) these systems adequately cover the positioning of sources in the horizontal plane. One of their basic shortcomings however is the inability to adequately control source height information. Even though the future development of these systems promises the inclusion of one or more extra channels dedicated to source height information, such a possibility is of course limited by the number of sources that can be practically accommodated in a standard installation.

### *1.1.6 The binaural technique*

Finally, an altogether different approach to the problem at hand, the approach considered in this thesis, is that of binaural reproduction (Møller 1992). The starting point of this approach lies in the fact that the auditory perception of the listener is solely determined by the pair of sound pressure signals at his/her eardrums. Hence, if this pair of sound pressure signals (binaural signals) that corresponds to a given sound

---

<sup>5</sup> For details on the various sound reproduction formats one can refer to the abundance of relative information on the Internet, see e.g. [http://www.dolbylaboratories.com/consumer/technology/tech\\_overview.html](http://www.dolbylaboratories.com/consumer/technology/tech_overview.html) or <http://www.digit-life.com/articles2/sacd-dvd-a/>

event is known and is locally reproduced at the position of the listener’s ears, then the listener will perceive the corresponding auditory event without the presence of the sound event itself. Evidently, the two basic issues arising in the context of the binaural technique are (i) the determination of the binaural signals corresponding to the given sound event and (ii) the accurate reproduction of these signals at the ears of the listener. The former of these issues is discussed in the next section. For the latter, two basic types of implementation exist; the first is reproduction over headphones (for which the term “binaural reproduction” is commonly used) and the second is reproduction over loudspeakers, for which the term “transaural reproduction” was coined by Cooper and Bauck (1989).

The principle advantage in using headphones is the guaranteed suppression of the cross-talk, i.e. of the fact that each one of the binaural signals only reaches the intended ear. Furthermore, the use of headphones alleviates the design task from the influence of the room acoustics in the reproduction space. Thus, the only processing required for the reproduction is the equalisation of the response between the headphones and the eardrums, something that is achieved rather easily (Møller 1992). On these grounds the headphone-reproduction approach indeed offers an efficient implementation option. It does however pose a number of challenges, the most prominent being those of the avoidance of the in-head localisation effect, of the robust control of images directly in front of the listener, of the avoidance of front-back confusions and of the achievement of a sense of natural hearing. It is generally agreed that such problems can be successfully addressed if the binaural material is determined on the basis of measurements made on the specific listener<sup>6</sup>. Of course such a requirement adds an extra parameter of complexity to the implementation which to a degree negates the simplicity mentioned above.

On the other hand, the main design challenge in systems where the binaural signals are delivered over loudspeakers is the suppression of the cross-talk in the contralateral loudspeaker-ear transmission paths and also the equalisation of the ipsilateral paths which contain the response of the loudspeakers, the room and the listener’s head and

---

<sup>6</sup> For further details on headphone reproduction issues and the related issue of the use of individualised and non-individualised HRTFs see (Asano et al. 1990), (Begault and Wenzel 1993), (Hartmann and Wittenberg 1996), (Kulkarni and Colburn 1998), (Kulkarni et al. 1999), (Martens 1997), (Middlebrooks et al. 1989), (Møller 1989), (Møller et al. 1995a), (Morimoto and Ando 1980), (Wenzel et al. 1993), (Wightman and Kistler 1989a) and (Wightman and Kistler 1989b).

body for the given reproduction geometry. This adds one further design and implementation stage in comparison to the case of reproduction over headphones, something that at first sight might be considered undesirable. Nevertheless there are a number of reasons that render the use of loudspeakers equally, if not more, attractive.

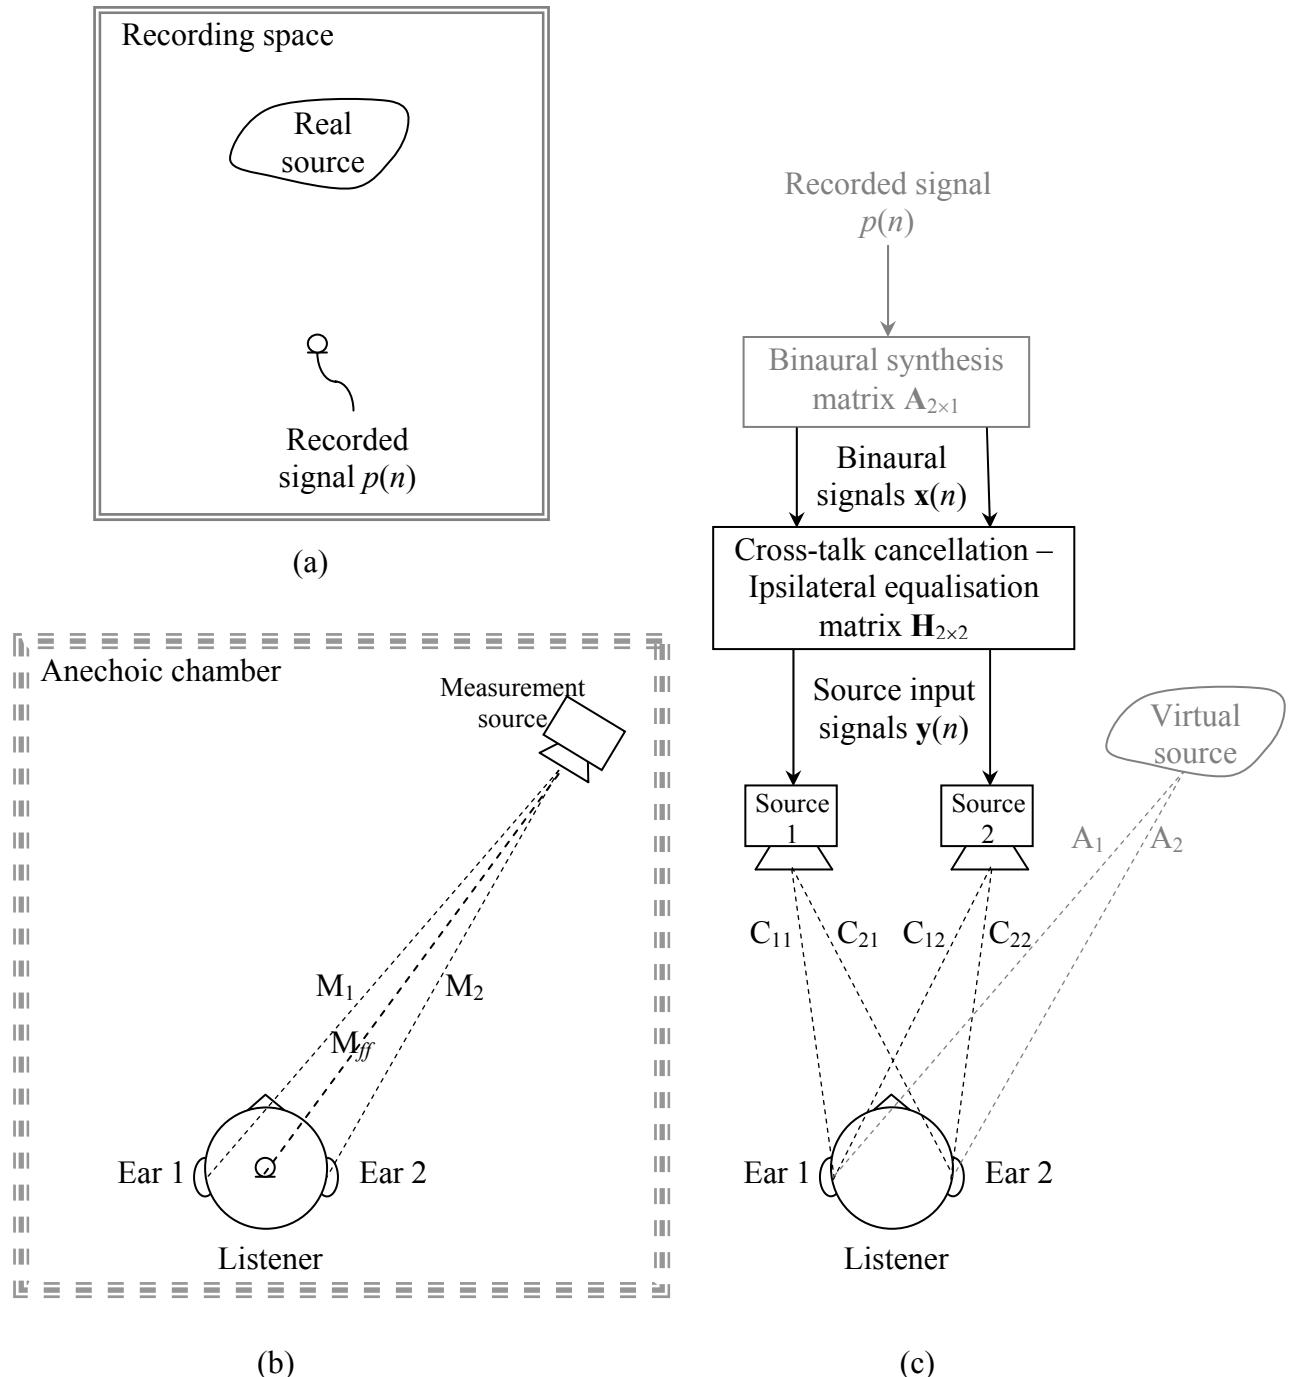
First, the listening experience when loudspeakers are used is effortlessly more natural than that achieved with headphones. Even when the cross-talk cancellation and ipsilateral equalisation fails to meet the required standard in order to create a completely convincing virtual acoustic environment, the auditory display degrades very gradually with the virtual images being concentrated on the front half of the horizontal plane or at worse collapsing to the location of the sources. As long as reasonable care is taken, the timbral characteristics of the display can be maintained at an acceptable quality. Conversely, imperfections in the design of headphone implementations can automatically lead to very unnatural and even discomforting auditory displays. This, in conjunction with the headphone-related “listener fatigue” effect, can prove completely prohibitive in applications such as multimedia applications and video gaming where long-time exposure is customary.

Furthermore, the technology and know-how gained over the years in the design and production of loudspeakers is greater than that of headphones. Loudspeaker unit/cabinets are generally considered to be better sources of sound in terms of their ability to cover the whole audio frequency range with good frequency response and adequate dynamic range. In this respect there is certainly a greater industrial/commercial interest in designs that use loudspeakers than in designs using headphones.

## ***1.2 The specific model discussed in the thesis – Local sound-field reproduction over loudspeakers***

In this section we discuss the two basic points mentioned above in the context of the binaural reproduction technique, namely that of obtaining the binaural signals and that of delivering them at the listener’s ears over loudspeakers. We give the basic analytical formulation of the inverse filtering problem arising from the need to

suppress the cross-talk and equalise the ipsilateral paths between the loudspeaker inputs and the output at the listener's ears.



**Figure 1-1:** The sound reproduction model discussed in the thesis. (a) Single-channel recording of a real source. (b) Measurement of HRTFs. (c) Reproduction of a pair of binaural signals over loudspeakers (black line) and creation of a single virtual acoustic image (grey line).

The model of virtual acoustics imaging systems discussed in this thesis is described in figure 1-1. Two cases are described: In the first, a pair of binaural signals is available (e.g. in the form of a binaural recording with microphones inserted at the ears of the listener or of a dummy-head) and it is to be reproduced using two loudspeakers at the ears of a listener positioned in a different physical space. This is illustrated in part (c) of figure 1-1 with the parts plotted in black line. In the second case, a monophonic recording signal is to be reproduced over the loudspeakers in order to create a virtual image at a specified point of the listener’s auditory space. This case is discussed immediately afterwards.

We denote with  $\mathbf{x}(n)$  the  $2 \times 1$  vector containing the pair of binaural signals  $[x_1(n) \ x_2(n)]$ , with  $\mathbf{y}(n)$  the  $2 \times 1$  vector containing the pair of source input signals  $[y_1(n) \ y_2(n)]$ , with  $\mathbf{d}(n)$  the pair of desired signals  $[d_1(n) \ d_2(n)]$  that we want to reproduce at the listener’s ears, with  $\hat{\mathbf{d}}(n)$  the  $2 \times 1$  vector containing the pair of signals  $[\hat{d}_1(n) \ \hat{d}_2(n)]$  that are actually reproduced at the listener’s ears and with  $\mathbf{e}(n)$  the pair of error signals corresponding to the difference between the desired and the reproduced signals. Furthermore, with  $\mathbf{H}(z)$  we denote the  $2 \times 2$  matrix of transfer functions that we use in order to create the necessary source input signals and with  $\mathbf{C}(z)$  the  $2 \times 2$  plant matrix that contains the transfer functions  $C_{ij}(z)$  relating the input to the  $j$ -th source to the pressure at the ear denoted with the index  $i$ . With this notation the schematic of figure 1-1 can be readily seen to translate to the block diagram of figure 1-2 (again the case where binaural signals are already available is shown with the black-line part of the figure).

Evidently, for the case of the reproduction of binaural signals over loudspeakers, the vector of desired signals  $\hat{\mathbf{d}}(n)$  is identical to the vector of binaural signals  $\mathbf{x}(n)$ . In that case then, and using then the  $z$ -transforms of the related signals and systems, the error signal becomes equal to:

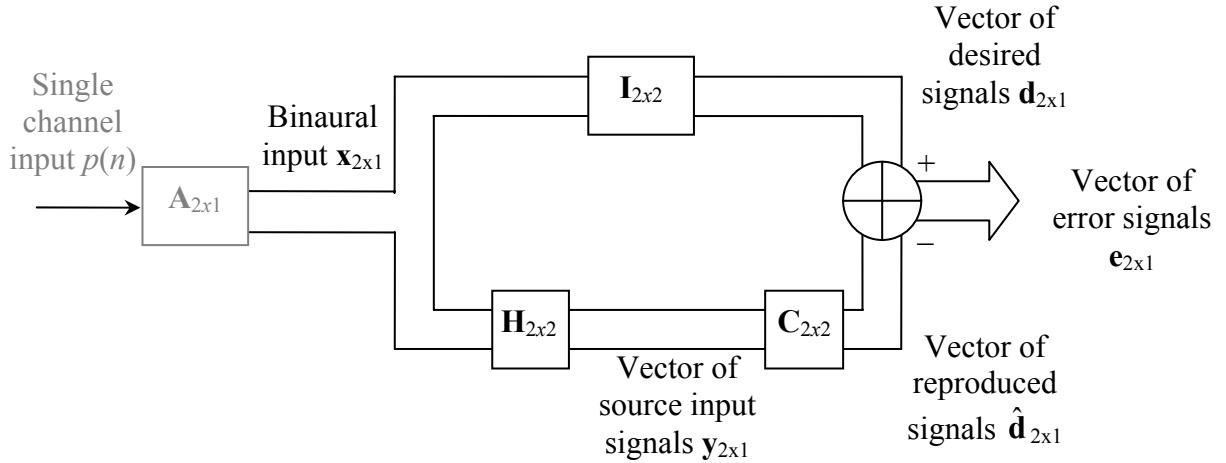
$$\mathbf{e}(z) = \mathbf{d}(z) - \hat{\mathbf{d}}(z) = \mathbf{x}(z) - \hat{\mathbf{d}}(z) \quad (1-1)$$

or equivalently

$$\mathbf{e}(z) = \mathbf{x}(z) - \mathbf{C}(z)\mathbf{H}(z)\mathbf{x}(z) \quad (1-2)$$

and with  $\mathbf{I}$  used to denote the  $2 \times 2$  identity matrix

$$\mathbf{e}(z) = [\mathbf{I} - \mathbf{C}(z)\mathbf{H}(z)]\mathbf{x}(z) \quad (1-3)$$



**Figure 1-2:** Block diagram of the sound reproduction model discussed in the thesis

Consequently, assuming that the binaural input signals  $\mathbf{x}(z)$  are “spectrally rich” across the audio frequency range, the ideal  $2 \times 2$  matrix  $\mathbf{H}$  that drives the error signal  $\mathbf{e}(z)$  to zero will be given by

$$\mathbf{C}(z)\mathbf{H}(z) = \mathbf{I} \quad \text{or equivalently} \quad \mathbf{H}(z) = [\mathbf{C}(z)]^{-1} \quad (1-4)$$

which in the time-domain becomes

$$\mathbf{c}(n) * \mathbf{h}(n) = \delta(n) \cdot \mathbf{I} \quad (1-5)$$

Henceforward we denote the matrix product  $\mathbf{C}(z)\mathbf{H}(z)$  that describes the equalisation and crosstalk cancellation with  $\mathbf{X}_{eq}(z)$ .

$$\mathbf{X}_{eq}(z) = \mathbf{C}(z)\mathbf{H}(z) = \begin{bmatrix} C_{11}H_{11} + C_{12}H_{21} & C_{11}H_{12} + C_{12}H_{22} \\ C_{21}H_{11} + C_{22}H_{21} & C_{21}H_{12} + C_{22}H_{22} \end{bmatrix} \quad (1-6)$$

or equivalently in the time-domain

$$\mathbf{x}_{eq}(n) = \mathbf{c}(n) * \mathbf{h}(n) = \begin{bmatrix} c_{11} * h_{11} + c_{12} * h_{21} & c_{11} * h_{12} + c_{12} * h_{22} \\ c_{21} * h_{11} + c_{22} * h_{21} & c_{21} * h_{12} + c_{22} * h_{22} \end{bmatrix} \quad (1-7)$$

Equations (1-4) and (1-5) form the analytical basis of the implementations considered in this thesis.

In the second case, that of positioning a virtual image at a given point in space, the formulation described above has to be extended to cover the creation of the binaural signals which are now not available (*binaural synthesis*). This is illustrated with the grey line parts in figures 1-1 and 1-2. The notion of the Head Related Transfer Function (HRTF) plays a central role in this process. This describes the linear transformation that is imposed on the free-field soundfield which is created by a source emitting sound from a given direction, by the presence of the listener's head and torso. It is defined (see eq. (1-8)) as the ratio of the Fourier transform of the sound pressure signal  $P_i(\omega, \varphi, \theta, r)$  at a given point along the listener's ear-canal (where the index  $i$  denotes the left or right ear of the listener) created by a source that is located at an azimuth angle  $\varphi$ , elevation angle  $\theta$  and distance  $r$ , to the Fourier transform of the free-field sound pressure signal  $P_{ff}(\omega, \varphi, \theta, r)$  that would be created by the same source at the position of the centre of the listener's head with the listener not present.

$$HRFT_l(\omega, \varphi, \theta, r) = \frac{P_l(\omega, \varphi, \theta, r)}{P_{ff}(\omega, \varphi, \theta, r)} \quad HRFT_r(\omega, \varphi, \theta, r) = \frac{P_r(\omega, \varphi, \theta, r)}{P_{ff}(\omega, \varphi, \theta, r)} \quad (1-8)$$

We note that the definition given here corresponds to the *free-field equalised* HRTF while other definitions have also been proposed. For a detailed discussion on the issue see (Blauert 1996), (Gardner 1998) (where the term *measurement equalised* HRTF is used for the definition given here) and (Møller 1992). It should also be noted that, despite the fact that the above general model describes the HRTF as a function of distance  $r$ , the usual practise is to assume that the source is located far enough from the listener so that the sound wave reaches the listener as a plane wave. In that case the dependence on distance simplifies to a scaling following the inverse square law. Hence, in all published HRTF databases, the source distance is kept constant (typically in the range of 1-2m) across the different values of azimuth and elevation

and does not appear as a variable. Finally, given the spherical symmetry of the formulation as regards the free-field pressure measurement, it is easily conceived that the free-field sound pressure signal  $P_{ff}$  should be considered independent of the angle of incidence and hence the variables  $\varphi$  and  $\theta$  become redundant. In this case the formulae of (1-8) become:

$$HRFT_l(\omega, \varphi, \theta) = \frac{P_l(\omega, \varphi, \theta)}{P_{ff}(\omega)} \quad HRFT_r(\omega, \varphi, \theta) = \frac{P_r(\omega, \varphi, \theta)}{P_{ff}(\omega)} \quad (1-9)$$

The basic arrangement of an HRTF measurement is depicted in part (b) of figure 1-1. A measurement loudspeaker fed with a stimulus signal<sup>7</sup> is positioned at a given direction with reference to the centre of the head of the listener (or a dummy-head) who has a pair of microphones placed at his/her ears<sup>8</sup>. The responses  $M_1(\omega, \varphi, \theta)$  and  $M_2(\omega, \varphi, \theta)$  from the input to the measurement loudspeaker to the output at the listener's ears are measured as well as the response  $M_{ff}(\omega)$  from the input to the loudspeaker to a (omni-directional) microphone positioned at the position of the centre of the listener's head with the listener absent. The responses  $M_i(\omega, \varphi, \theta)$  are measured on a grid sampling the sphere that surrounds the listener. The free-field response  $M_{ff}(\omega)$  has to be measured only once unless more than one measurement loudspeakers are used in which case it should be measured for each one of them. It then follows directly that the HRTF corresponding to the left and right ear (denoted henceforward with  $A_1(\omega, \varphi, \theta)$  and  $A_2(\omega, \varphi, \theta)$ ) will be as in equation (1-10). The measurement is typically made in an anechoic chamber or a room that is adequately sound-insulated and large enough so that the reflections due to the walls occur well

---

<sup>7</sup> Different measurement techniques exist for this type of electroacoustic plant response measurement. In all cases, a stimulus that is spectrally rich in the frequency range of interest is used to drive the source and the output signal is captured at the sensors. Subsequently, using a process that varies according to the type of the stimulus, the response of the stimulus is deconvolved from the output response to obtain the response of the system under test. For a detailed review see (Muller and Massarani 2001) and the references therein.

<sup>8</sup> The position along the ear-canal where the microphone is placed is a free parameter in this type of measurement. Ideally the microphone should be detecting the pressure output at the listener's eardrum. This is very easily achievable in dummy-head measurements but not quite trivial in human listener's measurements where it can only be achieved with probe microphones. Such transducers are typically of inferior specification compared to standard measurement-type microphones. This can be addressed however with the placement of the microphone higher-up along the ear-canal or even at the entrance of a blocked ear canal. This is because the contribution of the sound propagation along the ear-canal to the HRTF, despite being highly individualised across humans, has been shown by Hammershoi and Møller (1996), Mehrgardt and Mellert (1977), Middlebrooks et al. (1989) and Møller et al. (1995b) to be directionally independent.

after the measured impulse responses and can thus be removed by a “windowing” operation.

$$\begin{aligned} A_l(\omega, \varphi, \theta) &= HRFT_l(\omega, \varphi, \theta) = \frac{P_l(\omega, \varphi, \theta)}{P_{ff}(\omega)} = \frac{M_l(\omega, \varphi, \theta)}{M_{ff}(\omega)} \\ A_r(\omega, \varphi, \theta) &= HRFT_r(\omega, \varphi, \theta) = \frac{P_r(\omega, \varphi, \theta)}{P_{ff}(\omega)} = \frac{M_r(\omega, \varphi, \theta)}{M_{ff}(\omega)} \end{aligned} \quad (1-10)$$

The pair of left and right ear HRTFs for a given source location in relation to the listener, contains the information that is used by the listener’s auditory system in order to localise the position of the source. That is, the magnitude of the ratio of these two transfer functions, termed the Interaural Transfer Function (ITF), encodes the Interaural Level Difference (ILD) localisation cue that describes the head-shadowing effect and is known to be the basic localisation cue for sounds with frequency content above 1.5kHz. In the same manner, the phase delay of the ITF encodes the Interaural Time Delay (ITD)<sup>9</sup> localisation cue that has been shown to be the dominant source of information for the localisation of sounds with frequency content up to 1.5kHz. The time delays between the envelope of the signals at the two ears, encoded at the group delay of the ITF, have also been shown by Blauert (1996) to provide a localisation cue at frequencies above 1.5kHz.

Furthermore, the magnitude of each of the left/right HRTFs encodes information related to the monaural spectral cues that represent the basic source of information for the localisation in the median plane (Asano et al. 1990), (Blauert 1996) and (Wenzel et al. 1993). We note that the human auditory system is known to also use other cues that occur in situations where the source-listener geometry changes dynamically and which are thus, strictly speaking, not covered by the above time-invariant model of the HRTF. Dynamic cues that occur from movements of the listener’s head are primarily used in order to resolve situations where the binaural cues leave an ambiguity (front-back confusion and ambiguity about points on the cone of confusion).

Returning then to the process of binaural synthesis, we start from a single-channel signal that is available and which contains all the audible information of a given

---

<sup>9</sup> The term ITD is used in some sources to denote the frequency independent time delay between the HRTFs of the two ears (e.g. the time difference between the onset of the corresponding HRIRs) as opposed to the frequency dependent phase delay of the ITF (Blauert 1996).

sound event other than its spatial characteristics. In the most typical case this is a recording of the sound pressure  $P$  that is created at a given point  $M$  of a space, by a source  $S$  that is located at another point  $N$  of the same space as depicted in part (a) of figure 1-1. It is then readily seen from equation (1-9) that the sound pressures that would have been created at the ears of a listener positioned with the centre of his/her head at point  $M$  if the sound source  $S$  (virtual source) was emitting sound from a direction with a given azimuth and elevation, say  $(\varphi, \theta)$ , would be given by the product of the corresponding left and right HRTFs with the spectrum of the recorded sound pressure  $P(\omega)$ . This formulation, which is more commonly implemented in the time-domain as the convolution of the recorded signal  $p(n)$  with the time-domain version of the HRTFs, termed Head Related Impulse Responses (HRIRs), forms the basis of the binaural synthesis part of the design. This is illustrated in figures 1-1 and 1-2 with the parts plotted in grey lines<sup>10</sup>. Concluding, we see that making use of the linearity of the binaural synthesis process, this can be extended to give the binaural signals corresponding to a number of sources positioned at a number of different locations. In that case, the binaural signals will be obtained by the addition of the binaural signals corresponding to each one of the sources.

A few notes can be made regarding the model described above: First it should be clear that, unlike the binaural synthesis matrix  $\mathbf{A}$  the elements of which are HRTFs, the electroacoustic responses  $C_{ij}$  in the plant matrix  $\mathbf{C}$  are not HRTFs. That is to say, the elements of  $\mathbf{C}$  are not equalised with respect to a free-field response and so they incorporate the response of the reproduction loudspeakers and also the response of the room where the reproduction takes place. In an exact implementation of the model these components of the responses in the plant matrix also have to be corrected by the

---

<sup>10</sup> A point that should be made here is that this formulation of the binaural synthesis part of the design cannot, in principle, take fully into account the individual reverberation characteristics of the space where the recording of the original sound event is made. This is because these characteristics depend not only on the temporal but also the spatial distribution of the arrival of reflections, the latter being clearly lost in a single-channel pressure recording. In a similar sense the information related to the directivity characteristics of the source (which would be prominent in situations where the listener is moving in relation to the source) is not covered by the model. That is to say, this model can only produce the binaural signals that correspond to the same relative rotation between the virtual source and listener as that between the recording transducer and the real source.

inverse filtering process<sup>11</sup>. This approach has been adopted in the study presented in this thesis.

Another point to be made, that is evident from figure 1-2, is that in the second case described above, i.e. that of the reproduction of a single-channel recording for the creation of a single virtual acoustic image, the matrices  $\mathbf{H}$  and  $\mathbf{A}$  can be consolidated into a single  $2 \times 1$  vector  $\mathbf{H}_v = \mathbf{C}^{-1} \mathbf{A}$ . This effectively adds zeros to the recursive transfer functions of the inverse matrix  $\mathbf{H}$  (see also §2.3.1) making the corresponding impulse responses decay faster. Thus fewer and shorter FIR filters are needed for the implementation. The downside of this approach is that, when multiple or moving virtual images are to be created, the inverse filtering has to be applied to each source or position. Furthermore, the separation of the binaural synthesis from the inverse filtering stage facilitates the investigation of the effectiveness of each stage alone. This approach has been taken in the present thesis in which the object of interest is the inverse filtering stage.

Finally, following previous findings regarding the effectiveness of the inverse filtering relative to the reproduction geometry (see also the review of §2.2), we use the Stereo Dipole geometry ( $10^\circ$  total loudspeaker angle span) as the exemplifying case in the measured results and the simulations presented in the thesis.

### ***1.3 The contribution of the thesis and the organisation of the presented material***

The initial motivation for the undertaking of this research project was to investigate the possibility of realising the inverse matrix  $\mathbf{H}$  of equation (1-4) with recursive filters. As is further elaborated below, the filters comprising  $\mathbf{H}$  are by nature recursive in the sense that they correspond to rational transfer functions. Hence, in principle, one would expect such a recursive realisation to result in computational cost savings. However, the denominator polynomial of these transfer functions is typically non-

---

<sup>11</sup> Strictly speaking, the reproduced signals will also contain the response of the recording microphone in part (a) figure 1-1 and of the microphones used to measure the plant. But, presenting closer approximations to the ideal response, the microphone responses introduce less distortion if left uncorrected.

minimum phase when measured electroacoustic responses are involved in the determination of  $\mathbf{H}$ . In that case the direct use of the coefficients of the associated rational transfer functions does not result in realisable recursive filters.

An approach, novel to the area of inverse filtering for virtual acoustic imaging applications, was taken to address this issue. This new approach decomposes the related transfer functions in two parts; one realisable in forward time and one realisable in backward time. At a certain stage in the course of this research it became apparent that such an implementation is indeed possible. Nonetheless, its merit in terms of the associated computational efficiency could not be properly evaluated unless a closer look was taken into the relation between the design parameters of  $\mathbf{H}$  and the corresponding inversion accuracy that is realistically achievable in the real-world implementation of a virtual acoustic imaging system. Such an investigation was thus undertaken which turned out to be of great research interest in itself. This is also reflected in the extent of the presentation of the related material in this thesis. Upon completion of the investigation of this newly added research topic, the results related to the forward-backward time recursive filtering technique were augmented and the technique’s computational efficiency was evaluated with reference to this real-world-implementation context. The presented material is therefore organised as follows:

In Chapter 2 we review the previously published work related to the inversion problem at hand. We discuss the conditions under which the inverse matrix  $\mathbf{H}$  is realisable and we give the formulation details of the specific FIR inverse design utilised in the thesis. We review a previously introduced method for the determination of this FIR inverse in the frequency-domain and we propose a minor modification in order for the method to be directly utilisable in the experimental procedure of the following chapters. The use of regularisation in the inversion is discussed and the formulae are given for the use of frequency-varying regularisation. Finally, the implications of the fact that the poles of the rational transfer functions in the exact inverse matrix  $\mathbf{H}$  are clustered close to the unit circle are discussed and an explanation is offered that is based on a simple analytical model.

In Chapters 3, 4 and 5 we use the inverse design discussed above in a series of objective experiments that quantitatively evaluate the effectiveness of the inverse filtering in a wide range of implementation scenarios. In Chapter 3 we explain the

motivation for the undertaking of the specific type of objective evaluation of the inverse filtering implementation. We give the details of the experimental arrangement that we used and of the methods employed for the acquisition of the presented results.

A detailed exposition is presented Chapter 4 in which we determine the absolute optimum of the inversion accuracy that can be realistically achieved with this specific design. The presentation continues by giving quantitative measures of the degradation imposed on the inversion by the introduction of a series of errors and approximations in the design of the inverse filter matrix. These include the cases when the *in situ* measurement of the plant matrix is not feasible, when regularisation is introduced as a means to increase the overall dynamic range of the processing, when the FIR approximation to the ideal inverse filters is of insufficient order, when the responses of the reproduction transducers are not included in the design of the inverse and when a stable and causal inverse is used based on the minimum-phase approximation of its recursive part. Even though the presentation is primarily interested in the objective-quantitative evaluation of the inversion accuracy, the expected perceptual weight of the findings is also briefly discussed in each case.

In Chapter 5 we expand the exposition to cover a number of errors that typically arise in the implementation of the inverse in real-world virtual acoustic imaging applications. These include the presence of reflections in the plant, positional deviations from the ideal reproduction geometry and the use of non-individualised HRTFs for the modelling of the plant. Highly realistic quantitative measures of the inaccuracy introduced in the inversion are given for these cases and, again, the expected influence on the subjective perception of the listeners is discussed.

In Chapter 6 we formulate the introduced forward-backward time inverse filtering method. Initially the method is illustrated in an artificial single-channel context. Subsequently, a number of two-channel variants are considered and a highly accurate numerical method for the determination of the involved transfer functions' decomposition in parts realisable in forward and backward time is presented. A number of weaknesses of the method are identified and ways to rectify them are proposed. These modifications to the method are exemplified on the basis of one of the two-channel variants and using the measured plant model of chapter 3. In this way, the inversion results obtained are compared with the corresponding FIR

inversion results of chapter 3 in terms of the realistically obtainable inversion accuracy and the associated computational cost. The possibility of realising the inverse with reasonably low order recursive filters implemented in forward-backward time is thus ascertained.

In Chapter 7 we present a block-processing type algorithm for the on-line implementation of the proposed forward-backward time inverse filtering method. We discuss the computational load associated with the on-line implementation of the algorithm and give the formulae that relate the computational load with the incurred input-output latency. The computational load vs. input-output latency performance of the algorithm is then compared with the corresponding performance of two representative frequency-domain block-processing filtering algorithms. The comparison shows that our algorithm’s performance is closely following that of the existing frequency-domain based algorithms and hence it may well represent an attractive alternative if further optimised.

Finally in Chapter 8 we summarise the conclusions of the research presented and propose directions for its further development.

# ***2 Inverse filtering in virtual acoustic imaging systems***

## ***2.1 Introduction***

In this chapter (in §2.2) we give a brief review of the various designs that have been proposed in the literature for the implementation of the inverse matrix  $\mathbf{H}(z)$  of equation (1-4). The review extends to cover a wider range of audio-related inverse filtering formulations than those strictly pertaining to equation (1-4) and it also covers a number of works that address single-channel inversion problems of associated interest. In §§2.3-2.4 we focus our interest on one specific method for the determination of the inverse and discuss its main properties. The FIR inverse models resulting from this particular method described in these two sections are those examined in the results of chapters 3, 4 and 5.

## ***2.2 Review of inverse filtering designs***

The work that practically initiated the research topic of binaural reproduction over loudspeakers and the related research into inverse filtering for multi-channel equalisation and cross-talk cancellation was that of Shroeder and Atal in the early 1960's (Atal and Schroeder 1962) and (Schroeder 1975). The basic form of a cross-talk cancellation and equalisation filtering network is presented in those works and the properties of its implementation are discussed. Early contributions to the field also include those of Bauer (1961) and Damaske (1971) in which an experimental procedure is described for the determination of an analogue filter of concentrated elements (termed “90° filter”) that is used for the implementation of a cross-talk cancellation network.

Cooper and Bauck (1989) build on Atal’s and Shroeder’s inverse filter design for the axisymmetric reproduction geometry case and propose a shuffler topology for the implementation of the two-channel inverse. They introduce the notion of *joint minimum phase* for a pair of transfer functions and argue that the shuffler they propose can be implemented with stable and causal recursive filters. The same authors (Bauck and Cooper 1996) extend the application range of the inverse filter design to cases of more than one listener, and more than two channels and sources. They investigate the possibility of also extending the two-channel shuffler topology to these cases by factorising the corresponding inverse (or pseudoinverse) matrices using the Singular Value Decomposition method. They also discuss the issue related to the total loudspeaker-input signal power needed for the generation of the required control signals. Similar issues are also investigated by Kirkeby (1995). Moller (1989) discusses the same equalisation/cross-talk cancellation network introducing free-field equalised responses (HRTFs) in the design. Thus his version of the design explicitly accounts for the equalisation of the reproduction loudspeakers’ free-field response as a separate filtering stage.

Sakamoto et al. in (1982) propose a slightly modified version of the inverse filtering network that essentially separates the cross-talk cancellation and overall equalisation stages. Using subjective experiments based on listener individualised plant models they show that the overall equalisation stage is necessary for the creation of virtual images in the median plane and the rear half of the horizontal plane. They also give subjective experimental results for the degradation of the system’s performance due to listener position misalignment or when the reproduction takes place in non-anechoic environments. Similar issues are considered by Köring and Schmitz (1993) where the performance of the inverse filtering in reverberant environments is evaluated again on the basis of subjective experiments. That study also investigates the use of non-individualised inverse models and evaluates the use of a “typical ear” model derived from spectral smoothing of measured HRTFs.

Utilising a feedforward control design previously used in active control of noise and vibration applications Elliott et al. (1987) and Nelson et al. (1992) discuss the implementation of the inverse matrix  $\mathbf{C}^{-1}$  with adaptive FIR filters. They derive the least squares solution to the inverse filtering problem and describe a multi-channel

version of the filtered- $x$  LMS algorithm for the determination of the inverse. Measured experimental results are provided for the two-channel case that verify the design’s performance while the ineffectiveness of the design in controlling the frequency range below 200Hz is identified and discussed<sup>12</sup>. The design is further investigated in (Nelson et al. 1995) where the case of using more sources than points of intended control is considered and it is shown that in this case the adaptive method yields the same result as the MINT method (reviewed below). The completely general case of more than two sources and more than two points of intended control (i.e. more than one listener) is examined in (Nelson 1994) where a formulation is also proposed for the extension of the adaptive design technique to recursive inverse filters by means of the output error and the equation error methods.

A computationally efficient method for the determination of the inverse matrix  $\mathbf{C}^{-1}$  in the frequency domain is described by Kirkeby et al. (1998c). The introduction of regularisation in the related optimisation formulation (a penalty in the control effort needed for the inversion) is proposed as a means to avoid the time-aliasing arising from the computation of the deconvolution in the frequency domain. An additional advantage in the use of regularisation is the moderation effected in the excessive control power associated with the exact (non-regularised) inversion at the low and high ends of the spectrum. This advantage is identified in (Kirkeby and Nelson 1999) and further exploited by the introduction of frequency-dependent regularisation<sup>13</sup>. The implementation of the inverse with lower order warped FIR filters is investigated by Kirkeby et al. (1999b).

Extensive research has been presented by the above authors (Kirkeby et al. 1998a), (Kirkeby et al. 1998b), (Kirkeby et al. 1999a), (Nelson et al. 1997), (Takeuchi et al. 1997), (Takeuchi et al. 1998), (Takeuchi et al. 2001) and (Vasileiadis and Nelson 2004) regarding the influence of the audio reproduction geometry on the effectiveness of the imposed control. On this basis the use of a system with the two loudspeakers positioned close together (termed the “Stereo Dipole”) is advocated. This principle has been further investigated by Takeuchi (2001), in which the use of a two-channel

---

<sup>12</sup> Due to limitations of the then available hardware, the measurements in that work were conducted at a sampling rate of 6.4kHz thus limiting the considered frequency range to below 2kHz. Furthermore, the results only consider the free-field case (no listener present) and are restricted to anechoic reproduction. The results presented in Chapters 3, 4 and 5 of this thesis provide an extensive elaboration of that work.

<sup>13</sup> A more detailed discussion on the use of regularisation is given below in §2.3.3.

multi-way system is proposed (termed the “Optimal Source Distribution” system). In that design, the author exploits the dependence of the effectiveness of the audio control on the relation between the frequency of interest and the reproduction geometry by assigning the control of different frequency intervals to pairs of sources positioned at different angle-spans relative to the listener.

Gardner (1998) implements a visually adaptive cross-talk cancellation system in which the position of the listener is tracked and the inverse filters are accordingly modified on-line. He builds on findings by Jot et al. (1995) according to which HRTFs can be represented up to 8-10kHz by a minimum-phase function cascaded with a pure delay element and implements a zero-delay band-limited crosstalk canceller operating up to 6kHz. A detailed investigation of the implementation properties and the perceived quality of his design is presented.

With the assumption that the squared magnitude of the cross-talk transmission terms is much lower than that of the direct transmission terms<sup>14</sup> Kyriakakis et al. in (1999) propose a method for the factorisation of the inverse matrix  $\mathbf{C}^{-1}$  in two terms, one of which performs cross talk cancellation and the other ipsilateral equalisation. They discuss the use of least squares inversion and the LMS adaptive method for the determination of the filters comprising their inverse design. Simulation results of the implementation are presented by Mouchtaris (2000). In these works the case of early energetic reflections in the plant (as for example the reflections due to the PC monitor screen or the table in the case of a PC desktop oriented application) and of the possibility of tracking the listener’s position and adapting the inverse model dynamically are also considered but no inversion results are provided.

Garas and Sommen (1998) describe an adaptive solution for the determination of a model of the  $2 \times 1$  matrix  $\mathbf{C}^{-1}\mathbf{A}$  (see discussion in p. 23). They propose its implementation with multiresolution filters as a means to address the issue of the non-robustness of the system to small listener position errors. Aarts et al. (1998) describe the implementation of the filtered-x adaptive algorithm using block processing in the frequency-domain. A model for the implementation of adaptive multi-channel

---

<sup>14</sup> We note that such an assumption is not always valid as, for instance, is the case with the Stereo Dipole geometry.

equalisation in subbands is proposed by Weiss et al. (1999) as a means to reduce the computational cost associated with the inverse filtering and the adaptation process.

Important aspects of the theory and implementation of the inverse filtering problem at hand have also been investigated by a number of studies associated with single-channel inverse designs in the context of room dereverberation and loudspeaker response correction problems. Neely and Allen (1979) use impulse responses determined on the basis of a room acoustics model to determine the conditions under which the response between a source and a receiver located inside a room of given characteristics is of minimum or mixed phase. They find that even though such a response in a typical room with the source and receiver located not very close to the walls should be expected to be minimum phase, maximum phase zeros can occur in more reverberant rooms and in cases where the source or the receiver moves closer to a boundary.

Miyoshi and Kaneda (1986) and the same authors (1988) formulate the “multiple-input/output inverse theorem” (MINT) on the basis of which they propose a method for the exact inversion of a single-channel mixed-phase plant. Unlike the commonly used FIR-inverse technique, in which the realisability issue of the inverse is addressed with the introduction of modelling delay, the authors propose the introduction of an additional audio channel and source in the plant. Two FIR filters (one for each channel) are then to be designed for the implementation of the control. With this arrangement the  $z$ -domain equation describing the inversion transforms to a Diophantine equation in the polynomial transfer functions of the two FIR filters to be designed. It is shown that, assuming that the responses of the two plant channels do not share any common zeros, an exact inversion (zero residual error) can be effected.

Work on the single-channel inversion of loudspeaker-to-listener responses in reverberant environments that extends back to the early 80's has been presented by Mourjopoulos et al. (1982), Clarkson et al. (1985), Mourjopoulos (1994) and Mourjopoulos (2003). Issues regarding the realistically achievable inversion quality in such environments are extensively investigated. The use of time-domain least-squares techniques for the determination of the inverse rather than the frequency-domain method used here is advocated for that type of application (see (Mourjopoulos 2003) and references therein). A comparative study of computational methods for the

determination of the inverse of a single-channel impulse response is presented by Irisawa et al. (1998). They discuss their properties in terms of the accuracy achieved and the computational load required by each.

Finally, a detailed analysis of the feasibility and the properties of the inversion of a single-channel room impulse response is presented by Fielder (2003). The aspects that bear the greatest importance in terms of the perceived audio quality are discussed and criteria for the inaudibility of inversion errors are given in the frequency and time domains. In a more recent study Norcross et al. (2004) investigate the use of inverse filtering for loudspeaker response correction. They evaluate different inverse filter design techniques in subjective experiments and directly relate the presence of specific types of imperfections in the inverse with the perceived quality of the achieved equalisation. The results and analysis of chapters 3, 4 and 5 of this thesis build on the results presented in those two works and extend them to the two-channel case of crosstalk cancellation and equalisation for binaural reproduction.

## ***2.3 Computation of the inverse in the frequency-domain***

### ***2.3.1 Impulse response of a rational transfer function***

In this section we give the form of the stable but not always causal impulse response  $g(n)$  corresponding to the rational transfer function  $G(z)$  of equation (2-1)<sup>15</sup>. We assume that  $G(z)$  is proper, i.e. that  $M < N$ , as is the case with the inverse transfer functions encountered in the context of this work.

$$G(z) = \frac{B(z)}{A(z)} = \frac{\sum_{i=0}^M b_i z^{-i}}{\sum_{j=0}^N a_j z^{-j}} \quad \text{with } a_j, b_i \text{ real numbers} \quad (2-1)$$

---

<sup>15</sup> The presentation of this is essentially a recompilation of what is presented by Proakis and Manolakis (1992 pp. 264-271). It is included here for the purposes of consistent referencing in the material presented in the following.

The impulse response  $g(n)$  will be given by the inverse  $z$ -transform of  $G(z)$  as in equation (2-2):

$$g(n) = Z^{-1}\{G(z)\} \quad (2-2)$$

Assuming that the denominator  $A(z)$  in equation (2-1) has no multiple-order roots and factorising it into  $N$  factors of the form  $(1-p_i z^{-1})$  there will be  $N$  numbers  $A_i$  so that the partial fractions expansion of equation (2-3) will hold<sup>16</sup>:

$$G(z) = \frac{1}{a_0} \left[ \frac{A_1}{1-p_1 z^{-1}} + \frac{A_2}{1-p_2 z^{-1}} + \dots + \frac{A_N}{1-p_N z^{-1}} \right] \quad (2-3)$$

Given that the coefficients in  $G(z)$  are real, all the zeros  $p_i$  of the polynomial  $A(z)$  will be either real or they will appear in complex conjugate pairs. That is, for every strictly complex  $p_i$  there will be a  $p_j$  for which  $p_j = (p_i)^*$ . Furthermore for every real pole  $p_i$ , the corresponding term  $A_i$  will be real and for every complex conjugate pair of poles  $p_j = (p_i)^*$ , the corresponding terms  $A_i$  and  $A_j$  will also be conjugate:  $A_j = (A_i)^*$ . Irrespective of  $p_i$ ,  $A_i$  being real or complex, each partial fraction in equation (2-3) can be written as the infinite sum of a geometric series either as in equation (2-4) where the region of convergence of the series is the exterior of the circle  $|z|=|p_i|$  or as in equation (2-5) where the region of convergence of the series is the interior of the circle  $|z|=|p_i|$ .

$$\frac{A_i}{1-p_i z^{-1}} = A_i \sum_{j=0}^{+\infty} (p_i)^j z^{-j}, \quad |z| > |p_i| \quad (2-4)$$

$$\frac{A_i}{1-p_i z^{-1}} = \left( \frac{z}{p_i} \right) \frac{-A_i}{1-p_i^{-1} z} = -A_i \sum_{j=1}^{+\infty} (p_i)^{-j} z^j, \quad |z| < |p_i| \quad (2-5)$$

Depending then on the modulus  $|p_i|$  of the real and complex poles appearing in equation (2-3) we can distinguish between the following cases:

---

<sup>16</sup> In the case where  $A(z)$  has poles of higher multiplicity then equation (2-3) does not hold but a similar expansion in terms of higher multiplicity can be written, see (Dyer 2003)

### Poles with modulus less than unity

In the case where  $|p_i| < 1$  it is easy to see that the right-sided time series<sup>17</sup> equal to the inverse  $z$ -transform of the sum in (2-4) is decaying exponentially in forward time. We can thus take the causal and stable sequence  $g_i(n)$  of equation (2-6) as the inverse transform of the terms of (2-3) for which  $|p_i| < 1$

$$g_i(n) = Z^{-1} \left\{ \frac{A_i}{1 - p_i z^{-1}} \right\} = u(n) A_i (p_i)^n, \quad |p_i| < 1 \text{ and } |z| > |p_i| \quad (2-6)$$

where  $u(n)$  is the unit-step function defined as:

$$u(n) = \begin{cases} 1, & n \geq 0 \\ 0, & n < 0 \end{cases} \quad (2-7)$$

In the case where  $p_i$ , and consequently  $A_i$ , are real then  $g_i(n)$  will also be real. If  $p_i$  is complex then we can see that the sum of the time-series  $g_i(n)$  and  $g_i(n)'$  corresponding to the fractions containing  $p_i$ ,  $A_i$ , and  $(p_i)^*$ ,  $(A_i)^*$  will again be real as is shown in equation (2-8):

$$\begin{aligned} g_i(n) + g_i(n)' &= Z^{-1} \left\{ \frac{A_i}{1 - p_i z^{-1}} \right\} + Z^{-1} \left\{ \frac{(A_i)^*}{1 - (p_i)^* z^{-1}} \right\} = \\ &= u(n) \left( A_i (p_i)^n + (A_i)^* ((p_i)^*)^n \right) = \\ &= u(n) |A_i| |p_i|^n \left( e^{i(\arg(p_i)n + \arg(A_i))} + e^{-i(\arg(A_i)n + \arg(A_i))} \right) = \\ &= 2u(n) |A_i| |p_i|^n \cos(\arg(p_i)n + \arg(A_i)), \quad |p_i| < 1 \text{ and } |z| > |p_i| \end{aligned} \quad (2-8)$$

### Poles with modulus higher than unity

In the case where  $|p_i| > 1$  we choose the left-sided sequence  $g_i(n)$  corresponding to the inverse  $z$ -transform of the sum in (2-5). As is shown in equation (2-9) this sequence is again decaying exponentially but in backward time and is thus stable but anti-causal.

$$g_i(n) = Z^{-1} \left\{ \frac{A_i}{1 - p_i z^{-1}} \right\} = -u(-n-1) A_i (p_i)^n, \quad |p_i| > 1 \text{ and } |z| < |p_i| \quad (2-9)$$

<sup>17</sup> For the terms right-sided and left-sided series see (Oppenheim and Schafer 1975).

Again, as in the previous case,  $g_i(n)$  will either be real or there will be a  $g_i(n)'$  so that the sum of equation (2-10) will be real:

$$\begin{aligned}
 g_i(n) + g_i(n)' &= Z^{-1} \left\{ \frac{A_i}{1 - p_i z^{-1}} \right\} + Z^{-1} \left\{ \frac{(A_i)^*}{1 - (p_i)^* z^{-1}} \right\} = \\
 &= -u(-n-1) \left( A_i (p_i)^n + (A_i)^* ((p_i)^*)^n \right) = \\
 &= -2u(-n-1) |A_i| |p_i|^n \cos(\arg(p_i)n + \arg(A_i)), \quad |p_i| > 1 \text{ and } |z| < |p_i|
 \end{aligned} \tag{2-10}$$

### Poles with modulus equal to unity

In the case where the modulus of one (or more) of the poles  $p_i$  in (2-3) is exactly equal to 1,  $|p_i|=1$ , it is easy to see that if  $p_i$  is real then the corresponding real impulse response  $g_i(n)$  will be as in equation (2-11) whereas if  $p_i$  is not real then the sum of the complex impulse responses  $g_i(n)$  and  $g_i(n)'$  corresponding to  $p_i$ ,  $A_i$ , and  $(p_i)^*$ ,  $(A_i)^*$  will be as in equation (2-12).

$$g_i(n) = Z^{-1} \left\{ \frac{A_i}{1 \pm z^{-1}} \right\} = \pm A_i u(n), \quad |z| > |p_i| \tag{2-11}$$

$$\begin{aligned}
 g_i(n) + g_i(n)' &= Z^{-1} \left\{ \frac{A_i}{1 - p_i z^{-1}} \right\} + Z^{-1} \left\{ \frac{(A_i)^*}{1 - (p_i)^* z^{-1}} \right\} = \\
 &= 2u(n) |A_i| \cos(n \arg(p_i) + \arg(A_i)) \quad p_i = 1 \text{ and } |z| > |p_i|
 \end{aligned} \tag{2-12}$$

Hence in both cases the corresponding impulse response (or sum of impulse responses) will be either of a constant finite value it will “ring” indefinitely forming a sinusoidal envelope.

Gathering now the three cases  $|p_i| < 1$ ,  $|p_i| > 1$  and  $|p_i|=1$  and grouping accordingly the terms of equation (2-3) we see that the impulse response  $g(n)$  of equation (2-2) will be a sum of three types of time-series, one decaying in forward time, one decaying in backward time and one extending indefinitely. It is also easy to see that the rate of decay of these series will be controlled by the magnitude  $|p_i|$  of the poles, the closer each  $|p_i|$  being to unity the slower the decay while in the case where one (or more) has a magnitude exactly equal to 1 the series does not decay but extends indefinitely in time.

$$g(n) = \frac{1}{a_0} \left[ \sum_i g_i(n) + \sum_j g_j(n) + \sum_k g_k(n) \right] \quad (2-13)$$

### 2.3.2 Realisability of the inverse filter matrix

As was described in §1.2, the design objective for the virtual acoustic imaging system under discussion translates to the design and the realisation of the inverse filter matrix  $\mathbf{H}$  according to equation (1-4). The form of the exact inverse matrix  $\mathbf{H}$  will then be as in equation (2-14) where the responses  $C_{ij}(z)$  are those defined in §1.2.

$$\begin{aligned} \mathbf{H}(z) &= [\mathbf{C}(z)]^{-1} = \frac{1}{\det[\mathbf{C}(z)]} \text{adj}[\mathbf{C}(z)] \\ &= \frac{1}{\det[\mathbf{C}(z)]} \begin{bmatrix} C_{22}(z) & -C_{12}(z) \\ -C_{21}(z) & C_{11}(z) \end{bmatrix} \\ &= \begin{bmatrix} \frac{C_{22}(z)}{C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z)} & \frac{-C_{12}(z)}{C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z)} \\ \frac{-C_{21}(z)}{C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z)} & \frac{C_{11}(z)}{C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z)} \end{bmatrix} \end{aligned} \quad (2-14)$$

It can be seen from equation (2-14) that the inverse filter matrix  $\mathbf{H}(z)$  will comprise 4 recursive filters  $H_{ij}(z)$  the transfer functions of which all share a common denominator equal to the determinant of the plant matrix  $\mathbf{C}(z)$ . In general, as was explained in the previous section (see §2.3.1), a rational transfer function  $G(z) = B(z)/A(z)$  corresponds to a causal and stable (and thus realisable) impulse response  $g(n)$  if and only if all the (generally complex) zeros of the polynomial  $A(z)$  have a modulus of less than unity (i.e. define minimum phase poles). If on the other hand one, at least, of the zeros of  $A(z)$  has a modulus greater than unity (making it a maximum phase pole) then the impulse response  $g(n)$  will be either unstable or non-causal and thus, in both cases, unrealisable<sup>18</sup>.

The issue of the phase characteristics of measured HRTFs and electroacoustic plant responses has been extensively covered in the literature. It is well established that

<sup>18</sup> As is explained in 2.3.1, in the event where the magnitude of a zero of  $A(z)$  is exactly equal to one, the corresponding impulse response extends indefinitely in time. We can note that, even though this “pathological” case has never been observed in realistic situations (Kirkeby et al. 1996), it can be addressed with the introduction of regularisation (see §2.3.3)

single-channel room responses can be minimum phase if the energy of the early reflections and the reverberation is kept low (Neely and J. B. Allen 1979) and that HRTFs measured in anechoic conditions are generally non-minimum phase but with the maximum phase zeros appearing at higher frequencies (see also the discussion in §5.2). As has already been reviewed in §2.2, inverse filtering designs have been proposed that explore the fact that band limited HRTFs are minimum-phase thus resulting in causal and stable realisations of the inverse filtering network. In order to achieve that, those designs reformulate equation (2-14) as a function of the ratio of the contralateral to the ipsilateral HRTFs  $C_{ji}(z)/C_{ii}(z)$  which corresponds to a stable and causal impulse response if  $C_{ii}(z)$  is minimum-phase<sup>19</sup>.

In the work presented here we have not adopted such an approach. That is because, first, one of the objectives defined at the outset of this research project was the examination of the effectiveness of the particular full-audio range inverse filtering design of equation (2-14) and an investigation of the possibility of implementing this design with recursive filters. As can be seen in equation (2-14), even if the responses  $C_{ij}(z)$  are minimum-phase this does not necessarily hold for the expression  $\det[\mathbf{C}(z)]$ . In fact, the relative delay between the two ear responses results to a rising initial part in  $\det[\mathbf{C}(z)]$  (see e.g. figure 2-2) which can indeed enhance its non-minimum phase character. Hence the minimum-phase approach is not trivially applicable. Second, the need to address the problem without the non-minimum phase assumption is also imposed by the use of regularisation. The motivation for introducing regularisation in the inversion is discussed in §2.3.3 and also in Chapter 4. As will be seen there, the regularised inverse matrix is bound to be of mixed phase. Finally, in cases where the inverse is designed to also invert reflections in the plant (more in §5.2), the minimum-phase assumption fails altogether.

With the denominator of the transfer functions  $H_{ij}(z)$  in (2-14) being mixed-phase, the ideal inverse impulse responses  $h_{ij}(n)$  will consist of a causal and an anti-causal part both extending infinitely in forward and backward time respectively. Given, however, that none of the zeros of  $\det[\mathbf{C}(z)]$  has a magnitude equal to one, after a certain number of samples both the forward and backward parts of the series will have

---

<sup>19</sup> Note that ratios of this form are independent of the method of HRTF equalisation, or indeed of whether equalisation is applied at all, as the equalisation response cancels out (Gardner 1998). Hence for such formulations of the inverse the plant matrix can comprise HRTFs rather than responses from the input to the loudspeaker at the output of the listener's ears.

decayed to insignificantly low values:  $h_{ij}(n) \approx 0$  for  $n < -\Delta$  or  $n > N$  with  $\Delta$  and  $N$  positive integers. Under that assumption a finite length approximation to the ideal inverse impulse responses  $h_{ij}(n)$  can be obtained by truncating the samples corresponding to the indices  $n < -\Delta$  and  $n > N$ . Shifting these truncated impulse responses by  $\Delta$  samples forward in time so that they become causal we obtain a causal FIR approximation to the ideal inverse of a total length of  $N_H = \Delta + N + 1$  coefficients as is described in equation (2-15).

$$h_{ij}^{\text{FIR}}(n) = w_{N_H}(n) Z^{-1} \{z^{-\Delta} H_{ij}(z)\} \quad (2-15)$$

where  $w_{N_H}(n)$  is the unit-height square window of length  $N_H$  samples:

$$w_{N_H}(n) = \begin{cases} 1, & n = 0, 1, \dots, N_H - 1 \\ 0, & \text{otherwise} \end{cases} \quad (2-16)$$

It is then easy to see that the matrix  $\mathbf{h}^{\text{FIR}}(n)$  having the FIR filters  $h_{ij}^{\text{FIR}}(n)$  as elements will satisfy equation (2-17), i.e. a time-delayed version of the original cross-talk cancellation equations (1-4) and (1-5).

$$\begin{aligned} \mathbf{c}(n) * \mathbf{h}^{\text{FIR}}(n) &\approx \delta(n - \Delta) \cdot \mathbf{I} && \text{in the time-domain or} \\ \mathbf{C}(z) \mathbf{H}^{\text{FIR}}(z) &\approx z^{-\Delta} \cdot \mathbf{I} && \text{in the } z\text{-domain} \end{aligned} \quad (2-17)$$

The approximate equality in equation (2-17) is due to the fact that the filters  $h_{ij}^{\text{FIR}}$  represent a truncated version of the ideal infinite-length responses  $h_{ij}$ . Hence a truncation error will appear at the two ends of the time sequences in the matrix  $\mathbf{c} * \mathbf{h}^{\text{FIR}}$  in the form of a deviation from the ideal zero value. As is further discussed in §4.2 this type of error can be very important in terms of the perceptual impression of the listener and should thus be suppressed by the use of an adequately long window  $w_{N_H}(n)$ . Another way to address this issue could be the use of a window that gradually tapers the filters  $h_{ij}^{\text{FIR}}$  towards zero at their two ends rather than sharply truncating them as does the square window. Such an approach has not, however, been considered in this thesis. A frequency-domain method for the computation of the FIR inverse model of equation (2-15) is described by Kirkeby et al. (1998c) and the FIR models used in the present thesis were obtained by a modified version of that method (see §2.3.4).

### 2.3.3 Regularisation

The introduction of regularisation was originally proposed in Kirkeby et al. (1998c) as a means to address the problem of time-aliasing that arises from the computation of the inverse in the frequency-domain (see also §2.3.4). This is done by modifying the exact inverse modelling problem of equations (1-1)–(1-3) to the optimisation problem of minimising the quadratic cost function of equation (2-18) which is now formulated using the frequency domain expression of the quantities defined in §1.2<sup>20</sup>.

$$J(e^{j\omega}) = \mathbf{e}^H(e^{j\omega})\mathbf{e}(e^{j\omega}) + \beta\mathbf{y}^H(e^{j\omega})\mathbf{y}(e^{j\omega}) \quad (2-18)$$

where the real valued and positive regularisation parameter  $\beta$  is introduced. In this cost function the minimisation of the squared error  $\mathbf{e}^H\mathbf{e}$  is penalised by the term  $\beta\mathbf{y}^H\mathbf{y}$ , which is proportionate to the magnitude of the control effort represented by the source input signals  $\mathbf{y}$ . The unique minimum of the cost function of equation (2-18) can be shown (Nelson 1994) to be obtained when the frequency response of the inverse filter matrix  $\mathbf{H}$  becomes equal to:

$$\mathbf{H}(e^{j\omega}) = \left[ \mathbf{C}^H(e^{j\omega})\mathbf{C}(e^{j\omega}) + \beta\mathbf{I} \right]^{-1} \mathbf{C}^H(e^{j\omega}) \quad (2-19)$$

As is readily seen in equation (2-19), the introduction of regularisation effectively penalises the power of the control effort  $\mathbf{y}$  that is needed for the exact inversion. That is, for a value of  $\beta$  equal to zero the optimal solution of equation (2-19) coincides with the exact solution of equation (1-4) while for positive values of the regularisation parameter  $\beta$  a trade-off is effected between the accuracy in the inversion and the required control effort. As is shown by Kirkeby et al. (1996), this translates to the replacement of each pole in the responses  $H_{ij}(z)$  (i.e. each zero in the response  $\det[\mathbf{C}(z)]$ ) with one zero and a pair of poles placed on either side of the unit circle in the  $z$ -plane<sup>21</sup>. These poles are further away from the unit circle than the original (non-regularised) pole of  $H_{ij}(z)$  and consequently the impulse responses  $h_{ij}(n)$  will decay faster in time, thus reducing the required length of the inverse DFTs needed to

---

<sup>20</sup> The superscript  $H$  denotes the conjugate transpose.

<sup>21</sup> Obviously, this would also address the hypothetical case whereby a zero of  $\det[\mathbf{C}(z)]$  lies exactly on the unit circle.

suppress the time-aliasing effect. This was the basic principle for the original introduction of regularisation as described Kirkeby et al. (1998c).

On the other hand, as is demonstrated in further detail with the results of Chapter 4, by moderating the applied control power the introduction of regularisation effectively increases the dynamic range of the inversion. As will become apparent in Chapter 4 this is a more desirable property of the introduction of regularisation, since the dynamic range of the non-regularised inversion can be prohibitively low. When used with this objective, however, an extension to the formulation of regularisation seems appropriate. This is because, as was also identified by Kirkeby and Nelson (1999), the control power needed for the exact inversion becomes excessively high at the low and high end of the spectrum where the application of control is either prohibitively inefficient or of no interest. The increase at the low end of the spectrum in the magnitude response of the inverse is due to the roll-off of the loudspeakers' response<sup>22</sup> at this region and also due to the acoustical properties of the reproduction<sup>23</sup>. In this low frequency region, the control is bound to be ineffective even when no regularisation is used. Similarly, when the plant matrix comprises measured responses, it will have to exhibit a roll-off at the high end of the spectrum (close to the Nyquist frequency) because of the antialiasing filters used for the measurement. Obviously, the correction of this roll-off to a flat response is of no interest. Hence it would be desirable to be able to apply a stricter optimisation penalty in these two frequency regions and penalise less (or even not at all) the control in the region in-between which we both wish to and are able to apply control more efficiently. This can be straightforwardly achieved by replacing the regularisation parameter with the frequency variable parameter  $\beta(\omega)$  in equation (2-18) which then results in the form of the inverse matrix described by equation (2-20)<sup>24</sup>.

$$\mathbf{H}(e^{j\omega}) = \left[ \mathbf{C}^H(e^{j\omega}) \mathbf{C}(e^{j\omega}) + \beta(\omega) \mathbf{I} \right]^{-1} \mathbf{C}^H(e^{j\omega}) \quad (2-20)$$

---

<sup>22</sup> It should be noted that this is not true for the case where the plant matrix is modelled with HRTFs which by definition should converge to unity at the DC limit (Algazi et al. 2001). It is however true in the case where the plant matrix also contains the responses of the loudspeakers used for the reproduction, in which case the inverse filtering will have to correct these responses also.

<sup>23</sup> As is discussed in detail in (Takeuchi 2001), the inverse control at the low frequency end of the spectrum becomes ill-conditioned, a characteristic that is further amplified by small loudspeaker-span angles as is the case with the Stereo Dipole.

<sup>24</sup> Another approach, formulated in the time-domain is taken by Kirkeby and Nelson (1999). We believe the method proposed in equation (2-20) to be simpler.

The application of this form of frequency varying regularisation is further discussed in §4.3.

### 2.3.4 Modified method for the determination of the inverse in the frequency-domain

In this section we describe the details of the computation method used to determine the FIR inverse models that were used for the results presented in Chapters 3, 4 and 5. The method is based on the frequency-domain based method that is described in (Kirkeby et al. 1998c) but incorporates a slight modification which allows the chosen length of the inverse filters to be independent of (but always lower than) the length of the DFTs used. A simplified-case example is given that illustrates the justification of such a modification to the method of Kirkeby et al. (1998c).

The objective here is to find the, say  $L$ , coefficients of an FIR filter  $H_{mod}$  that models the complex response (magnitude and phase up to a constant delay factor) of the reciprocal of a given (plant) response  $C(z)$ <sup>25</sup> as in equation (2-21). The term  $z^{-\Delta}$  introduces a frequency independent delay of  $\Delta$  samples in order to address the case where the transfer function  $C(z)$  is non-minimum phase (see §2.3.1)

$$H_{mod}(z) \simeq \frac{1}{C(z)} \cdot z^{-\Delta} \quad (2-21)$$

The existence of an FIR model that satisfies equation (2-21) up to an arbitrary level of accuracy is guaranteed (Claerbout and Robinson 1964) as long as the number of coefficients  $L$  in the FIR model and the number of modelling delay samples  $\Delta$  are sufficiently high.

The method proposed by Kirkeby et al. (1998c) for the determination of  $H_{mod}$  consists of:

- Choosing the intended number of coefficients of the inverse model  $H_{mod}$ , say  $L$ .

---

<sup>25</sup> The discussion is presented in a single-channel context but is straightforwardly applicable to the multi-channel models presented in the thesis. This is done by applying what is presented here for  $1/C(z)$  to the rational forms  $C_{ij}(z)/\det[C(z)]$  of the inverse matrix  $\mathbf{H}(z)$  of equation (2-14) or their regularised equivalent of equation (2-19).

- Taking a  $K$ -point DFT of  $c(n)$ ,  $C(k)=\text{DFT}[c(n)]$  with  $K=L$  (choosing  $K=L$  to be a power of 2 maximises the computational efficiency of the FFT).
- Computing the ratio  $H(k)=1/C(k)$  for  $k=0,\dots,K-1$  that samples the frequency response  $H(\omega)$  at  $K$  equidistant points.
- Taking the  $K$ -point inverse DFT,  $\hat{h}(n)=\text{IDFT}[H(k)]$ .
- Applying a circular shift of  $K/2$  points to  $\hat{h}(n)$  to obtain  $h(n)$  which is the sought model of the inverse of length  $K$  and with the modelling delay  $\Delta$  set to  $K/2$  points.

It is then obvious that the length of the model  $H_{mod}$  is directly determined by the length of the DFTs that are used to get to the frequency domain and back ( $L=K$ ). However the length of the DFT used and the length of the model are, in principle, separate concepts. The former essentially determines the degree to which time-aliasing is present in the response  $\hat{h}(n)$  that results from the IDFT while the latter controls the length up to which the model  $H_{mod}$  represents the exact inverse<sup>26</sup>. It could, for instance, be the case that one would desire to set  $L$  at a relatively low value, hence choosing a short model of the inverse that is truncated early, but at the same time would want to guarantee that the  $L$  samples kept in the model are virtually free of time-aliasing distortion, in which case a high value of  $K$  would be required. Moreover the amount of modelling delay introduced with the above method is fixed to half the length of the filter ( $\Delta=K/2$ ). This could easily be set to different values by appropriately changing the amount of circular shift but this does not alleviate the issue of the time-aliasing.

The obvious way to address the above issue is to compute the DFT and IDFT involved with an excess of zero-padding. This way the aliasing replicas present in  $\hat{h}(n)$  can be pushed further way and hence their contribution can be suppressed down to an arbitrary level. Rectangular windowing can then be applied on  $\hat{h}(n)$  to keep a chosen length of the causal and the anti-causal decay. This was the approach followed for the determination of the FIR models used in the thesis. More specifically, the method used can be described as follows:

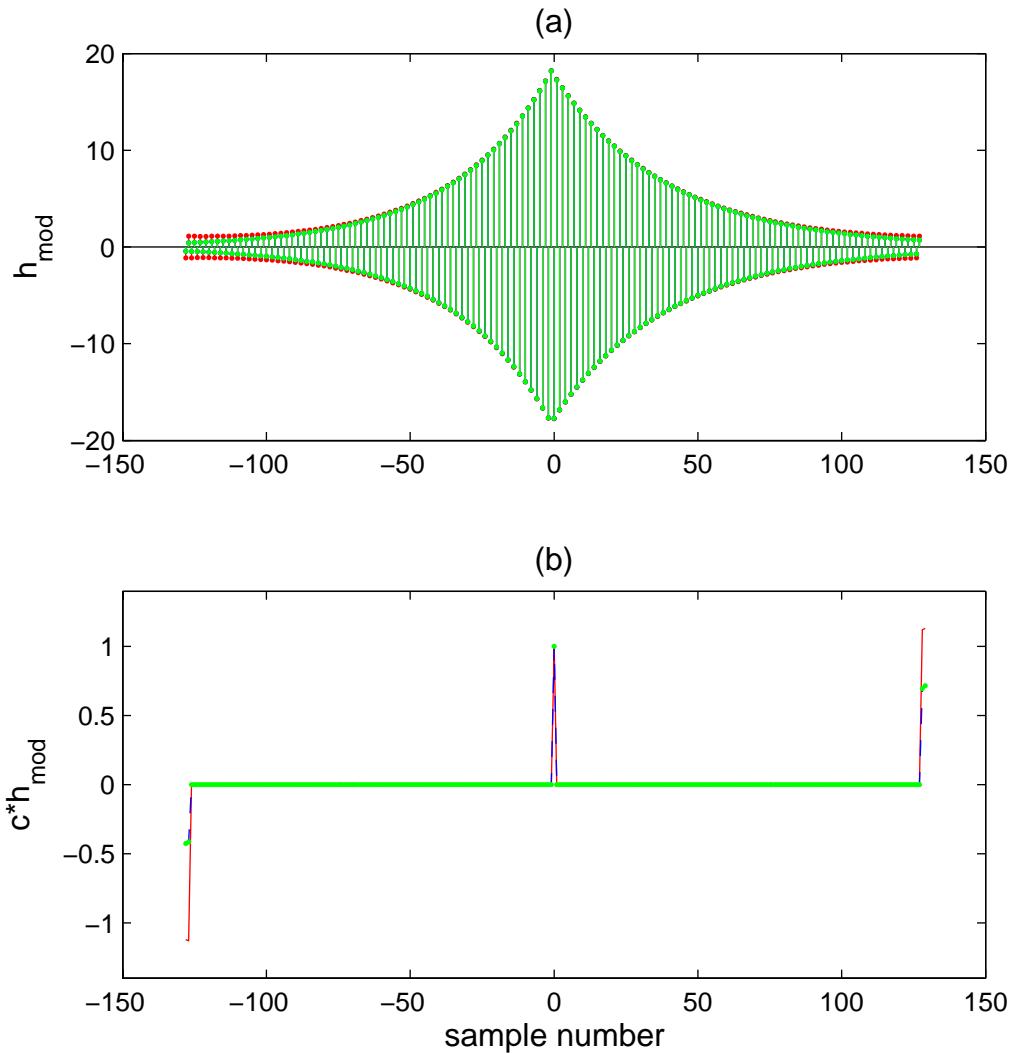
---

<sup>26</sup> Obviously, being the time-domain version of a rational transfer function, the exact inverse will extend infinitely in positive time and, if  $C(z)$  is of non-minimum phase, it will also extend infinitely in negative time.

- Choose the intended number of coefficients of the inverse model  $H_{mod}$ , say  $L$ , and how many of them will correspond to negative time indices (number of samples of modelling delay), say  $\Delta$ .
- Take a  $K$ -point DFT of  $c(n)$ ,  $C(k)=\text{DFT}[c(n)]$ , with  $K$  sufficiently large for the time-aliasing effect to be negligible up to the index  $\Delta$  in negative time and  $L-\Delta$  in positive time (in the cases examined in this thesis we typically used for  $K$  the first power of 2 for which  $K > 2L$  and this was always found to be sufficient).
- Compute the ratio  $H(k)=1/C(k) \cdot e^{-j2\pi k \Delta/K}$ . This applies the circular shift needed in order to bring  $\Delta$  points corresponding to negative time indices followed by  $L-\Delta$  points corresponding to positive time indices at the beginning of the time-sequence obtained with the following IDFT. (For the multi-channel cases that are actually addressed in the thesis, this translates to computing the elements of the matrix  $\mathbf{H}(k)=e^{-j2\pi k \Delta/K} \cdot [\mathbf{C}^{-1}(k)]$ ).
- Compute the inverse DFT  $\hat{h}(n)=\text{IDFT}[H(k)]$  (or of the elements of matrix  $\hat{\mathbf{h}}(n)=\text{IDFT}[\mathbf{H}(k)]$  for the multi-channel case).
- Keep the first  $L$  points of  $\hat{h}(n)$  to obtain the sought FIR model  $h^{\text{FIR}}(n)$ .

The points discussed above are illustrated with the aid of the following simplified-case example. We take here the plant  $C(z)=1+2.005z^{-1}+1.00425z^{-2}$  which has one maximum phase zero at  $z = -1.03$  and one minimum phase zero at  $z = -0.975$ . For such a short plant, the impulse response  $h(n)$  of the inverse  $1/C(z)$  can be straightforwardly computed up to a given length by taking partial fractions expansion (see §2.3.1). The three cases corresponding to the exact inverse computed by partial fractions expansion (no time-aliasing), the modified method presented above (time-aliasing independent of the filter length) and the method of Kirkeby et al. (1998c) are plotted in part (a) of figure 2-1 with blue, green and red colours respectively. For the results of the figure the length of the inverse is set to  $L=2^8$ , the modelling delay to  $\Delta=2^7$  and for the modified method the length of the DFT was set equal to  $K=2^{10}$ . The exact inverse results (blue colour) are not visible because they exactly coincide with the results of the modified method (green colour). However, the effect of the time-aliasing is visible on the results obtained with the method of Kirkeby et al. (1998c)

(red colour). The effect of the time-aliasing becomes more obvious in part (b) of figure 2-1 where we plot the convolution of the three inverse models with  $c(n)$ . Even though all three inverse models are truncated at that same point, the truncation end-effects are more severe in the inverse computed with the method of Kirkeby et al. (1998c) (red line) than those of the modified method (green dotted line) which exactly coincide with the exact inverse computed with a partial fraction expansion (blue dashed line).



**Figure 2-1:** (a) Inverse models of the two-zero mixed-phase plant computed with the modified method (green), the partial fractions expansion (blue, not visible) and the method of Kirkeby et al. (1998c) (red). (b) Convolution of the inverse models with the plant. Modified method (green dotted line), partial fractions expansion method (blue dashed line) and method of Kirkeby et al. (1998c) (red solid line).

It must be said that the results shown in figure 2-1 are somewhat arbitrary; with appropriate placing of the poles of the plant in relation to the unit circle and appropriate choices of the lengths  $L$  and  $K$  the performance of the method of Kirkeby et al. (1998c) can be shown to range from insignificantly to exceedingly worse than the modified method. It has been found that in realistic cases where the plant results from measured impulse responses the difference between the two methods is very rarely significant. But in any event it is evident that if the length of the DFT is equal to the length of the filter, either the choice of the filter length will be exaggerated or the model will be influenced by time-aliasing. The same holds even when regularisation is used even though in that case the effect is further minimised because the inverse responses decay much faster and a low choice of  $L=K$  in the method of Kirkeby et al. (1998c) is enough to yield short models of the inverse that are practically free of time-aliasing. In fact, the use of regularisation was introduced in (Kirkeby et al. 1998c) mainly as a means to minimise the time-aliasing effect. However, since the objective of Chapters 3, 4 and 5 of the present work is to determine the actual inversion accuracy for any combination of the inverse's length and the regularisation applied, a modification had to be made so as to give us the ability to separately adjust these two design parameters of the inversion.

## ***2.4 Notes on the implementation of the inverse***

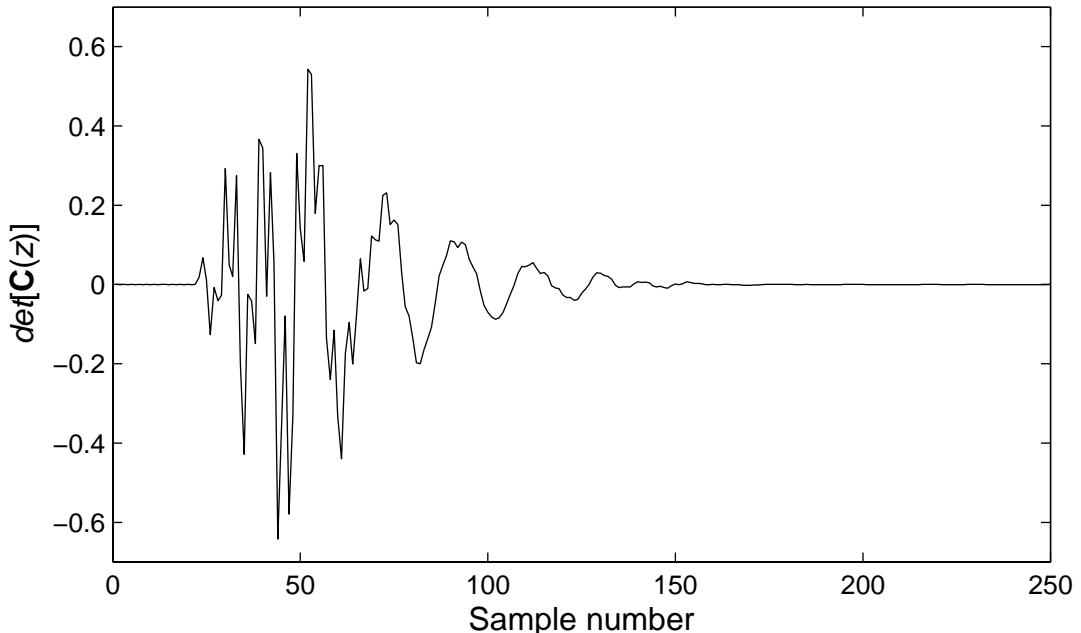
### *2.4.1 Distribution of the zeros of the measured electroacoustic transfer functions*

As can be seen from equation (1-4), the form of the inverse filters  $H_{ij}$  is principally determined by the zero distribution of the polynomial  $\det[\mathbf{C}(z)]$ . The presence of zeros of  $\det[\mathbf{C}(z)]$  outside the unit-circle in the  $z$ -plane necessitates the introduction of modelling delay in order to obtain a causal time-domain representation of each  $H_{ij}$ . Equally, the proximity of the zeros of  $\det[\mathbf{C}(z)]$  to the unit-circle results in a slowly decaying time-response of each  $H_{ij}$ <sup>27</sup>. Both of these characteristics consistently occur when the plant matrix  $\mathbf{C}(z)$  is formed on the basis of measured electroacoustic

---

<sup>27</sup> This is unless each one of these poles of  $H_{ij}$  is cancelled by a zero.

responses and even though the former has been interpreted to some extent (see also the discussion in §5.2), the latter remains somewhat obscure. The presence of high- $Q$  resonances at certain frequencies has been put forth as an explanation by Fielder (2003) in cases where  $\mathbf{C}(z)$  is measured in enclosures. However, in the cases we are concerned with, the relevant electroacoustic responses are measured in anechoic conditions. Most importantly, the zeros of  $\det[\mathbf{C}(z)]$  are clustered very close to the unit circle in an apparently frequency-uniform way. This is shown in figure 2-3 where we plot the zero distribution of the polynomial  $\det[\mathbf{C}(z)]$  computed using the HRTFs corresponding to the  $10^\circ$  Stereo Dipole geometry. The time-domain version of the polynomial  $\det[\mathbf{C}(z)]$  is depicted in figure 2-2<sup>28</sup>.



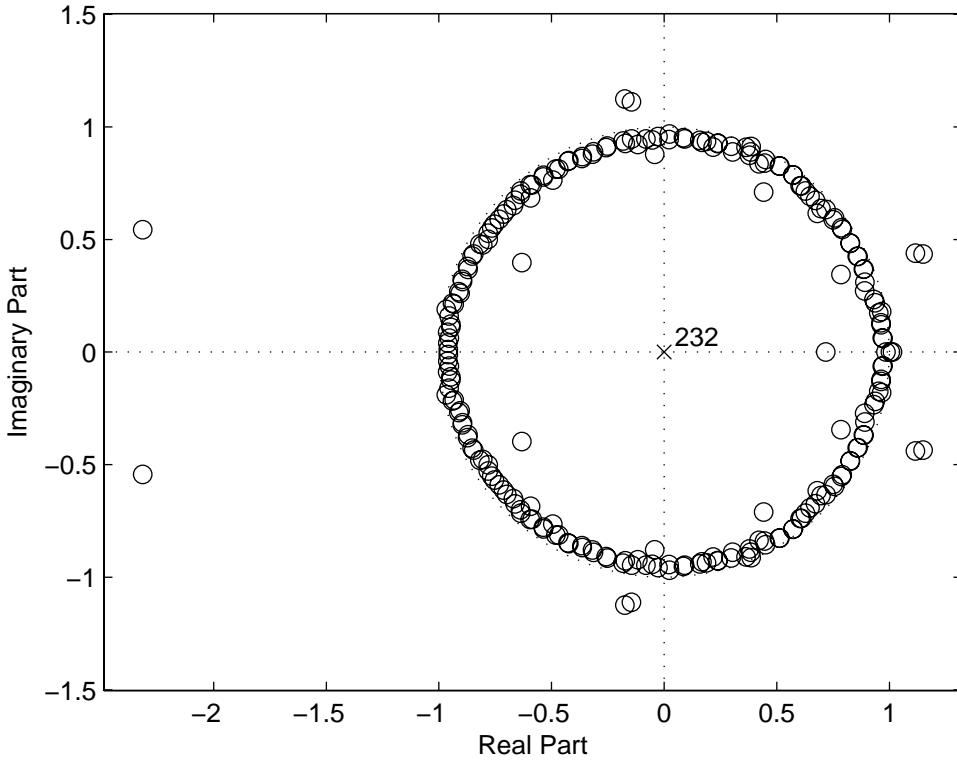
**Figure 2-2:** Impulse response of the determinant polynomial of the plant matrix.

We point out here that, apart from the possible presence of sharp resonances or notches in the measured system, this clustering of the zeros of  $\det[\mathbf{C}(z)]$  close to the unit circle is inherent in the modelling of a system that may contain poles and zeros with an FIR model resulting from measurement. This effect should thus be present (as

---

<sup>28</sup> Note that for the results of figure 2-3 we have excluded the initial delay part of the first 22 samples seen in figure 2-2

is the case) in any measured electroacoustic response of the type considered here. This can be illustrated by means of the following simplified example.



**Figure 2-3:** Zero distribution of the determinant polynomial of the plant matrix.

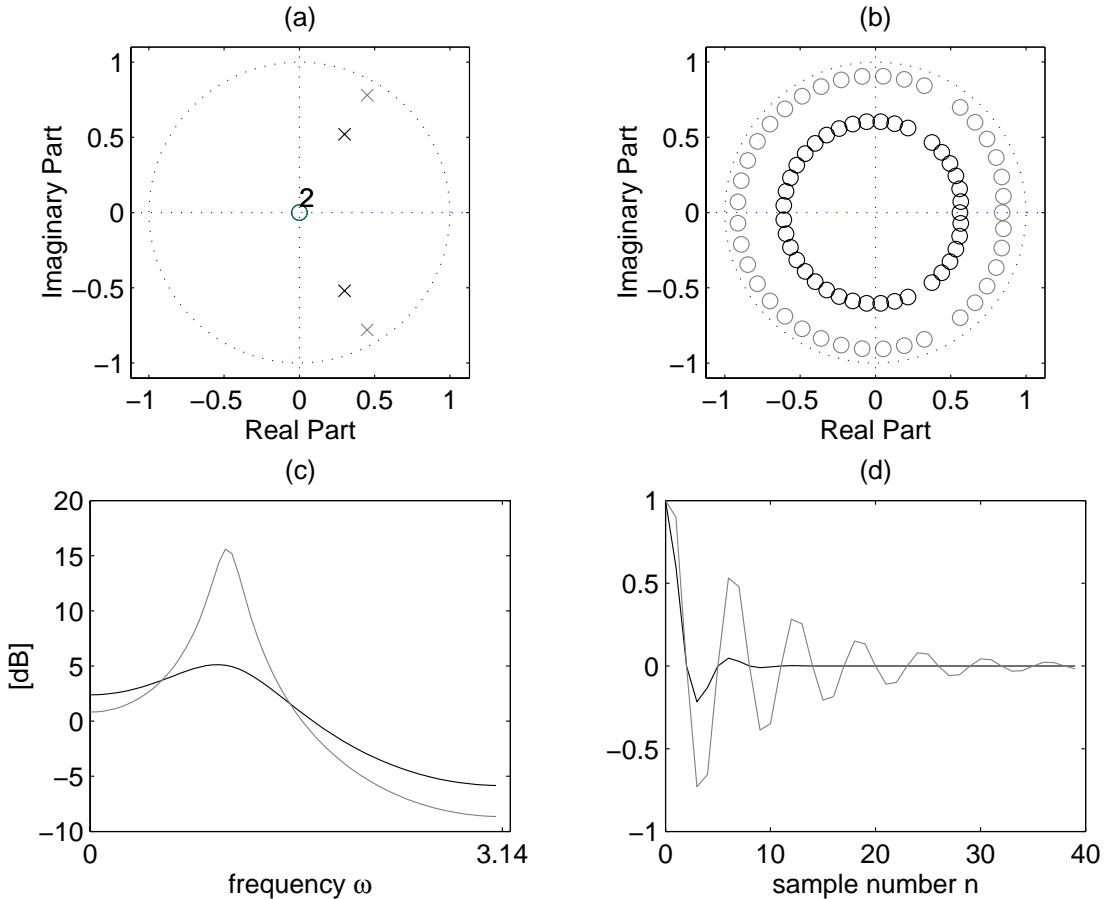
Let us consider the simple case of a second order resonator (Proakis and Manolakis 1992) with transfer function  $C_{exact}(z)$  and impulse response  $c_{exact}(n)$  as defined in equations (2-22) and (2-23) respectively. The pole locations of  $C_{exact}(z)$  with the parameter  $r$  taking the values 0.6 and 0.9 and the parameter  $\omega$  equal to  $\pi/3$  are shown in part (a) of figure 2-4 in part (c) of which we plot the frequency response of the system.

$$C_{exact}(z) = \frac{1}{(1 - re^{j\omega_0}z^{-1})(1 - re^{-j\omega_0}z^{-1})}, \quad \begin{matrix} 0 < r < 1 \\ 0 < \omega < 2\pi \end{matrix} \quad (2-22)$$

$$c_{exact}(n) = Z^{-1}\{C_{exact}(z)\}, \quad n = 0, 1, \dots, \infty \quad (2-23)$$

Now when the system is modelled by the non-recursive transfer function of equation (2-24) having the (finite length) impulse response shown in part (d) of figure 2-4, we see that the corresponding zero distribution will be as in part (b) of the same figure.

$$C_{model}(z) = \sum_{n=0}^{39} c_{exact}(n) z^{-n} \quad (2-24)$$



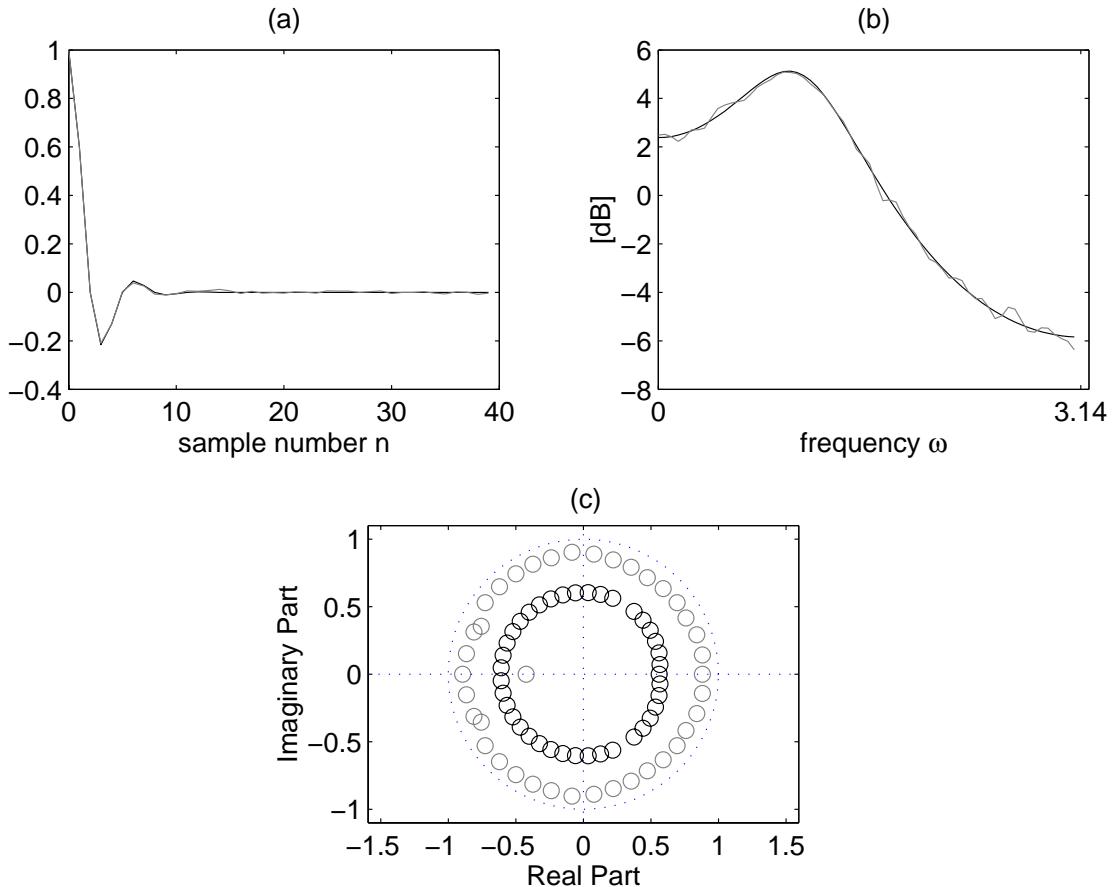
**Figure 2-4:** (a) System's pole locations (b) Model's zero locations (c) System's frequency response and (d) Model's impulse response for  $\omega=\pi/3$ ,  $r=0.6$  (black) and  $r=0.9$  (grey).

It is then seen that, as the model has no poles, the sharpness of the resonance peak is indeed controlled by the proximity of the zeros to the unit circle, in other words, the lower the damping of the system the closer the zeros of the model to the unit circle. However, the location of the zeros of the FIR model is also significantly influenced by the, inevitable, presence of measurement noise in the model as is seen below. Assume that the measured impulse response  $c_{measured}(n)$  is as in equation (2-25) where

we assume the presence of additive noise modelled in this case as a white sequence of zero mean and variance<sup>29</sup>  $\sigma^2=5\times10^{-2}$ .

$$c_{\text{measured}}(n) = c_{\text{model}}(n) + w(n) \quad (2-25)$$

The finite impulse response and the frequency response of the model  $c_{\text{measured}}(n)$  are plotted in parts (a) and (b) of figure 2-5. We then see in part (c) of the same figure that, even in the moderately damped case ( $r=0.6$ ), the addition of noise in the measured impulse response has “pushed” the zeros of the actual model further closer to the unit circle.



**Figure 2-5:** (a) Impulse response (b) Frequency response and (c) Zero-map of error-free (black) and noisy (grey) model with  $r=0.6$  and  $\omega=\pi/3$

<sup>29</sup> With the given variance the energy of the additive noise would then be roughly 40dB below the energy of the model’s impulse response  $c_{\text{exact}}(n)$ . This is higher than the amount of noise typically present in measurements of the type considered here (see §3.3). A higher level of additive noise is used here in order to illustrate the point of discussion.

### 2.4.2 Filtering through the common denominator

Another simple result is presented in this section that is called upon in later parts of the thesis. As was discussed in §2.4.1, the computational load in the implementation of the inverse matrix  $\mathbf{H}$  is mainly due to the low rate of decay of the time-domain representation of the denominator polynomial of equation (2-14). Hence, given that this polynomial is common to all four inverse responses  $H_{ij}$  one might expect that the separate implementation of this stage could result in some computational advantage. That is, instead of filtering each one of the binaural signals  $x_1(n)$ ,  $x_2(n)$  through the elements  $h_{ij}(n)$  of the inverse matrix  $\mathbf{H}$ , one could filter  $x_1(n)$  and  $x_2(n)$  through a model  $h_{CD}(n)$  of the transfer function  $1/\det[\mathbf{C}(z)]$  and then filter the two resulting intermediate signals through the elements of the matrix  $\text{adj}[\mathbf{C}(z)]$ .

This amounts to replacing the direct computation of the inverse filters' output<sup>30</sup>

$$\mathbf{y}(n) = \begin{bmatrix} h_{11}(n) & h_{12}(n) \\ h_{21}(n) & h_{22}(n) \end{bmatrix} * \begin{bmatrix} x_1(n) \\ x_2(n) \end{bmatrix} = \begin{bmatrix} h_{11}(n)*x_1(n) + h_{12}(n)*x_2(n) \\ h_{21}(n)*x_1(n) + h_{22}(n)*x_2(n) \end{bmatrix} \quad (2-26)$$

by the separate stages

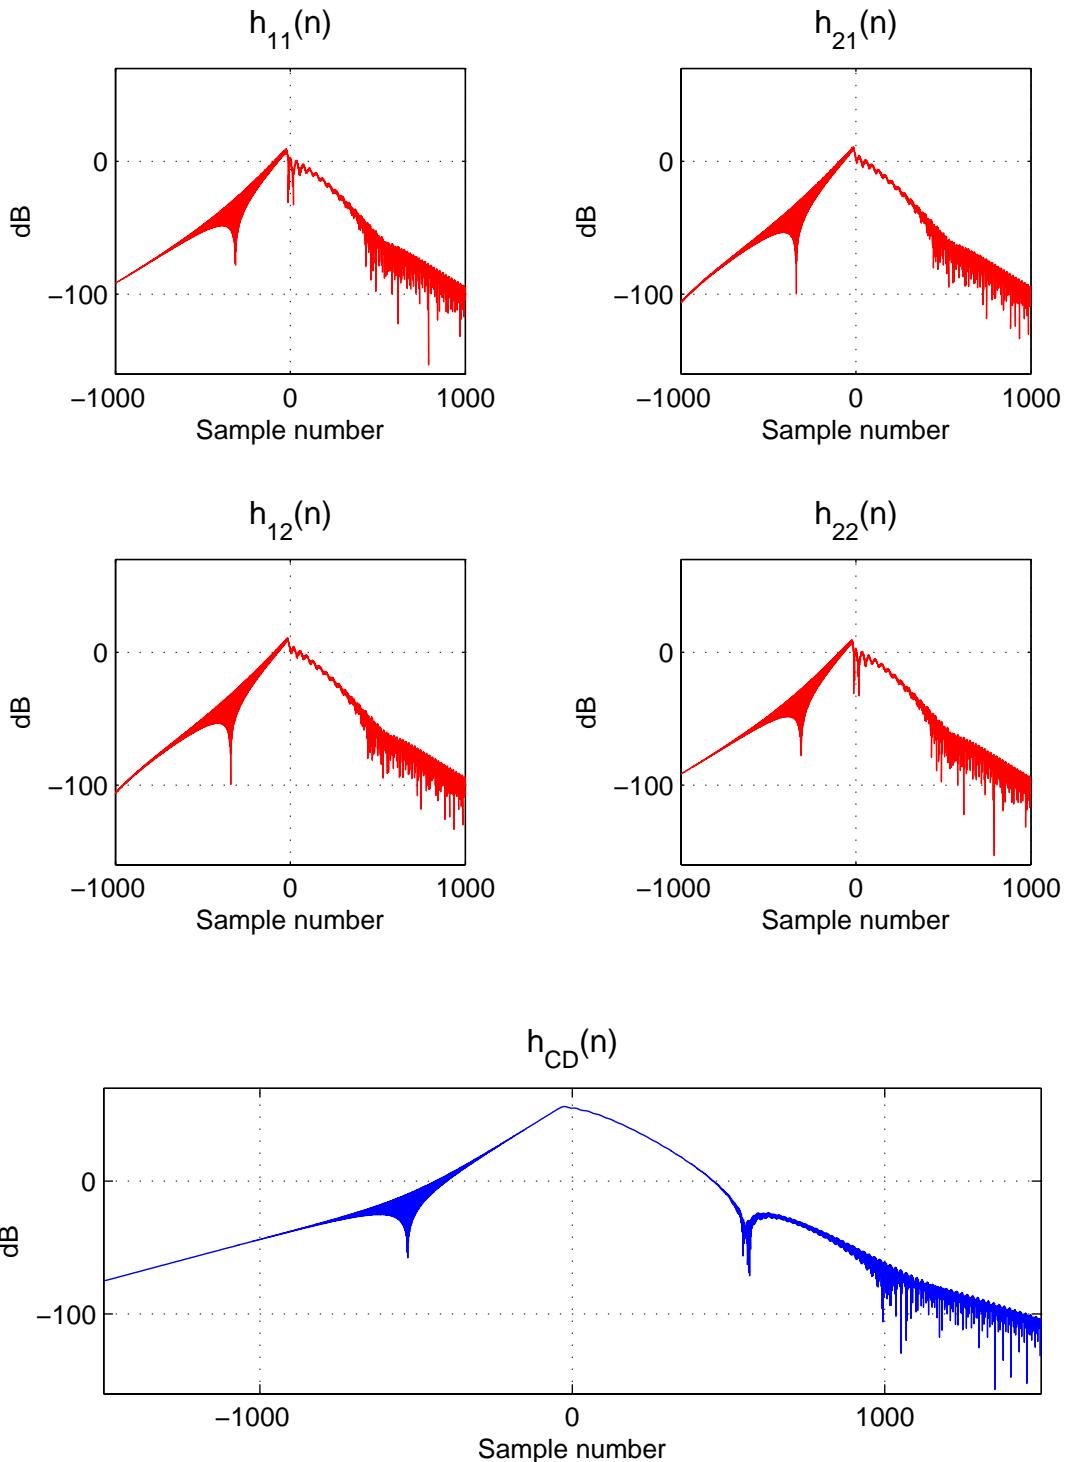
$$\mathbf{y}(n) = \begin{bmatrix} c_{22}(n) & -c_{12}(n) \\ c_{21}(n) & c_{11}(n) \end{bmatrix} * \left( h_{CD}(n) * \begin{bmatrix} x_1(n) \\ x_2(n) \end{bmatrix} \right) \quad (2-27)$$

We denote with  $N_H$  the chosen order of the FIR filters  $h_{ij}(n)$ , with  $N_{CD}$  the chosen order of the FIR filter  $h_{CD}(n)$  used to model the common denominator reciprocal transfer function  $1/\det[\mathbf{C}(z)]$  and with  $N_C$  the order of the measured responses of the plant. We see then that the computation of equation (2-26) requires four convolutions with the impulse responses  $h_{ij}(n)$ , i.e.  $4N_H$  multiplications and additions for the computation of each pair of output samples in direct sample-by-sample filtering. Alternatively, the computation of equation (2-27) requires two convolutions with the impulse response  $h_{CD}(n)$  and four convolutions with the impulse responses  $c_{ij}(n)$ , i.e. a total of  $2N_{CD} + 4N_C$  multiplications and additions for the computation of each pair of output samples. Since the order of the measured model of the plant  $N_C$  is typically much lower than the order of the inverse models  $N_H$  it would be expected that the total

---

<sup>30</sup> We use here the signals and systems notation introduced in p. 17.

cost of the latter process could be lower than that of the former. However, this is not always the case as is illustrated with the following set of simulation results.



**Figure 2-6:** Impulse response of the elements of the inverse filter matrix  $\mathbf{H}$  (red line) and of the reciprocal expression of the common denominator (blue line). In both cases the regularisation is set to zero.

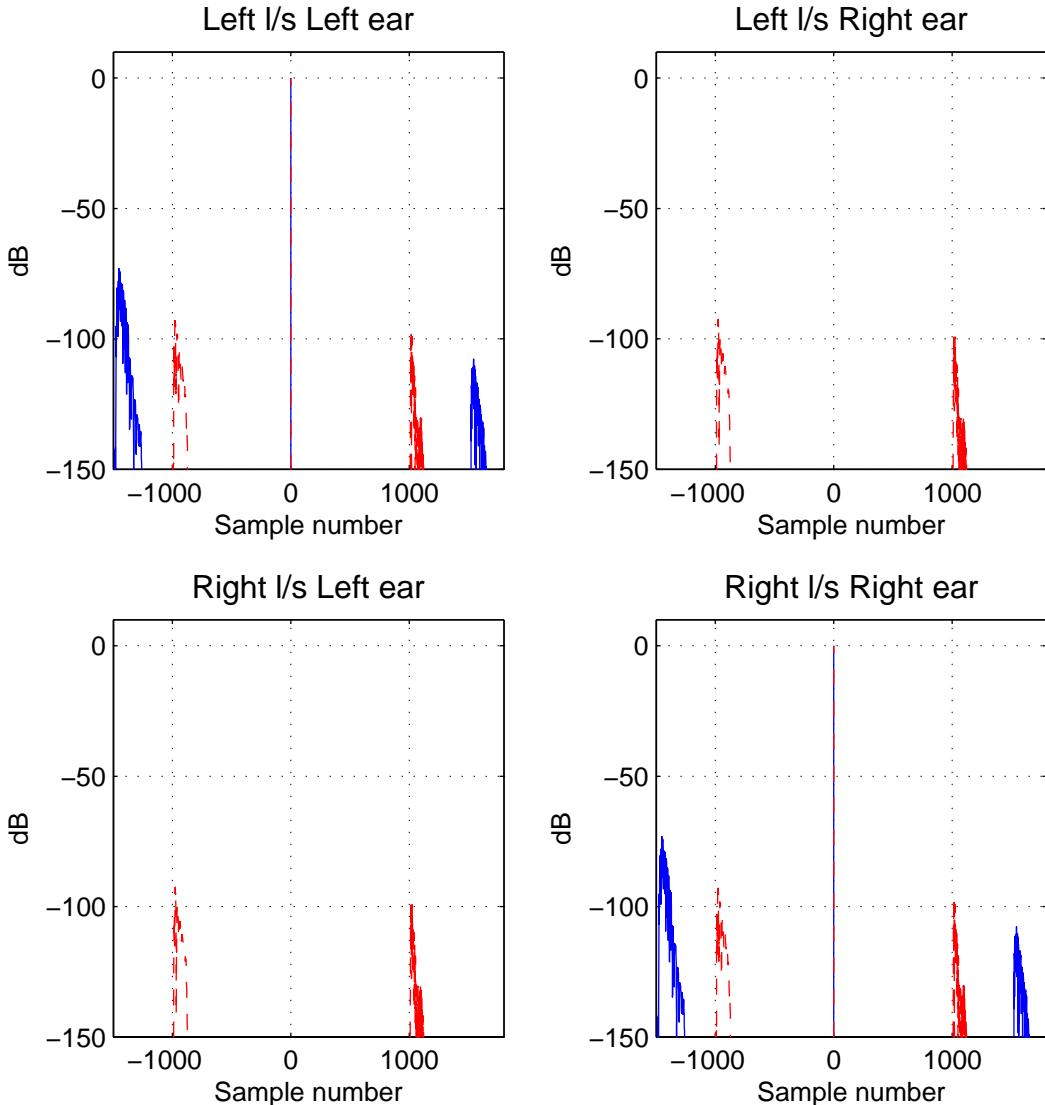
In figure 2-6 we plot the impulse responses of the inverse matrix  $\mathbf{H}$  defined from the MIT database of HRTFs (Gardner and Martin 1994) for the Stereo Dipole geometry with the regularisation set equal to zero. In the red line we plot the four FIR filters  $h_{ij}(n)$  that model the exact inverse matrix  $\mathbf{H}$  with 1000 anti-causal and 1000 causal coefficients<sup>31</sup>. In the blue line we plot the FIR filter  $h_{CD}(n)$  that models the single-channel transfer function  $1/\det[\mathbf{C}(z)]$  of the common denominator with 1500 anti-causal and 1500 causal coefficients. Hence the total cost for the computation of each sample pair of the output would be 8000 multiplications and additions using the direct arrangement of equation (2-26) and 6500 multiplications and additions using the common denominator arrangement of equation (2-27).

In figure 2-7 we compare the inversion results obtained with these two types of inverse models. With the red dashed line we plot the results obtained with the direct arrangement and with the blue line the results of the common denominator arrangement. As is made apparent by these results, the 1500 coefficients used to model the anti-causal part of the reciprocal common denominator transfer function result in a truncation of the model  $H_{CD}$  at a higher level of its decay tail than is the case with the 1000 coefficients used for the modelling of the anti-causal delay of the filters  $H_{ij}$ . This is because the presence of zeros in the transfer functions  $H_{ij}$  (see equation (1-4)) results in a faster decay of the corresponding time-responses than is the case for the all-pole transfer function  $1/\det[\mathbf{C}(z)]$ <sup>32</sup>. Hence, for an equal suppression of these truncation artefacts one would have to use a longer model of the reciprocal common denominator transfer function and thus practically eliminate any cost saving. Additional to that, such a realisation would also introduce more delay in the inversion.

---

<sup>31</sup> All models discussed here are computed using the frequency-domain method described in §2.3.

<sup>32</sup> This “all-pole effect” is restricted to the exact inversion case. As was discussed in §2.3.3, with the introduction of regularisation each pole due to the common denominator is replaced by a pair of a pole and zero. Hence the effect described in figure 2-7 should not be present in regularised inversion cases. As is also discussed in chapter 8, a complete investigation of the potential presented by the separation of equation (2-27) is left as further work.



**Figure 2-7:** Simulated inversion results. The plant matrix  $\mathbf{C}$  is convolved with a model of the inverse  $\mathbf{H}_D$  (red dashed line) and with the matrix corresponding to the common denominator filtering process (blue solid line).

Note that in figure 2-7 the cross-talk appears to be completely suppressed in the realisation utilising the common denominator. This should be considered to be an artificial effect. In effect, the common denominator arrangement effectively separates the cross-talk stage, applied by the  $\text{adj}[\mathbf{C}(z)]$  part from the overall equalisation stage applied by the  $1/\det[\mathbf{C}(z)]$  part<sup>33</sup>. As is readily seen in equation (1-4), the multiplication of the matrices  $\mathbf{C}(z)$  and  $\text{adj}[\mathbf{C}(z)]$  results in a matrix with zero off-diagonal terms whereas the term  $1/\det[\mathbf{C}(z)]$  effectively equalises the diagonal terms

<sup>33</sup> As was reviewed in §2.2, such a separation has appeared in previously proposed inverse filtering implementations but it has not been investigated in the context of the specific inverse filtering design discussed here.

to unity. When the inverse is directly realised as in equation (2-26), any error in the computation of the inverse (in this case due to the truncation of the infinitely long time responses  $H_{ij}$ ) will be present in all four elements of the resulting matrix  $\mathbf{X}_{eq} = \mathbf{CH}$ . On the other hand, in the simulation example presented here, the matching of the model of  $\text{adj}[\mathbf{C}(z)]$  with the plant model  $\mathbf{C}(z)$  is perfect, so the cross-talk part will be apparently implemented perfectly. The truncation is only applied to the model of the  $1/\det[\mathbf{C}(z)]$  term and is hence only visible in the diagonal terms. However, in realistic cases where the available model of the plant is not perfect (see the analysis in chapter 4) such a perfect suppression of the cross-talk terms is by no means guaranteed.

## 2.5 *Conclusions*

In this chapter we have reviewed the available literature relevant to the inverse filtering problem at hand. A specific type of inverse design was presented in detail and some of its basic properties highlighted. This type of inverse design, implemented in the manner described in §§2.3.3-2.3.4, is utilised in the following chapters where a detailed evaluation of its performance is presented.

## ***3 Objective evaluation of inverse filtering –***

### ***Motivation and methodology***

#### ***3.1 Introduction***

In this chapter we describe the experimental arrangement used to obtain the results presented in chapters 4 and 5. First, in §3.2 we discuss the original motivation for the design of such an experimental procedure and we outline the questions we set out to answer. The experimental arrangement itself is described in §3.3. Finally, in §3.4 and §3.5, we elaborate on certain aspects of the experimental procedure that are of importance in regard to the results presented in the following two chapters.

#### ***3.2 Motivation for the undertaking of the chosen objective evaluation experimental procedure***

On the basis of the specific inverse filtering design presented in §1.2 and §2.3 (and equally on the basis of most of the other designs reviewed in §2.2) lies a common set of assumptions regarding the conditions for the implementation of the inverse filtering. These are: (i) the assumption that we have exact knowledge of the plant and (ii) the assumption that the transducers involved in the control behave perfectly linearly and are transparent i.e. that they have a flat magnitude response and linear phase across the audio frequency range. For the first of these assumptions to be completely satisfied the following conditions have to hold:

- The reproduction has to take place in an anechoic environment or the plant has to be measured in-situ.

- The individual HRTFs of the listener for the specific reproduction geometry have to be known.
- The reproduction geometry has to be fixed or, if the listener is allowed to move, the inverse filters need to adapt on-line. In this last case the individual HRTFs of the listener for the whole set of possible listener-source geometries have to be known.

Obviously not all of the conditions listed above are fully satisfied in everyday audio reproduction situations. On the contrary, more often than not, none of them holds. Usually, the inverse filter matrix  $\mathbf{H}$  is designed using a model of the plant that has been measured in anechoic conditions. However, in most realistic cases the actual plant contains reflections as is the case in all home entertainment applications. Furthermore, hardly ever are individualised HRTFs used and even in specifically designed 3D-audio experiments the plant is usually modelled using HRTFs measured with a dummy head. Finally, the responses corresponding to the exact source/listener geometry are usually not available. In that case the plant is approximated using the responses of the closest available position (or an interpolation between the closest available positions) among a spatially sampled set of positions.

Similarly, the second assumption mentioned above can be only approximately satisfied when the non-ideal properties of real transducers are considered. This is especially true for medium-priced loudspeakers which, being the weak link of the audio reproduction chain, exhibit non-linear behaviour when over-driven, respond poorly at the low and high end of the audio spectrum and have non-flat response at their crossover region(s) and off-axis.

These deviations from the ideal model, along with any imperfections or approximations inherent in the design of the inverse filtering network, will of course degrade the performance of the implementation. This fact is readily verified even in simple listening tests of virtual acoustic imaging systems of this type. Front-back reversals, increased localisation blur, reduced quality of the presented audio material, colouration and artefacts, erratic behaviour of virtual sound images placed at the median plane are the most common types of errors encountered in applications of the design where one or more of the aforementioned elements of error are present.

Numerous studies have been previously published that examine the influence of those elements of error on the effectiveness of a number of different inverse filtering designs among the ones reviewed in §2.2. The performance of the inverse filtering when implemented in a non-anechoic environment (with the inverse filters determined either on the basis of anechoic models of the plant or of models that include early reflections and reverberation) is examined by Farina and Ugolotti (1997), Gardner (1998), Köring and Schmitz (1993), Nelson et al. (1996) and Takeuchi (2001). The influence of listener positioning errors is examined by Gardner (1998), Nelson et al. (1997), Rose (2004) and Takeuchi (2001). The feasibility of using non-individualised plant models for the computation of the inverse is discussed and examined by Gardner (1998), Köring and Schmitz (1993), Martens (1997) and Takeuchi (2001). The use of non-individualised HRTFs (measured or otherwise synthesised) in the binaural synthesis stage is investigated and discussed by Asano et al. (1990), Evans et al. (1998), Evans et al. (2000), Kistler and Wightman (1992), Kulkarni and Colburn (1998), Kulkarni et al. (1999) and Wenzel et al. (1993). Finally, issues related to the effectiveness of the specific inverse filtering design of §2.3 in comparison to other inverse filtering designs are examined by Farina and Ugolotti (1997). The same issues are also discussed by Fielder (2003), Mourjopoulos (1994), Mourjopoulos (2003) and Norcross et al. (2004) in the context of the problems of single channel equalisation and room dereverberation. An important feature of these latter studies in relation to our study, despite the fact that they are confined to a single channel context, is that they offer significant insight into the relation between the design parameters of the inverse (i.e. the length of the inverse filters and the amount of regularisation introduced) and the expected perceived quality of the inversion.

Nevertheless, despite this wealth of existing literature, we believe that a certain group of questions relevant to the research topic of this thesis remain to a great extent unanswered. By this we mean that there does not exist a definite analysis for the optimal design parameters of the inverse filter matrix  $\mathbf{H}$  that takes into account the imperfections and approximations inherent in real-world reproduction situations. In such real-world situations, the presence of errors and approximations, like the ones described above, renders the use of an exact realisation of the inverse matrix  $\mathbf{H}$  unnecessary. Sub-optimal realisations resulting, for instance, from the use of higher regularisation values and shorter inverse filter lengths could yield the same (and in

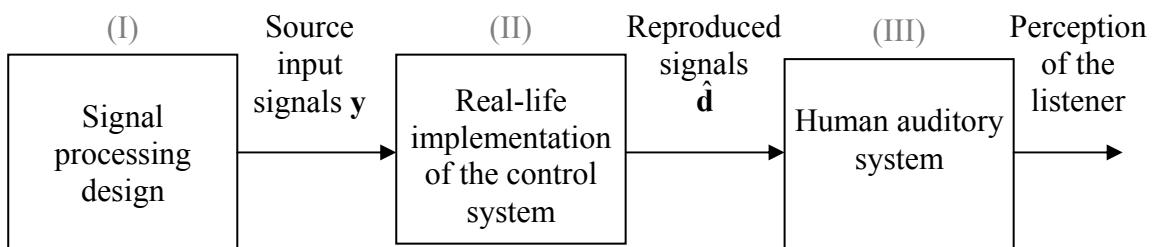
some cases even better) inversion results as those obtained by the exact realisation of  $\mathbf{H}$ . The reason why such an investigation of the realistically optimal design parameters of  $\mathbf{H}$  is not presented in the aforementioned studies is that all those studies share as a common characteristic the fact the evaluation of the inverse filtering is based either on computer simulations of the inversion process or on subjective experiments.

Computer simulations indeed offer a versatile and easy-to-implement method for the investigation of the isolated influence and the interrelation of each one of the inverse filter design parameters on the system's performance. Their downside, however, is that they typically overestimate the efficacy of the inversion as they rely on implicit assumptions regarding the perfect linearity and ideal response of the sources, the absence of noise or other sources of error in the measurement of the plant etc. Especially when the question of the most accurate, realistically achievable level of inversion is considered, computer simulations (as will be seen in the following) can lead to exceedingly optimistic results and can hence offer misguiding information regarding the optimal design parameters.

Conversely, evaluation results that are based on subjective listening tests are guaranteed to exactly reflect the overall efficacy of the inverse filtering design in terms of the perceived impression of the listener. On the other hand, the involvement of human subjects restricts the versatility of the experimental design. The exact control of factors like the listener's positioning or the use of individualised plant measurements becomes non-trivial. Hence the inadvertent presence of such types of error in the implementation of the inversion can conceal the influence of the design parameter under investigation. Furthermore, the quantifiable results of such experiments are inevitably coupled with the response of the human auditory system thus further obscuring the association of the specific choice of the design parameters with the actual effectiveness of the imposed control.

The above are summarised schematically in figure 3-1 which helps describe the rationale of the work presented in this and the following two chapters. In the figure, the block labelled (I) stands for the specific signal processing design (along with the choice of the related design parameters) that is used to derive the source input signals  $\mathbf{y}(n)$  of figure 1-1. The block labelled (II) represents the actual implementation of the

design and incorporates any possible deviations from the ideal-implementation assumptions listed above (see p. 55). The output of this stage is the signals reproduced at the listener's ears. The block labelled (III) represents the human auditory system that includes the human sound localisation mechanism and all the other known or unknown attributes and mechanisms (various masking effects, the precedence effect, thresholds of perceptibility of the specific characteristics of the reproduced signal, etc.). The output of this stage is of course the auditory perception of the listener.



**Figure 3-1:** Conceptual stages in the realisation of a virtual acoustic imaging system.

As discussed above, many studies exist that investigate and quantify the effectiveness of the design on the basis of the outputs of the first and third blocks of figure 3-1. Hardly any have appeared however to provide insight into the output of the second block. The results presented in the following chapter aim to provide a closer look at the influence of what may be incorporated in the block (II) of figure 3-1 and hence to relate the specific choices made in the design of the inverse filtering network to the realistically achievable audio control. Even though some discussion is provided regarding the expected influence of such findings on the auditory perception of the listener, the complete investigation of this issue is left outside the scope of this research project.

The obvious interest in undertaking the analysis outlined above is that its outcome could show a significant reduction of the required specifications for the implementation of  $\mathbf{H}$ , the most important of these specifications being the length of the inverse filters in  $\mathbf{H}$ . A second point of interest in such an analysis, one that is even more important for the development of the research presented in this thesis, is that its results are used as a realistic “benchmark” of accuracy in view of the alternative reduced-order recursive designs proposed in chapter 6.

The goals of the analysis presented in the following two chapters are therefore the following:

- (i) To verify the extent to which the inversion of a given measured plant with the inverse model of §2.3 is successful in ideal conditions of implementation and to give a detailed investigation of the role that the parameters related to the design of the inverse play in the actual accuracy of the inversion. This is done in chapter 4 of the thesis.
- (ii) To examine to what extent the full audio frequency range inversion of a measured model of the plant of §2.3 (with the design challenges it entails) is justified in non-ideal conditions of implementation. This is done in chapter 5

### ***3.3 Experimental arrangement***

The experimental arrangement that was used to obtain the results presented in the following two chapters is described in figure 3-2. The arrangement was based on a PC equipped with the HURON Audio Workstation (Lake 2003). A deterministic pseudo-random stimulus with pink spectrum was fed through HURON's audio I/O to a stereo power amplifier<sup>34</sup> that drove a pair of bookshelf two-way loudspeakers<sup>35</sup>. The loudspeakers were positioned with the centre of their units spaced 26.5cm apart. KEMAR dummy-head, fitted with the DB-60 and DB-61 (small) right and left pinna models<sup>36</sup>, was placed on the axis of symmetry of the two loudspeakers with the centre of its head 150cm from the face of the loudspeakers. This arrangement gives a total angle span of 10° between the lines joining the centre of KEMAR's head with the two loudspeakers (the Stereo Dipole geometry). The loudspeakers and KEMAR were placed inside the large anechoic chamber of ISVR. The grid floor was not removed from the chamber but the whole area between the loudspeakers and KEMAR as well as the area surrounding the arrangement by approximately 1m was covered with

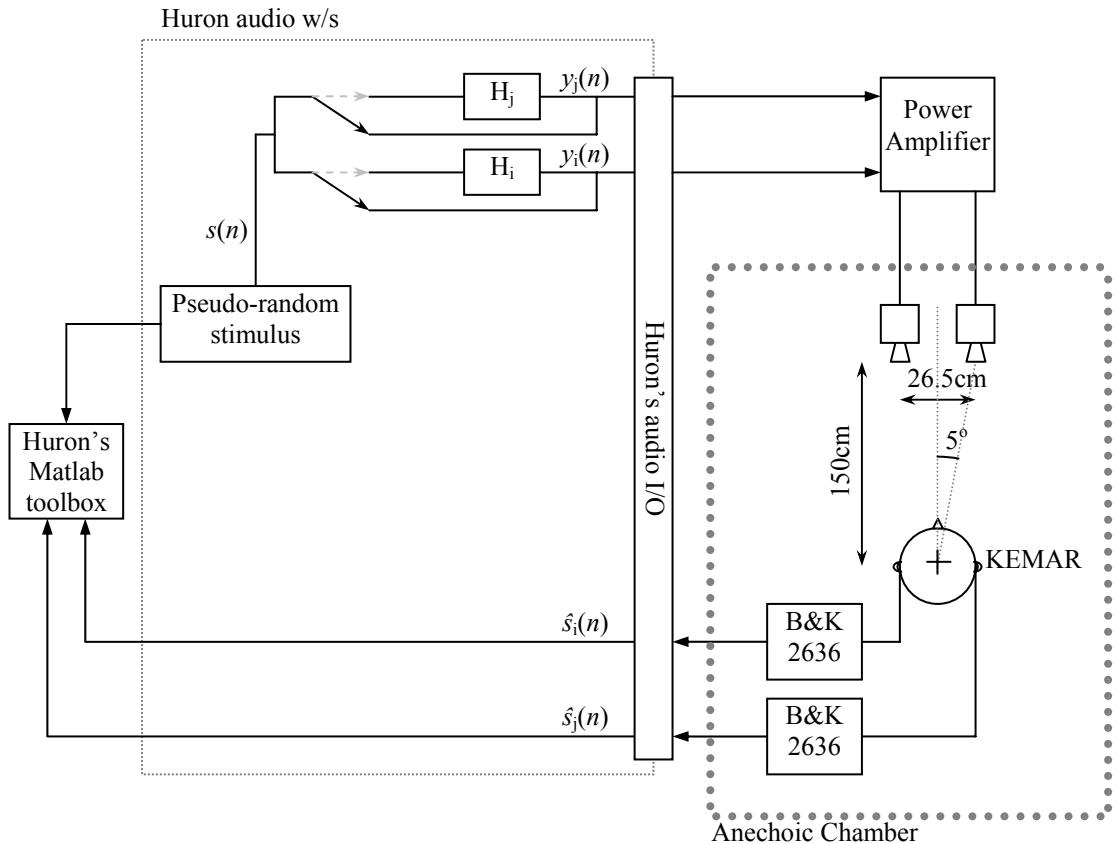
---

<sup>34</sup> The power amplifier used was Yamaha H5000. This model offers a gain control knob that can be adjusted in the range from -∞dB to 0dB in discrete values. For more details see *Yamaha Power amplifier H7000/H5000/H3000, Owner's manual*. Hamamatsu Japan. (available online at [http://www.yamaha.co.jp/english/product/proaudio/products/power\\_amps/h7000/index.html](http://www.yamaha.co.jp/english/product/proaudio/products/power_amps/h7000/index.html)).

<sup>35</sup> The “Celestion 1” loudspeakers were used for the measurements. They comprise a 5" midrange unit and a 1 1/8" and use bass reflex.

<sup>36</sup> See (Burkhard and Sachs 1975), (Burkhard et al. 1985) and (Maxwell and Burkhard 1979).

absorbing foam. Two  $\frac{1}{2}$  inch microphones mounted at the position of KEMAR's eardrums were used to pick up the sound pressure signals at KEMAR's ears (i.e. an open ear canal measurement). The microphone signals were input back to HURON's audio I/O through a pair of measuring amplifiers<sup>37</sup>.



**Figure 3-2:** Measurement apparatus.

All signals obtained from the ADCs and reaching the DACs of HURON's I/O were sampled at 48kHz and quantised at 24bit word-length. The input signals from the ADCs were uploaded to Matlab together with the pseudo-random stimulus and the response of the electroacoustic system was computed in Matlab's 64bit double-precision arithmetic by use of the *tcdeconv* function included in Matlab's HURONMAT toolbox (Lake 2003).

<sup>37</sup> The B&K 4134 pressure-type microphones were used with B&K 2619 preamplifiers and B&K 2636 measuring amplifiers.

With the (notional) switch in the position drawn in figure 3-2 with solid black line, the measured response corresponds to the electroacoustic plant to be inverted. This response includes KEMAR's HRTFs for this specific geometry and also the responses of the power and measuring amplifiers, the microphones, the I/O stage of HURON and the responses of the loudspeakers. The pink pseudo-random stimulus used for the measurement of the plant was 5.46s long ( $2^{18}$  samples at 48kHz sampling rate) and the microphone signals were averaged 10 times.

The inverse filtering was realised in real-time with FIR filters implemented in HURON as noted in the schematic of figure 3-2 with the grey dashed-line notional switches. As can be seen in equation (3-1), this arrangement was used to measure the ipsilateral equalisation in one channel and the cross-talk cancellation from that channel to the opposite ear. This is done by considering the case where the binaural input signal vector  $\mathbf{x}$  and the desired output signal vector  $\mathbf{d}$  is equal to  $[s_1 \ 0]$ , i.e. the case where a signal  $s_1$  is to be reproduced in the left channel with no signal in the right channel. The signals  $\hat{s}_{11}$  and  $\hat{s}_{21}$  captured at the two microphones will then be given by:

$$\begin{aligned} \begin{bmatrix} \hat{s}_{11} \\ \hat{s}_{21} \end{bmatrix} &= \mathbf{CH} \begin{bmatrix} s_1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} s_1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} (C_{11}H_{11} + C_{12}H_{21})s_1 \\ (C_{21}H_{11} + C_{22}H_{21})s_1 \end{bmatrix} \end{aligned} \quad (3-1)$$

Clearly then, with  $s_1$  being again a pseudo-random stimulus with pink spectrum<sup>38</sup> and with FIR models of  $H_{11}$  and  $H_{21}$  implemented in the filters  $H_i$  and  $H_j$  of figure 3-2, the response between the input  $s_1$  and the output  $\hat{s}_{11}$  corresponds to the (1,1) element of the equalisation and crosstalk cancellation matrix  $\mathbf{X}_{eq}$  of equation (1-6). Similarly, the response between the input  $s_1$  and the output  $\hat{s}_{21}$  corresponds to the (2,1) element of the same matrix. In the same manner, substituting  $H_{12}$  and  $H_{22}$  for  $H_i$  and  $H_j$  and measuring the two responses between the stimulus  $s_1$  reproduced by the second (right)

---

<sup>38</sup> A different stimulus of the same characteristics as above, but of 2.73s duration was used for the measurement of the equalised responses. The measured outputs were again averaged 10 times.

channel and the microphone outputs  $\hat{s}_{12}$  and  $\hat{s}_{22}$  we measure the elements (2,1) and (2,2) of  $\mathbf{X}_{eq}$ .

Before being loaded onto HURON, the inverse filters were scaled so that both the coefficients of the four filters  $h_{ij}(n)$  and the four output signals  $y_i(n)$  were kept in the range [-1,1]. This scaling is hereafter referred to as normalisation. We note that, since the stimulus signal  $s_1(n)$  is scaled to occupy the full range [-1,1], the amount of normalisation introduced is equal to the loss of dynamic range in the digital part of the reproduction setup.

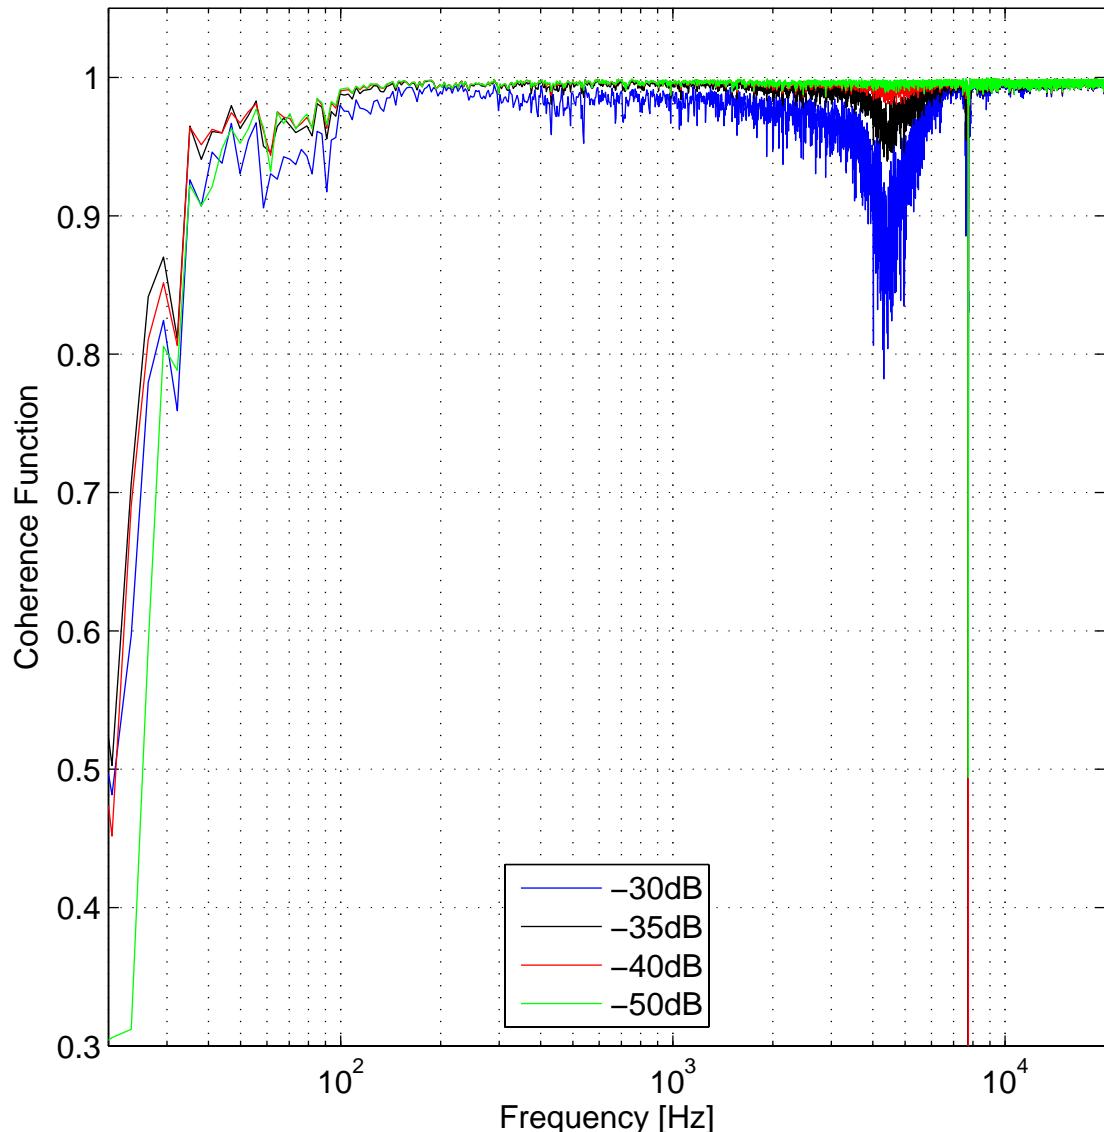
A choice needed to be made regarding the level of amplification of the pseudo-random stimulus before it was fed to the loudspeakers. This choice reflects a compromise between (i) the desired increase in the measurement's dynamic range achieved when the stimulus is further amplified by the power amplifier and (ii) the undesired excitation of significant non-linear behaviour in the audio reproduction chain (especially the loudspeakers) when a given level of amplification is exceeded. A measure of the presence of non-linearities and extraneous noise in the measurement is the coherence function between the input and output signals (Bendat and Piersol 1993). This is plotted in figure 3-3 for four different levels of amplification of the input signal corresponding to the values -50dB, -40dB, -35dB and -30dB of the power amplifier gain. The results of figure 3-3 were obtained for the left side ipsilateral response measurement (i.e. stimulus reproduced from the left loudspeaker and measured from the left ear microphone). The pink noise pseudo-random stimulus was 5.46s long ( $2^{18}$  samples at 48kHz sampling rate) and the microphone output signal was averaged 10 times<sup>39</sup>.

As can be seen in the figure, the coherence drops significantly in the region of 2-7kHz (loudspeaker crossover region) for the amplification levels of -35dB and -30dB. This loss of coherence should be attributed to non-linear behaviour. For the levels of -50dB and -40dB the coherence is similar, with the -40dB measurement showing a slightly better coherence below 50Hz but slightly worse in the region 4-6kHz. The coherence drop at 8kHz is due to a notch of the left ipsilateral response at this frequency (see also figure 3-5). Similar results were obtained for the remaining three responses

---

<sup>39</sup> The coherence was computed using Matlab's *cohere* function (MathWorks 2000).

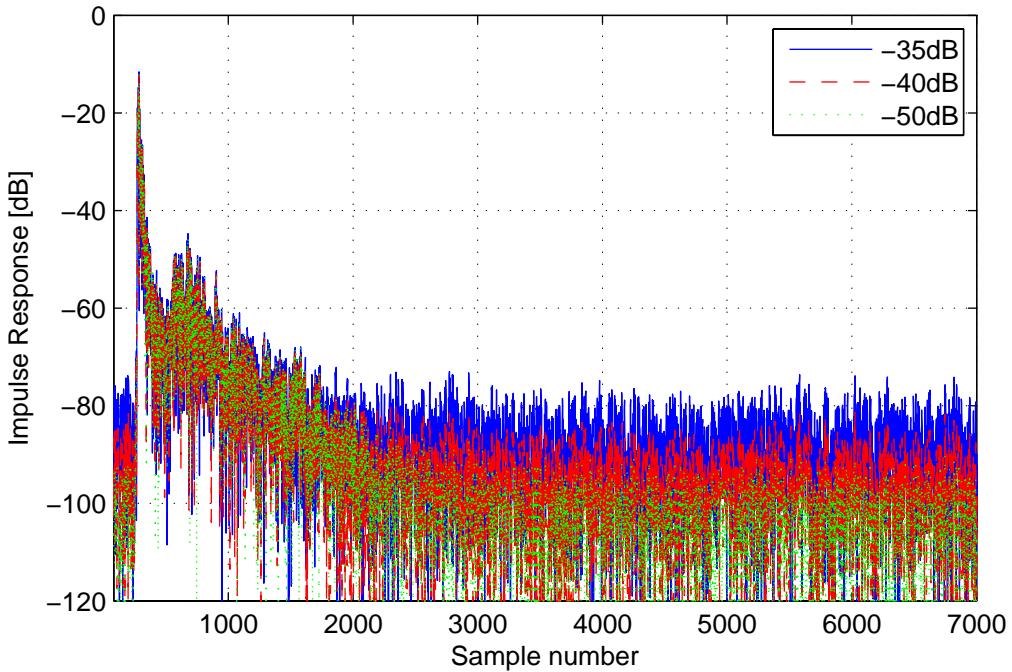
comprising the plant (i.e. the right-side ipsilateral and the two contralateral responses).



**Figure 3-3:** Coherence function of the plant measurement (left loudspeaker to left ear) with the power amplifier level set to -50dB (green line), -40dB (red line), -35dB (black line) and -30dB (blue line).

The reduction of the usable dynamic range of the plant measurement as the level of the power amplification is increased above the -50dB value, can also be seen in the time-domain version of the measured response. This is shown in figure 3-4 where we plot the impulse responses obtained with the power amplification set to -50dB, -40dB and -35dB. The -40dB and -35dB measurements are shifted down by 10dB and 15dB respectively and are thus directly comparable to the -50dB measurement. The

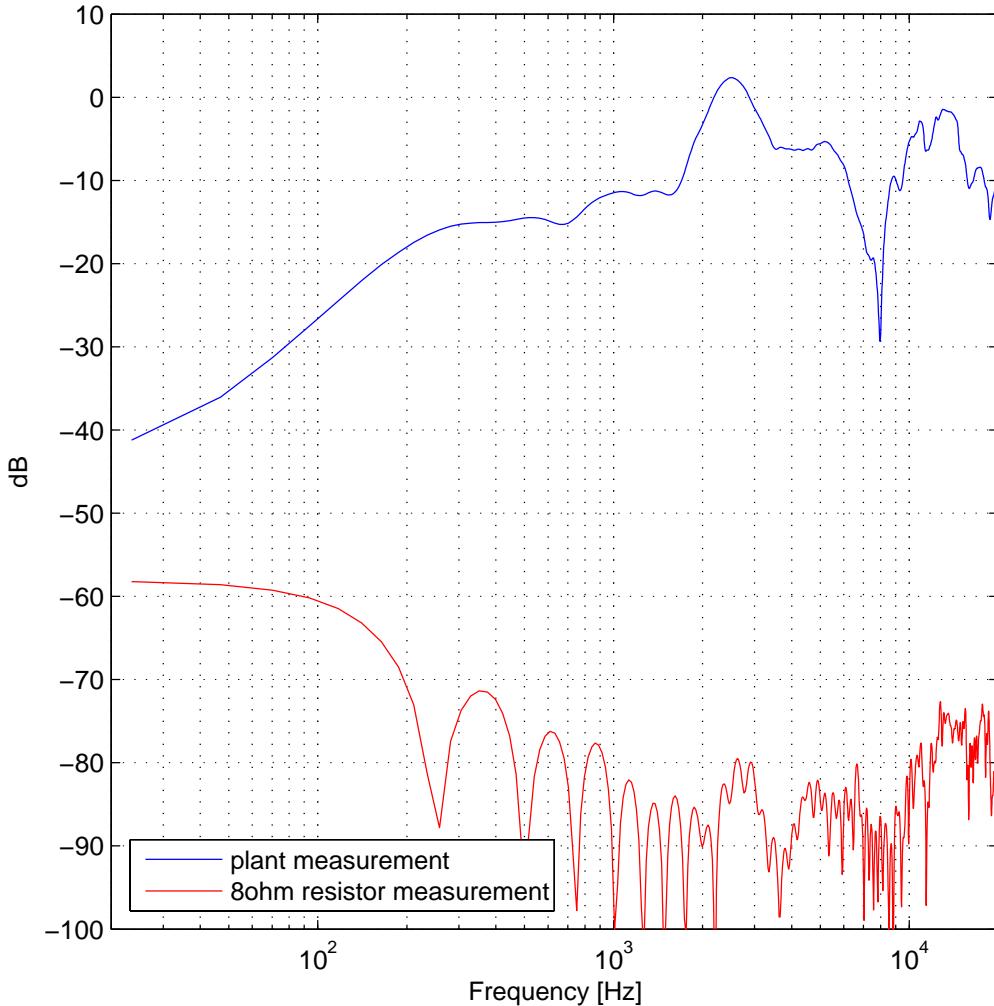
difference in the dynamic range of the -50dB and -40dB measurement is clearer in this plot as the level of the noise-floor in the latter has risen relatively higher than that of the former due to increased non-linear behaviour (Muller and Massarani 2001).



**Figure 3-4:** Squared impulse responses of the left ipsilateral side on a logarithmic scale. The amplifier amplification level is set to -35dB (blue solid line), -40dB (red dashed line) and -50dB (green dotted line). The -40dB and -35dB responses are shifted downwards by 10dB and 15dB respectively.

The sound pressure level at the dummy-head's ear microphones was measured at 71dB SPL (70.5dBA) for the -50dB setting. This level is close to the levels typically used in subjective experiments and listening tests (Møller 1995a) and hence we chose to use the -50dB gain level throughout the experiments. The measurement dynamic range with these settings compares favourably with the two previously published measurements that present similar data. Namely, Gardner and Martin (1994) estimate the dynamic range as the ratio of the total energy in 100 samples of the measured impulse response centred on the energetic part of the measurement to the energy in 100 samples of the initial delay that should ideally be zero. The quoted SNR is 65dB,

while the same estimate in the left ipsilateral impulse response of the measurements presented here is 73dB<sup>40</sup>.



**Figure 3-5:** Plant response measurement (blue line) and measurement of response with 8ohm resistor connected at the place of the loudspeaker (red line).

A different type of estimate for the dynamic range of the measurement arrangement is presented by Møller (1995a). In that case the dynamic range is estimated as the headroom between the transfer function of a representative measured response and the transfer function of a measurement with the output of the power amplifier connected to an 8-ohm resistor. The same procedure was followed with our arrangement and the

<sup>40</sup> In (Gardner 1998) the SNR estimate is calculated using the impulse response measured with the source directly ahead of KEMAR while in our measurements the estimate is taken on the measurement with the source 5° to the left of KEMAR. The difference should be considered negligible.

results are presented in figure 3-5. A headroom of more than 60dB between the measurement and the noise floor can be seen for the frequency range above 200Hz, a result equivalent to that in (Møller 1995a).

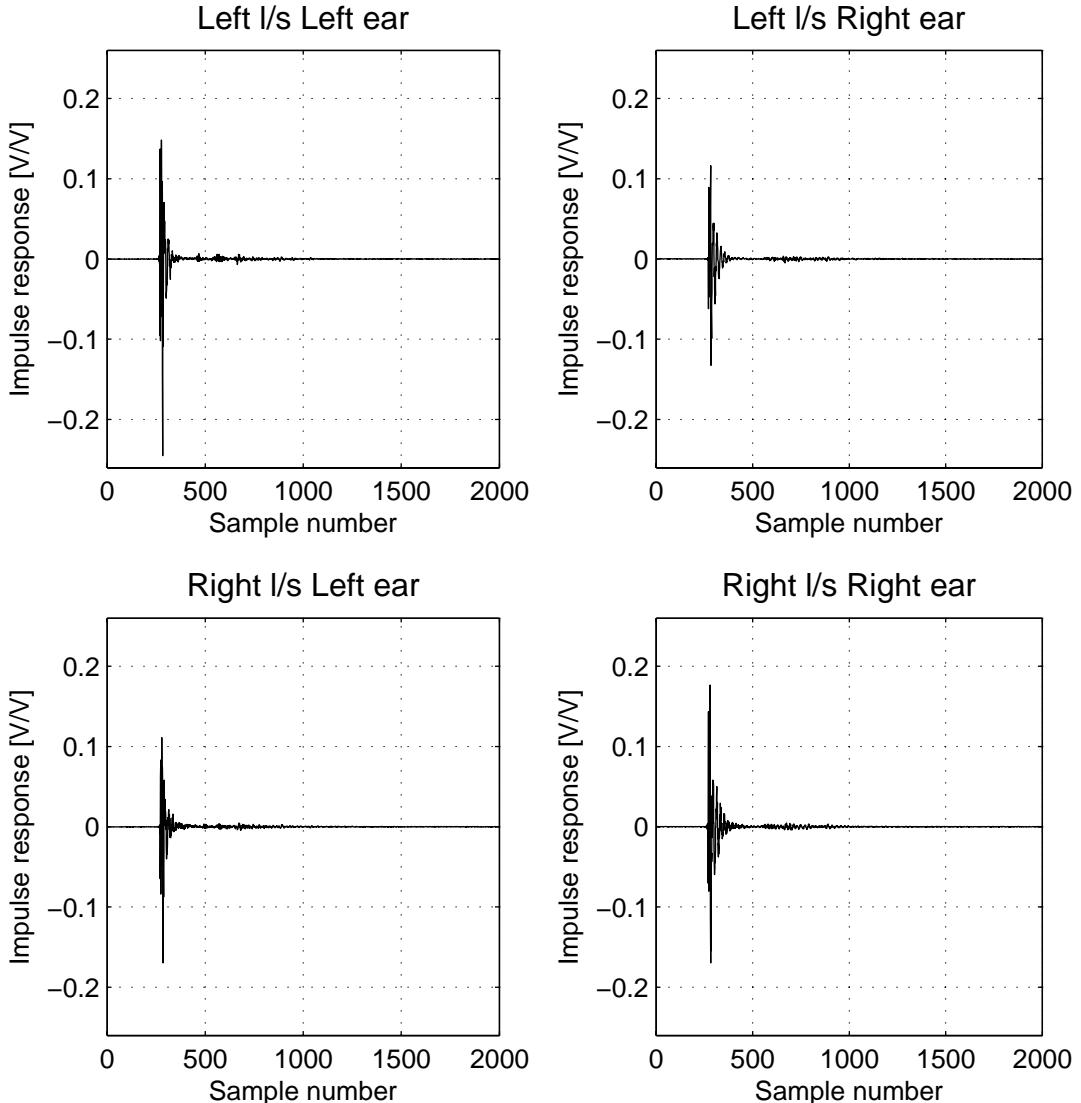
### **3.4 The measured model of the plant matrix**

Using the experimental arrangement of figure 3-2 (with the notional switches in the lower, solid black line position) we measured the four plant impulse responses of figure 3-6. As can be seen in the figure, the responses start with a delay of approximately 260 samples that corresponds to the flight time from the loudspeakers to the ears of KEMAR and the delay in the measurement system. Their initial part persists until about the 450<sup>th</sup> sample, i.e. a time-duration of approximately 200 samples which at the chosen sample rate of 48kHz corresponds to approximately 4ms. This time-duration is in agreement with previously published measurements of anechoic impulse responses of KEMAR and human subjects by Algazi et al. (2001), Gardner and K. Martin (1994) and Møller et al. (1995b) and with reviews on the implementation of the binaural technique (see Jot et al. (1995) and references therein).

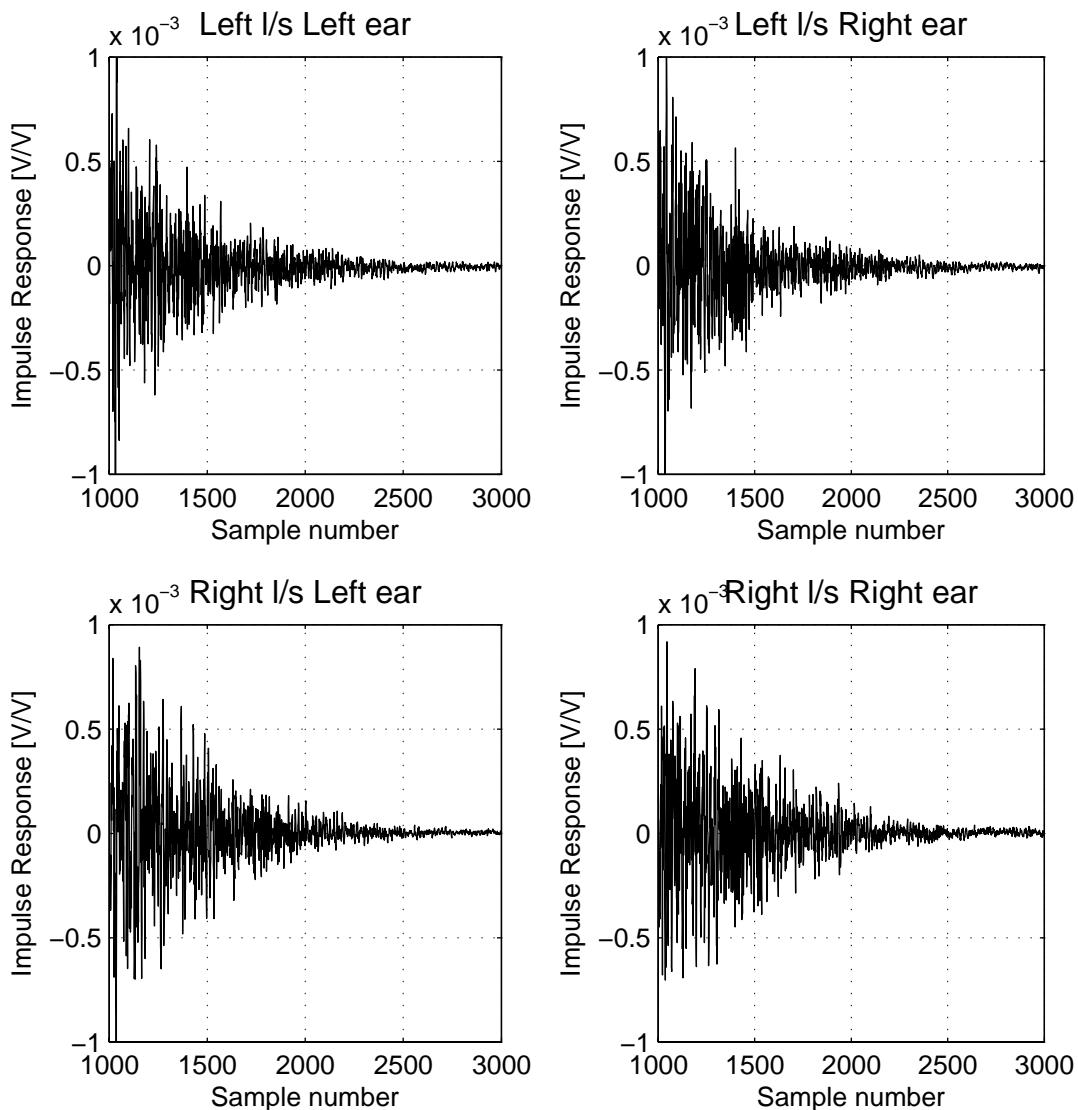
A series of individual reflections, caused by the stands of KEMAR and the loudspeakers and possibly the part of the grid floor that was not covered with foam, starts at approximately the 450<sup>th</sup>-500<sup>th</sup> sample. These individual reflections persist until about the 1000<sup>th</sup> sample of the impulse responses and are followed by a reverberation-like tail that, as can be seen in figure 3-7, decays down to the level of the noise floor between the 2000<sup>th</sup> and 2500<sup>th</sup> sample.

We introduce here the following notation to refer to models of the plant matrix  $\mathbf{C}$  corresponding to different lengths of the measured impulse responses of figure 3-6. With  $\mathbf{C}_{190}$  we denote the matrix containing the impulse responses  $c_{ij}(n)$  with the sample index  $n$  in the range  $260 \leq n \leq 450$ . Hence, the  $\mathbf{C}_{190}$  version of the plant model contains only the strictly anechoic part of the measured impulse responses. With  $\mathbf{C}_{740}$  we denote the matrix containing the impulse responses  $c_{ij}(n)$  with the sample index  $n$  in the range  $260 \leq n \leq 1000$ . This second model of the plant matrix also includes the

first individual reflections present in the measured impulse responses. Finally, with  $\mathbf{C}_{1740}$  we denote the matrix containing the impulse responses  $c_{ij}(n)$  with the sample index  $n$  in the range  $260 \leq n \leq 2000$ . In this third case the late decay part of the measured impulse responses down to the noise level is also included.



**Figure 3-6:** Measured plant impulse responses.



**Figure 3-7:** Decay tail of the measured plant responses of figure 3-6.

It is evident that a different realisation of the inverse filter matrix  $\mathbf{H}$  should result when each one of these different models of the plant is used in the method for the determination of a model of the inverse described in §2.3. We denote these different cases for the inverse with  $\mathbf{H}_{190}$ ,  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$ . It should be clear of course, that each of these symbols for the inverse matrix  $\mathbf{H}$  does not correspond to a unique realisation but to a whole group of different realisations varying with respect to the remaining inverse design parameters, i.e. the length of the inverse and the amount of regularisation introduced. The effectiveness of each of these models of the inverse is examined in the results of the following chapter.

### 3.5 On the accuracy of the computer-simulation of the inversion process

In this section we discuss the computer simulation of the inversion process and we describe how such properly designed simulations are used in the results presented in the following two chapters. More specifically, we show that when the inversion process is simulated by convolving a properly chosen model of the plant matrix  $\mathbf{C}$  with the model of the inverse matrix  $\mathbf{H}$  under evaluation, the outcome is indeed in good agreement with the measured results of the inversion process. On the other hand, this is not the case when the strictly anechoic model of the plant matrix is used for the convolution.

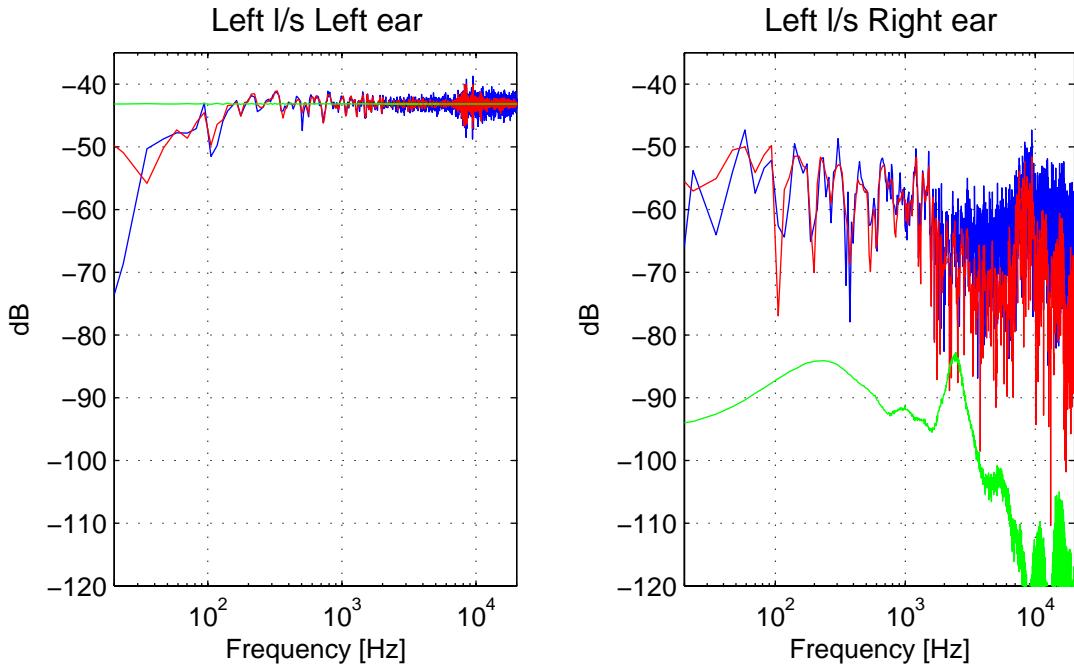
The comparison of the measured inversion results with the two different cases of simulation results is plotted in figure 3-8. The same model of the inverse matrix  $\mathbf{H}$  is used in all three cases depicted in the figure. This is  $\mathbf{H}_{190}$ , i.e. the inverse computed by substituting the plant model  $\mathbf{C}_{190}$  in the method of §2.3. In this instance the regularisation is set to the value  $\beta=0$ , the inverse length is set to 4000 coefficients and the modelling delay  $\Delta$  is set to 2000 samples.

With blue line in figure 3-8 we plot the inversion results measured by means of the experimental arrangement of figure 3-2 (with the notional switches in the upper, dashed grey line position). This is then the actual, directly measured, performance of the inversion against which the simulation process should be compared. With red line in the figure we plot the results obtained when the model of the plant containing the measured impulse responses  $c_{ij}(n)$  with the sample index  $n$  in the range  $1 \leq n \leq 3000$ <sup>41</sup> is convolved with the properly normalised (see p. 63) chosen model of the inverse  $\mathbf{H}_{190}$ . Finally, with green line we plot the results obtained with the strictly anechoic model of the plant  $\mathbf{C}_{190}$  is convolved with  $\mathbf{H}_{190}$ . We note that this latter method is the one typically used when computer simulations are used for the evaluation of the inversion process. In such cases, a model of the plant that corresponds to the strictly anechoic part of the plant responses is used for the determination of the inverse and this inverse is then convolved with the same anechoic model of the plant. Obviously

---

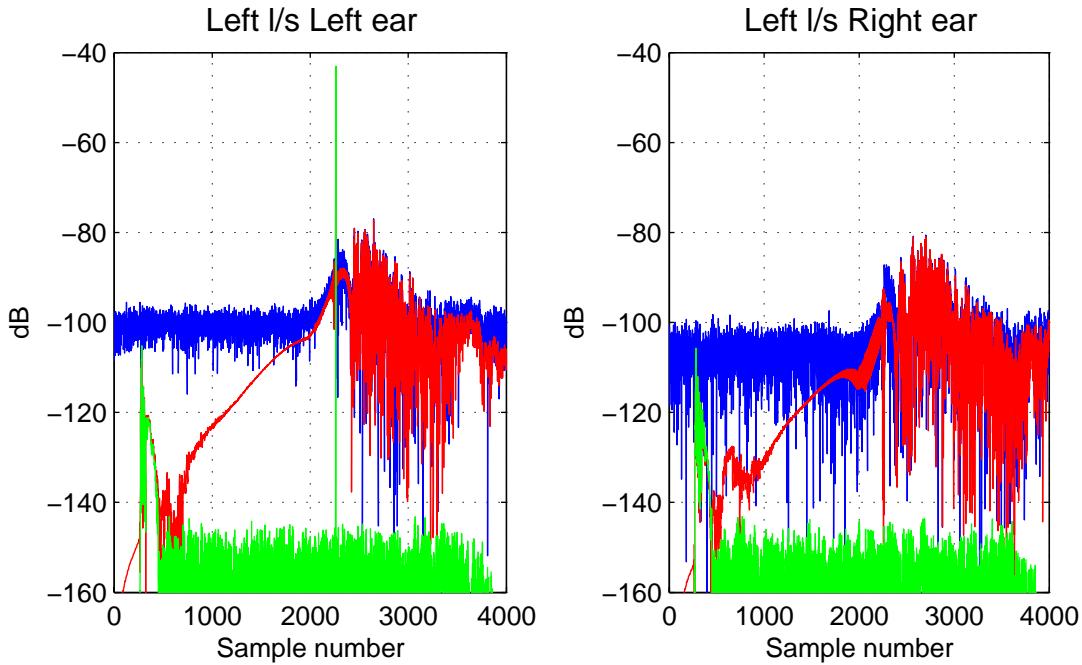
<sup>41</sup> Henceforward we refer to this model of the plant as the “full-length model of the plant”.

then, as is confirmed with the green line results of figure 3-8, when no regularisation is used and the length of the inverse is chosen to be sufficiently high, it is straightforward to (erroneously) obtain perfect inversion results.



**Figure 3-8:** Measured inversion results for  $\mathbf{H}_{190}$  (inverse computed using  $\mathbf{C}_{190}$  as the plant model,  $\beta=0$ , inverse length set to 4000 coefficients) (blue line). Simulation results obtained by convolving the 3000-coefficients model of the plant with  $\mathbf{H}_{190}$  (red line). Simulation results obtained by convolving  $\mathbf{C}_{190}$  with  $\mathbf{H}_{190}$  (green line).

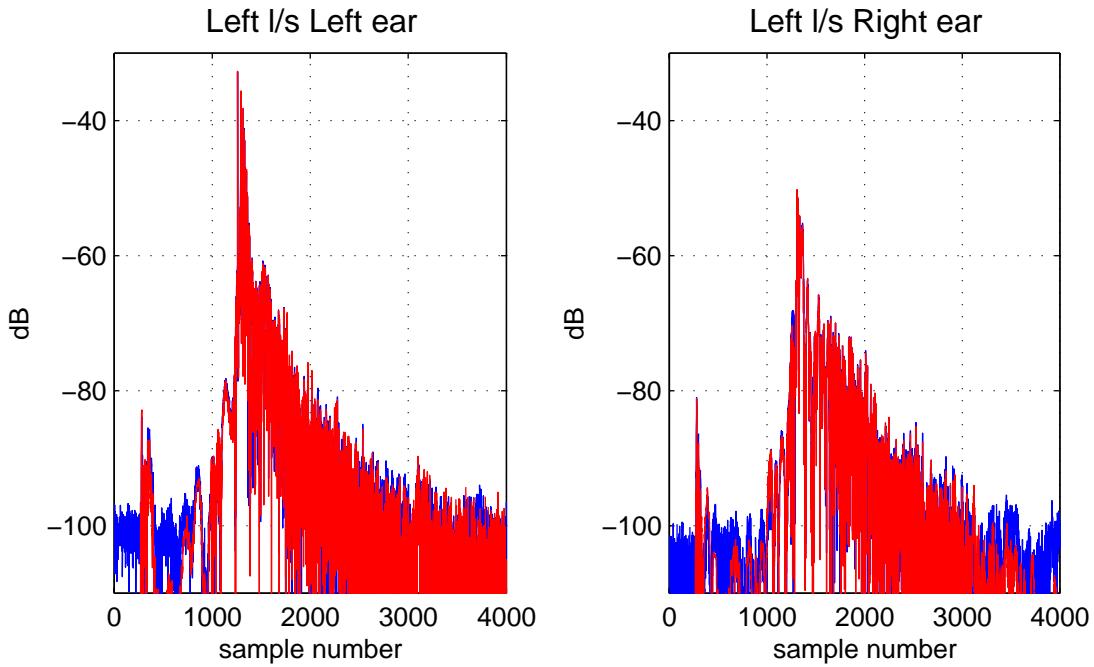
On the other hand, the comparison of the measured results (blue line) with the results obtained when the full-length model of the plant is used for the convolution (red line) can be seen in the figure 3-8 to validate this properly designed simulation process. As can be seen in the time-domain plot of figure 3-9, the simulated and measured impulse responses are identical above the measurement noise-floor. A small deviation of the simulation results from the measured results in figure 3-8 is due to the fact that, for the non-regularised case chosen here, the response level is rather close to the noise-floor. Hence the difference below the noise-floor between the simulation and the measurement becomes significant. In regularised cases where the level of the response rises higher above the noise-floor, the agreement between measurement and simulation is virtually perfect as is seen in the following example.



**Figure 3-9:** Measured inversion results for  $\mathbf{H}_{190}$  (inverse computed using  $\mathbf{C}_{190}$  as the plant model,  $\beta=0$ , inverse length set to 4000 coefficients) (blue line). Simulation results obtained by convolving the 3000-coefficients model of the plant with  $\mathbf{H}_{190}$  (red line). Simulation results obtained by convolving  $\mathbf{C}_{190}$  with  $\mathbf{H}_{190}$  (green line).

A second case of comparison between measurement and simulation is presented in figures 3-10 and 3-11 in the time-domain and frequency-domain respectively. In this case, the plant contains a strong reflection<sup>42</sup> but the inverse is computed on the basis of a model of the plant that does not contain this reflection. In this case, the regularisation is set to the value  $\beta=10^{-4}$ , the total length of the inverse to 2000 coefficients and the modelling delay  $\Delta$  to 1000 samples. Again with blue line we plot the directly measured results and with red line the results obtained when the inversion process is simulated by convolving the full-length model of the plant with the chosen inverse model.

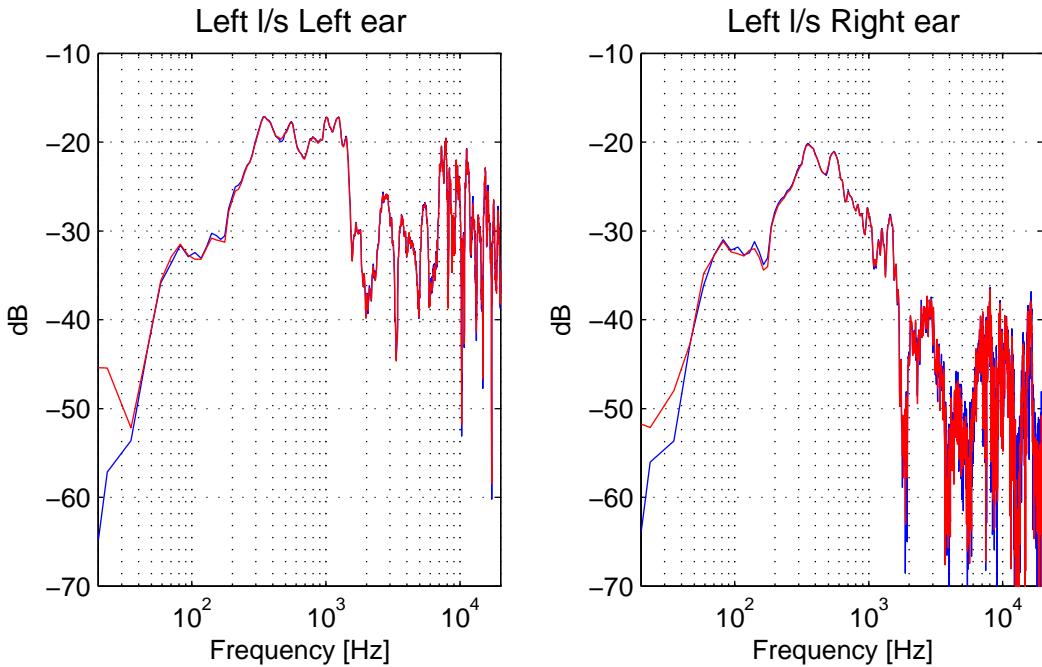
<sup>42</sup> This is the case with the wall reflection described in §5.2. A detailed description of that experimental arrangement that contains strong early reflections is given in that section. In this instance we are only interested in the comparison between the measured and the properly simulated results.



**Figure 3-10:** Comparison of measured (blue line) and simulated (red line) inversion results. The regularisation is set to  $\beta=10^{-4}$  and the inverse length is set to 2000 coefficients. The inverse is computed using the anechoic model of the plant.

As discussed in further detail in chapter 4, when regularisation is introduced, the level of the inversion rises higher above the noise-floor. This can be seen when the results of figure 3-10 are compared with those of figure 3-9. In that case then, the disagreement between the measured and simulated results below the noise-floor becomes less significant. Hence the agreement between measurement and simulation can be seen to be virtually perfect also in the frequency-domain results of figure 3-11.

Concluding then, we see that the two comparison cases presented in this section confirm the validity of the simulation of the inversion process when this is implemented by convolving the 3000-samples-long measured model of the plant with the chosen model of the inverse. More specifically, in that case the simulation results are virtually identical with the directly measured results above the measurement noise-floor in the time-domain. As long as the level of the inversion is reasonably higher than the noise-floor, the simulation results are also identical with the measured results in the frequency-domain.



**Figure 3-11:** Comparison of measured (blue line) and simulated (red line) inversion results. The regularisation is set to  $\beta=10^{-4}$  and the inverse length is to set to 2000 coefficients. The inverse is computed using the anechoic model of the plant.

### 3.6 Conclusions

In this chapter we described the basics of the inverse filtering objective evaluation procedure that is used in the following two chapters of the thesis. A detailed justification for the specific choice of evaluation procedure was given as well as a description of the experimental arrangement that used for the acquisition of the presented results. The measured model of the plant matrix  $\mathbf{C}$  was presented that is inverted in a number of different implementation cases in the next chapters. More specifically, it was shown that even in controlled anechoic conditions the strictly anechoic part of the plant model is followed by a series of reflections of low energy and that decay slowly to the level of the noise floor. When substituted into equation (2-20) that describes the inverse matrix  $\mathbf{H}$ , these different models of the plant matrix  $\mathbf{C}$  should result to more or less varying models of the inverse. The examination and the comparison of the effectiveness of these inverse models are presented in the next chapter. Finally, it was demonstrated that the use of the strictly anechoic part of the measured model of the plant for the simulation of the inversion process is highly

inaccurate. Conversely, when the full-length model of the plant is used for that simulation, the results were shown to be practically identical to the directly measured inversion results. This provides us with a versatile simulation method that is extensively used in the analysis of the following two chapters but also validates the measurement procedure that was described above.

## ***4 Objective evaluation of inverse filtering –***

### ***Inversion efficacy under ideal conditions***

#### ***4.1 Introduction***

In this chapter we present the results obtained when the inverse matrix  $\mathbf{H}$  is realised with a number of different FIR models. The inversion is implemented under “ideal” conditions as all the assumptions listed on p. 55 regarding the perfect knowledge of the plant are satisfied and, furthermore, the condition of “transparency” of the audio reproduction chain is circumvented by the fact that the transducers used for the measurement of the effectiveness of the audio control are the same as those used for the measurement of the plant. Hence any possible deviations from the flat spectrum and linear phase are taken into account in the inversion.

The only remaining sources of error in the inversion are then the imperfections inherent in the design of the inverse matrix  $\mathbf{H}$  and the presence of non-linear behaviour in the audio reproduction chain. As was already seen above, the non-linear behaviour in the reproduction chain appears in the measurement as a constant noise-floor of the measured impulse responses which for the -50dB power amplification setting appears at approximately -100dB below the [-1,1] full-scale range of our measurement apparatus<sup>43</sup> (see figure 3-4). This is of course far from the -144dBFS noise-floor of the 24-bit quantisation used here. Nevertheless, it exceeds the -96dBFS noise-floor of the 16-bit arithmetic of the standard PCM format that is commonly used in existing home audio installations (e.g. in Compact Disc equipment) and is thus the basic implementation platform of the virtual acoustic imaging system design under examination.

---

<sup>43</sup> Henceforward we denote the dB value relative to the maximum absolute signal value of 1 by the symbol dBFS (dB Full Scale).

As regards the imperfections introduced by the design parameters of the inverse, the parameter that is commonly identified as that primarily impeding the effectiveness of the inversion is the required length of the FIR filters in the inverse matrix  $\mathbf{H}$ . However, the results and analysis presented in the following show that the influence of the inverse filters' length on the actual effectiveness of the inversion is in turn determined by two other parameters. These are the length of the responses in the model of the plant matrix  $\mathbf{C}$  that is used for the determination of the inverse (see the discussion in §3.4) and the amount of regularisation introduced in the inversion. It is shown in the following that the chosen values of these two parameters affect the expected perceived quality of the inversion as much as that of the inverse filters' length and that for different values of these parameters the required length of the inverse filters varies dramatically. The exact relationships between these three parameters and their combined influence on the effectiveness of the imposed control are the objective of this chapter.

## ***4.2 Length of the model of the plant matrix***

As was shown in §3.4, even when the plant measurement is carried in anechoic conditions, the initial strictly anechoic part of the measured responses that corresponds to the head and body of the listener is followed by a low-energy decaying part that persists up to approximately the 2000<sup>th</sup> sample of the responses. Even though this late decay part is, strictly speaking, of deterministic nature, it will obviously not be repeatable across measurements of the same reproduction arrangement in different (even anechoic) spaces. Hence the only case when this part could be taken into account in the design of the inverse would be if an *in situ* measurement of the plant was available. This part of the experiment focuses then on the following two questions:

- (i) What is the highest level of cross-talk cancellation and equalisation that can be realistically achieved when only the strictly anechoic (i.e. approximately 4ms long) part of the plant is used for the determination of the inverse? This should be considered the absolute optimum for practically every real-world implementation of

the system, i.e. for any implementation that excludes the *in situ* measurement of the plant.

(ii) Can the exact knowledge of the whole length of the plant be used to improve the inversion results and under which conditions would that be possible? The answer to this second question would apply in cases where the whole length of the plant is known, e.g. in subjective experiments where the plant containing the individual HRTFs of the subjects is measured, or in real-world implementations where the plant pertaining to the specific reproduction geometry is measured *in situ*.

In order to answer these questions, different realisations of the inverse matrix  $\mathbf{H}$  were computed based on the three different models  $\mathbf{C}_{190}$ ,  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  of the plant matrix  $\mathbf{C}$  that were described in §3.4. In the first case examined, the exact inverse models (i.e. with the regularisation parameter  $\beta$  set uniformly equal to zero)  $\mathbf{H}_{190}$ ,  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  of the three models of the plant are realised with inverse filters that are 2000 samples long<sup>44</sup>.

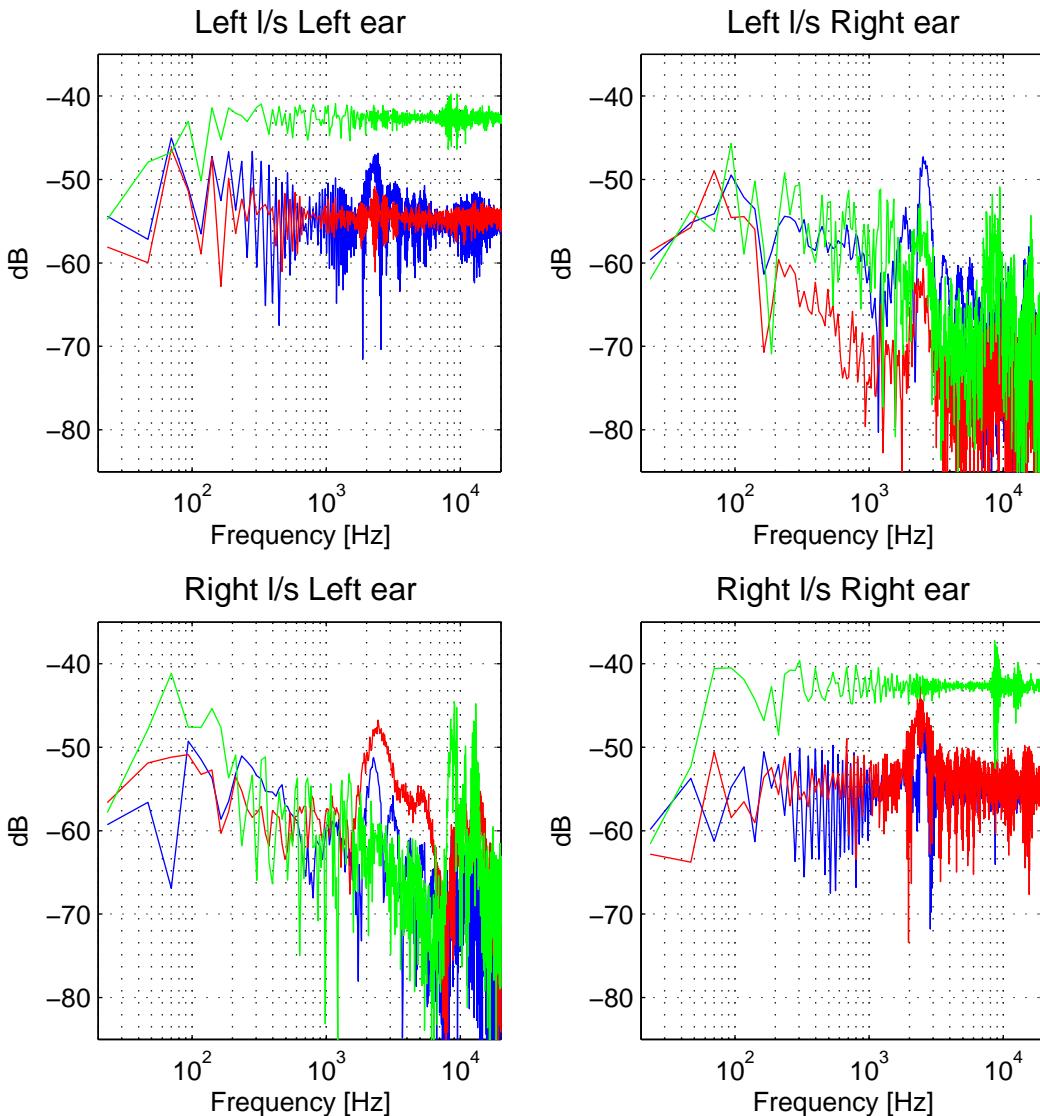
The measured ipsilateral equalisation and cross-talk cancellation results are plotted in figures 4-1 and 4-2. The frequency response magnitude results of figure 4-1 show that the inversion of  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  fails for these parameters of the inversion. The ipsilateral equalisation deviates from the flat line target by 15dB or more for most of the spectrum while the cross-talk cancellation reaches a maximum of 10dB with the contralateral channel being reinforced above the level of the ipsilateral channel in certain frequency regions. Compared to these results, the results obtained when the inverse is calculated using the  $\mathbf{C}_{190}$  model of the plant are far better. The cross-talk cancellation succeeds consistently from 100Hz up to 20kHz and ranges from 10dB to more than 20dB and the equalisation error is confined to less than 10dB for that frequency region.

Apart from the poor inversion results, the inverse filters  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  also impose a more stringent normalisation requirement, with the equalisation target being nearly 15dB lower than  $\mathbf{H}_{190}$ . This incurs a significant reduction in the overall dynamic range of the reproduction, something that can be seen more clearly in figure 4-2 where we

---

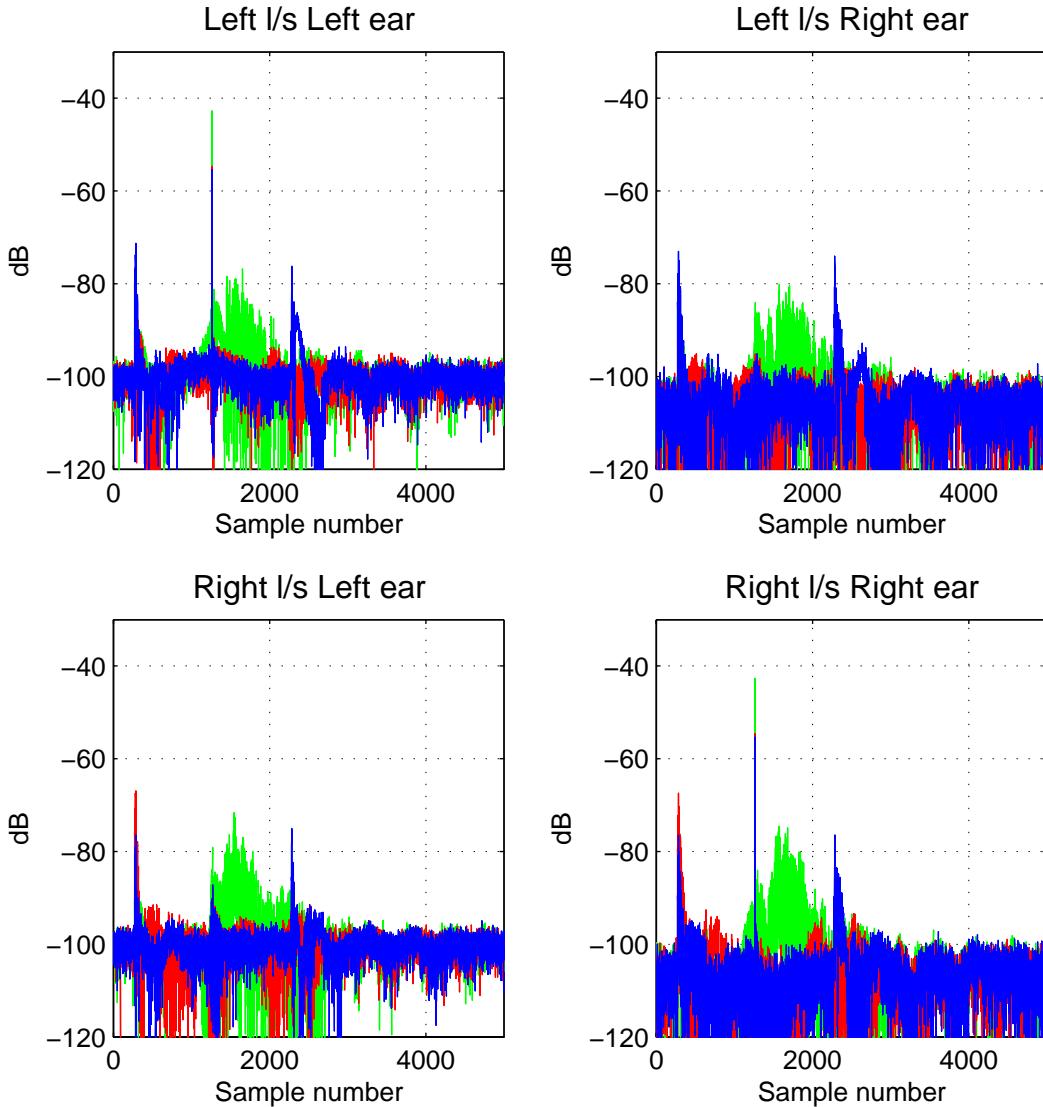
<sup>44</sup> Hereafter, unless otherwise noted, the modelling delay is set to half the length of the inverse (in this case 1000 samples), thus resulting to equal causal and anti-causal parts.

plot the time-domain equalised responses for the same three realisations of the inverse<sup>45</sup>. The measurement noise floor at about -100dBFS, constant across our measurements, can clearly be seen in this figure. As is then depicted in figure 4-2, the normalisation required for the implementation of  $\mathbf{H}_{190}$  is -42dB while the normalisation of the implementation of  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  is -54dB and the effective dynamic range of the whole processing scheme reduces from around 58dB to 44dB.



**Figure 4-1:** Measured inversion results. The regularisation is set to  $\beta=0$  and the inverse length to 2000 coefficients. The plant is modelled with  $\mathbf{C}_{190}$  (green line),  $\mathbf{C}_{740}$  (red line) and  $\mathbf{C}_{1740}$  (blue line).

<sup>45</sup> Unless otherwise stated throughout the thesis, in the time-domain plots of the equalised responses we plot the logarithm of the squared impulse response  $20\log_{10}(|x_{eq}(n)|)$ .



**Figure 4-2:** Measured inversion results. The regularisation is set to  $\beta=0$  and the inverse length to 2000 coefficients. The plant is modelled with  $\mathbf{C}_{190}$  (blue line),  $\mathbf{C}_{740}$  (red line) and  $\mathbf{C}_{1740}$  (green line).

The time-domain representation of figure 4-2 offers further information on the effectiveness of the inversion with respect to the expected perceived audio quality. More specifically, three types of error, i.e. deviations from the ideal delta-spike form, can be identified in the ipsilateral equalisation plots. Two of them are more prominent in the  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  results (red and blue lines respectively). They are visible as sharp rises of the equalized responses above the noise-floor at positions 1000 samples before and after the delta-spike. Evidently these are due to the fact that the inverse model used is a truncated version of the exact inverse with the truncation occurring at the samples of index  $n=-1000$  and  $n=999$  of the exact (non-causal) inverse. The

inverse computed on the basis of the shorter model of the plant  $\mathbf{C}_{190}$  decays faster than the other two inverses and hence the error due to the truncation at the same indices (shown by the green line) is buried below the noise floor. It should be noted however that the error due to the truncation of the anti-causal parts of the inverse filters in  $\mathbf{H}_{190}$  does rise above the noise floor but it is lower than in the  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  cases and is thus hidden behind the blue and red lines in the figure. A third type of error is on the other hand visible in the  $\mathbf{H}_{190}$  results. This appears as a broad rise starting 190 samples after the delta spike and is due to the fact that  $\mathbf{C}_{190}$  is a truncated version of the measured plant responses.

The correlation of deviations from the ideal equalized response such as those of figures 4-1 and 4-2 with the subjective evaluation of the inversion results is investigated by Fielder (2003) and in a more recent study by Norcross et al. (2004)<sup>46</sup>. As is discussed in (Fielder 2003), depending on the width of the corresponding peaks or notches and their position along the frequency axis, variations from the flat response as small as 0.5dB can be audible, i.e. distinguishable from the ideal flat frequency response. Based on that it is clear that binaural material presented through a plant equalised with  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  is likely to be perceived as significantly distorted. The absence of isolated peaks and notches in the results of  $\mathbf{H}_{190}$  should make the reproduction through  $\mathbf{H}_{190}$  more transparent but again, there is considerable room for improvement as the deviations in the order of 10dB from the flat response should be audible.

The deviations from the ideal delta-spike form in the time-domain results are also investigated in (Fielder 2003) and (Norcross et al. 2004). The principle psychoacoustic effect associated with such deviations is that of the temporal masking (also commonly referred to as “the precedence effect” or “the law of the first waveform”) which is in turn related to the temporal integration function of the ear. In simple terms this states that diminished-in-level replicas of a sound event that appear shortly before or after a main sound event are integrated onto the main sound event by the auditory system and thus appear as a single auditory event of increased loudness. A detailed exposition of the theory related to this effect is outside the scope of this

---

<sup>46</sup> As is also discussed in the review of §2.2 those studies are restricted to single-channel inversion cases but the findings presented there apply to the ipsilateral equalisation part of the two-channel inversion presented here.

thesis<sup>47</sup> but its main properties of interest here are the following: First, the integration of sounds arriving before the main sound event extends very briefly, in the order of 5-20ms and the level of sounds that are masked inside this interval falls very fast. Thus sounds that arrive more than 20ms before the main sound event will definitely not be masked and will be audible as separate auditory events or as distortion of the main event (Fielder (2003) sets this limit where backward masking disappears at 15ms). Second, the integration of sounds arriving after the main sound event extends for a substantially longer interval, up to 100-200ms. Providing an approximate mean of previously published studies, Fielder (2003) stipulates that sounds arriving up to 4ms after the main sound will be masked if their level does not exceed the main sound while this threshold of masking drops by 35dB per decade (in ms) for later arrivals thus falling down to approximately -50dB 100ms after the main sound. These stipulations are verified in subjective experiments by Norcross et al. (2004) in the context of the loudspeaker-response equalisation problem. As is shown in that study, truncation effects as low as 50dB below the level of the delta-spike can significantly degrade the perceptual performance of the inversion.

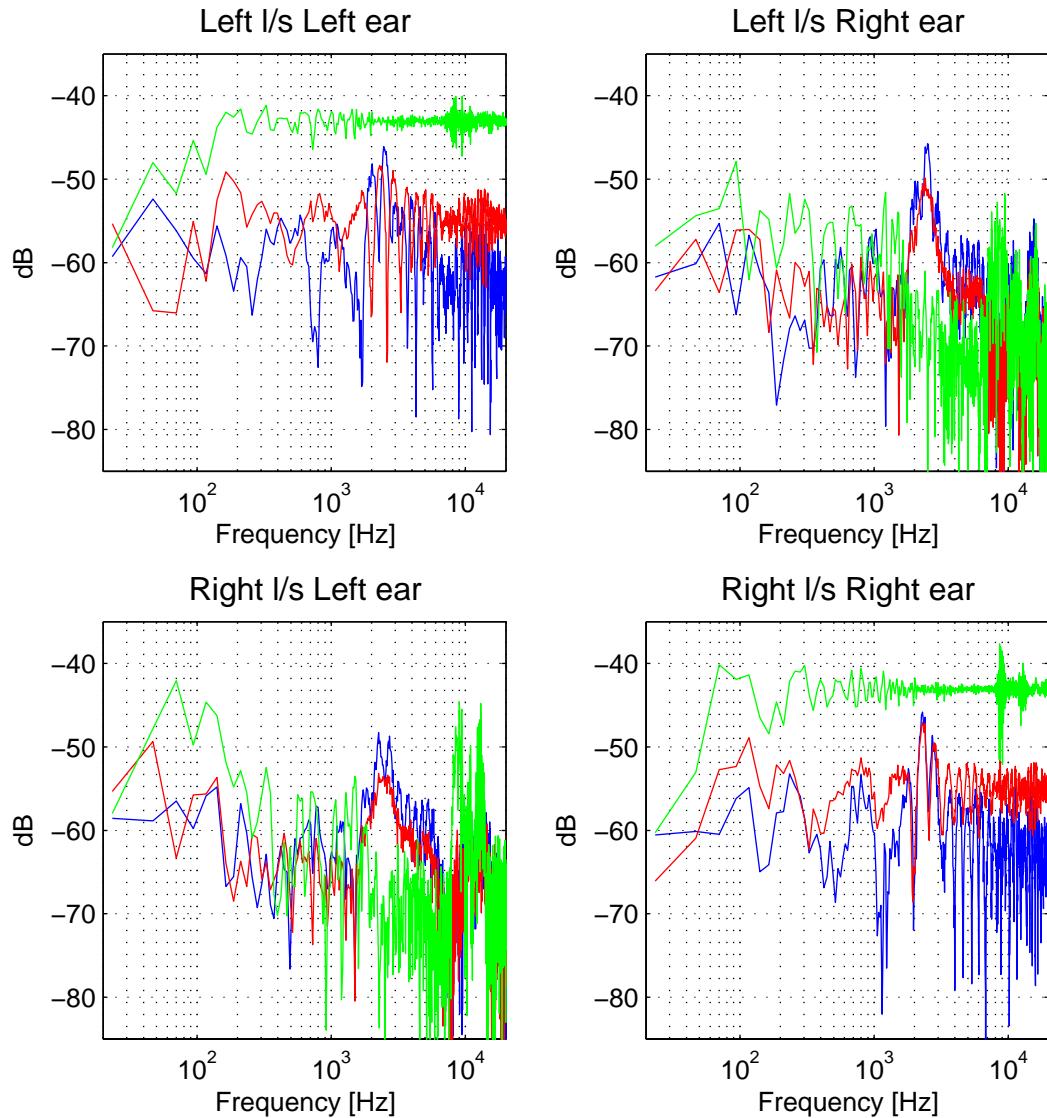
It is evident then, that the transients that appear in the equalised impulse responses of figure 4-2 will result in significant degradation of the perceived quality of the delivered binaural material. These transients are more severe for  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  and can be seen to be only around 15dB below the delta spike. This is normal as these inverse filters attempt to invert longer versions of the plant and should hence be expected to decay more slowly. On the other hand the truncation transients are much lower in the  $\mathbf{H}_{190}$  inversion results, their level being 40dB below the delta spike in the anti-causal end of the inverse filter and below the noise floor in its causal end. This decrease in the truncation error comes at the expense of a rise in the inversion error immediately after the delta spike that is due to the imperfect modelling of the plant by  $\mathbf{C}_{190}$ . However, as it appears immediately after the delta spike (extending up to approximately 15ms) this artefact should be totally masked by the preceding delta-spike. It should be clear of course that this temporal masking effect does not alleviate the influence of this error element in the equalised frequency response. This can be seen in figure 4-1 (green line) as a deviation from the ideal flat form which, as was argued above, should be perceptibly significant.

---

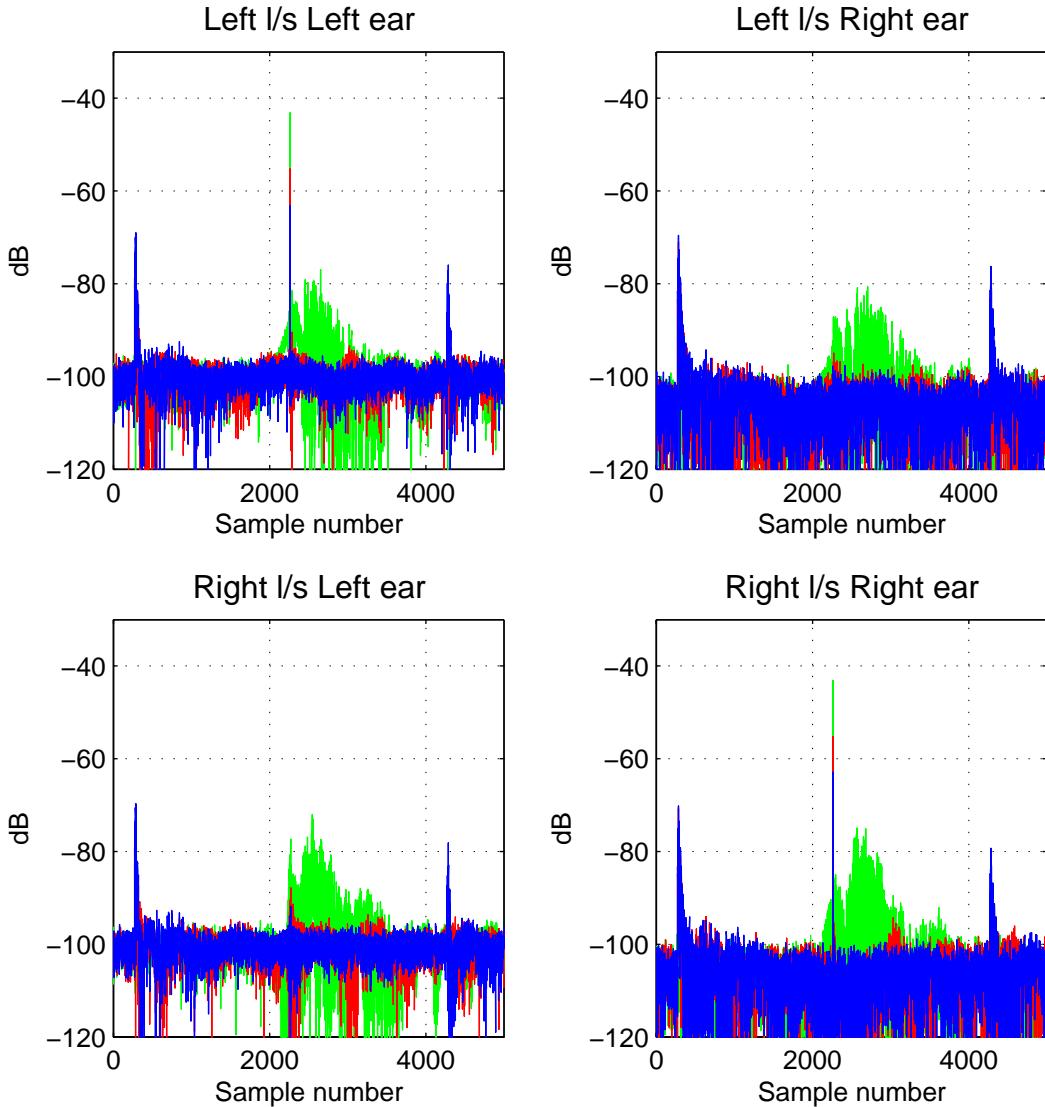
<sup>47</sup> For review of the current knowledge on the effect see (Fielder 2003) and references therein.

In the next set of results, shown in figures 4-3 and 4-4, we plot the equalisation and crosstalk cancellation obtained when  $\mathbf{C}_{190}$ ,  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  are inverted with inverse filters that are 4000 samples long and again with the regularisation set uniformly equal to zero (again we denote these realisations of  $\mathbf{H}$  with  $\mathbf{H}_{190}$ ,  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$ ). The frequency response results of figure 4-3 are practically the same as those of the 2000-samples-long realisations of  $\mathbf{H}$  presented above, with  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  failing both to equalise and suppress the cross talk and  $\mathbf{H}_{190}$  showing a reasonable success in both goals but again with a significant loss of dynamic range. In the time-domain plot of figure 4-4, the transients due to the truncation are still present in the results obtained with  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$ , but reduced in the  $\mathbf{H}_{740}$  case in comparison with the case of 2000 inverse coefficients shown in figure 4-1. They are now totally absent in the results obtained with  $\mathbf{H}_{190}$  as the 2000 taps of anticausal part of the filters in  $\mathbf{H}_{190}$  are now enough to suppress them below the noise floor. As should be expected, the error due to the imperfect modelling of the plant with  $\mathbf{H}_{190}$  is exactly the same as in the previous case of figure 4-2.

Referring then to the questions posed in p. 77, the point made by the results presented up to now is that: (i) the use of the typical 4ms duration plant results in an error rising above the noise floor due to the imperfect modelling of the plant, which is however of small perceptual importance and (ii) the inversion of the longer versions of the plant significantly diminishes the total dynamic range of the processing and results in end-effects that persist even when the inverse matrix  $\mathbf{H}$  is realised with 4000-samples-long filters. We expect the first of these points to be consistently verified as the length of the inverse increases even further (as long as the regularisation is kept to zero). As regards the second point, there will of course be a length of the inverse after which the end-effects are suppressed below the noise-floor. On the other hand, the issue regarding the loss reduction of dynamic range should remain in that case.



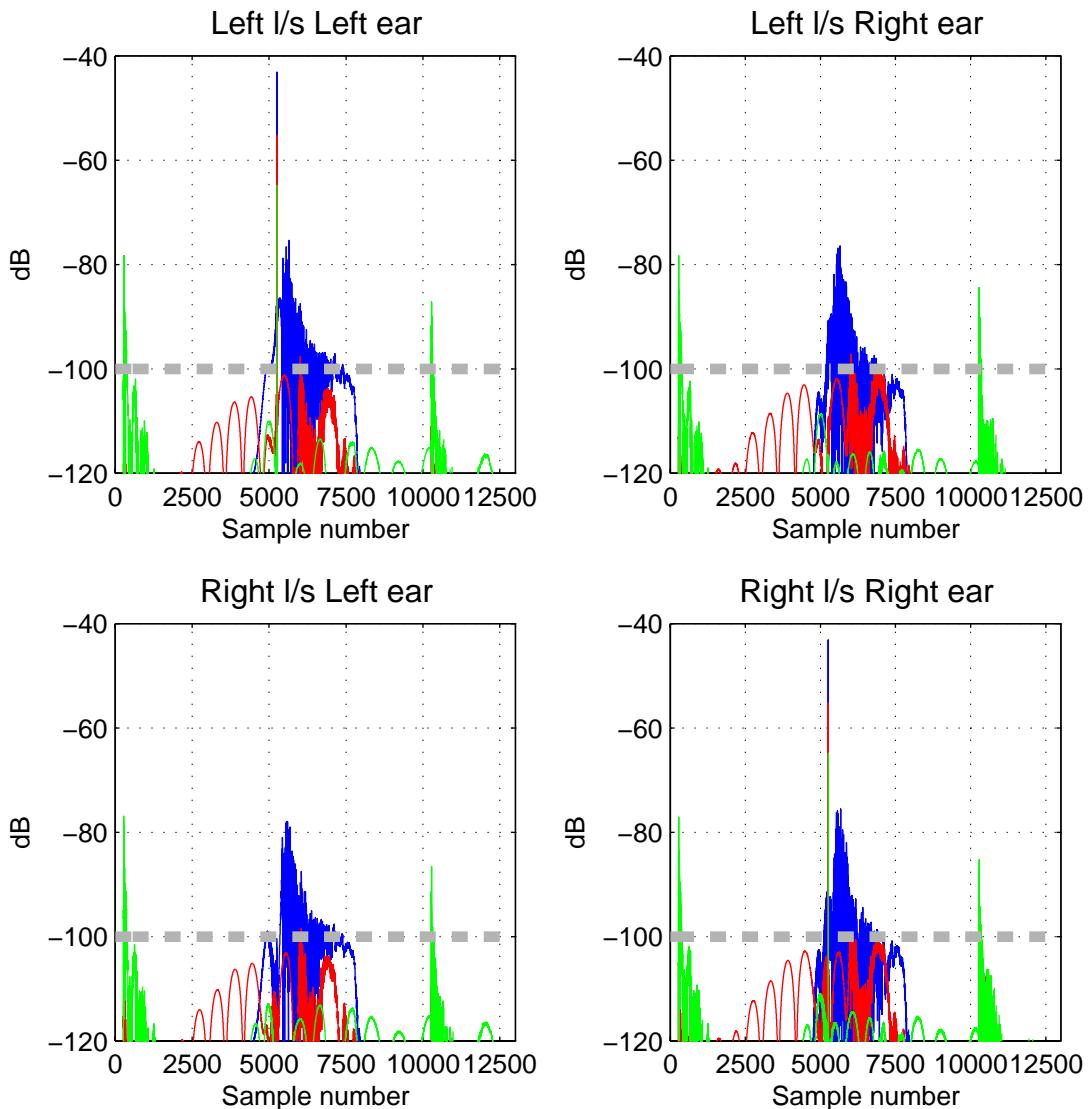
**Figure 4-3:** Measured inversion results. The regularisation is set to  $\beta=0$  and the inverse length to 4000 coefficients. The plant is modelled with  $\mathbf{C}_{190}$  (green line),  $\mathbf{C}_{740}$  (red line) and  $\mathbf{C}_{1740}$  (blue line).



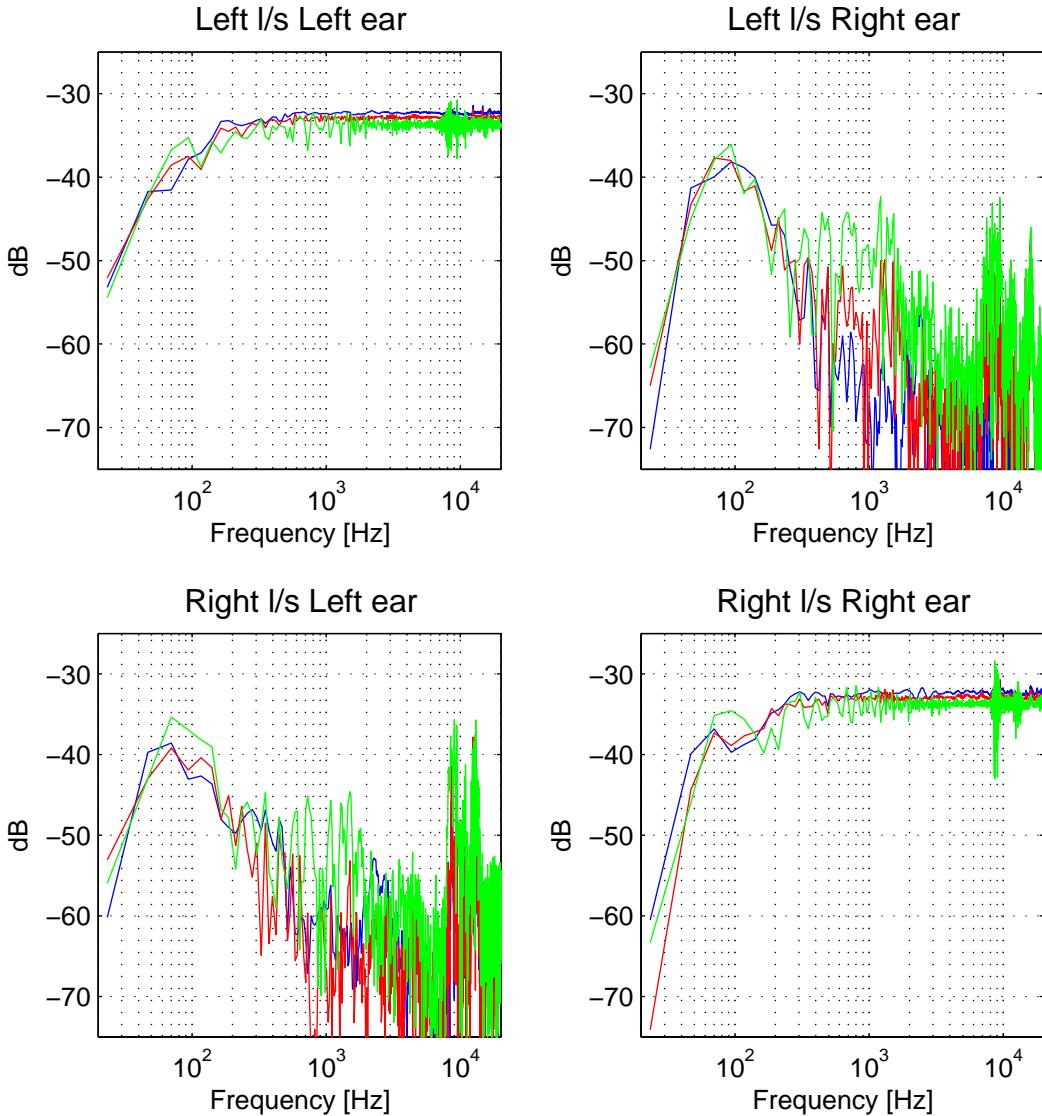
**Figure 4-4:** Measured inversion results. The regularisation is set to  $\beta=0$  and the inverse length to 4000 coefficients. The plant is modelled with  $\mathbf{C}_{190}$  (green line),  $\mathbf{C}_{740}$  (red line) and  $\mathbf{C}_{1740}$  (blue line).

The 4000-samples-long realisation presented in figures 4-3 and 4-4 was the longest set of inverse filters tried in the experiments, so the validity of the above statement cannot be verified with directly measured inversion results. Instead one can use the simulation process described in §3.5. This is done in figure 4-5 where we plot the inversion results when the three models of the plant  $\mathbf{C}_{190}$ ,  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  are inverted with 10000-samples-long inverse filters (denoted again as  $\mathbf{H}_{190}$ ,  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  respectively) and again with the regularisation parameter set uniformly to zero. As expected, the simulation results verify the fact that, even when realised with 10000-samples-long filters,  $\mathbf{H}_{190}$  demonstrates the error immediately after the delta-spike.

Similarly, the transients due to the truncation are present in  $\mathbf{H}_{1740}$  but they are now suppressed further compared to the previous cases. In the inversion with  $\mathbf{H}_{740}$  the truncation transients are suppressed below the noise-floor but of the overall dynamic range of the processing is again approximately 45dB as with the previous cases.



**Figure 4-5:** Simulated inversion results. The regularisation is set to  $\beta=0$  and the inverse length to 10000 coefficients. The plant is modelled with  $\mathbf{C}_{190}$  (blue line),  $\mathbf{C}_{740}$  (red line) and  $\mathbf{C}_{1740}$  (blue line). The results are obtained by convolving the 3000-coefficients model of the plant with the three inverse models. The grey dashed line shows the noise-floor of the corresponding measured results.



**Figure 4-6:** Measured inversion results. The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{190}$  and the inverse length is set to 2000 coefficients (green line). The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{740}$  and the inverse length is set to 4000 coefficients (red line). The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{1740}$  and the inverse length is set to 4000 coefficients (blue line).

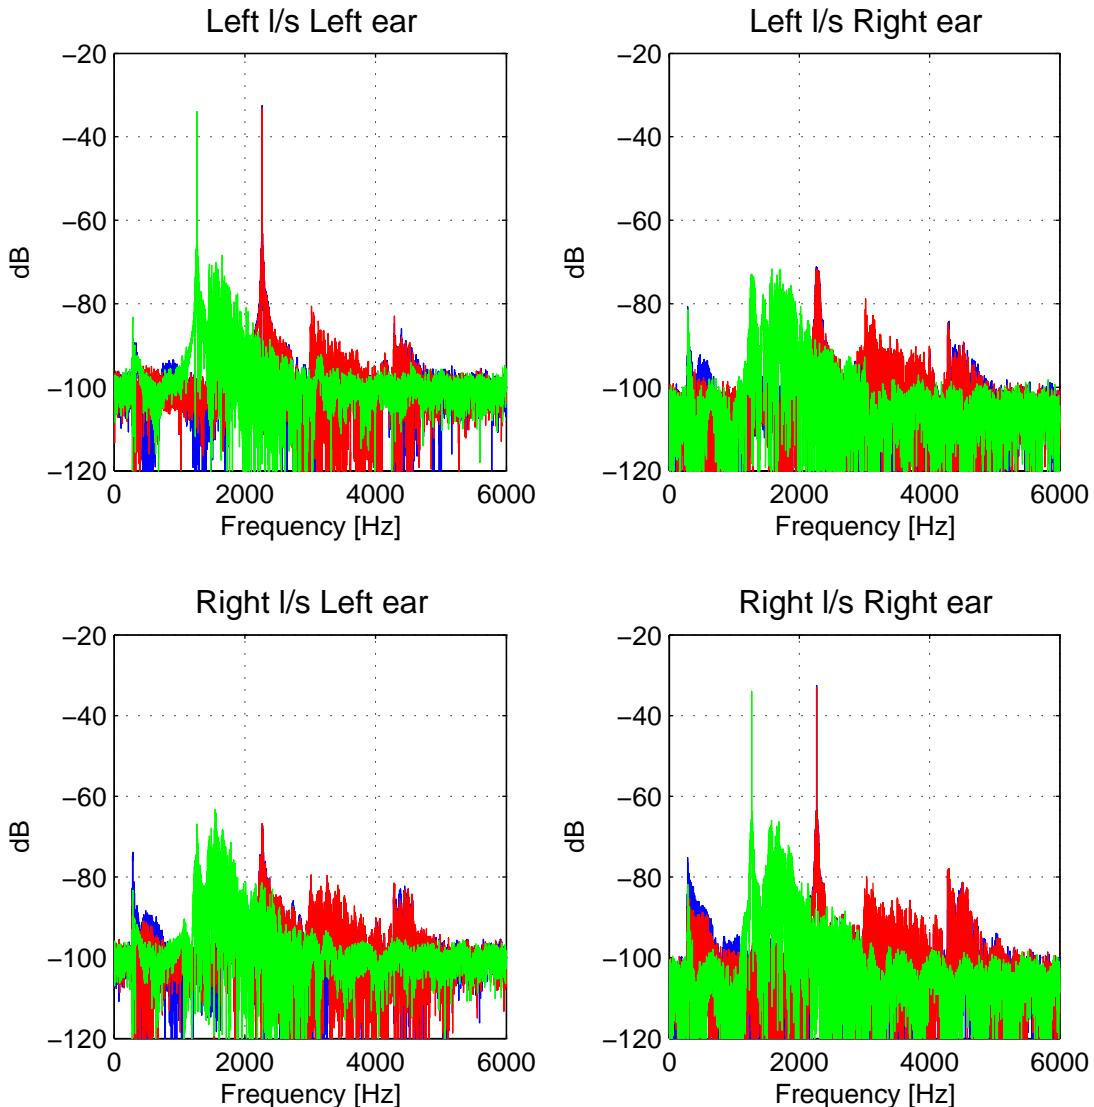
In the discussion up to now we established the extent of feasibility of the inversion for the three different cases of the plant model and for different lengths of the inverse models but keeping the regularisation parameter equal to zero. We now introduce regularisation into the solution. In figures 4-6 and 4-7 we plot the measured results obtained when  $\mathbf{C}_{190}$  is inverted with a 2000-samples-long inverse and  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  with 4000-samples-long inverse filters. All three realisations of the inverse (denoted again as  $\mathbf{H}_{190}$ ,  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$ ) are computed with the regularisation parameter set uniformly equal to the value  $\beta=10^{-4}$ . The frequency-domain results of figure 4-6

appear rather similar to each other with  $\mathbf{H}_{1740}$  giving the tightest ipsilateral equalisation confined to no more than  $\pm 1$ dB in the region above 200Hz and cross-talk cancellation of 15-30dB in that region. The normalisation requirement is also similar among the three cases, with the equalisation target level being at approximately -33dBFS. This is great improvement compared to the non-regularised cases presented above. Even at the low-frequency end, where the introduction of regularisation prevents the application of the necessary (but excessive) control needed for the inversion, the erratic behaviour of the non-regularised inversion is now replaced by a gentle roll-off. Regarding the issue of the length of the plant model that should be used for the determination of the inverse, the results show that in cases where the whole length of the impulse responses that comprise the plant is accurately known, the inversion of the whole length of the plant can indeed be advantageous. This however is shown to be true only with the extra condition that the proper amount of regularisation is introduced in the determination of the inverse.

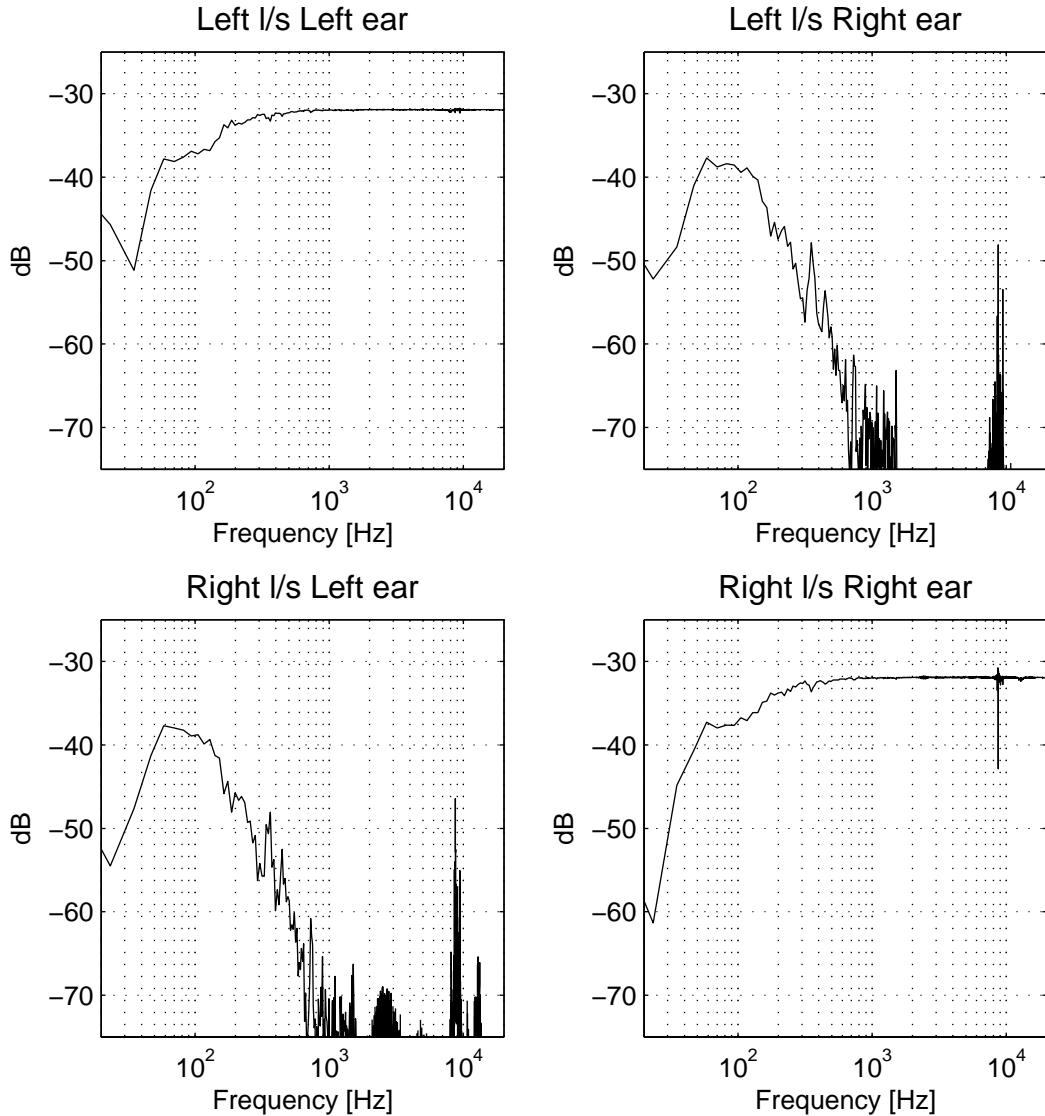
A point that needs to be made here is that the superiority of the inversion results obtained with  $\mathbf{H}_{1740}$  and  $\mathbf{H}_{740}$  over the results of  $\mathbf{H}_{190}$  is not due to the fact that  $\mathbf{H}_{1740}$  and  $\mathbf{H}_{740}$  are 4000-samples-long while  $\mathbf{H}_{190}$  is 2000-samples long. This becomes evident from the time-domain version of the results shown in figure 4-7. As can be seen in the figure, the basic element of error in the  $\mathbf{H}_{190}$  inversion results is the artefact following immediately after the delta-spike. As it was argued above, this is due to the imperfect model of the plant used for the determination of the inverse and is not affected by the length of the inverse. Conversely, the end-effects that are due to the truncation of the inverse can be seen to be of much lower level and overall power even in this 2000-samples-long realisation. Hence, the use of a longer realisation of the inverse would not improve the inversion results.

On the other hand, as should be expected, the error in the results of  $\mathbf{H}_{1740}$  is solely due to the end-effects caused by the truncation of the inverse. The results of  $\mathbf{H}_{740}$  are somewhat in between, as they demonstrate artefacts both due to the truncation of the inverse (occurring at the 260<sup>th</sup> sample and after the 4000<sup>th</sup> sample) and due to the truncated model of the plant (these appear at about the 3000<sup>th</sup> sample, further away from the delta-spike than in the  $\mathbf{H}_{190}$  case because  $\mathbf{C}_{740}$  is truncated at a later point

than  $\mathbf{C}_{190}$ ). The latter would be diminished by a longer realisation of the inverse but this clearly does not apply to the former.

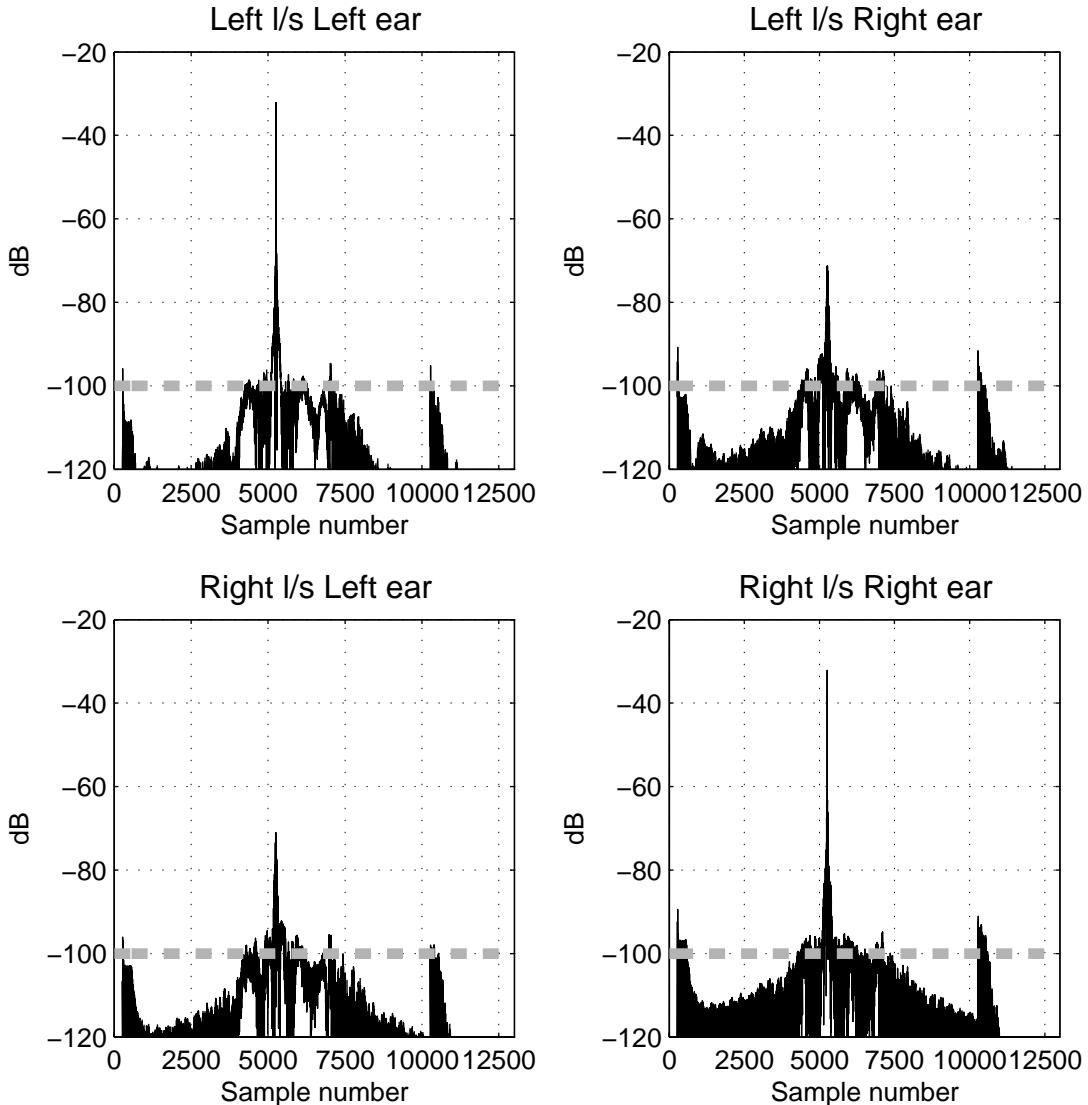


**Figure 4-7:** Measured inversion results. The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{190}$  and the inverse length is set to 2000 coefficients (green line). The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{740}$  and the inverse length is set to 4000 coefficients (red line). The regularisation is set to  $\beta=10^{-4}$ , the plant is modelled with  $\mathbf{C}_{1740}$  and the inverse length is set to 4000 coefficients (blue line). The introduction of different amounts modelling delay (1000 samples in the first case and 2000 samples in the second and third cases) is made apparent by the different positioning of the delta spike.



**Figure 4-8:** Simulated inversion results. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 coefficients. The plant model used for the determination of the inverse is  $\mathbf{C}_{1740}$ . The results are obtained by convolving the 3000-coefficients model of the plant with the inverse model.

As before, these 4000-samples-long realisations of the inverse were the longest inverse filters tried in the measurements. Hence the determination of the length of the inverse filters necessary for the suppression of the truncation effects below the noise-floor is left to be determined with a simulation. This is done in figures 4-8 and 4-9.



**Figure 4-9:** Simulated inversion results. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 coefficients. The plant model used for the determination of the inverse is  $\mathbf{C}_{1740}$ . The results are obtained by convolving the 3000-coefficients model of the plant with the inverse model. The grey dashed line shows the noise-floor of the corresponding measured results.

Comparing the results of figure 4-9 with the non-regularised 10000-samples-long inverse  $\mathbf{H}_{1740}$  (green line in figure 4-5) we see that, as should be expected, the introduction of regularisation increases the rate of decay of both the causal and non-causal parts of the inverse filters. This becomes evident by the fact that the regularised 10000-samples-long realisation is just enough to suppress the truncation transients below the -100dBFS noise-floor something that was not true for the non-regularised realisation of figure 4-5. Furthermore, the normalisation requirement is significantly relaxed in the regularised inversion, thus yielding an overall processing dynamic

range of about 70dB, which is double the 35dB of the corresponding non-regularised case. These improved results come at the expense of a rise in the cross-talk level at the position of the ipsilateral channel's delta-spike and also a slight deviation from the ideal delta-spike form in the ipsilateral channel. Still, both these effects should be expected to have minimal perceptual significance. The former does not significantly influence the level of the crosstalk suppression above 200Hz as is seen in figure 4-8 while the latter can be seen to fall inside the temporal masking intervals described above (see pp. 81-82).

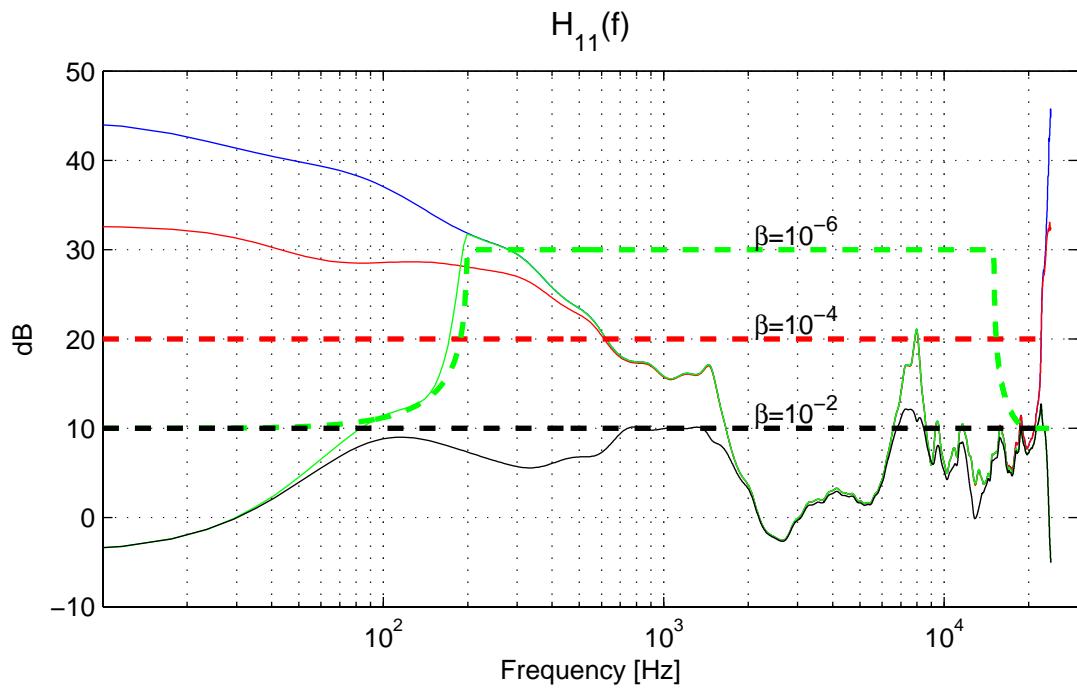
### **4.3 Regularisation**

The parameter of regularisation has already been introduced in the discussion of §4.2 where it has been shown that the regularised inversion of an adequately long model of the plant yielded better results compared to the non-regularised inversion. More importantly, the results of §4.2 brought up the role of regularisation as a means of relaxing the normalisation requirement and thus effectively increasing the dynamic range of the inversion. In this section we focus our interest on the parameter of regularisation and we examine its influence in the effectiveness of the inversion.

A clear depiction of the role of regularisation in the determination of the inverse is given in figure 4-10 where we plot (with a solid line) the magnitude of the frequency response of the element in the first row and first column of the inverse matrix  $\mathbf{H}$  with different values of regularisation. The responses plotted in the figure are those prior to normalisation, i.e. they correspond to the total control power necessary for the inversion of the plant in order for the input signal to be reproduced at the ears of the listener at its original level. In the same figure we use dashed thick line of the same colour to indicate the introduced amount of regularisation in each case. Note that, being a dimensionless scalar, the value of regularisation is not directly comparable to the frequency response magnitude values plotted in figure 4-10 and that, for illustrative purposes, the regularisation parameter value is depicted in the figure via the quantity  $-5\log_{10}(\beta)$ .

We see in the figure that, with  $\beta$  set to zero, most of the energy of the inverse filter is concentrated in the low frequency region below 500Hz and in the region above

20kHz. We saw however in the previous section that both the cross-talk cancellation and the ipsilateral equalisation fails in the region up to 200Hz, even in ideal conditions. This is normal if we consider that the control is applied with small, two-way loudspeakers that are virtually non-responsive in this frequency region. As for the high frequency region above 20kHz (which we obviously do not intend to control) the increase in the required control power is a result of the anti-aliasing filter roll-off present in the measured plant responses.



**Figure 4-10:** Frequency response of the element  $H_{11}$  of the inverse matrix  $\mathbf{H}$  for  $\beta=0$  (blue solid line),  $\beta=10^{-4}$  (red solid line),  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  (green solid line) and  $\beta=10^{-2}$  (black solid line). The quantity  $-5\log_{10}(\beta)$  is also plotted for the cases  $\beta=10^{-4}$  (red dashed line),  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  (green dashed line) and  $\beta=10^{-2}$  (black dashed line).

By penalising the control effort, as achieved by the introduction of regularisation, we see in figure 4-10 that the level of the applied control is suppressed. This suppression starts at the low and high ends of the spectrum for the value  $\beta=10^{-4}$  leaving the shape of the control unaltered for the region from around 600Hz up to about 18kHz. This is of course a desirable effect as it effectively (with the necessary introduction of normalisation) redistributes the control power from the low and high end of the spectrum (that we cannot, or do not wish to control) to the audible and controllable

middle region. Nevertheless, we see in the figure that this effect is not fully exploited with  $\beta$  set to the value  $\beta=10^{-4}$  as a significant amount of control is still exerted on the ends of the spectrum. By increasing the regularisation penalty to the value  $\beta=10^{-2}$  the control is completely suppressed at the two ends of the spectrum but at the same time the shape of the control is also significantly affected in the frequency region of interest.

The above rationale justifies the introduction of regularisation that varies with frequency so that the control effort at two ends of the spectrum is heavily penalised and at the same time allowing the exertion of the necessary control at the frequency region of interest. Such a design is also plotted in figure 4-10. In this case, the regularisation is set to the value  $\beta=10^{-2}$  in the frequency intervals from DC to 100Hz and from 18kHz up to the Nyquist frequency and to the value  $\beta=10^{-6}$  in the interval from 100Hz to 18kHz. In order to avoid an abrupt, “brick-wall” type, change in the frequency response of the inverse filters that would result in excessive ringing in their time response we have shaped the transition from  $\beta=10^{-2}$  to  $\beta=10^{-6}$  and back to  $\beta=10^{-2}$ , with two half-Hanning windows extending from 20Hz to 200Hz and from 15kHz to 20kHz<sup>48</sup>. As can be seen in the figure, the frequency response of this filter follows that of the  $\beta=10^{-2}$  inverse up to 20Hz and, after a transition region extending from 20Hz to 200Hz, it meets the frequency response of the non-regularised inverse<sup>49</sup> and follows it up to 18kHz where the opposite transition takes place.

These different realisations of the inverse are compared in the following measured results. First, we examine the typical case where the inverse is determined on the basis of the strictly anechoic model of the plant  $\mathbf{C}_{190}$ . Figures 4-11 and 4-12 compare the measured results obtained when  $\mathbf{C}_{190}$  is inverted with the regularisation parameter  $\beta$  set to the uniform values  $\beta=0$ ,  $\beta=10^{-4}$  and  $\beta=10^{-2}$  and the variable value  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  described above. The length of the inverse is set to 2000 taps for the  $\beta=10^{-4}$  and  $\beta=10^{-2}$  cases and 4000 taps for the  $\beta=0$  and  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  cases. The frequency-domain plot of figure 4-11 clearly shows the increase in the level of the signal reaching the listener’s ear as the value of  $\beta$  is increased. This desirable increase neither improves nor degrades the effectiveness of the equalisation and the cross-talk

---

<sup>48</sup> Henceforward we denote this variable regularisation value with  $\beta=[10^{-2} 10^{-6} 10^{-2}]$ .

<sup>49</sup> The regularisation value of  $\beta=10^{-6}$  results to no penalty in the region from 200Hz to 18kHz and is thus practically equivalent to the value  $\beta=0$ .

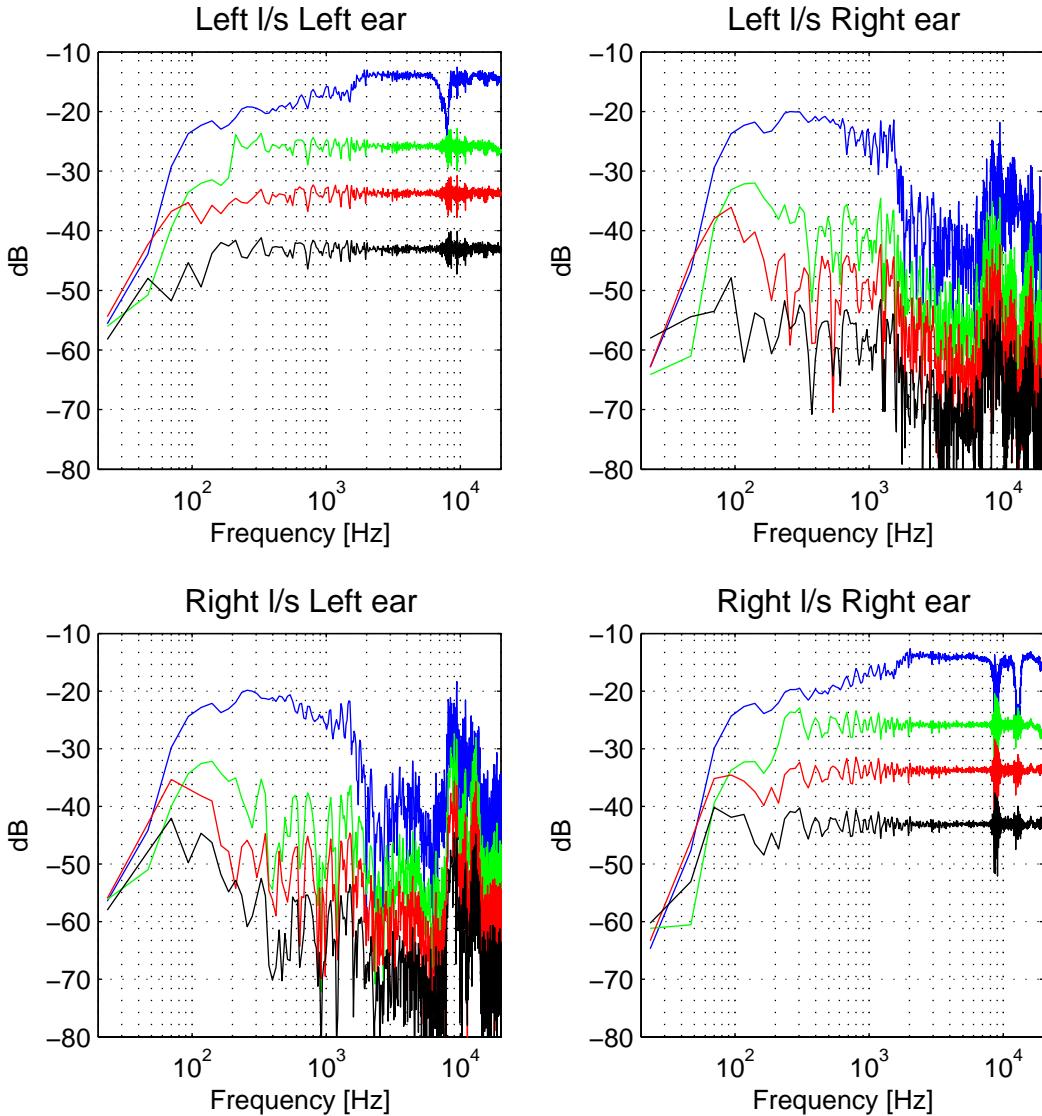
cancellation as the regularisation increases from  $\beta=0$  to  $\beta=10^{-4}$ . However, the degradation is visible in the comparison of the  $\beta=10^{-4}$  and  $\beta=10^{-2}$  cases. The variable regularisation case can be seen to succeed in increasing the equalisation target level compared to the  $\beta=10^{-4}$  case preserving the same inversion quality.

Further insight into the properties of these results is again given by the time-domain plot of figure 4-12. First, we see in the figure that the introduction of regularisation results in the smearing of the delta spike into a tight pulse that is symmetrically distributed around the position of the original, non-regularised, delta-spike (this symmetrical smearing is more evident by comparison of figures 4-5 and 4-9). The time-symmetry of this smeared pulse is imposed by the desired pure delay (i.e. linear phase) character of the sought ipsilateral equalisation and it is thus bound to appear in the inversion results when regularisation is introduced. This entails the presence of a pre-ringing artefact in the equalised response that<sup>50</sup> is more likely to be audible than its mirror image. For the regularisation values considered in this case, this smearing extends to no more than 5-10 ms and should thus be considered acceptable. Higher regularisation values could however render it audible. On the other hand we see in figure 4-12 that the smeared energy around the delta spike is also copied in the cross-talk paths where the absence of the masker should render it audible. The “spill-over” of this energy on the cross-talk paths increases with increasing values of the regularisation parameter, something that is also evident from the rise in the level of the cross-talk paths in the frequency-domain plot of figure 4-11.

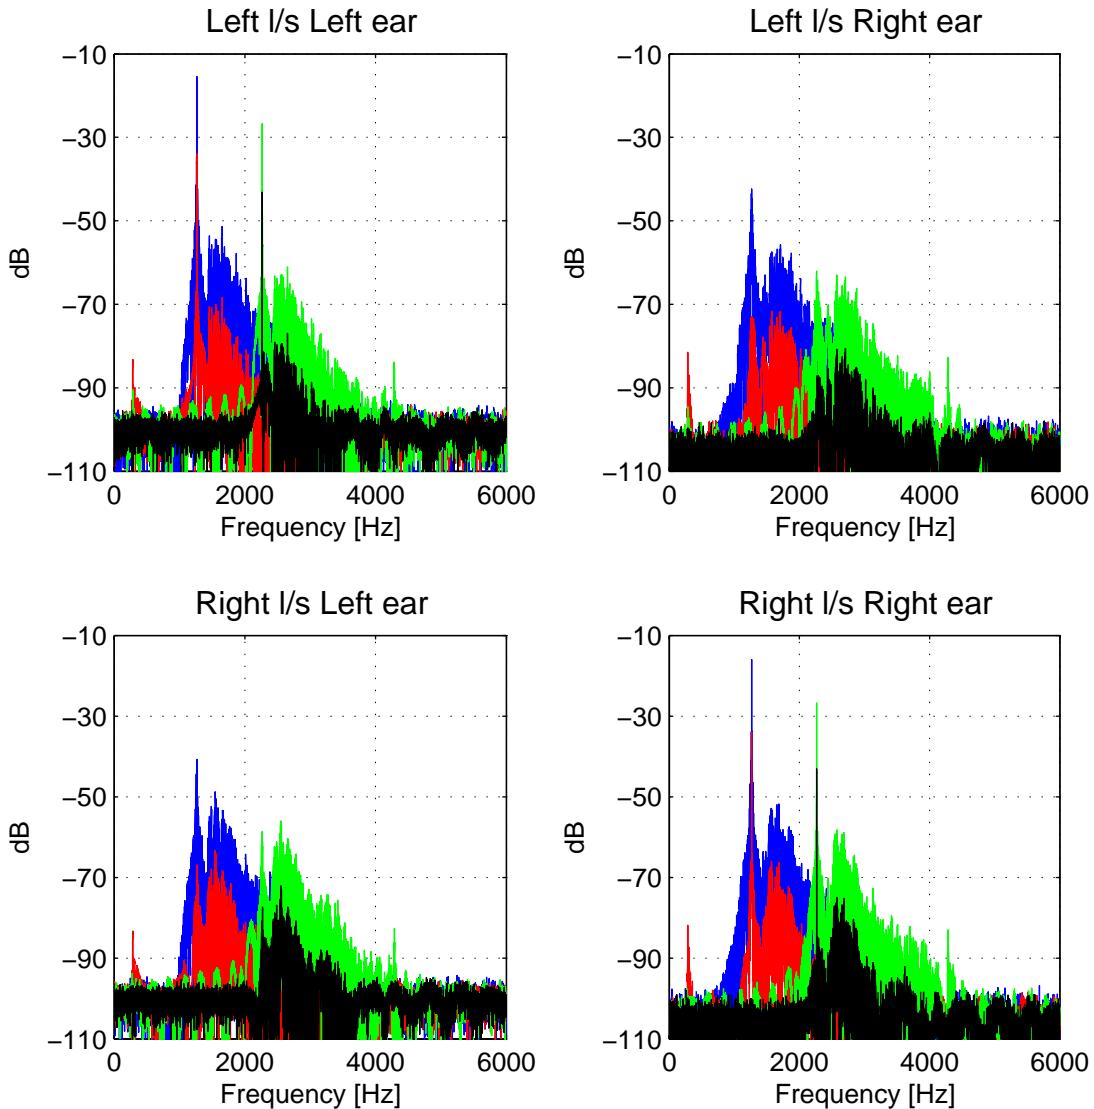
A similar influence of the regularisation can also be seen on the artefact that follows the pulse and, as was discussed in the non-regularised case of the previous section, is due to the imperfect modelling of the plant by  $\mathbf{C}_{190}$ . More specifically, with the introduction of regularisation and the consequent relaxation of the normalisation requirement, the inversion results are effectively scaled by a constant factor that is applied both to the equalisation pulse and to the residual error. Hence, this error is amplified by the same amount as the main pulse across the different regularisation values and is present in both the ipsilateral and the cross-talk paths.

---

<sup>50</sup> See also the discussion in pp. 81-82.



**Figure 4-11:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=0$  and the inverse length to 4000 coefficients (black line). The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 coefficients (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 2000 coefficients (blue line). The regularisation is set to  $\beta=[10^{-2} \ 10^{-6} \ 10^{-2}]$  and the inverse length to 4000 coefficients (green line).



**Figure 4-12:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=0$  and the inverse length to 4000 coefficients (black line). The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 coefficients (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 2000 coefficients (blue line). The regularisation is set to  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  and the inverse length to 4000 coefficients (green line).

To conclude the discussion of the role of regularisation we also examine the case where regularisation is used in the inversion of the longer model of the plant,  $\mathbf{C}_{740}$ . We saw in §4.2 that the introduction of uniform regularisation with  $\beta=10^{-4}$  made the inversion of this model possible and indeed advantageous over the inversion of  $\mathbf{C}_{190}$ . Here we are interested in establishing the conditions under which the introduction of the frequency varying regularisation penalty  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  can further increase the

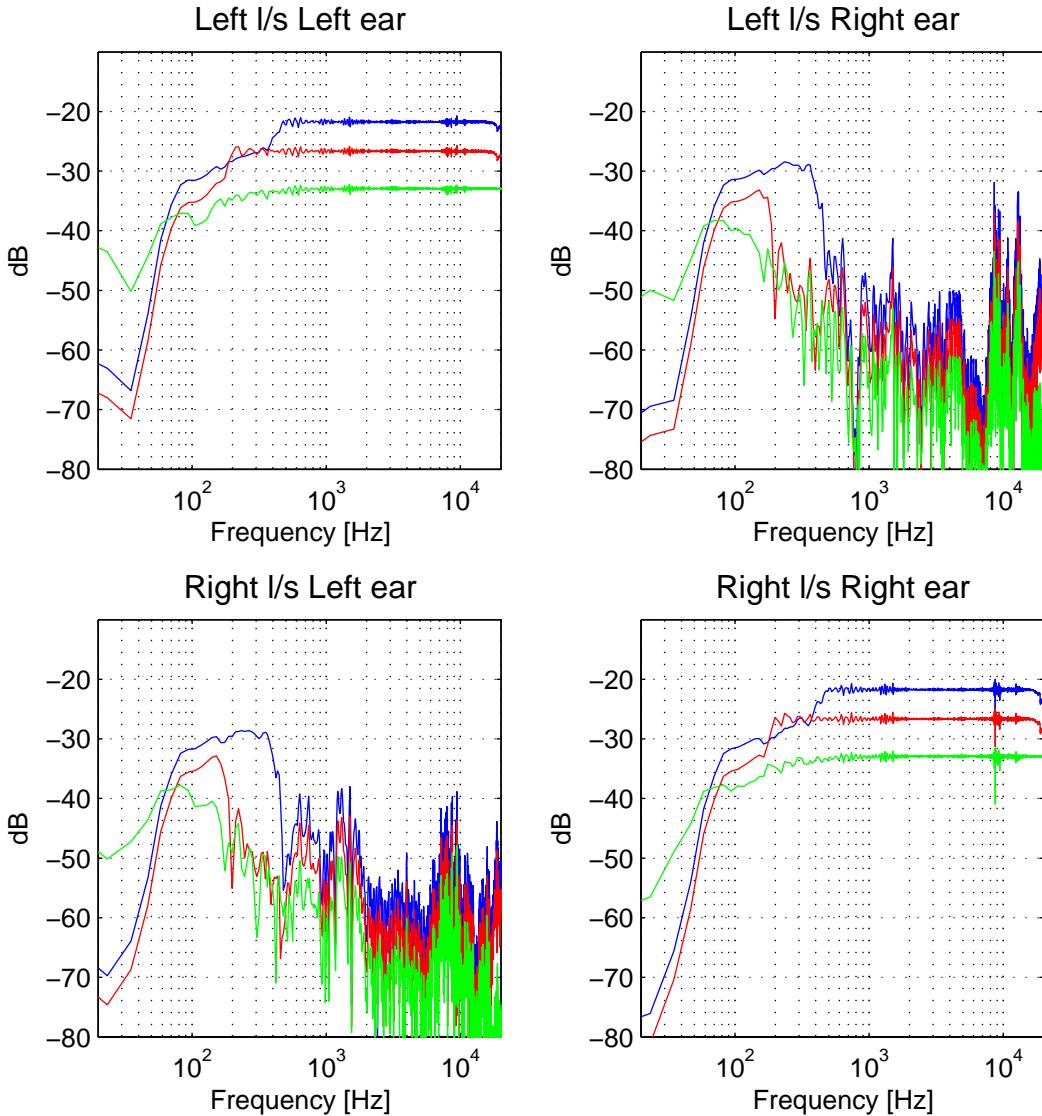
target level of the equalisation of the  $\beta=10^{-4}$  case shown in figures 4-6 and 4-7 preserving at the same the inversion accuracy of those results.

The simulation results<sup>51</sup> for the inversion of  $\mathbf{C}_{740}$  with  $\beta$  set to the value  $[10^{-2} \ 10^{-6} \ 10^{-2}]$  are depicted in figures 4-13 and 4-14 and compared with the results of the uniform value  $\beta=10^{-4}$  in the frequency and time domains respectively. The frequency varying regularisation is implemented in two ways, one with the penalty of  $\beta=10^{-2}$  extending up to 100Hz and the transition shaping with the half Hanning window extending from 20Hz to 200Hz (red line) and the other with the penalty of  $\beta=10^{-2}$  extending up to 200Hz (blue line) and the transition shaping with the half Hanning window extending from 20Hz to 500Hz. In both cases the  $\beta=10^{-6}$  penalty value extends up to 18kHz with the transition back to the  $\beta=10^{-2}$  penalty value being shaped by a half Hanning window extending from 15kHz to 20kHz. The filters that implement the uniform regularisation case are set to have 8000 taps (3000 taps for the non-causal and 5000 taps for the causal part) while the filters that implement the varying regularisation case are set to have 25000 taps (2000 taps for the non-causal and 23000 taps for the causal part). In all cases, the FFTs used for the computation of the inverse are set a length of  $2^{17}$  points.

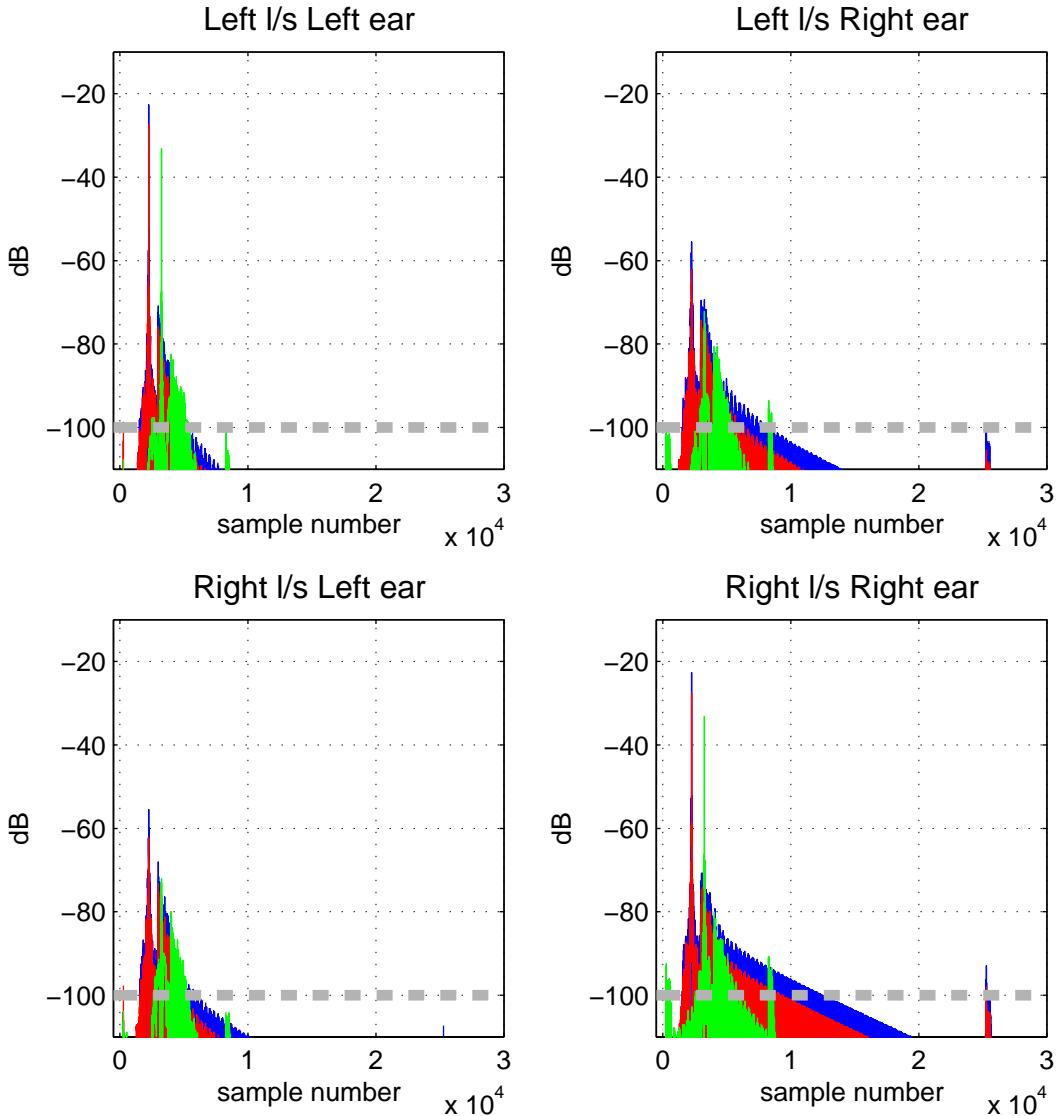
The increase in the target level of the inversion can be seen in the figures to be around 5dB from the  $\beta=10^{-4}$  to the first varying regularisation case and another 5dB to the second varying regularisation case. This is done without loss of the inversion accuracy in the region of low regularisation penalty. The reduction of the region that is effectively controlled is minimal in the first varying regularisation case with only the region from 150Hz to 200Hz being lost compared to the  $\beta=10^{-4}$  case. Naturally, the loss of control extends higher for the second varying regularisation case with the control starting to take effect after 400Hz.

---

<sup>51</sup> The results plotted in figures 4-13 and 4-14 are simulation results following the procedure described in §3.5.



**Figure 4-13:** Simulated inversion results. The plant model used for the determination of the inverse is  $\mathbf{C}_{740}$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 8000 coefficients (green line). The regularisation is set to  $\beta=[10^{-2} \ 10^{-6} \ 10^{-2}]$  (sharp transition) and the inverse length to 25000 coefficients (red line). The regularisation is set to  $\beta=[10^{-2} \ 10^{-6} \ 10^{-2}]$  (gentle transition) and the inverse length to 25000 coefficients blue line)



**Figure 4-14:** Simulated inversion results. The plant model used for the determination of the inverse is  $\mathbf{C}_{740}$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 8000 coefficients (green line). The regularisation is set to  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  (sharp transition) and the inverse length to 25000 coefficients (red line). The regularisation is set to  $\beta=[10^{-2} 10^{-6} 10^{-2}]$  (gentle transition) and the inverse length to 25000 coefficients (blue line). The grey dashed line shows the noise-floor of the corresponding measured results.

The most important difference between the uniform and varying regularisation cases can be seen in the time-domain plot of figure 4-14 which shows that for the same reduction of the end-effects, the frequency varying regularisation cases require 25000 taps compared to the 8000 taps required by the uniform regularisation case. This is of course due to the sharp roll-off of the frequency response at the point of change of the regularisation penalty. We note here that the computation of an inverse of  $\mathbf{C}_{1740}$  that

equally suppresses the end-effects of the inversion was impossible for these two choices of varying regularisation even when filters of 50000 taps were computed using FFTs of  $2^{18}$  points. It should be clear of course that the comparison between the inverse filters' length needed to suppress the end-effects below the noise floor for the regularised and non-regularised cases (figures 4-14 and 4-5 respectively) should take into account the fact that in regularised cases the headroom between the equalisation target level and the noise floor is 22-33dB wider than that of the non-regularised case.

Overall then we see that the introduction of uniform regularisation results in a benefit as regards the total dynamic range of the inversion between the level of the equalised signal and the noise-floor (65dB and 85dB for the  $\beta=10^{-4}$  and  $\beta=10^{-2}$  cases compared to 55dB for the  $\beta=0$  case). It does not however improve the level of cross-talk cancellation (while in certain cases it can degrade it) and when introduced in excess it can result in audible colouration of the reproduced signals (this is already evident in figure 4-11 for the results obtained with  $\beta=10^{-2}$ ). The use of frequency-varying regularisation can be used as a compromise between these two cases, as it does achieve an increase in the control target level compared to the low regularisation  $\beta=10^{-4}$  case avoiding at the same time the colouration caused by the high regularisation  $\beta=10^{-2}$  case. This is of course done at the expense of the total loss of the signal information in the very low frequency region.

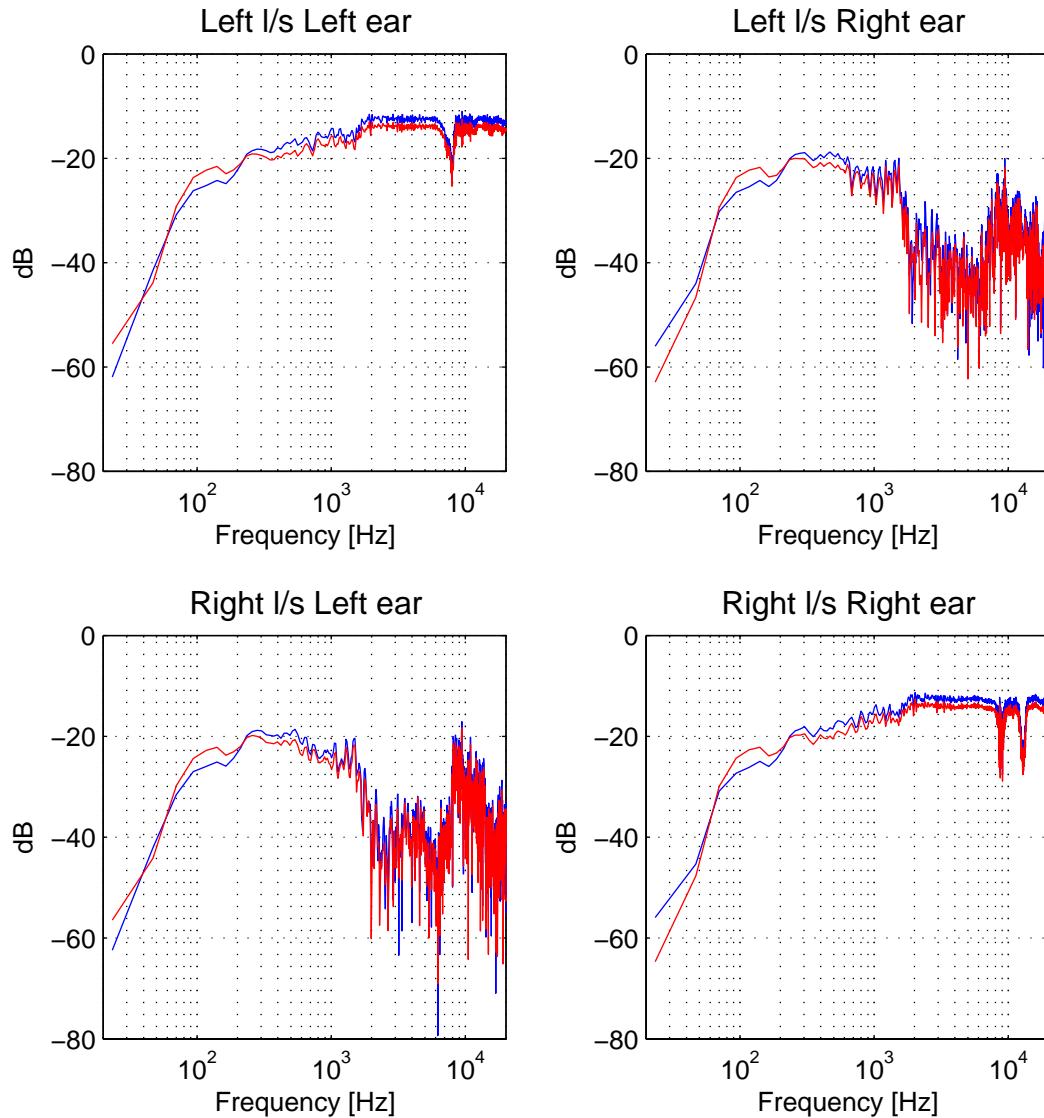
In fact, when used in this way, frequency varying regularisation effectively acts as a very sharp linear phase pass-band filter. In this sense it is normal for its time-domain response to extend for long, as would be the case for any frequency selective stage inserted in the reproduction chain (e.g. a digital crossover). The possible advantage of such a use of regularisation is that the delay needed for the linear-phase shaping of the control can only be applied once at this stage and effectively be combined with the modelling delay needed to compensate for the non-minimum character of the inversion. Obviously, the need for such a sharp frequency selection (and the entailed extended shape of the related inverse time-responses) is dictated by the fact that we desire to penalise the region below 100-200Hz but at same want the control to start being effective as close as possible to that lower limit. Such a problem does not exist

at the other end of the spectrum, where the rise in control effort is less steep and a less sharp frequency selection can be applied.

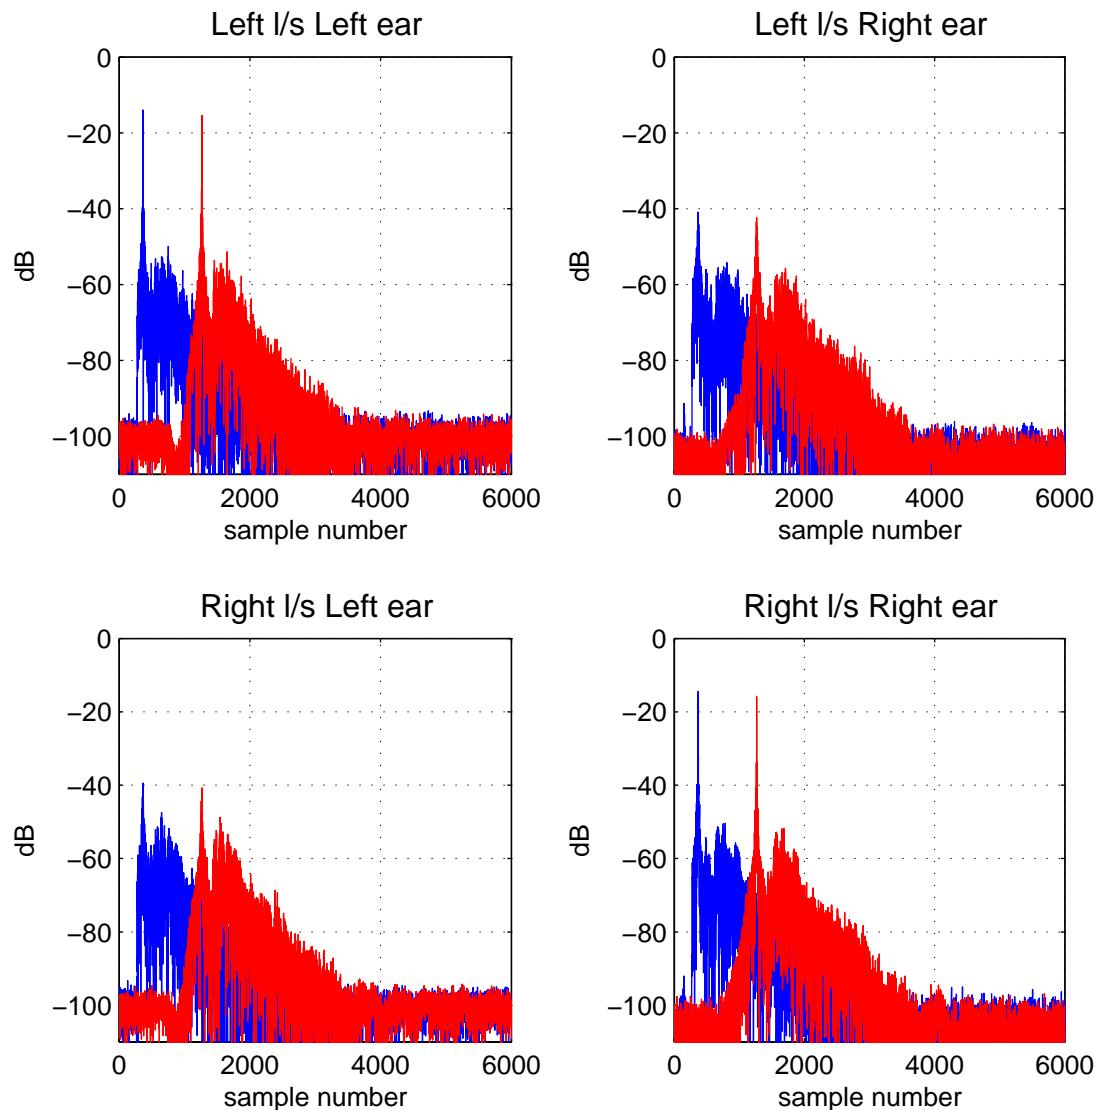
#### ***4.4 Length of the inverse when regularisation is applied***

The issue of the required length of the inverse filters was to some extent covered in the two previous sections. It was shown there that the length required for the inverse filters to adequately follow the decay of the ideal inverse was ranging from 1000 or 2000 taps (for the cases of the strictly anechoic model of the plant  $\mathbf{C}_{190}$  and a reasonable amount of regularisation) up to many thousands of taps (for the cases of the longer models of the plant  $\mathbf{C}_{740}$  and  $\mathbf{C}_{1740}$  and especially when frequency varying regularisation is implemented). This section completes the discussion on this issue by focusing on a number of cases where, as shown below, the choice of the remaining two parameters of the inversion limits the required length of the inverse filters to much lower values.

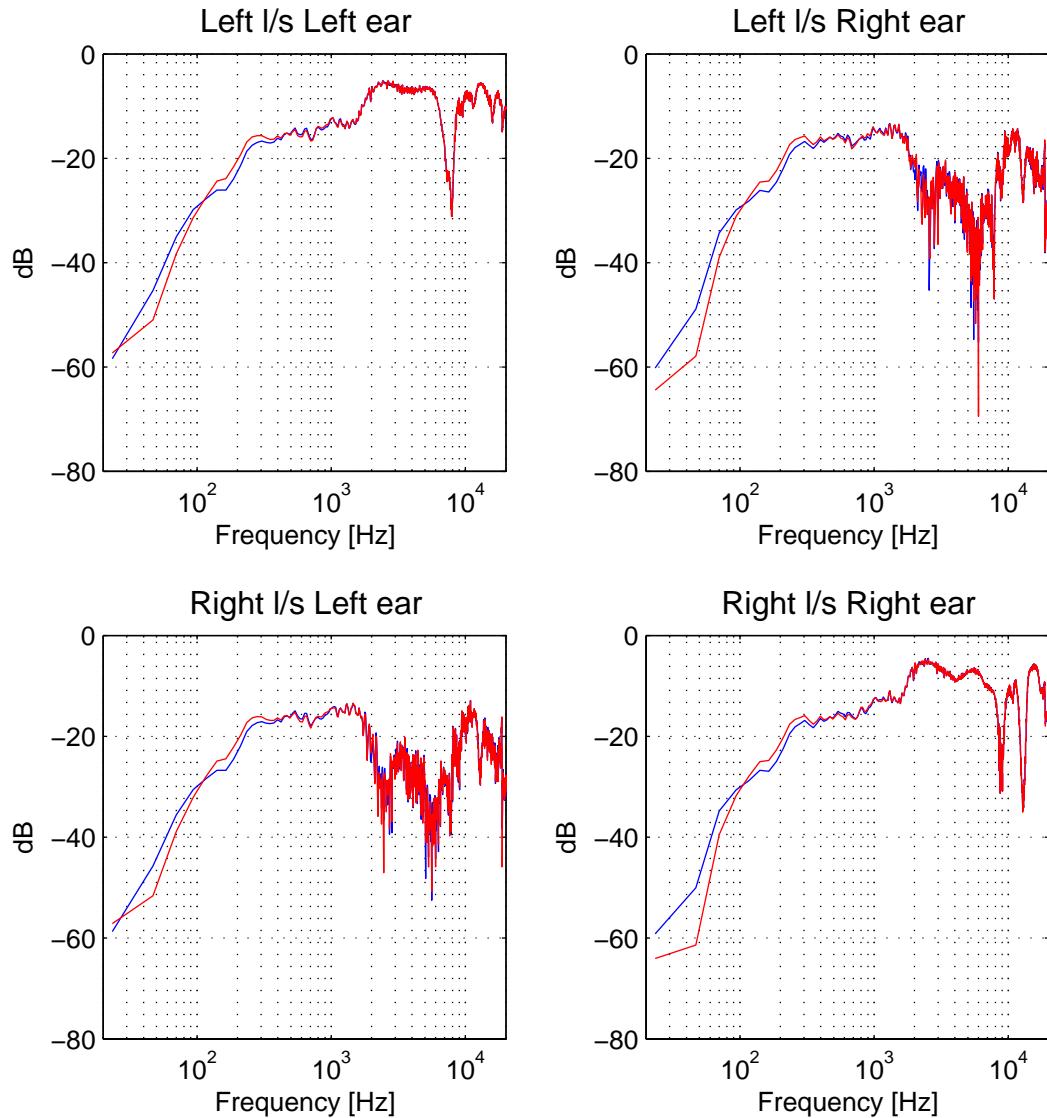
In figures 4-15 - 4-18 we plot the measured results of the inversion of  $\mathbf{C}_{190}$  with the regularisation parameter set to the value  $\beta=10^{-2}$  (figures 4-15 and 4-16) and  $\beta=10^{-1}$  (figures 4-17 and 4-18). The plots compare the results obtained when the inverse is realised with 200 tap and 2000 tap inverse filters. These measured results verify the fact that with the regularisation parameter set to these high values and with the plant modelled by the strictly anechoic part  $\mathbf{C}_{190}$ , the use of an inverse as short as 200 taps is enough to achieve the best obtainable inversion results for both these cases. The significant degradation of the ipsilateral equalisation and the reduction of the region of cross-talk cancellation as well as the further, but only marginal, increase in the dynamic range of the inversion can be seen in the  $\beta=10^{-1}$  results compared to those with  $\beta=10^{-2}$ . Furthermore, the similarity of the inversion results of figure 4-17 with the frequency response of the plant (see figure 4-22) demonstrate the fact that with the regularisation penalty  $\beta=10^{-1}$  the control is hardly applied at all.



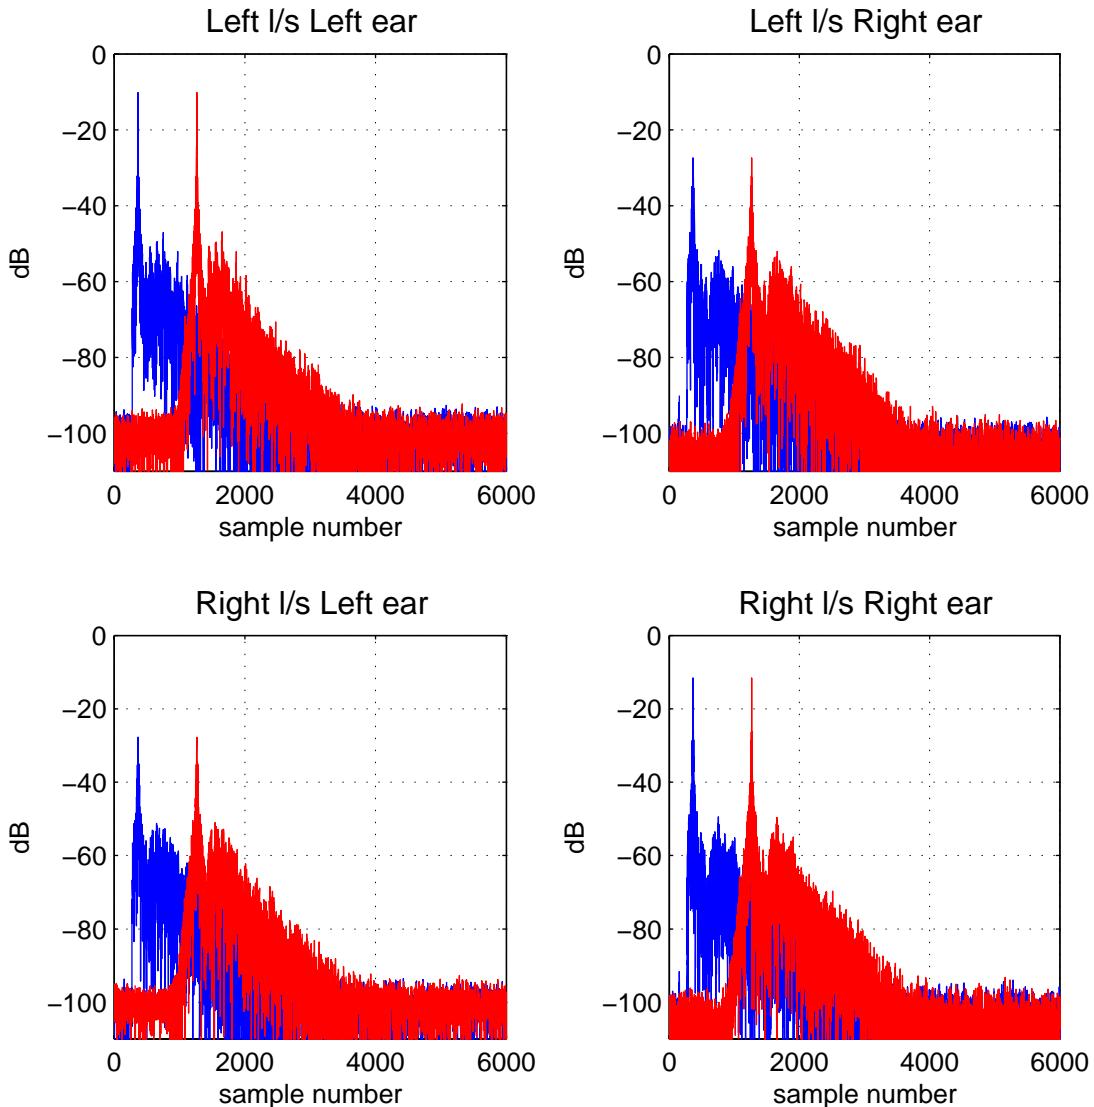
**Figure 4-15:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 2000 coefficients (blue line) and to 200 coefficients (red line).



**Figure 4-16:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 2000 coefficients (blue line) and to 200 coefficients (red line).

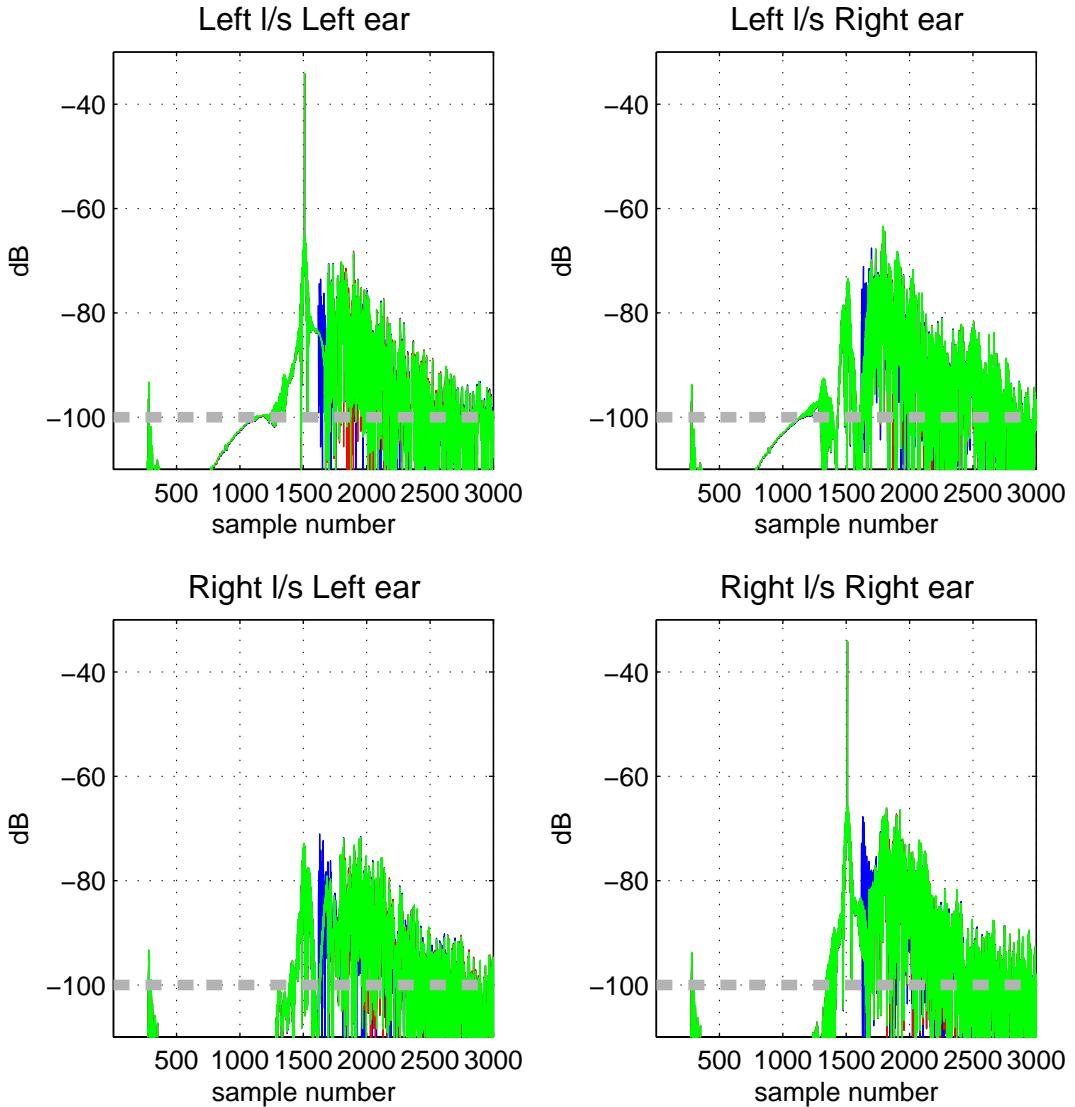


**Figure 4-17:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=10^{-1}$  and the inverse length to 2000 coefficients (blue line) and to 200 coefficients (red line).



**Figure 4-18:** Measured inversion results. The plant is modelled with  $\mathbf{C}_{190}$ . The regularisation is set to  $\beta=10^{-1}$  and the inverse length to 2000 coefficients (blue line) and to 200 coefficients (red line).

A further case of interest is presented in the results of figure 4-19. In this case we compare the inversion results obtained when  $\mathbf{C}_{190}$  is inverted with the regularisation parameter set to  $\beta=10^{-4}$  and with the length of the inverse set to three different values. The green line results correspond to an inverse length of 2500 coefficients (1250 causal and 1250 anti-causal), the red line to an inverse of 1550 coefficients (300 causal and 1250 anti-causal) and the blue line to an inverse of 1350 coefficients (100 causal and 1250 anti-causal).



**Figure 4-19:** Simulated inversion results. The plant model used for the determination of the inverse is  $\mathbf{C}_{190}$  and the regularisation is set to  $\beta=10^{-4}$ . The inverse length is set to 2500 coefficients (1250 causal and 1250 anti-causal, green line), 1550 coefficients (300 causal and 1250 anti-causal, red line) and 1350 coefficients (100 causal and 1250 anti-causal, blue line). The grey dashed line shows the noise-floor of the corresponding measured results. The impulse responses are time-shifted relative to each other so that the delta spike appears at the same point in time in all three cases.

As can be seen in the figure, the 1250 anti-causal coefficients are just enough to suppress the negative-time truncation effect down to the noise floor. However, even though a similar number of causal coefficients is needed to achieve the same suppression for the positive-time truncation effect, this is not reflected in the results. This is because at the point along the equalised responses where the truncation effect occurs, these responses exhibit the error due to the truncation of the plant model  $\mathbf{C}_{190}$ .

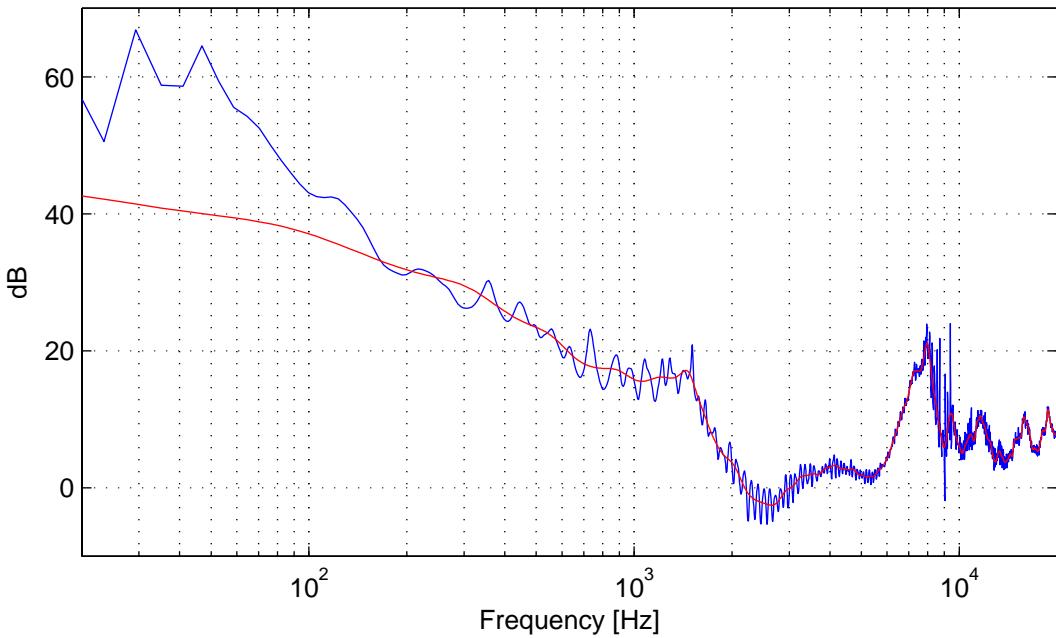
that is used for the determination of the inverse (see also the discussion in pp. 81-83). This error exceeds the error due to the truncation of the inverse filter. As a result, exactly the same inversion results obtained with 1250 causal coefficients are also obtained with 300 coefficients. A further reduction of the causal length to 100 coefficients becomes visible as it places the effect due to the truncation of the inverse further up along the equalised response, less than 190 samples after the delta spike. At this point, the error due to the truncated model  $\mathbf{C}_{190}$  that is used for the determination of the inverse is insignificant and hence the error due to the truncation of the inverse is revealed. It should be stressed however that even in this case, this concentrated increase in the inversion error is susceptible to temporal masking and would almost certainly be imperceptible. In that event we see that, given the practically unavoidable element of error introduced by the determination of the inverse using the strictly anechoic version of the plant matrix, as few as 100 coefficients are enough for the modelling of the causal part of the inverse.

## ***4.5 Comments on inversion results under ideal conditions***

We conclude this chapter by discussing a number of points related to the results presented above. The first of these points is the complete ineffectiveness of the control at the very low frequency region below 200Hz, something that was systematically verified in all the results presented. As has already been identified in the literature, this result has to be attributed to three basic reasons being: (i) the limited frequency resolution resulting from the truncation of the measured plant responses (Gardner 1998), (ii) the inherent properties of the active sound control problem that are primarily determined by the geometry of the reproduction (Takeuchi 2001) and (iii) the poor low-frequency response of the control transducers. In what follows we discuss the implications of these issues for the results presented above and we argue that the control of this low frequency region is effectively impossible with a sound reproduction arrangement like that examined here, even in ideal conditions of implementation.

Starting with the first of the aforementioned issues and considering the case where the inverse is computed using the strictly anechoic model of the plant  $\mathbf{C}_{190}$ , we see that the frequency resolution of the inverse for the 48kHz sampling rate that was used will be limited to approximately 250Hz. Hence any spectral detail finer than that will be inevitably left uncorrected by the inverse. Such fine spectral detail is conceptually irrelevant in terms of the HRTF notion which, by its very definition, is only related to the strictly anechoic part of the plant. But when the inversion of the actual plant is considered, the late decay of the plant's responses does indeed introduce such fine spectrum detail which is lost in the  $\mathbf{C}_{190}$  model.

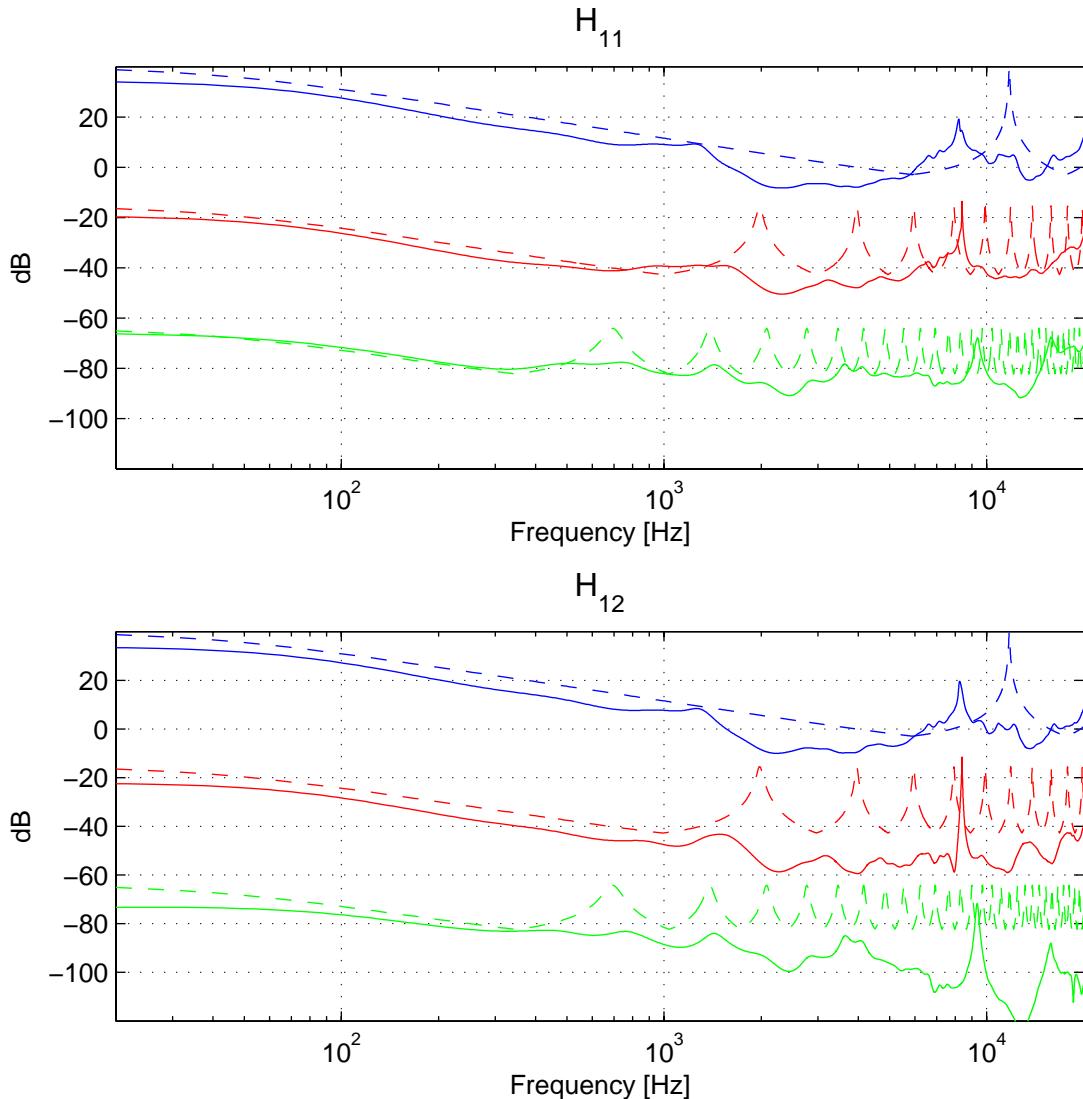
On the other hand, the use of a longer model of the plant for the determination of the inverse and the corresponding refinement of the inverse's frequency resolution entails other, more severe, difficulties and limitations. As was seen in figures 4-1 and 4-3, the equalisation applied by  $\mathbf{H}_{740}$  and  $\mathbf{H}_{1740}$  is indeed visible down to below 200Hz but is practically destroyed anywhere else due to the early truncation of the inverse imposed by the 2000 and 4000 tap inverse models. Most importantly, the equalisation target level is significantly reduced leading to a severe reduction of the dynamic range of the processing. Unlike the early truncation problem which can be addressed with the use of longer inverse models, the loss of dynamic range cannot be addressed as straightforwardly. In fact, this loss of dynamic range should be directly related to the refinement of the frequency resolution as the inverse now follows the boosted frequency response magnitude of the exact inverse (see figure 4-10) more closely and thus imposes a higher demand in terms of the amplification required for the exact inversion. This is illustrated in figure 4-20 where we plot the element  $H_{11}$  of the non-regularised inverse matrices  $\mathbf{H}_{190}$  and  $\mathbf{H}_{1740}$  (i.e. the inverse matrices computed on the basis of the  $\mathbf{C}_{190}$  and  $\mathbf{C}_{1740}$  plant models respectively). As can be seen in the figure, the refinement of the inverse's spectrum above 200Hz is confined to the correction of minimal details. Conversely, below 200Hz,  $\mathbf{H}_{1740}$  follows the power boost of the exact inverse, a feature that is absent in  $\mathbf{H}_{190}$ .



**Figure 4-20:** Element  $H_{11}$  of the non-regularised matrices  $\mathbf{H}_{190}$  (red line) and  $\mathbf{H}_{1740}$  (blue line) prior to normalisation.

Now this boost in the required control effort at the low end of the frequency spectrum is a result of the combination of the two remaining factors listed above. As is investigated in detail by Takeuchi (2001), the angle spanned by the two control transducers in symmetric geometries determines the position along the frequency axis for a number of regions where the inversion becomes ill-conditioned and the magnitude of the inverse increases. Using a simplified free-field model describing the acoustic propagation from two ideal monopole sources to two receivers in the free-field (no head present), Takeuchi (2001) shows that the position of these ill-conditioned frequencies will exhibit a periodicity with frequency and will be determined by the interleaving maxima of the two singular values of the inverse matrix  $\mathbf{H}$ . The first ill-conditioned frequency always appears at zero frequency. As the transducer span becomes smaller the ill-conditioned frequencies are pushed further apart but are associated with an increase in the magnitude of the inverse. In this sense, the  $10^\circ$  span Stereo Dipole geometry used here, yields a wide frequency interval between the zero frequency ill-conditioned region and the first ill-conditioned frequency to appear after that (at approximately 11kHz with nominal values of the distances describing the geometry), Kirkeby et al. (1998a). This, however, comes at

the expense of an increase in the associated inverse magnitude in the first ill-conditioned frequency region.



**Figure 4-21:** Magnitude of the elements  $H_{11}$  and  $H_{12}$  of the non-regularised inverse matrix for different loudspeaker span angles ( $10^\circ$  Stereo Dipole blue line,  $60^\circ$  Stereo red line and  $180^\circ$  green line). Magnitude predicted by the free-field ideal monopole transducer model (dotted lines). The results for the  $60^\circ$  and  $180^\circ$  cases are plotted with a vertical offset of  $-40\text{dB}$  and  $-80\text{dB}$  respectively.

This effect is depicted in figure 4-21 where we plot the magnitude of the elements  $H_{11}$  and  $H_{12}$  of the non-regularised inverse matrix for three different symmetric geometries, namely that of  $10^\circ$ ,  $60^\circ$  and  $180^\circ$  total angle spanned between the two loudspeakers. The results are obtained using HRTFs from the MIT database (Gardner

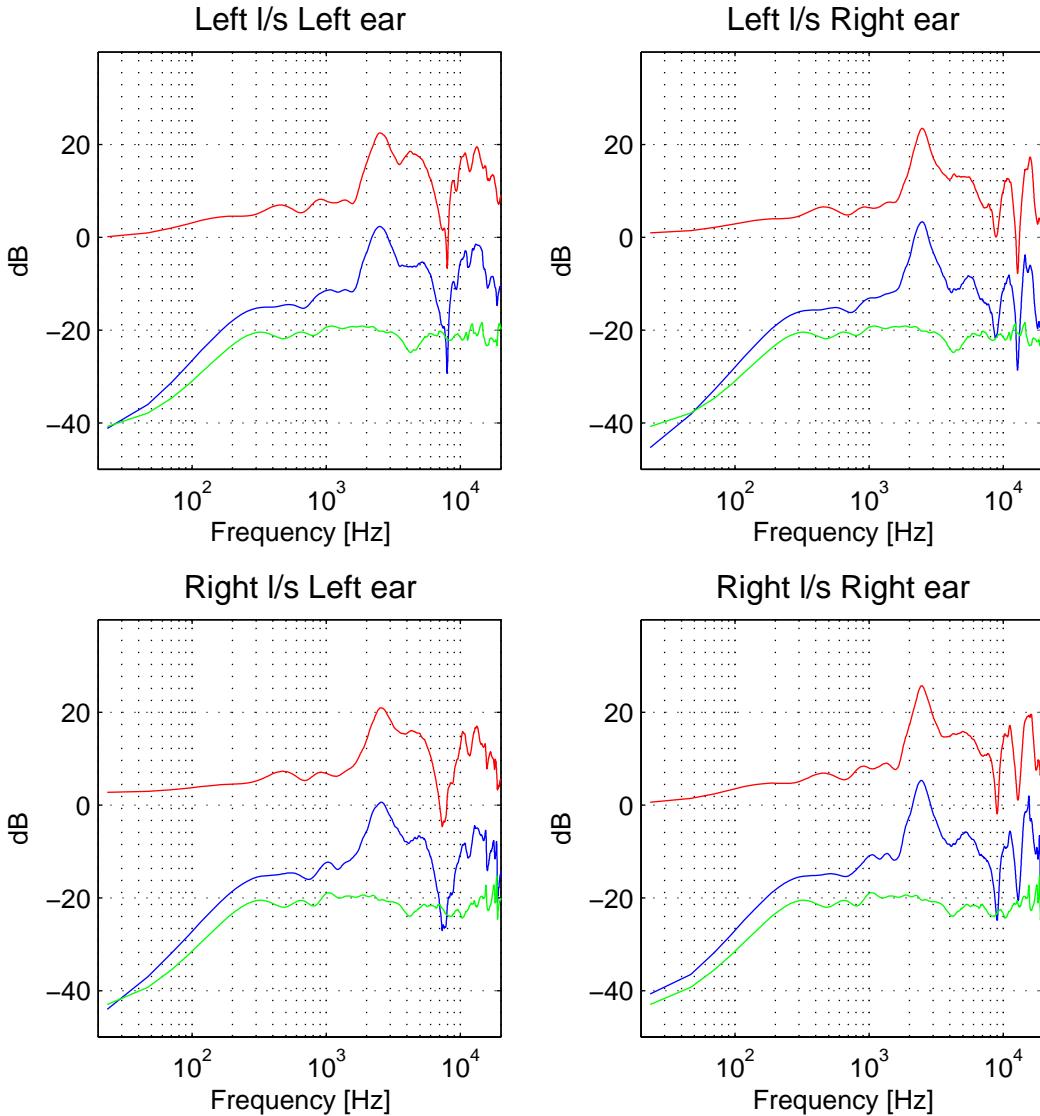
and K. Martin 1994)<sup>52</sup>. In the same plot, with a dotted line, we plot the maximum power amplification versus frequency as estimated by the Singular Value Decomposition of the inverse matrix  $\mathbf{H}$ . In this case,  $\mathbf{H}$  is derived from the plant corresponding to the free-field reproduction using ideal monopole sources for the same angle spans (Takeuchi 2001).

As can be seen in the figure, the increased low frequency power requirement estimated by the free-field model for the small angle span is very well replicated by the measured HRTF data. At higher frequencies, above 1kHz, a slight (but systematic) divergence of the estimated ill-conditioned frequencies can be observed for the 60° geometry. This is due to the fact that, as implemented here, the model does not take into account the frequency dependent interaural delay due to the presence of the head<sup>53</sup>. In the 180° span case, the estimated peaks due to the ill-conditioning are clustered very close together and appear in the measured results as a slight ripple. Furthermore, the free-field model is seen to overestimate the power output needed at higher frequencies as it does not take into account the natural shadowing of the head at higher frequencies which facilitates the cross-talk cancellation. This overestimation is obviously more severe in the 180° span case where the head shadowing is greater. The results of figure 4-21 also show that the sharp peak in the inverse's magnitude response at approximately 9kHz is not due to ill-conditioning of the inverse but due to a notch present in the measured single-channel HRTFs (pinna-notch, see (Lopez-Poveda and Meddis 1996) and references therein). The presence of this notch and its slow upwards shift in frequency as the source moves towards the side of the listener seen in figure 4-21 is also visible in similar results by Takeuchi (2001). It should be noted however that the combination of the two effects, namely the boost due to ill-conditioning and the solitary peak at around 9kHz can lead to a much sharper peak. This can be seen in the 60° span results and is due to the fact that in that case the 9kHz peak coincides with one of the ill-conditioned frequencies.

---

<sup>52</sup> In the MIT HRTF database KEMAR is assumed to be symmetrical and so when those data are used  $H_{ij}=H_{ji}$  and  $H_{ii}=H_{jj}$ .

<sup>53</sup> The results shown in figure 4-21 assume a frequency constant interaural time delay that is estimated here as the maximum of the interaural cross-correlation function between the left and right ear HRTFs.



**Figure 4-22:** Measured plant responses (blue line), measured free-field loudspeaker responses (green line) and corresponding HRTFs (red line) for the  $10^\circ$  geometry.

Finally we note that the increase in the required low-frequency power amplification of the inverse is further exaggerated in the case where the plant is modelled with measured responses that include the response of the loudspeaker rather than with HRTFs as was the case in the results of figure 4-21. This is made apparent in figure 4-22 where we plot the magnitude of the measured plant responses used for the results presented in the thesis and compare them with the corresponding free-field equalised HRTFs. As is seen in the figure, the measured responses display a roll-off below 200Hz following the low-frequency roll-off of the free-field loudspeaker responses. Evidently this leads to a further increase of the low-frequency content of the inverse matrix  $\mathbf{H}$ . The free-field equalised HRTFs on the other hand, converge to a 0dB level

at the low frequency end of spectrum. This is the typical low-frequency shape of the HRTF as at this frequency region the wavelength of the sound field is much larger than the physical dimensions of the listener's head and body and thus the listener does not disturb the soundfield. In the HRTF results of figure 4-22 note a significant asymmetry between the left and right ear responses of KEMAR. This asymmetry is not contained in the MIT HRTF database responses but it is verified in data from other databases (Algazi et al. 2001), (Mannerheim et al. 2004) and is specific to the small KEMAR pinnae (DB-60 and DB-61).

The results presented above then verify the argument put forth above regarding the inevitable inefficiency of the control at this very low frequency region. As was discussed, the acoustic control properties of the inversion problem result in a boost of the power output required by the control sources in order to control this region. This boost is maximal for the  $10^\circ$  as predicted by the free-field model of the problem but is also evident in the  $60^\circ$  degrees geometry. The remaining peaks predicted by the model are indeed present in the measured responses case, but at much lower amplitude than predicted due to the effect of the head shadowing. Hence, for the  $10^\circ$  geometry, the overall required power amplification is dictated by the low frequency content of the inverse. Furthermore, the exact correction of this low frequency region requires the knowledge of the full decay of the impulse responses comprising the plant and of the low frequency response of the control transducers both of which further increase the required low frequency amplification. The time-domain equivalent of that is a long ringing in the inverse impulse responses which, as was seen above, need to be implemented with several thousands of coefficients to avoid audible truncation effects. Even in that case, given that the plant response is measured using the full dynamic range of the audio reproduction chain (determined by the quantisation noise level, the analogue amplification capability, the sensitivity and maximum level of linear response of the loudspeakers, etc.) such an amplification requirement leads to a severe loss of dynamic range that would definitely be audible in the delivered audio material. In simple terms, the attempt to correct the region below 200Hz would inevitably destroy the control of the rest of the spectrum. Hence, the use of regularisation becomes necessary and (as was discussed in §4.3) the consequent loss of control over the 200Hz region becomes inevitable.

On the other hand, the presented results suggest that under certain conditions the proper use of regularisation can achieve a gentle transition between the controlled (above 200Hz) and the non-controlled (below 200Hz) region and even extend the controlled region down to 100Hz. This makes possible the reproduction of the very low frequency content by a separate transducer by-passing the inverse filtering network. Such a design has an already existing commercial counterpart in a number of “2.1” commercially available systems.

A second point to be discussed regards the approach taken in the thesis for the modelling of the plant matrix  $\mathbf{C}$  using measured responses that contain the response of the loudspeakers rather than the use of equalised HRTFs (see also the discussion in §1.2). Assuming that the responses of the reproduction loudspeakers are matched in magnitude and time-aligned, the difference between the use of non-equalised responses and HRTFs in  $\mathbf{C}$  essentially amounts to the choice between correcting and not correcting the free-field response of the loudspeakers. In simple terms, determining the inverse matrix  $\mathbf{H}$  using an HRTF-based model of  $\mathbf{C}$  should, in principle, result to the listener perceiving the virtual source from the intended position in space but as if being emitted by the reproduction loudspeaker rather than the original sound source.

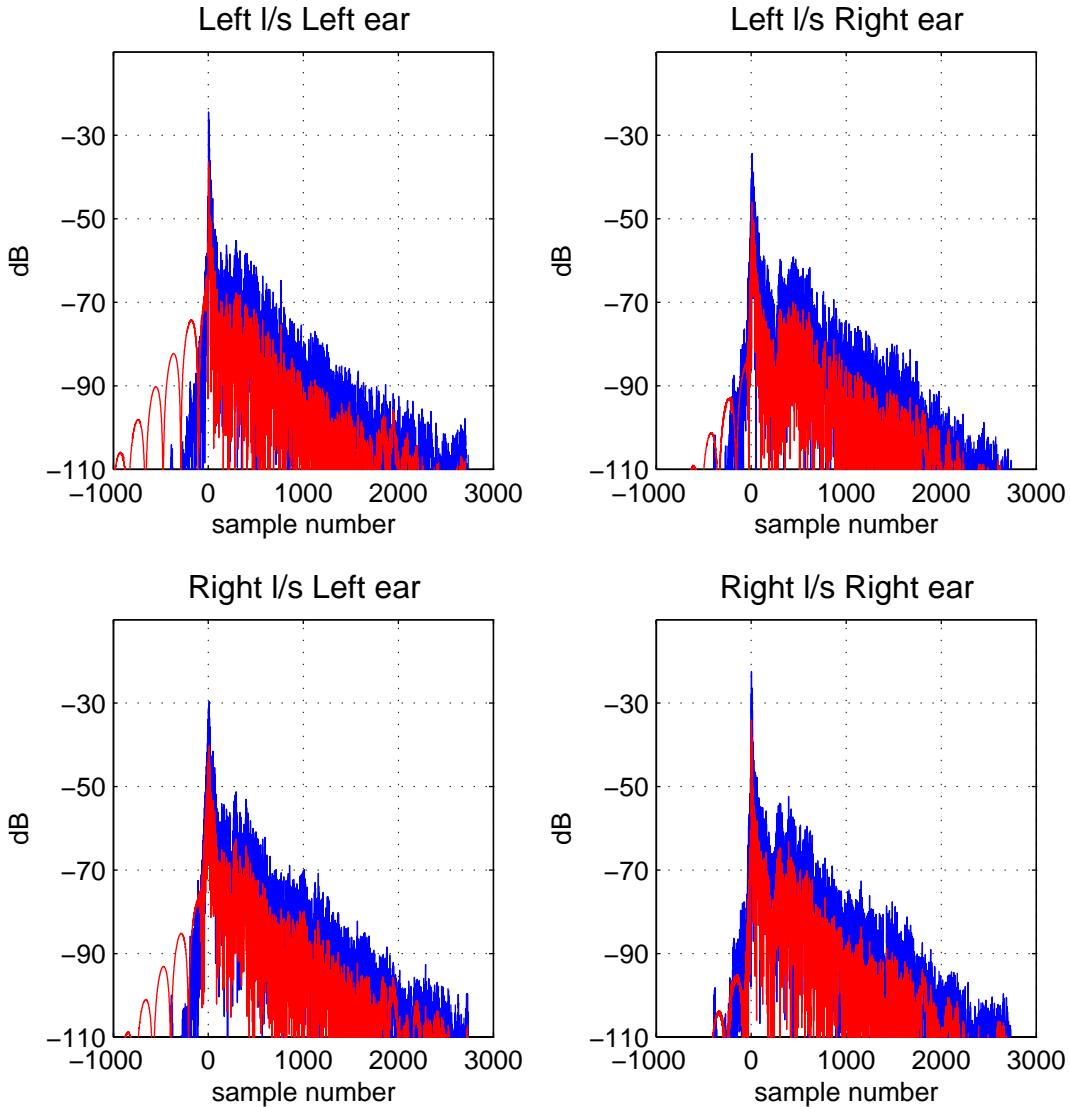
Even though this is an effect one would indeed desire to alleviate, the use of the inverse  $\mathbf{H}$  as a means to also correct the individual response of the reproduction loudspeakers introduces restrictions to the design. Apart from the fact that an *in situ* measurement of  $\mathbf{C}$  would be needed for such a design, the most prominent of these restrictions is the dynamic range loss imposed in the inverse filtering. The point is clarified with the results of figure 4-22. As can be seen in the figure, the free-field response of the loudspeakers presents a very good approximation to the ideal flat response for frequencies above 200Hz as is typically the case for on-axis responses of loudspeakers of reasonable quality. A similarly agreeable behaviour is demonstrated as regards the phase of the loudspeakers as is further discussed in §5.2. It would be thus reasonable to argue that, as far as the loudspeaker responses are concerned, the region above 200Hz could be left uncorrected while the region below 200Hz a correction should not be attempted due to the considerations discussed above.

Indeed the correction of the loudspeaker response was not included in the research interest of this project. In fact, similar issues regarding the depletion of the available dynamic range associated with the equalisation of loudspeaker responses have been mentioned since the early days of inverse filtering DSP applications (Mourjopoulos et al. 1982). Such issues are still considered to be one of the most significant inhibitory factors in the use of such techniques for the improvement of loudspeaker response (Stuart and Howard 2005). Instead, the reason we chose to use a plant matrix that contains the response of the loudspeakers was that (as is discussed in detail in §3.2) our objective is to determine the realistically optimal results that can be produced by a design like that discussed here and to investigate the influence of the associated design parameters in the achievement of these results. Such an investigation could obviously not be conclusive if an element of error was already present in the inversion as becomes apparent in the following set of results.

The effect of the use of HRTFs for the determination of  $\mathbf{H}$  is displayed in figures 4-23, 4-24 and 4-25<sup>54</sup>. In all these results the length of the inverse (1500 coefficients for the non-regularised case and 600 coefficients for the regularised case) is set by trial-and-error to the minimum value that gives the optimum frequency domain results and suppresses the truncation effects below the 100dBFS level. As seen in the figures, with the plant model used for the determination of the inverse being now effectively normalised by the level of the free field sound, the target level of the equalisation rises in comparison with the equivalent results of figures 4-3 and 4-4 where the inverse was determined using a plant model that contained non-equalised measurements. The -30dBFS equalisation target level of figures 4-23, 4-24 and 4-25 (red line) can thus be regarded as the actual loss of dynamic range of the non-regularised inversion. Furthermore, with the low frequency magnitude content of the plant matrix increased, the low frequency boost in the inverse filters is reduced. Consequently the decay of the inverse responses in time-domain is increased and hence a shorter model of the inverse is required in order to avoid the truncation effects.

---

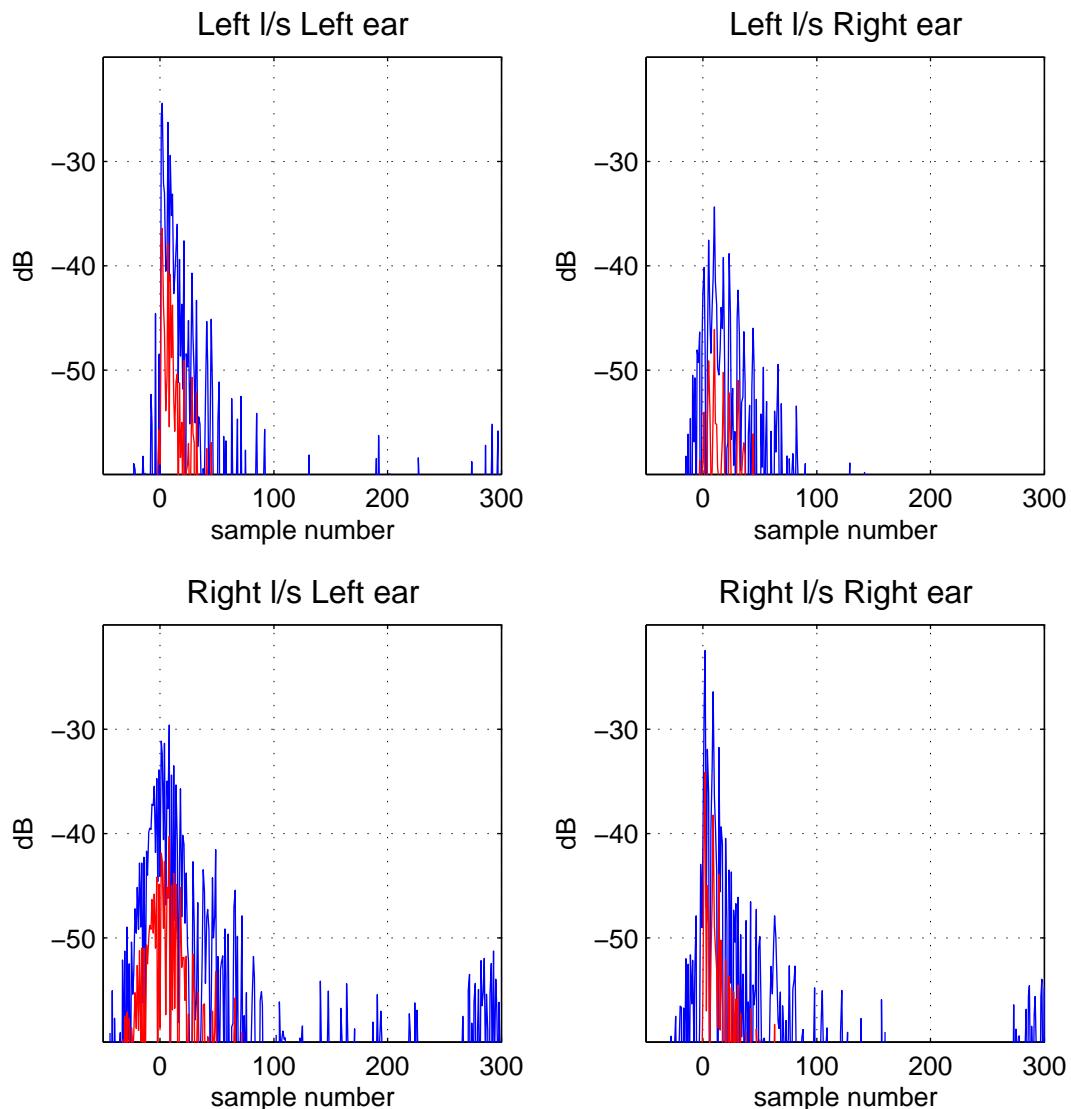
<sup>54</sup> These cases were not included in the measurements and the results plotted in figures 4-23, 4-24 and 4-25 are simulation results following the procedure described in §3.5.



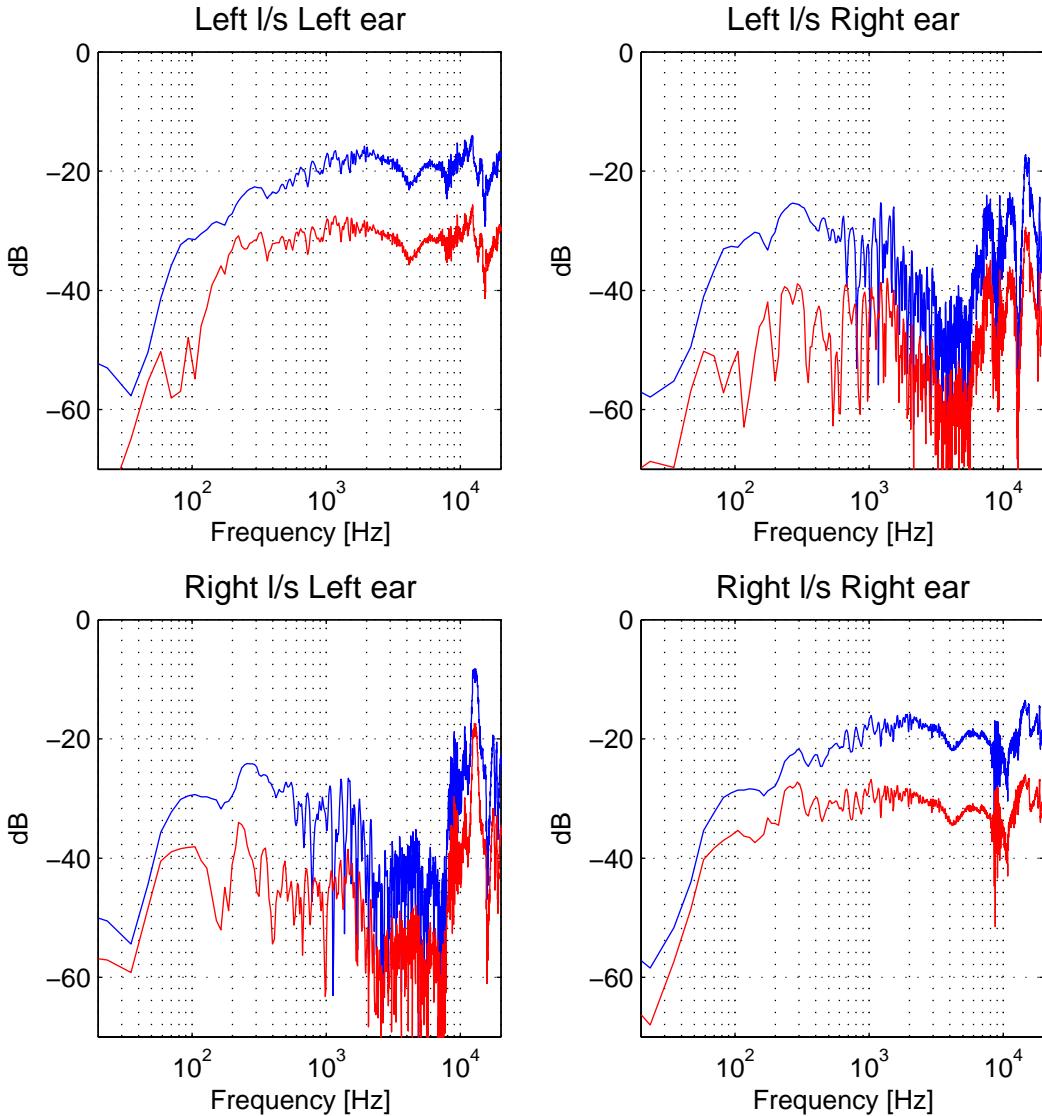
**Figure 4-23:** Simulated inversion results. The plant model used for the determination of the inverse contains the free-field equalised HRTFs of figure 4-22. The regularisation is set to  $\beta=0$  and the inverse length to 1500 coefficients (300 causal and 1200 anti-causal, red line). The regularisation is set to  $\beta=1$  and the inverse length to 600 coefficients (200 causal and 400 anti-causal, blue line).

This relaxed requirement for the length of the inverse filters is however predominantly due to the presence of a significant element of error following the initial equalisation spike. As is depicted in more detail in figure 4-24, the equalisation spike is now replaced by a smeared pulse with this forward-time smearing corresponding to the late arriving energy of the uncorrected loudspeaker responses. This decay is physically masking any truncation effect that would otherwise be visible (and separately audible) in the equalised responses thus allowing for the much earlier truncation of the inverse. Following the discussion in pp. 81-82 regarding the

perceptual masking of such late arriving energy, we should again expect this effect to not be audible as a time-domain artefact. In fact, given that this energy is arriving immediately following the equalisation spike and considering the time-integration properties of the human auditory system, we should expect it to enhance the perceived loudness of the delivered material.



**Figure 4-24:** Zoomed-in depiction of the results of figure 4-23.



**Figure 4-25:** Simulated inversion results. The plant model used for the determination of the inverse contains the free-field equalised HRTFs of figure 4-22. The regularisation is set to  $\beta=0$  and the inverse length to 1500 coefficients (300 causal and 1200 anti-causal, red line). The regularisation is set to  $\beta=1$  and the inverse length to 600 coefficients (200 causal and 400 anti-causal, blue line).

On the other hand, moving to the frequency-domain inversion results, we see that both the ipsilateral equalisation and the cross talk cancellation are significantly inferior to the optimal results presented in the previous sections. We note that the non-regularised results of figures 4-23 – 4-25 (red line) present the inversion's absolute optimum when the inverse matrix  $\mathbf{H}$  is computed using a plant model  $\mathbf{C}$  comprising HRTFs. As should be expected in this case, the ipsilateral equalisation results correspond to the ideal flat response multiplied by the loudspeakers' free-field

response with the characteristic roll-off at the low-frequency end and the shallow dip being clearly visible at the frequency interval of the loudspeakers' crossover.

More importantly, as can be seen by comparison of the non-regularised with the regularised results (red and blue line respectively), the error introduced by the imperfect modelling of the plant with HRTFs creates a large margin for reduction of the applied control effort that has effectively no cost on the quality of the inversion results. That is, when regularisation is introduced the target level of the equalisation increases by more than 10dB and the required length of the inverse reduces to less than half with virtually no further degradation of the inversion results. Note that given the increased magnitude of the plant matrix model, the absolute value of the regularisation parameter  $\beta$  resulting in an equally significant penalisation of the effort is also increased in magnitude.

We can then conclude that the choice of including the loudspeakers' response in the plant matrix to be inverted for the results presented here was indeed justified in terms of conclusions drawn. The correction of these responses does impose extra requirements compared to the inversion of the HRTF-based plant. However, when those requirements are met the inversion result can be significantly improved. More importantly, when the correction of the loudspeakers' response is not included in the inverse design, the inversion results (even in ideal conditions of implementation) should be expected to reach a plateau of accuracy. This plateau is met with much more modest values of the inverse design parameters (i.e. higher regularisation and lower filter length) than would be estimated by results where the inversion process is simulated by convolving through the same HRTF-based matrix  $\mathbf{C}$  that is used for the determination of the inverse.

Finally, a third point to be discussed is that of the phase of the equalised responses. In the results presented throughout the thesis we present the magnitude of the frequency response and the impulse response of the inversion results. Obviously this pair constitutes a full quantitative description of the results, assuming linearity of the system, and it also contains practically all information that can be straightforwardly interpreted in terms of its expected perceptual significance. Phase information on the other hand is not as easy to interpret. Even the choice of the phase-related quantity that bears the perceptually important information -be it phase, phase delay or group

delay- to a great extent remains an open issue. Even greater uncertainty exists as to the appropriate perceptual indices and perceptibility thresholds related to the phase distortion in a system or signal.

One thing being certain though is that the “Ohm’s law” assertion of the early-days of psychoacoustics (see the review of Schroeder (1975) and references therein), namely that the human auditory system is only sensitive to the spectrum amplitude of audio stimuli and is completely “phase insensitive”, does not hold. A significant amount of experimental evidence now exists that proves that the human auditory system can indeed detect specific types of phase distortion, e.g. (Blauert and Laws 1978), (Lipshitz et al. 1982), (Schroeder 1975) and much more recently (Flanagan et al. 2005). That said, no complete model of the perceptual weight of phase distortion has been yet convincingly put forward. Open issues remain as to whether findings based on highly specialised artificial test stimuli also apply to natural sound signals such as speech or music. Similarly, audibility thresholds determined using headphone reproduction or anechoic reproduction over loudspeakers and trained listeners as subjects, seem to significantly rise when the phase distorted signals are presented to either untrained listeners or through loudspeakers or in more reverberant environments. Furthermore, most of the references mentioned above seek to predict and model the absolute threshold of perceptibility of phase distortion. This should be expected to be lower than the level of distortion that is perceived as actually degrading the audio presentation. This in turn would be separately influenced by some of the parameters mentioned above such as the level of training of the listener and the type of material presented.

In this view, the detailed characterisation of the obtained results with respect to their phase characteristics falls outside the scope of this research project. Nevertheless, a simple finding regarding the phase of the obtained results can be shown. This is done in figure 4-26 where we plot the group delay of the equalised ipsilateral paths for four different inverse designs. These include the results<sup>55</sup> obtained with the non-regularised inverse  $\mathbf{H}_{190}$ , the regularised inverse  $\mathbf{H}_{190} (\beta=10^{-4})$ , the regularised inverse  $\mathbf{H}_{1740} (\beta=10^{-4})$  and a strictly causal and stable version of the non-regularised inverse  $\mathbf{H}_{190}$  denoted with  $\mathbf{H}_{min}$ . In this last realisation, the common denominator of equation

---

<sup>55</sup> The results of figure 4-26 are simulation results following the procedure described in §3.5.

(2-14) has been replaced by its minimum phase reconstruction as in equation (4-1)<sup>56</sup>. The four transfer functions  $H_{ij}$  of equation (4-1) will thus correspond to causal and stable time-responses without the need to introduce modelling delay. In all cases the length of the inverse filters is set to be long enough to eliminate any significant truncation effects.

$$\mathbf{H}_{min}(\omega) = \begin{bmatrix} \frac{\hat{C}_{22}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} & \frac{-\hat{C}_{12}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} \\ \frac{-\hat{C}_{21}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} & \frac{\hat{C}_{11}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} \end{bmatrix} \quad (4-1)$$

The equalisation matrix  $\mathbf{X}_{eq}$  resulting from the use of this approximation to the inverse  $\mathbf{H}_{min}$  will then be as in equation (4-2).

$$\begin{aligned} \mathbf{X}_{eq}(\omega) &= \mathbf{C}(\omega) \cdot \mathbf{H}_{min}(\omega) \\ &= \begin{bmatrix} \frac{C_{11}(\omega) \cdot \hat{C}_{22}(\omega) - C_{12}(\omega) \cdot \hat{C}_{21}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} & \frac{C_{12}(\omega) \cdot \hat{C}_{11}(\omega) - C_{11}(\omega) \cdot \hat{C}_{12}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} \\ \frac{C_{21}(\omega) \cdot \hat{C}_{22}(\omega) - C_{22}(\omega) \cdot \hat{C}_{21}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} & \frac{C_{22}(\omega) \cdot \hat{C}_{11}(\omega) - C_{21}(\omega) \cdot \hat{C}_{12}(\omega)}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} \end{bmatrix} \end{aligned} \quad (4-2)$$

As can be seen in equation (4-2) the effectiveness of the cross-talk cancellation in this case is determined by the numerator of the off-diagonal terms of  $\mathbf{X}(\omega)$ . But this expression is equal to that obtained by the mixed-phase inverse of equation (2-14). Hence the cross-talk cancellation achieved by this minimum-phase approximation to the inverse should be identical to that obtained with the equivalent mixed-phase inverse models. The same holds for the magnitude of the ipsilateral equalisation. This is solely determined by the ratio of the magnitudes of the numerator and denominator of the diagonal terms of  $\mathbf{X}_{eq}(\omega)$  and is not influenced by the minimum phase approximation as is seen in the derivation of equation (4-3). Hence the only element of error introduced in the inversion due to the approximation by  $\mathbf{H}_{min}$  is the amount of

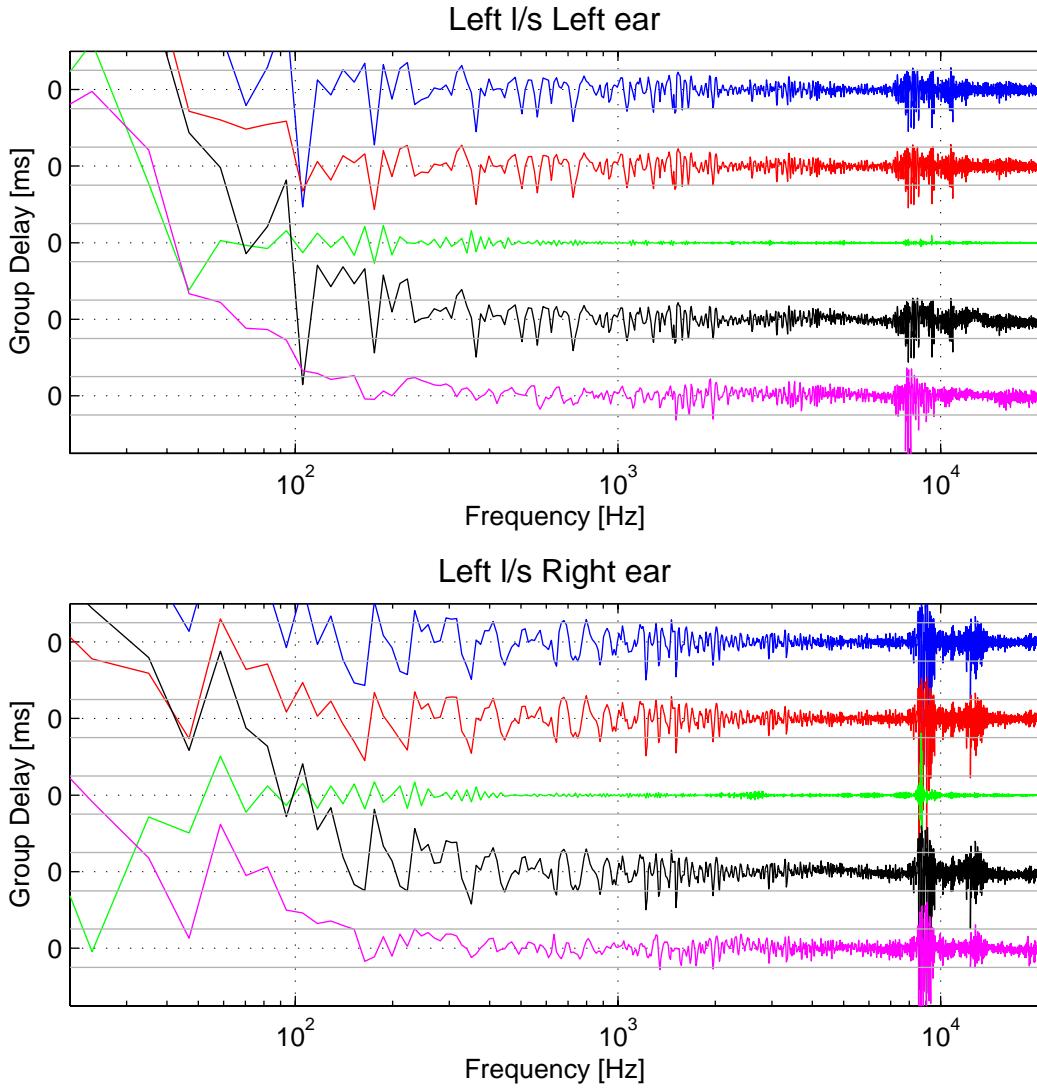
<sup>56</sup> In the following equations with  $\hat{C}_{ij}$  we denote the chosen model of the  $(i,j)$  element of the plant matrix used for the determination of the inverse (be it  $\mathbf{C}_{190}$  or  $\mathbf{C}_{1740}$ ). Similarly, with  $\det[\hat{\mathbf{C}}]$  we denote the determinant of the chosen model of the plant matrix  $\det[\hat{\mathbf{C}}] = \hat{C}_{11} \cdot \hat{C}_{22} - \hat{C}_{12} \cdot \hat{C}_{21}$ .

excess phase of the numerators of the diagonal terms that is left uncorrected by the minimum phase approximation to the denominators.

$$\begin{aligned}
 |\mathbf{X}_{ii}(\omega)| &= \frac{|C_{ii}(\omega) \cdot \hat{C}_{jj}(\omega) - C_{ij}(\omega) \cdot \hat{C}_{ji}(\omega)|}{\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}} \\
 &= \frac{|C_{ii}(\omega) \cdot \hat{C}_{jj}(\omega) - C_{ij}(\omega) \cdot \hat{C}_{ji}(\omega)|}{\left|\text{minphase}\{\det[\hat{\mathbf{C}}(\omega)]\}\right|} \\
 &= \frac{|C_{ii}(\omega) \cdot \hat{C}_{jj}(\omega) - C_{ij}(\omega) \cdot \hat{C}_{ji}(\omega)|}{\left|\det[\hat{\mathbf{C}}(\omega)]\right|} \\
 &= \frac{|C_{ii}(\omega) \cdot \hat{C}_{jj}(\omega) - C_{ij}(\omega) \cdot \hat{C}_{ji}(\omega)|}{\left|\hat{C}_{ii}(\omega) \cdot \hat{C}_{jj}(\omega) - \hat{C}_{ij}(\omega) \cdot \hat{C}_{ji}(\omega)\right|}
 \end{aligned} \tag{4-3}$$

In figure 4-26 the group delay of the equalised ipsilateral paths for the different realisations of the inverse is compared with the group delay of the measured plant responses  $C_{11}$  and  $C_{22}$  (magenta coloured line). These are computed using the responses that contain the full decay (first 3000 samples of the measured impulse responses) and represent the amount of phase distortion introduced by the uncorrected plant's ipsilateral paths, i.e. by the loudspeaker and the listeners HRTF. In all five cases, the initial frequency-independent delay (due to air propagation, modelling delay etc.) has been removed.

As can be seen in the figure the uncorrected plant's ipsilateral responses themselves introduce only a minimum amount of phase distortion. We can take the results reviewed by Preis (1980) as a rough guide. According to those results, for natural sound signals reproduced over loudspeakers in a reverberant environment, deviations of up to approximately  $\pm 2\text{ms}$  in frequencies above 1kHz should be imperceptible. This tolerance further increases at lower frequencies. Hence a signal reproduced in normal conditions over one of the loudspeakers used here would reach the listener's ears with imperceptible phase distortion (barring maybe an increase in the group delay deviation at 8-9kHz). This of course does not account for the effect of the cross-talk paths. The question is what happens when an inverse filtering network is introduced aiming at the suppression of the cross-talk paths and the equalisation of the ipsilateral paths.



**Figure 4-26:** Group delay of the ipsilateral paths for various inverse design parameters: Non-regularised  $\mathbf{H}_{190}$  (blue line),  $\mathbf{H}_{190}$  with  $\beta=10^{-4}$  (red line),  $\mathbf{H}_{1740}$  with  $\beta=10^{-4}$  (green line), causal and stable minimum phase reconstruction of the inverse (black line) and uncorrected plant ipsilateral responses (magenta line). In all cases the initial constant delay is removed. The grey-lined limits correspond to a  $\pm 2$ ms deviation from the ideal frequency-independent group delay.

In principle (Lipshitz 1982), the introduction of an inverse filtering correcting stage in a mixed-phase reproduction chain can result in a degradation of the system's phase response. As is again seen in figure 4-26, whatever cross-talk cancellation and ipsilateral equalisation is achieved with the non-regularised or regularised versions of the  $\mathbf{H}_{190}$  inverse (blue and red lines) does indeed introduce a slight extra phase distortion compared to the (single-channel) uncorrected plant's phase but still remaining below the perceptibility limit. Most importantly however, the same amount

of phase distortion is achieved by  $\mathbf{H}_{min}$ , at least above 200Hz. Recall that the only difference in inverting with  $\mathbf{H}_{min}$  compared to inverting with the non-regularised version of  $\mathbf{H}_{190}$  is that  $\mathbf{H}_{min}$  does not correct the excess phase part (i.e. additive to the minimum phase part) of the denominator  $\det[\mathbf{C}(\omega)]$  (equations (4-1) – (4-3)). It becomes then apparent that the excess phase of  $\det[\mathbf{C}(\omega)]$  is due to the late decay of the plant responses. The error in the determination of the inverse when this late decay is not taken into account (as is the case with the  $\mathbf{H}_{190}$  inverse models) is already much more significant than the phase discrepancy between  $\mathbf{H}_{190}$  and  $\mathbf{H}_{min}$ . This is also in agreement with the anechoic responses' weakly mixed-phase character demonstrated and discussed in §5.2.

Finally, as is seen in the green line results of figure 4-26, the inclusion of the late decay of the plant's responses in the determination of the inverse (by means of  $\mathbf{H}_{1740}$ ) can further reduce the group delay deviation. Even though this corresponds to no perceptual improvement in terms of phase detection, it is by-product of a perceptually significant improvement in terms of amplitude ipsilateral equalisation and cross-talk cancellation (see figure 4-6). Note that such a result is only obtainable with the use of regularisation (see the relevant discussion in §4.2). In that case, the denominator response in  $\mathbf{H}$  is replaced by the expression  $\det[\mathbf{C}^H(\omega)\mathbf{C}(\omega)+\beta\mathbf{I}]$  which, as was explained in §§2.3.1–2.3.3, will always be of strongly mixed-phase character. A minimum phase approximation should thus not be expected to be as easy to implement.

## 4.6 Conclusions

An extensive evaluation was presented in this chapter of the inverse filtering design described in §2.3. The evaluation was based on a set of objective measurements using the KEMAR manikin. The first set of results in the presentation established and quantified the optimal inversion accuracy that can be realistically achieved with the inverse design under consideration. It was thus demonstrated that under ideal conditions of implementation highly accurate inversion results are possible. In such a case it is possible to obtain ipsilateral equalisation confined to  $\pm 1\text{dB}$  in the region above 200Hz and cross-talk cancellation of 15-30dB in that region. It was explained,

however, that these results are practically restricted to the case where the plant can be measured *in situ*. Furthermore it was shown that this degree of inversion accuracy is only possible if a small degree of sub-optimality is allowed in advance by the introduction of a small amount of regularisation and, additionally, if the inverse matrix is implemented by FIR filters comprising several thousands of coefficients. In that case, the influence of regularisation is restricted to the very low frequency range which is anyway uncontrollable in designs of this type.

As a first element of error introduced in the inversion, the case was considered where only the strictly anechoic part of the otherwise exact model of the plant is available. This corresponds to the case where the plant is measured for the individual listener and reproduction system in a given (possibly anechoic) space but the reproduction takes place in a different anechoic space. It was demonstrated that even this marginal amount of disturbance introduces a perceptually significant amount of error. The fundamental perceptual indices relevant to this element of error were discussed and it was explained that even though the indices pertaining to the time-domain description of the reproduced signals should be of minor perceptual weight, those related to the frequency-domain description are indeed significant.

An issue of central importance throughout the presentation of this thesis was highlighted, namely the loss of dynamic range incurred by the application of the audio control. Consequently, the use of regularisation was advocated as a means to address this issue. The role of regularisation was examined in detail and it was shown that in certain cases its use can considerably enhance the overall dynamic range of the inversion without significantly, if at all, degrading its accuracy. The use of a frequency-varying regularisation penalty was also examined. In cases where the very low frequency content is to be reproduced by a separate transducer circuit this frequency-varying regularisation can achieve a further increase of the dynamic range by acting as linear-phase frequency separation stage. This however entails an increased requirement of the inverse filters' length. We do not know if this is indeed advantageous compared to the use of a separate frequency-selective network acting prior to the inversion stage.

On completion of the investigation of the aforementioned cases, the issue of the required length of the inverse filters was revisited and it was demonstrated that when

such parameters as the above are taken into consideration, the required length of the inverse reduces to no more than one or two thousand coefficients. Furthermore, it was shown (figure 4-19) that the element of error that is directly related to the use of the strictly anechoic part of the plant for the determination of the inverse, physically masks the error due to the use of a much shorter realisation of the inverse than would be otherwise expected.

Subsequently, the issue of the ineffectiveness of the design under investigation in controlling the frequency region below 200Hz was discussed. It was explained that this is the combined effect of more than one reasons, namely of the inherent acoustic properties of the reproduction geometry and of the determination of the inverse on the base of the strictly anechoic part of the plant which in turn is modelled using HRTFs. It was thus argued that the control of this region is probably impossible with the inverse design under discussion.

The most typical implementation case, in which HRTFs are used for the determination of the inverse, was also examined. When this is the case the responses of the reproduction loudspeakers are not corrected. Again, a significant amount of error was seen to be inherent in this type of implementation which in practice further relaxes the requirement of accuracy in the inversion. Thus higher regularisation can be applied and, consequently, shorter inverse filters can be used.

Finally, the issue of the phase response of the inverted plant was also considered. It was shown that, barring the case where an *in situ* measurement of the plant is available, a minimum phase approximation to the recursive part of the inverse can give practically equivalent results as the exact inverse. It is not certain if this can be used to obtain computational savings in the inversion but it could certainly be used to eliminate the need for the introduction of modelling delay.

## ***5 Objective evaluation of inverse filtering –***

### ***Inversion efficacy under non-ideal conditions***

#### ***5.1 Introduction***

In the previous chapter we established the extent of the optimum inversion under ideal conditions and presented in detail the role of the design parameters of the inverse in achieving this optimum. In the following sections of this chapter we return to the issues discussed in §3.1 and examine the effect that the deviation from this ideal inversion conditions has on the effectiveness of the inversion. The exposition does not intend to exhaustively cover this issue but rather to put the ideal conditions results of the previous sections in the context of a more realistic implementation situation and to give a measure of the validity of the findings of the previous sections in such realistic implementation conditions.

The effect of the three basic elements of error identified in §3.2 is studied in the following three sections. These elements of error are the presence of reflections in the plant, the misalignment of the position of the listener and the determination of the inverse on the basis of a plant measurement conducted on a different listener (typically a dummy-head). They are studied in sections §5.2, §5.3 and §5.4 respectively.

#### ***5.2 Presence of reflections in the plant***

In realistic applications the audio reproduction system is situated in a non-anechoic space, e.g. a living room or a study with a personal computer on top of a desk. In such cases, the impulse responses comprising the plant would represent the linear distortion in the sound field due to the head and the body of the listener followed by one or more

early reflections caused by objects and surfaces close to the direct sound propagation paths and then by a late reverberation decay due to the overall acoustics of the room. Naturally, the signals delivered at the listener's ears after being processed by an inverse designed on the basis of an anechoic measurement of the plant will differ from the desired binaural signals. This should be expected to degrade the performance of the system.

Studies that evaluate the degradation of the perceived quality of the reproduced audio material when the reproduction takes place in a reverberant environment include those of Damaske (1971), Gardner (1998), Köring and Schmitz (1993), Sakamoto et al. (1982) and Takeuchi (2001). They are all based on subjective experiments with the exception of (Takeuchi 2001), in which the results of the subjective experiment follow a simulation-based objective analysis of the effect of a single reflection by an infinitely long surface on the side of the reproduction geometry. These studies cover a diverse range of inverse filtering designs<sup>57</sup> and different instances of reverberant environments<sup>58</sup> but their common finding is that virtual imaging in the horizontal plane remains to a great extent unaffected by the presence of discrete reflections and reverberation, with an increase being observed in the front-back confusion rate. On the other hand, virtual imaging in the median plane and in elevated directions is found to suffer in such conditions.

These results are in accordance with the author's experience from extensive informal listening tests using inverse realisations of the type discussed in this thesis. These tests consistently verify the fact that in reverberant conditions, imaging in the front half of the horizontal plane does indeed succeed in terms of the directional characteristics of the presentation but with increased localisation blur, erratic behaviour of the distance perception and audible colouration of the presented material. Conversely, presentations that succeed in creating imaging in the median plane, in the rear half of the horizontal plane and in elevated positions inside the anechoic chamber, fail in reverberant environments. The degree of failure seems to be graver when the plant

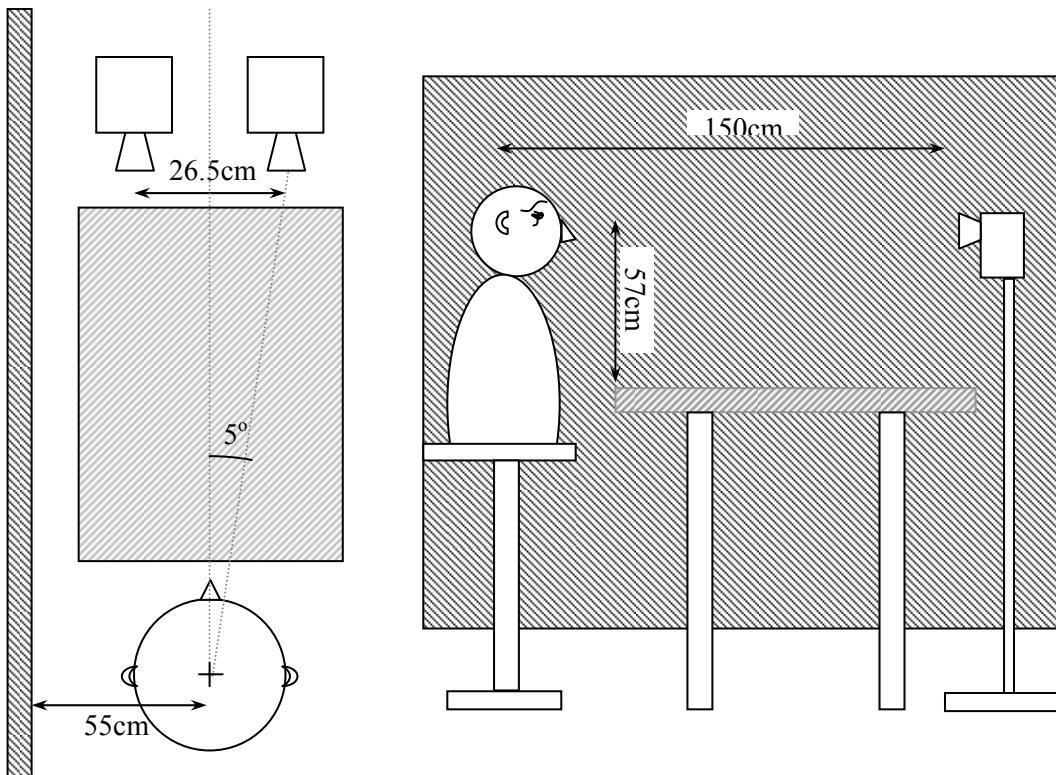
---

<sup>57</sup> From those listed above, the only study which uses the same inverse filter design with regularisation as the one used in this thesis is that of Takeuchi (2001).

<sup>58</sup> In (Damaske 1971), (Gardner 1998) and (Sakamoto et al. 1982) the subjective experiments were conducted in rooms of various reverberation times RT60. The same holds for (Köring and Schmitz 1993) in which information is also given about the influence of the energy of the reflections relative to the energy of the anechoic part of the plant. In (Takeuchi 2001) the role of a single reflection is evaluated by use of the mirror image model implemented inside the anechoic chamber.

contains strong distinct reflections close to the direct transmission paths, something that is in general agreement with the results of Köring and Schmitz (1993).

The aim of the discussion and the results presented in this section is to give an objective measure of the influence of reflections in the inversion accuracy. More specifically, we aim to determine the point after which the advantages of an accurately implemented inverse of the anechoic model of the plant  $\mathbf{c}(n)$  are lost due to the presence of the reflections. In this view, we seek to determine the degree to which the design parameter requirements can be relaxed without any loss in the inversion quality. Following that, we examine the case where an accurate model of the plant with the reflections, denoted here as  $\hat{\mathbf{c}}(n)$ , is known and verify the conditions under which such a plant can be effectively inverted<sup>59</sup>.



**Figure 5-1:** Plant with reflections, top and side view.

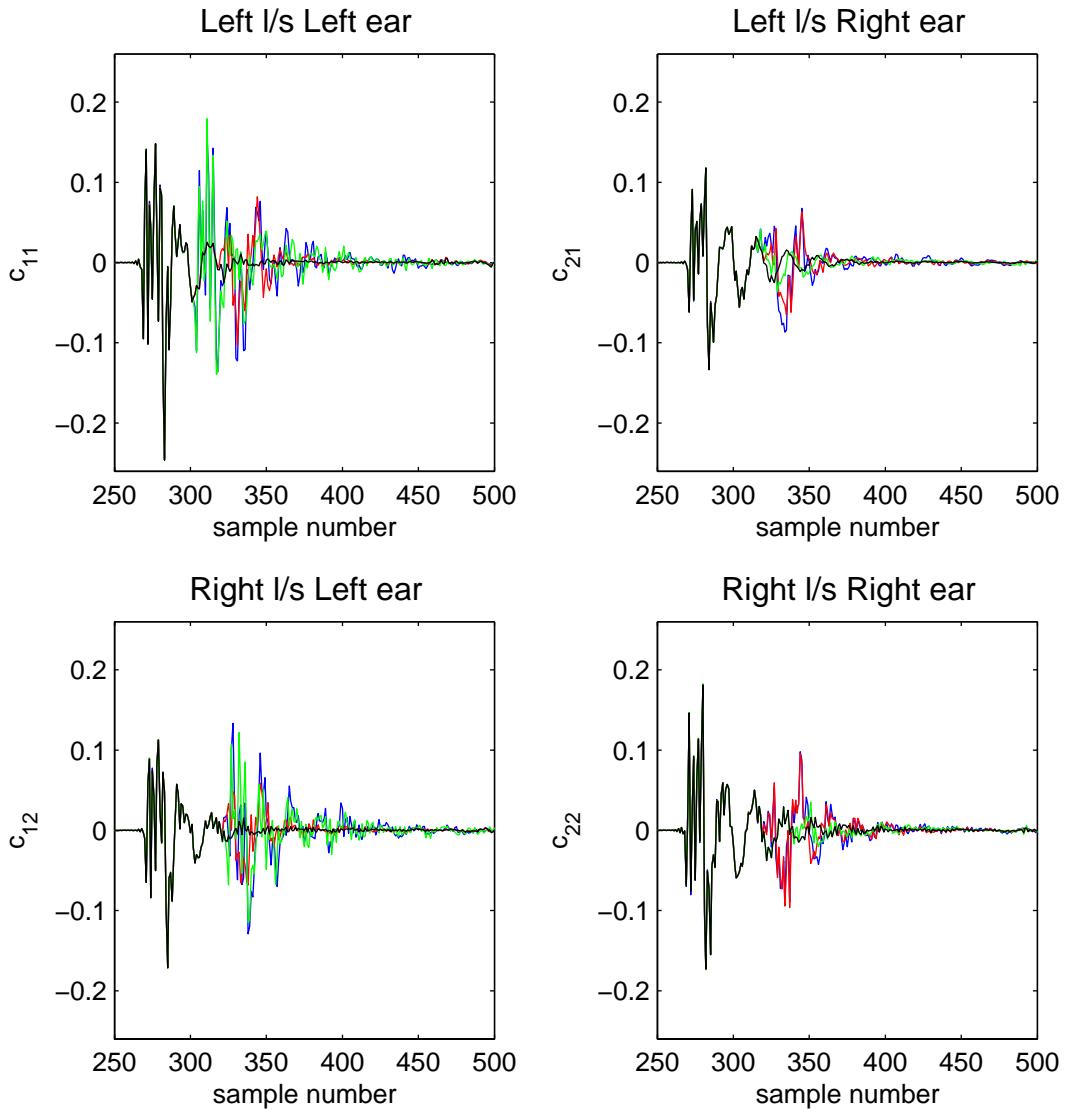
<sup>59</sup> The model of the plant used for the determination of the inverse throughout this section (§5.2) contains the measured impulse responses  $c_{ij}(n)$  for indices  $260 \leq n \leq 500$ . This length is the same both for the case where the anechoic model of the plant  $\mathbf{c}(n)$  is used for the determination of the inverse and for the case where the model  $\hat{\mathbf{c}}(n)$  containing the reflections is used for the determination of the inverse.

The experimental arrangement used to obtain the results of this section is the same as that described in §3.3, the only difference being that reflective surfaces are introduced in the plant. Three cases are implemented and are illustrated in figure 5-1. In the first case, a horizontal reflective surface is placed between the loudspeakers and KEMAR, corresponding to the case where the loudspeakers are on top of a table in front of the listener, eg on the sides of a pc monitor that is situated on a desk. In the second case, a vertical hard reflective surface is placed on the side of KEMAR correspond to the situation where the reproduction system is positioned next to a wall. The third case is implemented with both these surfaces in place<sup>60</sup>.

In figure 5-2 we plot the measured impulse responses corresponding to the four different cases (anechoic, reflection due to the wall, reflection due to the table and reflections due to both the wall and the table). The initial parts of the measured impulse responses are identical and are hence not visible in the plot as they overlay. The first reflection to appear is that due to the wall and it can be seen to be stronger in the left ear responses. A simple geometric calculation shows that the path-length difference between the direct path and the wall-reflection path is approximately 25cm for the left ipsilateral response and approximately 72cm for the right ipsilateral response. At the sampling rate of 48kHz and with a speed of sound of 340m/s these path-length differences correspond to 36 and 67 samples. The delay for the contralateral paths lies between these two limits being approximately 50 samples for the left loudspeaker to right ear path,  $\hat{c}_{wall,21}(n)$ , and 55 samples for the right loudspeaker to left ear path,  $\hat{c}_{wall,12}(n)$ . These delays can be seen in figure 5-2 where the reflection due to the wall can also be seen to be lower in the right ear responses due to the head shadowing. In the case of the reflection due to the table, due to the symmetry of the geometry and the small 10° loudspeaker span, the path-length difference between the ipsilateral and the contralateral sides is very small and the effect of the head shadowing minimal. The path-length difference for all four responses is approximately 38cm, i.e. 55 samples, as is verified by the plotted results. The case where both the table and the wall reflections are present can also be seen in the figure as a combination of the two previous cases.

---

<sup>60</sup> Where needed we specifically denote the plant model containing the wall, table or wall and table reflection with  $\hat{c}_{wall}(n)$ ,  $\hat{c}_{table}(n)$  and  $\hat{c}_{wall-table}(n)$  respectively.

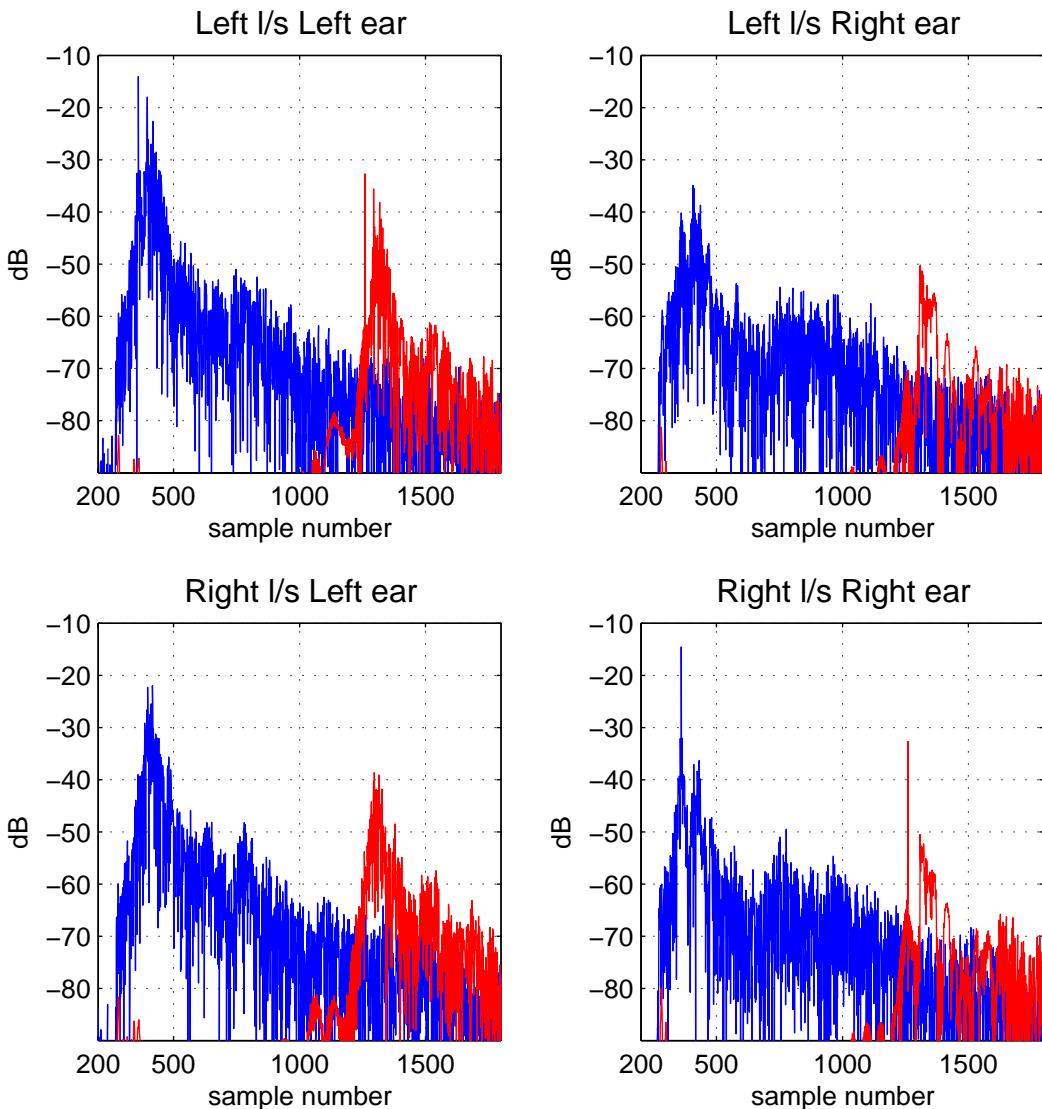


**Figure 5-2:** Measured impulse responses for the plants depicted in figure 5-1. Anechoic response (black line), response with the table in place (red line), response with the vertical reflective surface in place (green line), response with the table and the vertical reflective surface in place (blue line).

As was said above, the first set of results in this section aims at demonstrating the effect of the presence of the reflections in the effectiveness of the inversion, when the inverse is determined on the basis of the anechoic measurement of the plant  $\mathbf{c}(n)$ . In figure 5-3 we plot the inversion results<sup>61</sup> obtained using such an inverse and compare two cases: In the first (red line) the anechoic plant is inverted with a 2000-tap inverse and with the regularisation parameter set to the value  $\beta=10^{-4}$  whereas in the second

<sup>61</sup> All the results presented in the section are obtained using the simulation procedure described in §3.5

(blue line) the anechoic plant is inverted with a 200-tap inverse and with the regularisation parameter set to the value  $\beta=10^{-2}$ .



**Figure 5-3:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $\mathbf{c}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 taps (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 200 taps (blue line). In both cases the measured plant matrix describing the actual plant is the one containing the wall reflection.

The presence of the reflection in the plant but not in the model used for the computation of the inverse can be seen to result in a rise of the error in the impulse responses. This rise in the error appears with a delay of about 35 samples relative to the pulse for the left side direct path. The delay is higher for the other paths with a highest value of 55 samples for the right-side direct path. It can also be seen to be

higher in level in the left-ear responses, for both the direct and the cross talk paths, compared to the right-ear responses. These effects are readily interpreted by inspection of equation (5-1) that describes the convolution of the inverse matrix  $\mathbf{h}(n)$  determined on the basis of the anechoic plant matrix  $\mathbf{c}(n)$  with the matrix  $\hat{\mathbf{c}}(n)$  that represents the plant with the reflection:

$$\begin{aligned}\mathbf{x}_{eq}(n) &= \hat{\mathbf{c}}(n) * \mathbf{h}(n) \\ &= \begin{bmatrix} \hat{c}_{11}(n) * h_{11}(n) + \hat{c}_{12}(n) * h_{21}(n) & \hat{c}_{11}(n) * h_{12}(n) + \hat{c}_{12}(n) * h_{22}(n) \\ \hat{c}_{21}(n) * h_{11}(n) + \hat{c}_{22}(n) * h_{21}(n) & \hat{c}_{21}(n) * h_{12}(n) + \hat{c}_{22}(n) * h_{22}(n) \end{bmatrix} \quad (5-1)\end{aligned}$$

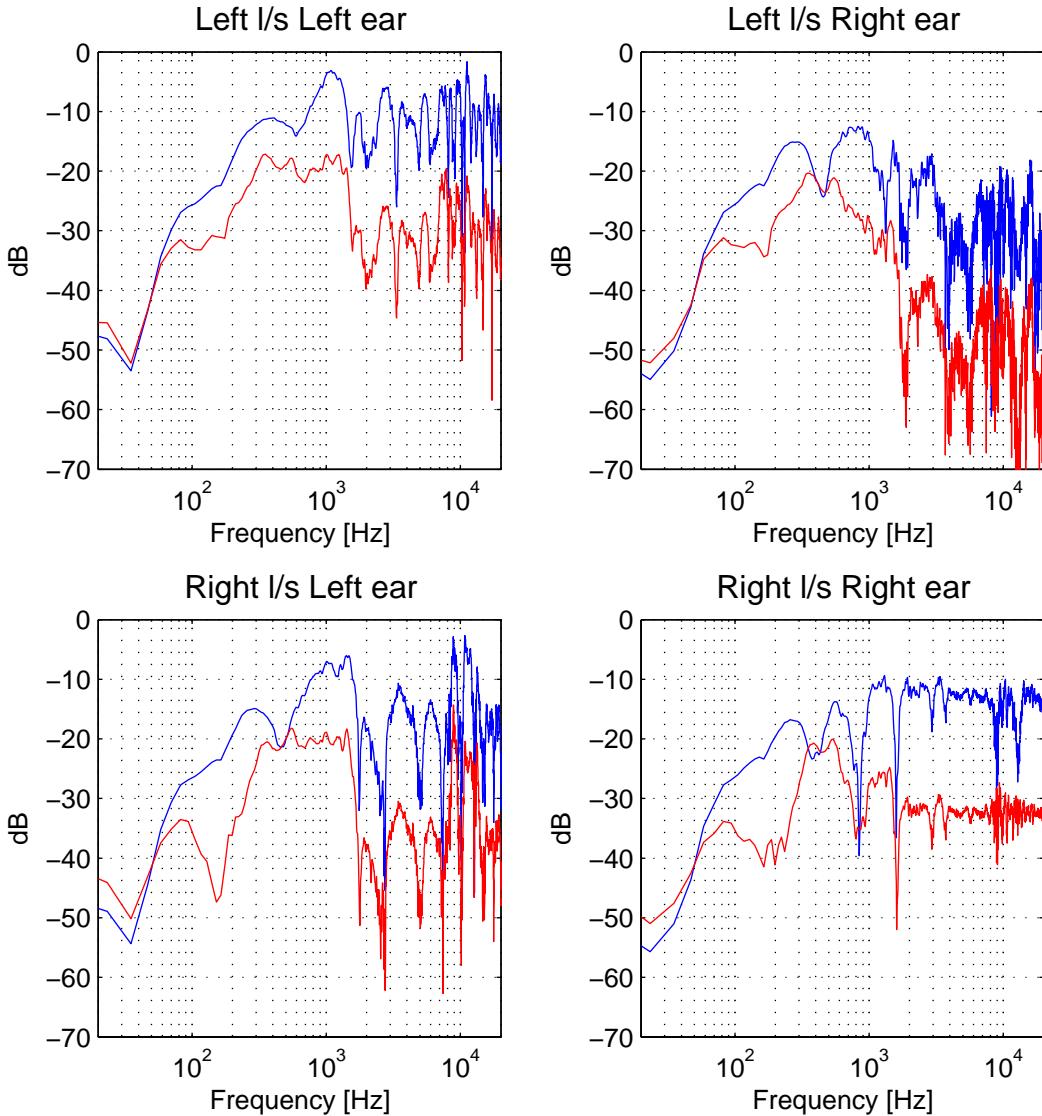
As was discussed above (see also figure 5-2), the disturbance from the wall reflection in the plant's impulse responses is more significant in the left-ear responses  $\hat{c}_{11}(n)$  and  $\hat{c}_{12}(n)$  whereas the deviation of the right-ear responses  $\hat{c}_{21}(n)$  and  $\hat{c}_{22}(n)$  from the anechoic responses  $c_{21}(n)$  and  $c_{22}(n)$  is much lower. Consequently the (1,1) and (1,2) elements of the matrix  $\mathbf{x}_{eq}(n)$  that correspond to the left-ear equalised responses will deviate significantly from the desired form (i.e. delta-spike and zero respectively). Conversely the disturbance in the right-ear (2,1) and (2,2) elements of  $\mathbf{x}_{eq}(n)$  will be lower. Furthermore, we see from the form of the elements of the matrix  $\mathbf{x}_{eq}(n)$  in equation (5-1) that the point at which the deviation from the desired form starts in the equalised responses, is determined by the point in which the terms  $\hat{c}_{ij}(n)$  deviate from the anechoic responses  $c_{ij}(n)$ . More specifically, the point where the rise in the inversion error should appear in the (1,1) and (1,2) elements  $\mathbf{x}_{eq}(n)$  will be the point at which  $\hat{c}_{11}(n)$  deviates from  $c_{11}(n)$ , since this point comes earlier than the point at which  $\hat{c}_{12}(n)$  deviates from  $c_{12}(n)$ . Similarly, the inversion error should appear in the (2,1) and (2,2) elements of  $\mathbf{x}_{eq}(n)$  at the point where  $\hat{c}_{21}(n)$  deviates from  $c_{21}(n)$ . This is in accordance with the 35 and 55 samples of delay between the pulse and the rise in the inversion error noted above.

Despite the significant increase in the inversion error, both in terms of the ipsilateral equalisation and in terms of the cross-talk cancellation, the characteristics of the time-domain results of figure 5-3 do not bear evidence for a major degradation of the perceptual impression of the listener. The rise in the error in the direct paths is physically bound to appear after the equalisation pulse. For the level and delay interval at which it appears in figure 5-3 it will be masked and integrated with the equalisation pulse as single auditory event occurring at the time of the pulse. We note

that the level of the error relative to the main pulse shown in the results of figure 5-3 is the highest we should expect to occur due to a side reflection since in this case the hard reflective surface was placed only 55cm from the listener's head. In realistic situations a side-wall would probably be further away. In such a case the error due to the reflection would be further delayed but also further attenuated and hence equally or more susceptible to temporal masking. The level of cross-talk cancellation suffers the severest degradation and can be seen to be around 10dB and 20dB for the two contralateral paths at the point of the equalisation pulse (the corresponding level for the anechoic inversion case of figure 4-12 was 30dB). Again, in realistic situations where the reflective surface is further away from the listener the cross-talk cancellation level would increase.

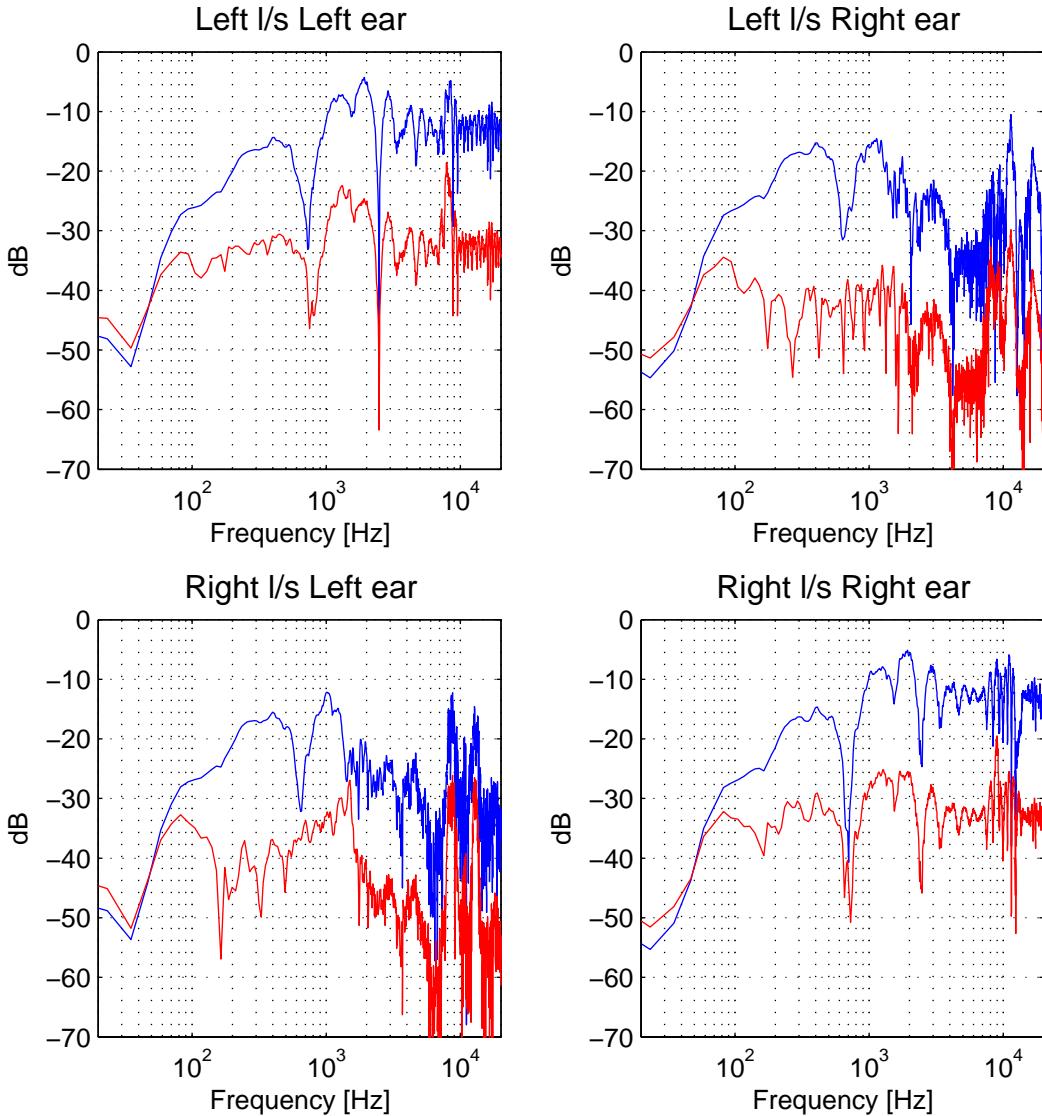
As regards the comparison between the  $\beta=10^{-4}$  and  $\beta=10^{-2}$  cases we see that, in agreement with the results of §4.3, the increase in the value of  $\beta$  results in an increase of the target level of equalisation and thus of the dynamic range of the inversion. Unlike the cases discussed in §4.3 however, the smearing of the equalisation pulse due to the increase in regularisation is in this case of minimal significance compared to the error components due to the presence of the reflection. Additionally, the 1000 taps corresponding the non-causal part of the inverse filters are not enough to completely suppress the truncation end-effect that can be seen to be approximately 50dB below the equalisation level for the  $\beta=10^{-4}$  case. Conversely, in the  $\beta=10^{-2}$  case 100 taps are enough to totally suppress the non-causal end-effect.

As was shown in §4.3, the effect of the introduction of regularisation is more visible in the frequency domain where the smearing of the equalisation pulse translates to gradual deterioration of the ipsilateral equalisation quality and rise of the cross-talk transmission in certain frequency regions. The comparison of the  $\beta=10^{-4}$  (2000-tap) and  $\beta=10^{-2}$  (200-tap) implementations in the frequency domain is shown in figures 5-4, 5-5 and 5-6 for the cases of the reflection due to the wall, the table and both the table and the wall respectively.



**Figure 5-4:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $\mathbf{c}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 taps (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 200 taps (blue line). In both cases the measured plant matrix describing the actual plant is the one containing the wall reflection.

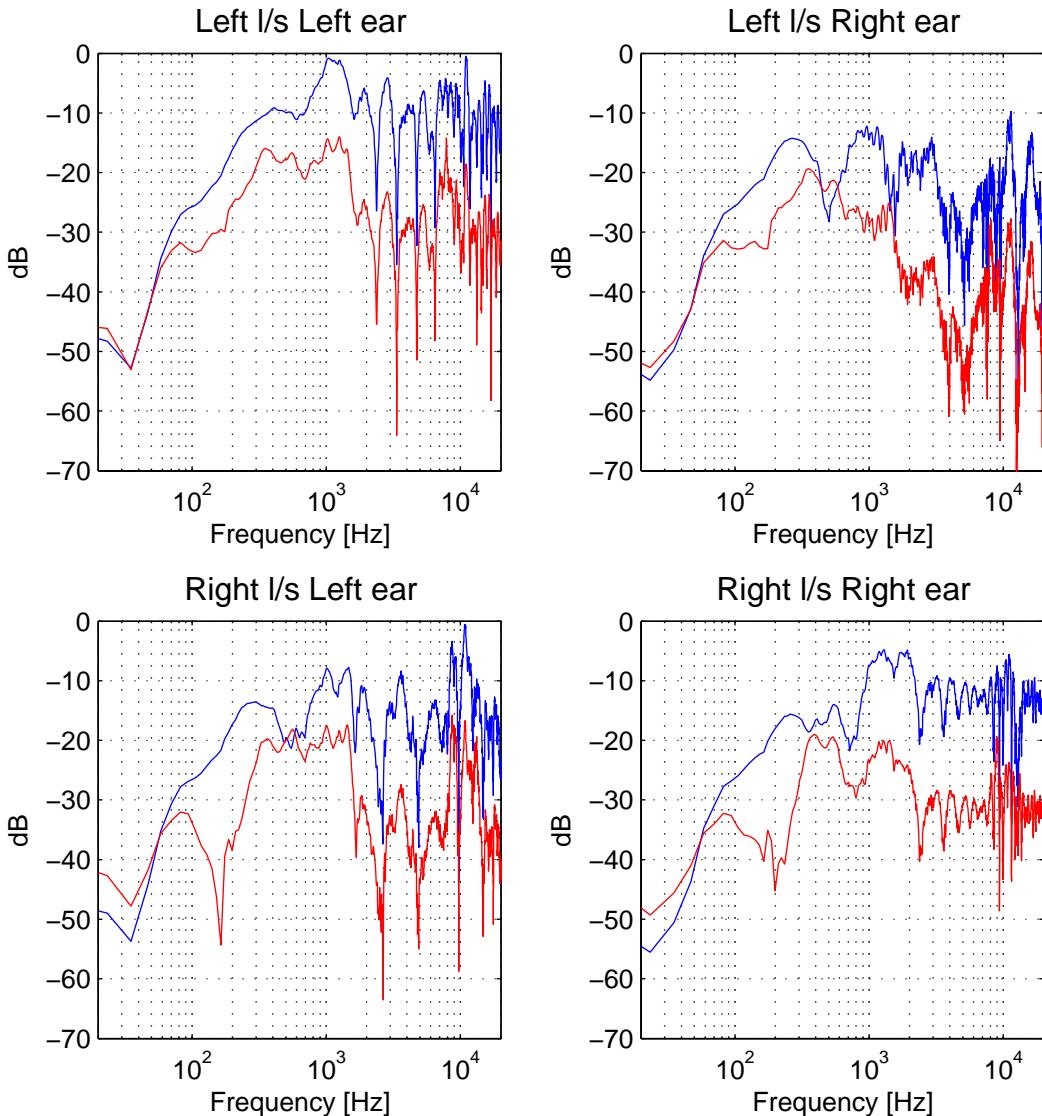
As should be expected, the common feature in all three cases is that the increased regularisation penalty of  $\beta=10^{-2}$  results in a gradual loss of control for the low frequency range below 200Hz compared to the  $\beta=10^{-4}$  case. This effect however should be considered to be of minimal importance compared to the general degradation of the control effectiveness especially in the case of the table and wall reflection of figure 5-6 in which the overall inversion results are very poor.



**Figure 5-5:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $\mathbf{c}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 taps (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 200 taps (blue line). In both cases the measured plant matrix describing the actual plant is the one containing the table reflection.

In all cases the ipsilateral equalisation can be seen to suffer and with the exception of the right ipsilateral path in the table reflection case, the degree of degradation is the same for both the  $\beta=10^{-2}$  and the  $\beta=10^{-4}$  cases. Similarly, the level of cross-talk cancellation is significantly reduced in comparison with the anechoic results of chapter 4 and confined to the frequency region above 1kHz for the table reflection case and in the interval from 2kHz to 7kHz for the wall and table-wall reflection cases. With the exception of a small interval around 150Hz for the wall and table-wall

reflection cases and of the region below 1kHz for the table reflection case, the level of cross-talk cancellation is again identical for the  $\beta=10^{-2}$  and the  $\beta=10^{-4}$  cases.



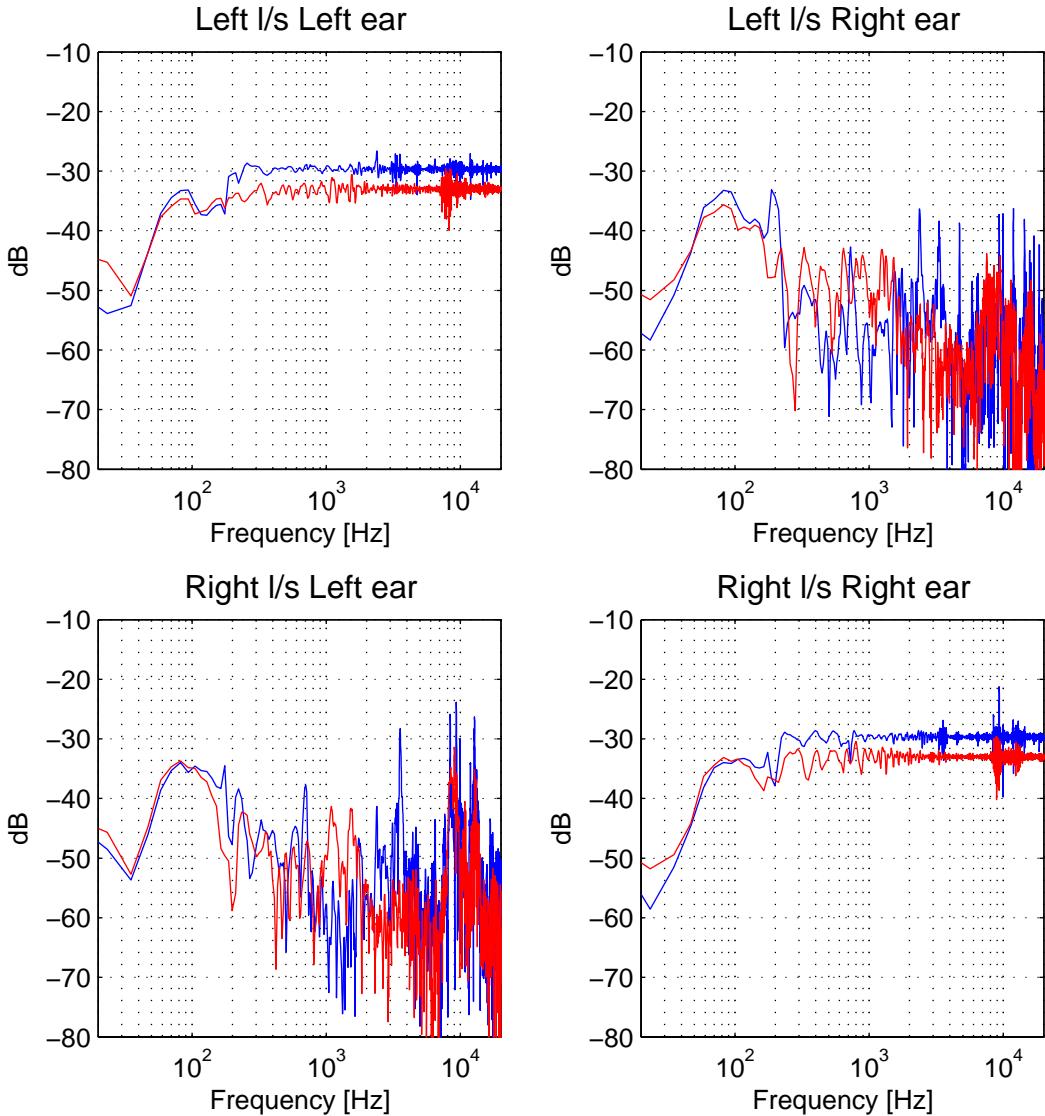
**Figure 5-6:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $\mathbf{c}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2000 taps (red line). The regularisation is set to  $\beta=10^{-2}$  and the inverse length to 200 taps (blue line). In both cases the measured plant matrix describing the actual plant is the one containing the reflections due to both the table and wall.

Overall then, the conclusions drawn from the results presented in figures 5-3 – 5-6 are as follows. The time-domain cues are delivered to a reasonable extent despite the presence of the reflections. Taking into account that these cues are the dominant source of information for localisation in the front half of the horizontal plane, the results of figure 5-3 offer a quantitative explanation of the fact that localisation in the

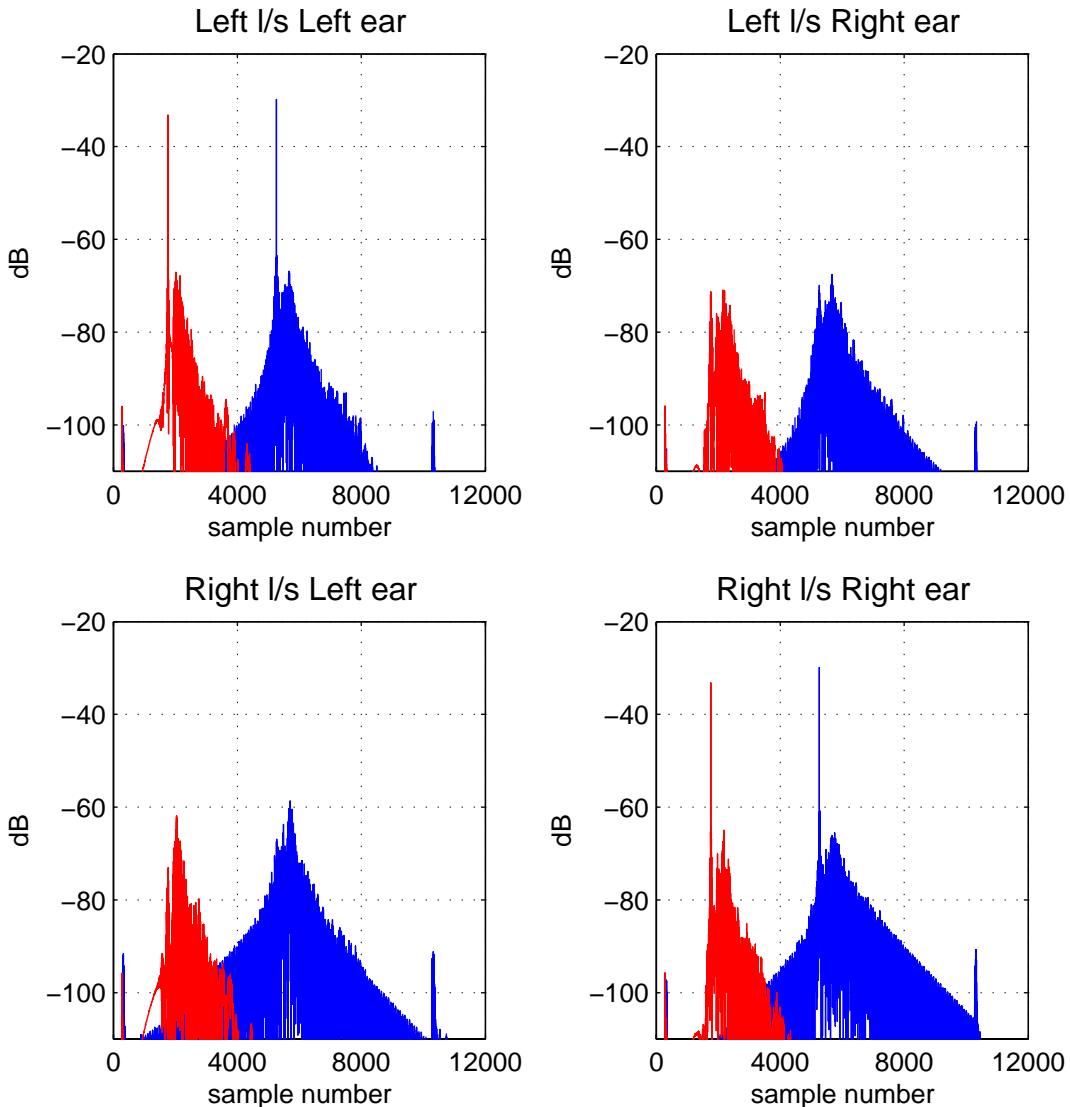
front horizontal plane is to a large extent unaffected by the presence of reflections. The level of cross-talk cancellation is affected to a greater extent. This amounts to a reduced ability to imitate the interaural level differences due to the head shadowing at directions well to the side of the listener. It should thus result in a shift of virtual images placed at the sides of the listener towards the median plane. Similarly, the spectral cues are severely distorted as can be seen in the ipsilateral equalisation results with clearly visible comb effects appearing in all three cases. It is worth taking into account the findings of Blauert (1996) regarding the “directional bands” and “boosted bands” and their role in front-back discrimination and localisation in the median plane. We expect this failure of the ipsilateral equalisation to be responsible for the aforementioned degradation of the virtual imaging in the median plane and the front-back reversals in the presence of reflections.

As far as the parameters of the design of the inverse are concerned, the results show that the degradation due to the presence of the reflections is such that the use of a more accurate inverse does not translate into any significant improvement of the inversion results. The heavier penalisation of the control effort effected by setting  $\beta$  equal to  $10^2$  yields the same inversion accuracy with the  $\beta=10^{-4}$  case, offering at the same time a dynamic range that is approximately 20dB higher and requiring less than 1/10 of the taps needed to adequately suppress the truncation end-effects.

Finally, we examine the case where an accurate model of the plant containing the reflections is available. In figures 5-7 and 5-8 we plot and compare the inversion results obtained for the anechoic plant and the plant with the table and wall reflection. Both plants are inverted with the regularisation parameter set to the value  $\beta=10^{-4}$ . The length of the inverse is set to 2500 taps (1000 causal and 1500 anticausal) for the anechoic case and 10000 taps (6000 causal and 4000 anticausal) for the case with the reflections. As can be seen in figure 5-7, when the plant model with the reflections is used for the determination of the inverse, equalisation and cross-talk cancellation equivalent to that of the anechoic case can be obtained. As is shown in figure 5-8 however, in order to obtain these equivalent inversion results and to suppress the end-effects at the same level as for the anechoic case, we have to use more than four times the number of coefficients for the inverse filters.



**Figure 5-7:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $c(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2500 taps. The measured plant matrix describing the actual plant is the anechoic one (red line). The inverse is determined using the model of the plant containing the reflections  $\hat{c}_{\text{wall-table}}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 taps. The measured plant matrix describing the actual plant is the one containing the reflections due to the table and the wall (blue line).



**Figure 5-8:** Simulated inversion results. The inverse is determined using the anechoic model of the plant  $\mathbf{c}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 2500 taps. The measured plant matrix describing the actual plant is the anechoic one (red line). The inverse is determined using the model of the plant containing the reflections  $\hat{\mathbf{c}}_{\text{wall-table}}(n)$ . The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 taps. The measured plant matrix describing the actual plant is the one containing the reflections due to the table and the wall (blue line).

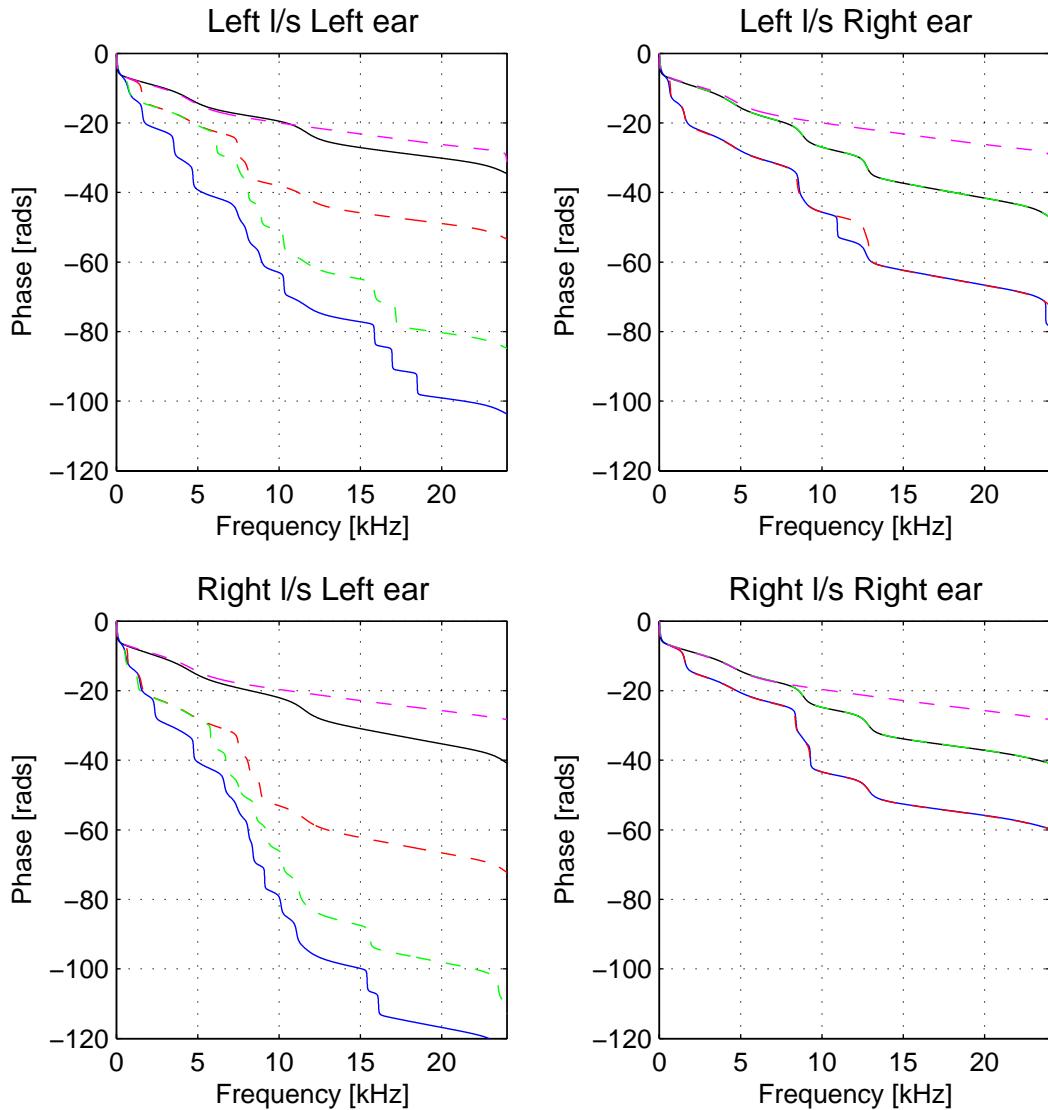
A last point that deserves to be made concerning the inversion of the plant with the reflections, is that the presence of the reflections adds maximum phase zeros in the responses comprising the plant. The relation of reflections and the phase characteristics of room responses was long ago investigated by Neely and Allen (1979) who by use of synthesised responses it showed that the maximum phase zeros start to appear when the reflections increase in strength and move closer to the direct

path between the source and the receiver. The presence of maximum phase zeros in KEMAR HRTFs measured in anechoic conditions was shown by Gardner (1998) to be restricted to frequencies above 6kHz and was attributed to reflections of the pinnae that become adequately energetic only in that frequency region. The non-minimum phase character of HRTFs of human subject was also examined by Møller et al. (1995b) who also found that, in general, human subjects' HRTFs are non-minimum phase. A few examples of the phase of the all-pass component of such responses is presented in that reference in which the maximum phase zeros indeed appear above 6kHz, but no mention is made as to if this holds for all the responses measured by those researchers.

In figure 5-9 we plot the unwrapped phase of the all-pass component of the measured responses of the anechoic plant and the three plants with the reflections. In the same figure we plot the unwrapped phase of the all-pass component of the free field responses of the left and right loudspeakers positioned at the same points used for the control and measured with KEMAR removed from the plant and a microphone placed at the position of the centre of its head. The results of figure 5-9 were obtained by decomposing the measured responses (after windowing out the initial delay part) into a minimum-phase and an all-pass component (Oppenheim and Schafer 1975, ch. 7). The maximum phase zeros appear thus as  $2\pi$  negative jumps in the phase curves of figure 5-9 (Proakis and Manolakis 1992, §5.4.5).

As can be seen in the figure, the maximum phase zeros of the anechoic plant that are not present in the free-field response (and are thus introduced by the measurement apparatus) indeed appear above 6kHz, at approximately 12kHz for the left ear responses and 8kHz and 13kHz for the right ear responses. More maximum phase zeros, spread across the spectrum, start to appear in the responses of the plants containing reflections. The maximum phase part of the right ear responses is unaffected by the presence of the wall as should be expected since the influence of the wall on the right ear responses not very energetic. Thus the maximum phase curve of the right ear responses is identical for the anechoic and wall reflection cases and also for the table and the wall and table reflections cases. On the other hand the presence of the wall introduces a number of maximum phase zeros on the left ear responses that are not present in the anechoic or the table reflection cases. The maximum phase

zeros due to the reflection off the table are roughly equal in number in all four responses.



**Figure 5-9:** Excess phase part of the measured plant responses. Anechoic responses (black line), responses with wall in place (dashed green line), responses with table in place (dashed red line), responses with wall and table in place (blue line), free-field responses of the loudspeakers (dashed magenta line).

### 5.3 Listener position misalignment

In this section we focus our interest on a second source of error in the inversion process, namely the imperfect positioning of the listener. The influence of the listener's position misalignment in the accuracy of the inversion has previously been

investigated by Gardner (1998), Rose (2004)<sup>62</sup> and Takeuchi (2001). The basic finding in these studies is that a displacement of the head by only a few centimetres is enough to severely degrade the accuracy of the inversion and the recreation of the localisation cues. This is shown on the basis of simulated inversion results in which the plant corresponding to the disturbed geometry is approximated by means of interpolation between the closest available measured responses. A positioning error of such magnitude is shown to be enough to significantly degrade the intended localisation properties of the perceived sound impression, something that is validated in those studies on the basis of subjective experiments. For symmetrical reproduction geometries such as the one discussed here, the degradation of the inversion is shown to be worse for misalignment along the left-right direction (i.e. along the interaural axis) and much more tolerable in the front-back direction. In (Takeuchi 2001), the results show an overall better performance for the Stereo Dipole geometry compared to the wider source-span geometries. Earlier studies that investigate the influence of listener position misalignment based solely on subjective experiments include those of Damaske (1971) and Sakamoto et al. (1982) while measured results of the influence of positioning misalignment in a free-field reproduction setup (no head present) are presented by Nelson et al. (1992). In a very recent study, Nelson and Rose (2005) offer a detailed analytical formulation that describes the effect of misalignment on the inversion accuracy which is again based on a free-field model of the reproduction arrangement.

Again, following the general objective of the analysis presented in this chapter, the results of this section aim to quantify the influence of this element of error in comparison to the inversion results under ideal conditions presented in the previous chapter. More specifically we aim to determine the extent of misalignment for which the utilisation of the full-frequency-range inverse under investigation is justified. To this end, the same experimental arrangement was used as that described in §3.3. The arrangement was slightly modified by placing the two loudspeakers on a small board that could be slid in the left-right direction on top of a frame. The frame was covered with foam to minimise reflections. The movement of the loudspeakers was (manually) controlled by means of low gear-ratio threaded shaft. The precision of the loudspeaker positioning was less than 0.5mm.

---

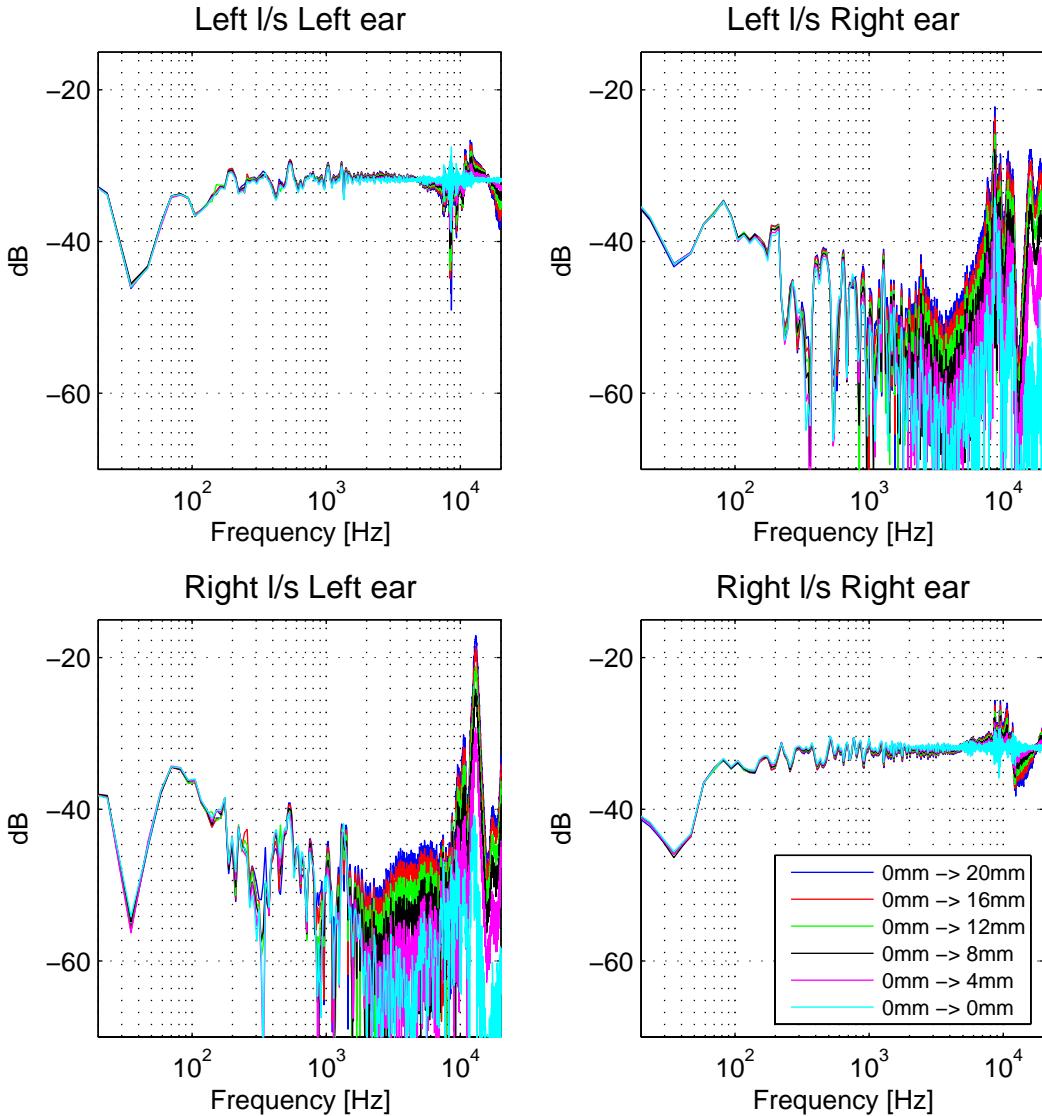
<sup>62</sup> Being published very recently, this reference has not been fully reviewed by the author.

In this way, measurements were made of the plant models corresponding to the same axisymmetric geometry as in chapter 4 as well as of the plant models resulting from the left and right displacements of the loudspeakers by 2mm up to a maximum of 20mm. The same measurement procedure was repeated for increments of 2cm up to a maximum of 10cm<sup>63</sup>. Inverse models of the properly aligned plant models  $C_{190}$  and  $C_{1740}$  were computed and the inversion results were then obtained by simulating the inversion process as described in §3.5, i.e. by convolving the inverse models the 3000-coefficient full-length measured model of the plant for the various misaligned positions. Following the results of the studies mentioned above, according to which the inversion accuracy is mostly influenced by the lateral misalignment along interaural axis, we restrict the investigation to this case of misalignment.

Figure 5-10 shows the degradation of the crosstalk cancellation and ipsilateral equalisation as the loudspeakers are moved from the 0mm to the 20mm position to the right in 4mm increments (i.e. the listener is moved in equal increments to the left). For these results, the inverse filters are computed on the basis of the  $C_{190}$  model of the plant with their length set to 4000 coefficients and the regularisation to the value  $\beta=10^{-4}$ . As seen in the figure, the degradation increases at higher frequencies where the corresponding wavelength becomes comparable with the amount of the path lengths' derangement in the misaligned plants. These results are in general agreement with the related analytical estimation of Gardner (1998). Overall, as can be seen in the figure, both the ipsilateral equalisation and the cross-talk cancellation are still effective for a 2cm displacement but only up to approximately 10kHz. Above that frequency a 2cm displacement renders the inversion practically pointless.

---

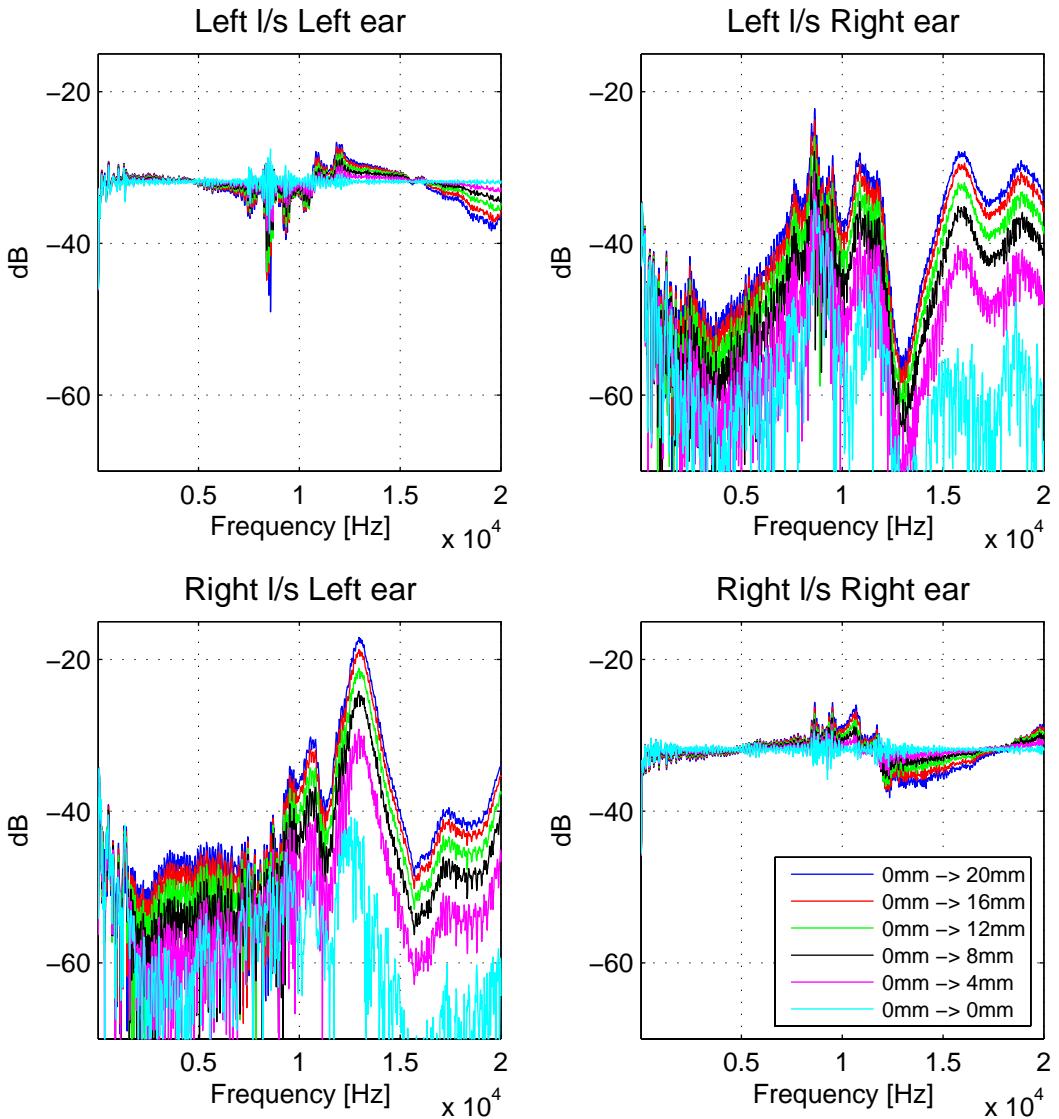
<sup>63</sup> The second series of measurements (0cm to 10cm in 2cm increments) was made with essentially the same setup but in the small anechoic of ISVR. This is a semi-anechoic chamber with rigid floor which can be covered with sound absorbing cones. For the measurements we covered the floor with sound absorbing cones. However the covering could not be perfect and hence a more energetic decay tail was present in the measured responses. This accounts for the slight difference between the zero displacement results of figures 5-17 and 5-18 shown below.



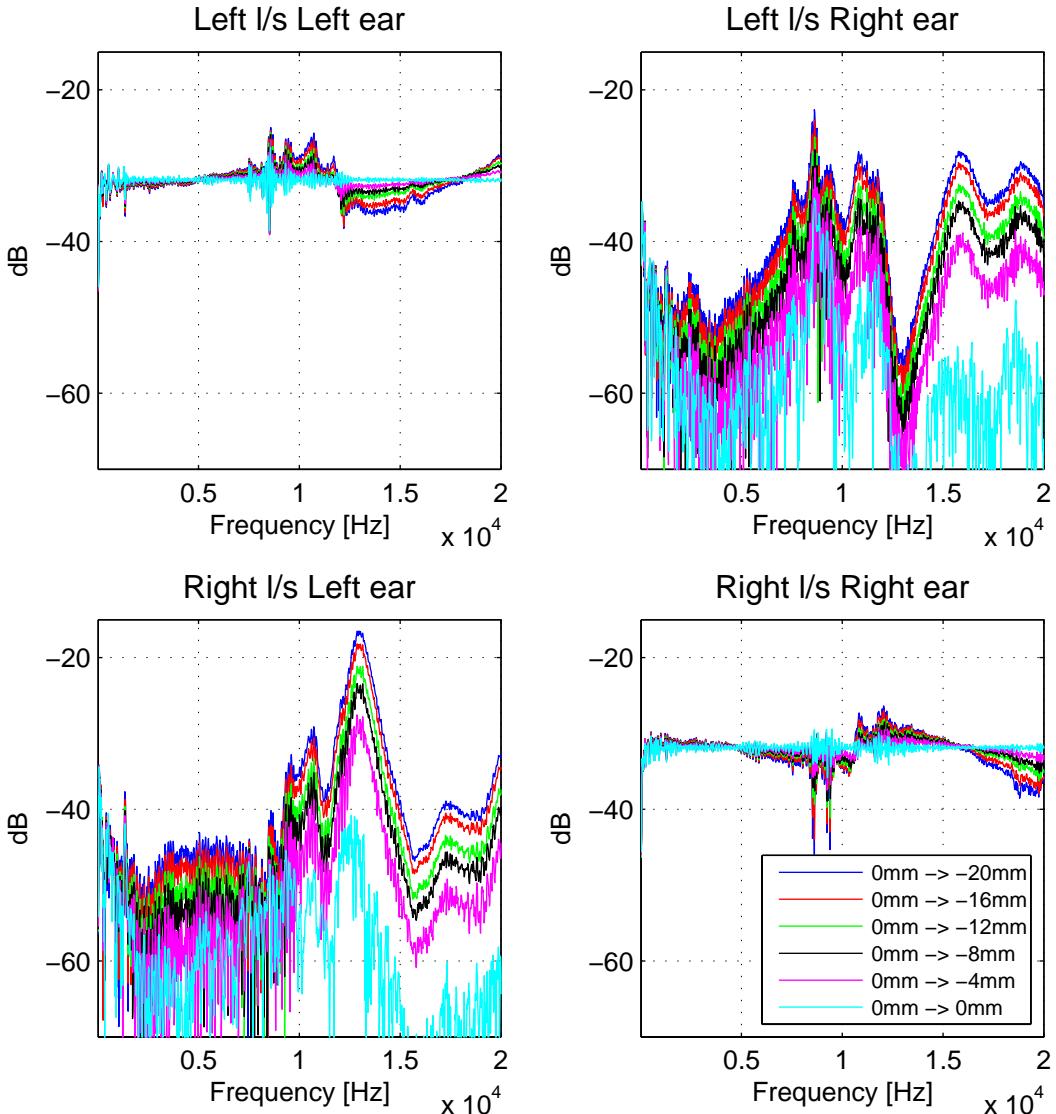
**Figure 5-10:** Simulated inversion results for listener misplacement up to 20mm to the left in 4mm increments. The inverse is determined using the strictly anechoic model of the plant  $\mathbf{C}_{190}$  for the 0mm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 4000 coefficients.

Note the systematic pattern of error in the misaligned results. This is more visible in the linear frequency axis plot of figure 5-11 and particularly in the equalisation results of the ipsilateral paths. These demonstrate a systematic fluctuation above and below the ideal equalisation line that increases with increasing misalignment. A number of points along the frequency axis can be identified for which the equalisation error is constantly kept equal to zero. These are at approximately 5kHz, 10kHz and 15kHz for the left-side ipsilateral path and at approximately 6kHz, 12kHz and 18kHz for the right-side ipsilateral path. Maximal deviations from the ideal occur between these

frequencies. The pattern is not easily identifiable in the cross-talk paths where the cross-talk cancellation reduction pattern is superimposed in the already fluctuating pattern of the non-misaligned result. However an increase of the error as frequency increases is clear visible.



**Figure 5-11:** Simulated inversion results for listener misplacement up to 20mm to the left in 4mm increments. The inverse is determined using the strictly anechoic model of the plant  $\mathbf{C}_{190}$  for the 0mm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 4000 coefficients. Linear frequency axis.



**Figure 5-12:** Simulated inversion results for listener misplacement up to 20mm to the right in 4mm increments. The inverse is determined using the strictly anechoic model of the plant  $\mathbf{C}_{190}$  for the 0mm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 4000 coefficients. Linear frequency axis.

The pattern of figure 5-11 is exactly reversed in figure 5-12 where we plot the inversion results for the same case but with the listener's position misaligned to the right this time. As is seen in the figure, the fluctuation around the ideal equalisation line is reversed between the left and right ipsilateral channels. The same is true for the distribution of the “node” frequency points, with the 5kHz, 10kHz, 15kHz sequence appearing now in the left-side ipsilateral path and the 6kHz, 12kHz, 18kHz sequence in the right-side ipsilateral path.

The error pattern of figures 5-11 and 5-12 is readily interpreted by use of the following analysis. When the listener moves away from the ideal position, the plant changes from the originally measured model  $\mathbf{C}$  to a disturbed model  $\hat{\mathbf{C}}$ . For small amounts of misalignment, each one of the responses  $\hat{C}_{ij}(z)$  in the disturbed model  $\hat{\mathbf{C}}$  should be approximately equal to a delayed (or advanced) and attenuated (or amplified) version of the corresponding response  $C_{ij}(z)$  in the original model  $\mathbf{C}$  depending on whether the specific path length is increased or decreased by the misalignment. Hence, denoting with  $d_{ij}$  the delay in (fractional) number of samples and with  $a_{ij}$  the attenuation ratio due to spherical spreading, we have

$$\hat{C}_{ij}(z) \approx a_{ij} z^{-d_{ij}} C_{ij}(z) \quad (5-2)$$

Using a typical value of 15cm for the interaural distance, a simple geometrical analysis shows that the maximum absolute path length change is equal to 2.87mm and occurs in the contralateral paths in the  $\pm 20$ mm misalignment cases. Even in that extreme case, the path length ratio between the non-misaligned and the misaligned case is equal to 1517.154mm/1514.284mm and corresponds to an attenuation (or amplification) by a factor of 1.0019. The corresponding ratios for the cases of smaller misalignment will be even lower. Hence, in equation (5-2), it is safe to approximate the geometrical spreading constants  $a_{ij}$  with unity and ignore them<sup>64</sup>.

Taking into account the symmetry of the geometry about the interaural axis and the small loudspeaker angle span considered here, an additional approximation can be made. That is, the ipsilateral path responses  $C_{11}(z)$  and  $C_{22}(z)$  can be considered to be approximately equal to each other

$$C_{11}(z) \approx C_{22}(z) \quad (5-3)$$

and the same will hold for the contralateral responses  $C_{12}(z)$  and  $C_{21}(z)$

$$C_{12}(z) \approx C_{21}(z) \quad (5-4)$$

Furthermore, the contralateral responses can be approximated by a scaled and delayed version of the ipsilateral responses

---

<sup>64</sup> The validity of this approximation is verified by the agreement of the results presented below with the measured results.

$$C_{ij}(z) \approx a_{sh} z^{-d_{ia}} C_{ii}(z) \quad (5-5)$$

In equation (5-5), the constant<sup>65</sup>  $a_{sh}$  corresponds to two physical effects; one being the attenuation of the contralateral response due to the longer path travelled from the source to the contralateral ear and the other being the shadowing effect of the head. The symbol  $d_{ia}$  denotes the delay corresponding to the path length difference between the contralateral and the ipsilateral paths, again in fractional number of samples.

Returning then to the inversion process we see that, in the case where the listener has moved from the correct position, the inversion process will be described by the convolution of the plant matrix  $\hat{\mathbf{C}}$  that corresponds to the misaligned geometry with the inverse matrix  $\mathbf{H}$  that is determined on the basis of the non-misaligned plant matrix  $\mathbf{C}$ . Hence, equation (1-6) becomes

$$\begin{aligned} \mathbf{X}_{eq} &= \hat{\mathbf{C}}(z) \mathbf{H}(z) \\ &= \begin{bmatrix} \hat{C}_{11}H_{11} + \hat{C}_{12}H_{21} & \hat{C}_{11}H_{12} + \hat{C}_{12}H_{22} \\ \hat{C}_{21}H_{11} + \hat{C}_{22}H_{21} & \hat{C}_{21}H_{12} + \hat{C}_{22}H_{22} \end{bmatrix} \end{aligned} \quad (5-6)$$

and substituting from equation (2-14) we obtain

$$\mathbf{X}_{eq} = \begin{bmatrix} \hat{C}_{11}(z)C_{22}(z) - \hat{C}_{12}(z)C_{21}(z) & -\hat{C}_{11}(z)C_{12}(z) + \hat{C}_{12}(z)C_{11}(z) \\ C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z) & C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z) \\ \hat{C}_{21}(z)C_{22}(z) - \hat{C}_{22}(z)C_{21}(z) & -\hat{C}_{21}(z)C_{12}(z) + \hat{C}_{22}(z)C_{11}(z) \\ C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z) & C_{11}(z)C_{22}(z) - C_{12}(z)C_{21}(z) \end{bmatrix} \quad (5-7)$$

Substituting the approximations of equations (5-2) – (5-5) into equation (5-7) and after some rearrangement we obtain

$$\mathbf{X}_{eq} = \begin{bmatrix} \frac{z^{-d_{11}} - a_{sh}^2 z^{-(d_{12}+2d_{ia})}}{1 - a_{sh}^2 z^{-2d_{ia}}} & \frac{-a_{sh} z^{-(d_{11}+d_{ia})} + a_{sh} z^{-(d_{12}+d_{ia})}}{1 - a_{sh}^2 z^{-2d_{ia}}} \\ \frac{a_{sh} z^{-(d_{21}+d_{ia})} - a_{sh} z^{-(d_{22}+d_{ia})}}{1 - a_{sh}^2 z^{-2d_{ia}}} & \frac{z^{-d_{22}} - a_{sh}^2 z^{-(d_{21}+2d_{ia})}}{1 - a_{sh}^2 z^{-2d_{ia}}} \end{bmatrix} \quad (5-8)$$

<sup>65</sup> In higher loudspeaker angle span geometries, the shadowing of the head varies significantly with frequency. In such a case the approximation of the head shadowing with a constant may not be well-suited. As is further discussed below, this is one of the issues related to the material presented in this section that is left as further work.

In figures 5-13 and 5-14 we plot the frequency response of the equalisation matrix  $\mathbf{X}_{eq}$  that is described in equation (5-8) for listener misalignment of 20mm to the left and right respectively. For the results of these figures, the values of the fractional delays  $d_{ij}$  are calculated by a simple geometric analysis using the typical value of 15cm for the interaural distance, the value of 340m/s for the speed of sound and the value of 48kHz for the sampling rate. The values of  $a_{sh}$  and  $d_{ia}$  on the other hand, are estimated by use of the measured model of the plant  $\mathbf{C}$  in the non-misaligned position.

The delay  $d_{ia}$  is estimated by upsampling the measured responses  $c_{ij}$  from one loudspeaker to the two ears by a factor  $r_{fs}$  to obtain the upsampled versions  $\tilde{c}_{ij}$  and subsequently computing the index of the maximum value of the cross-correlation function between the upsampled responses. That is, the estimated value of  $d_{ia}$  is the value of  $n_0$  that maximises the expression of equation (5-9) divided by the upsampling factor  $r_{fs}$ .

$$d_{ia} = \frac{n_0}{r_{fs}} \quad \text{with } n_0 \text{ such that} \quad \max_{n_0} \left\{ \sum_n \tilde{c}_{ij}(n+n_0) \tilde{c}_{jj}(n) \right\} \quad (5-9)$$

For the results presented here an upsample factor of 10 was used which gives a maximum resolution of 1/10 of a sample. The delay computed with this method was equal to 1.6 samples for the responses from the left loudspeaker and 2.6 samples for the responses from the right loudspeaker. The large difference between the two values is probably due to a deviation of our non-misaligned arrangement from the axisymmetric geometry but also due to the asymmetry present in the responses of the KEMAR dummy-head (see the discussion in p. 114). For the results of figures 5-13 and 5-14 we used the average of the two values,  $d_{ia} = 2.1$  samples.

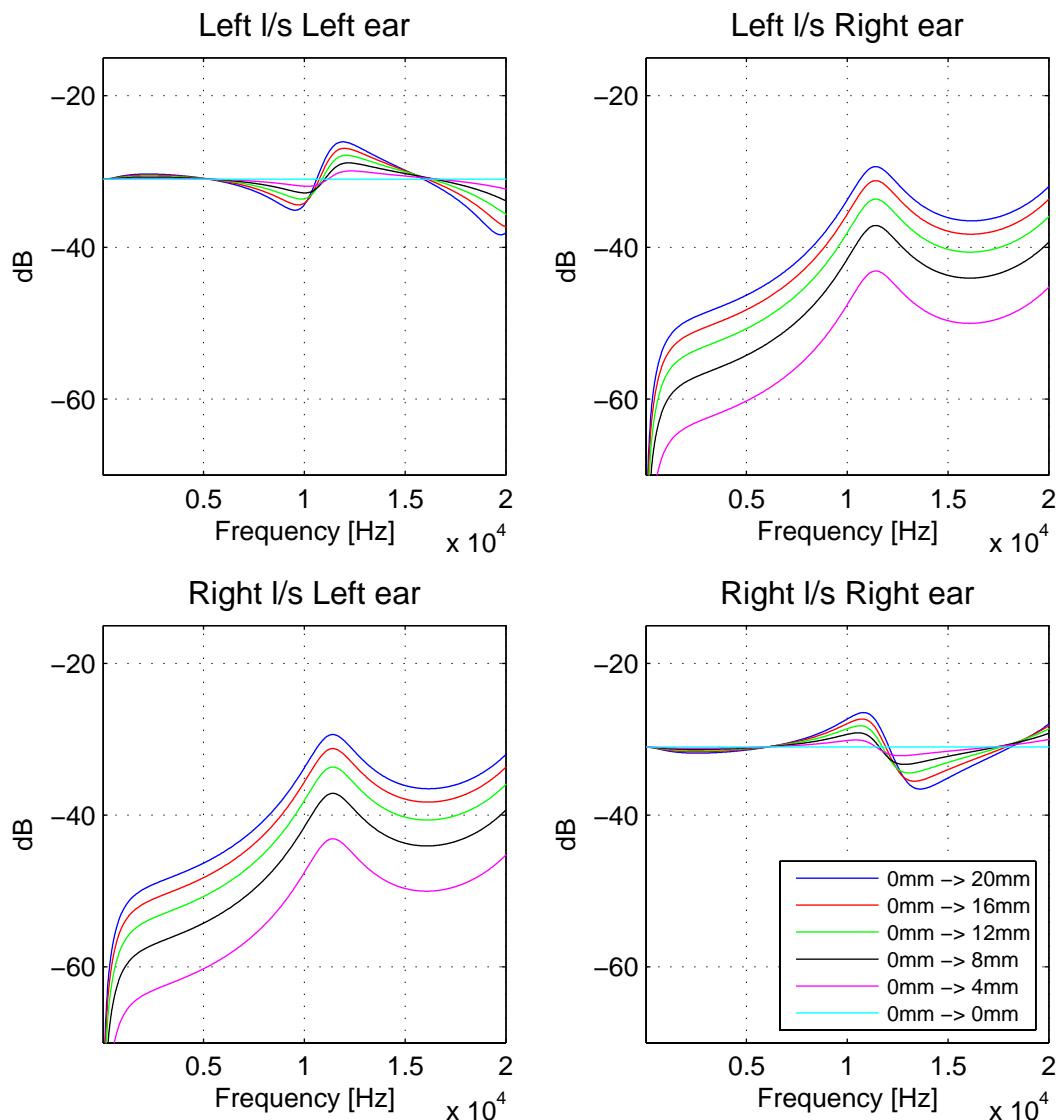
The shadowing factor  $a_{sh}$  was estimated as the square root of the ratio between energy in the strictly anechoic part of the responses for the two ears<sup>66</sup>.

---

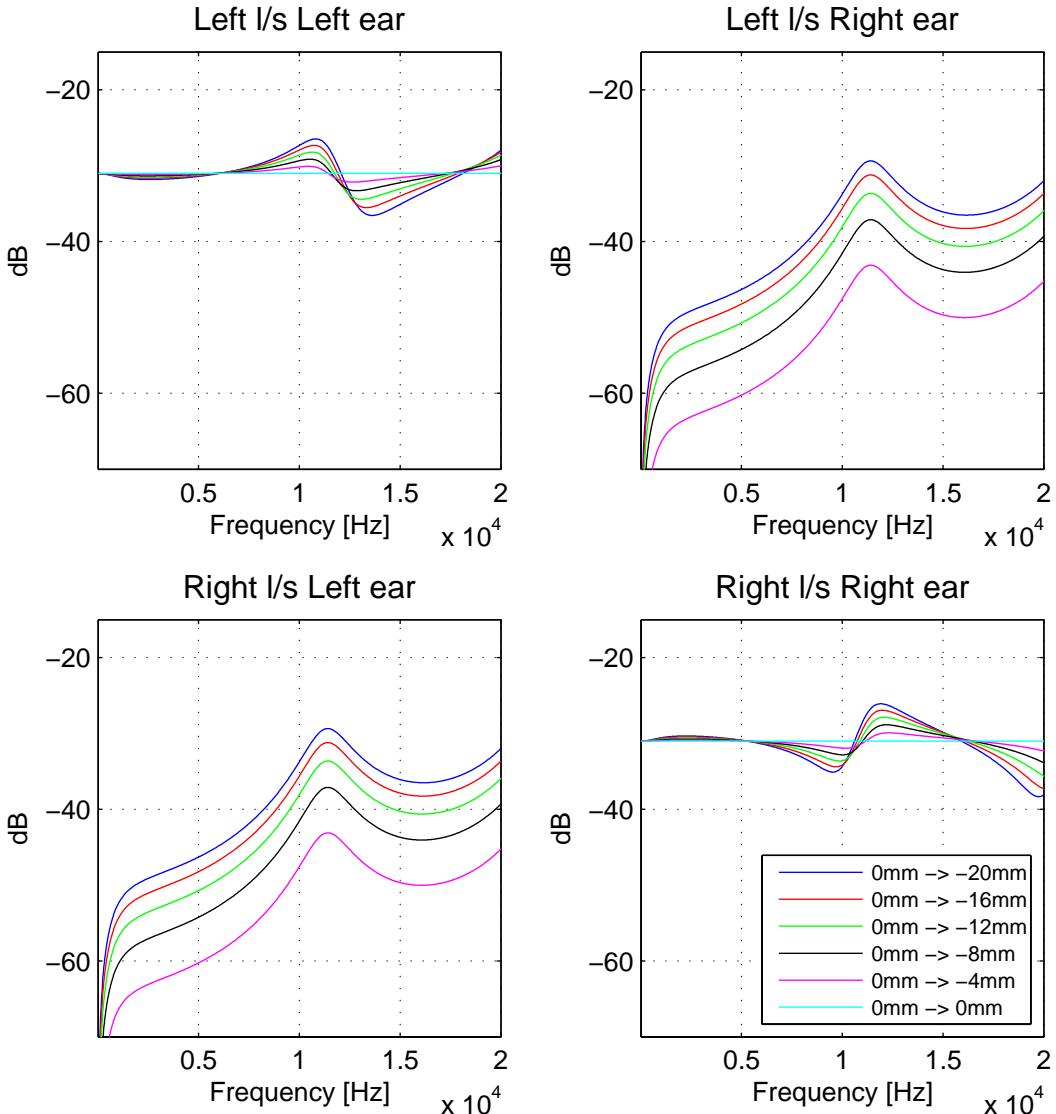
<sup>66</sup> Another method for the estimation of the shadowing factor would be to find the value of  $a_{sh}$  that minimises the quantity  $c_{ij}(n) - a_{sh}c_{jj}(n)$  in the least squares sense after  $c_{jj}(n)$  was (fractionally) delayed by  $d_{ia}$  samples. However, the results obtained with the value estimated by equation (5-10) are in good agreement with the measured results so the issue was not pursued further.

$$a_{sh} = \sqrt{\frac{\sum_{n=260}^{450} c_{ij}^2(n)}{\sum_{n=260}^{450} c_{jj}^2(n)}} \quad (5-10)$$

The value obtained by substituting the left-source binaural pair  $[c_{11}(n) \ c_{21}(n)]$  into equation (5-10) was  $a_{sh} = 0.7$  whereas the right-source binaural pair  $[c_{22}(n) \ c_{12}(n)]$  yielded the value  $a_{sh} = 0.75$ . The average value  $a_{sh} = 0.725$  was used.



**Figure 5-13:** Inversion results as estimated by equation (5-8) for listener misplacement up to 20mm to the left in 4mm increments.



**Figure 5-14:** Inversion results as estimated by equation (5-8) for listener misplacement up to 20mm to the right in 4mm increments.

The results of figures 5-13 and 5-14 can be seen to capture the error pattern of figures 5-11 and 5-12 very accurately<sup>67</sup> and hence completely validate the analysis presented above. More specifically, the characteristics of the fluctuation of the ipsilateral equalisation above and below the ideal, properly aligned, result that were described above are exactly replicated. As regards the cross-talk paths, a peak in the error can be seen at approximately 11kHz. This corresponds to a very well documented result

<sup>67</sup> Note that in order to make the analytical results of figures 5-13 and 5-14 directly comparable with the measured results of figures 5-11 and 5-12, we have introduced in the calculation of equation (5-8) a scaling factor equal to the normalisation (see p. 63) applied to the inversion results of 5-13 and 5-14. Note also that in the non-misaligned case, the analytical model predicts perfect cross-talk cancellation of  $-\infty$ dB and is thus not plotted in the results.

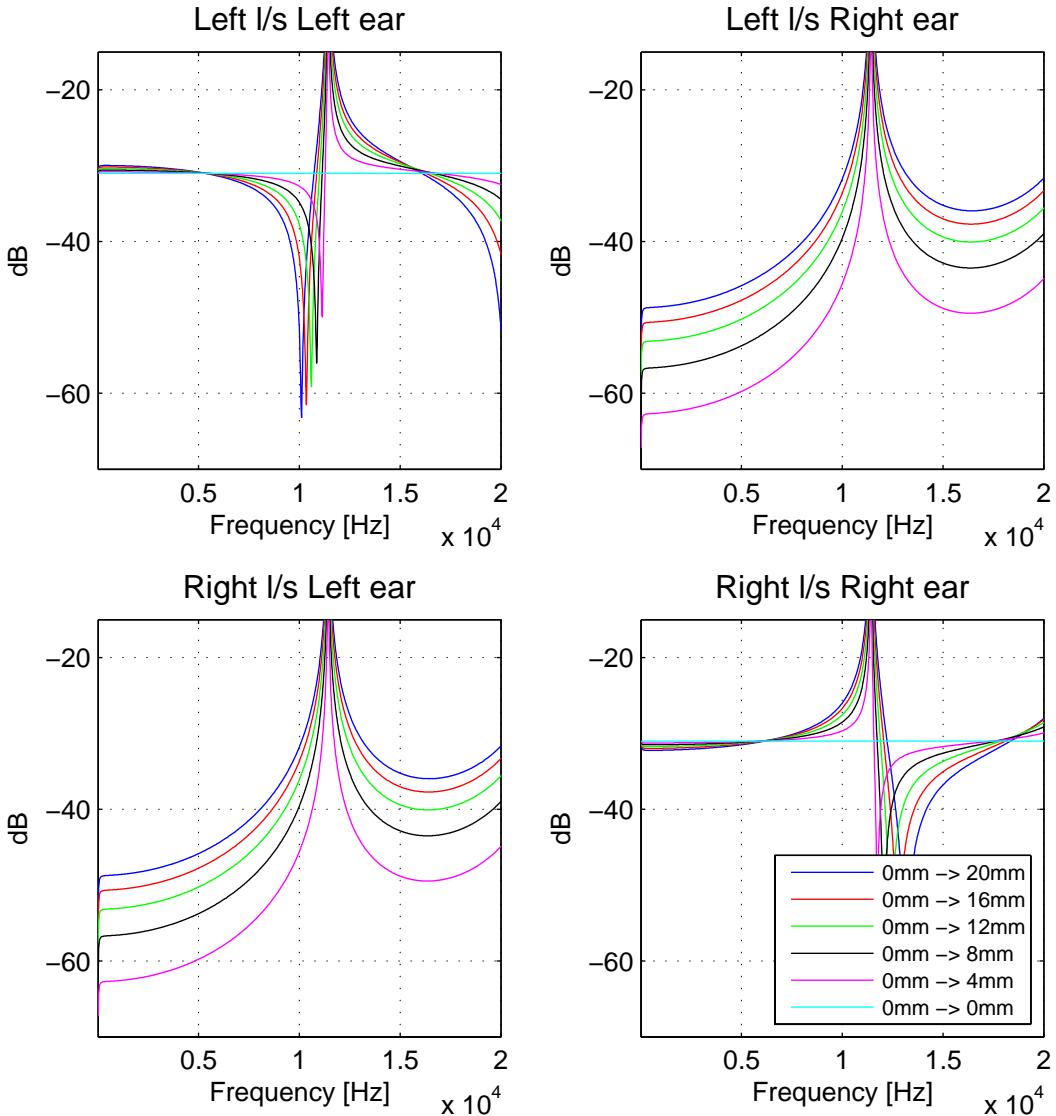
(Kirkeby et al. 1998a), (Nelson and Rose 2005), (Takeuchi et al. 2001) regarding the “ringing frequency” of the Stereo Dipole geometry, i.e. the first point along the frequency axis (after the zero frequency point) where the inversion for this particular loudspeaker angle span becomes ill-conditioned. As is explained in those works, the effect of the misalignment should be expected to have the severest impact at the frequency regions where the inversion becomes ill-conditioned. Hence, by placing the first ill-conditioned frequency of the audio range as high as 11kHz the Stereo Dipole geometry is found to be more robust to errors due to misalignment and hence yield a larger sweet-spot. This, of course, comes at the expense of a more severe impact of the zero frequency ill-conditioned region (see the discussion in §4.5) which, however, does not affect the system’s robustness to misalignment error.

Hence the measured and analytical results presented here are in very good agreement with the aforementioned works<sup>68</sup>. However, the results presented here offer an additional element of information to the relevant analysis. That is, unlike those previous works in which the analysis is founded on the free-field approximation of the inversion process, the analysis presented here incorporates the effect of the shadowing due to the presence of the listener. This was done by virtue of the  $a_{sh}$  parameter. Only with the inclusion of this element do the analytical results coincide so well with the measured results.

This is demonstrated with the results of figure 5-15 in which we plot the evaluation of equation (5-8) with the parameter  $a_{sh}$  set to the value predicted solely by the spherical spreading effect. In that case the path length ratio between the ipsilateral and contralateral paths (and hence the estimated value of  $a_{sh}$ ) is calculated geometrically to be equal to 0.99. The interaural sample delay is again set to the value  $d_{ia}=2.1$  samples as the results obtained with the value calculated geometrically were in worse agreement with the measurement results.

---

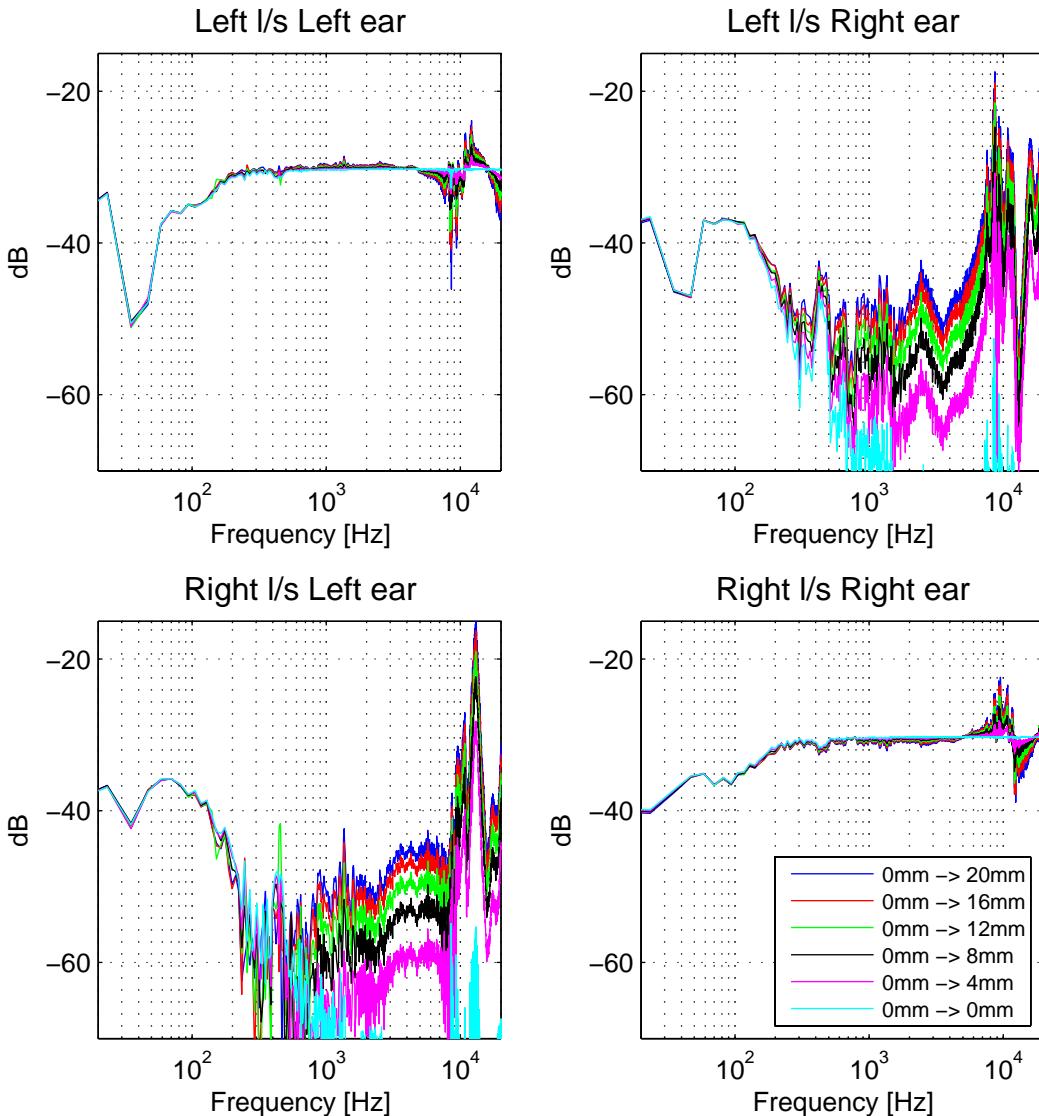
<sup>68</sup> In fact, the analytical approximation presented in equations (5-2) – (5-8) essentially amounts to an algebraic reformulation of the detailed analysis previously presented by Nelson and Rose (2005).



**Figure 5-15:** Inversion results as estimated by equation (5-8) for listener misplacement up to 20mm to the left in 4mm increments. The head-shadowing parameter  $a_{sh}$  is set to the value 0.99 and the interaural sample delay to  $d_{ia}$  is set to 2.1 samples.

Clearly, the results of figure 5-15 demonstrate the same physical effect as the ones of figures 5-11 – 5-14. However, they are seen to exaggerate the effect of the ill-conditioning to the extent that they imply the complete breakdown of the inversion. This, however, is not true since the effect of the shadowing of the head was seen above to considerably abate the undesirable effect of the ill-conditioning both in the ipsilateral equalisation and in the cross-talk cancellation. We believe that this observation could have interesting implications as regards the actual effect of the ill-

conditioning in other loudspeaker angle span geometries also. However the complete investigation of this issue is left outside the scope of this thesis.

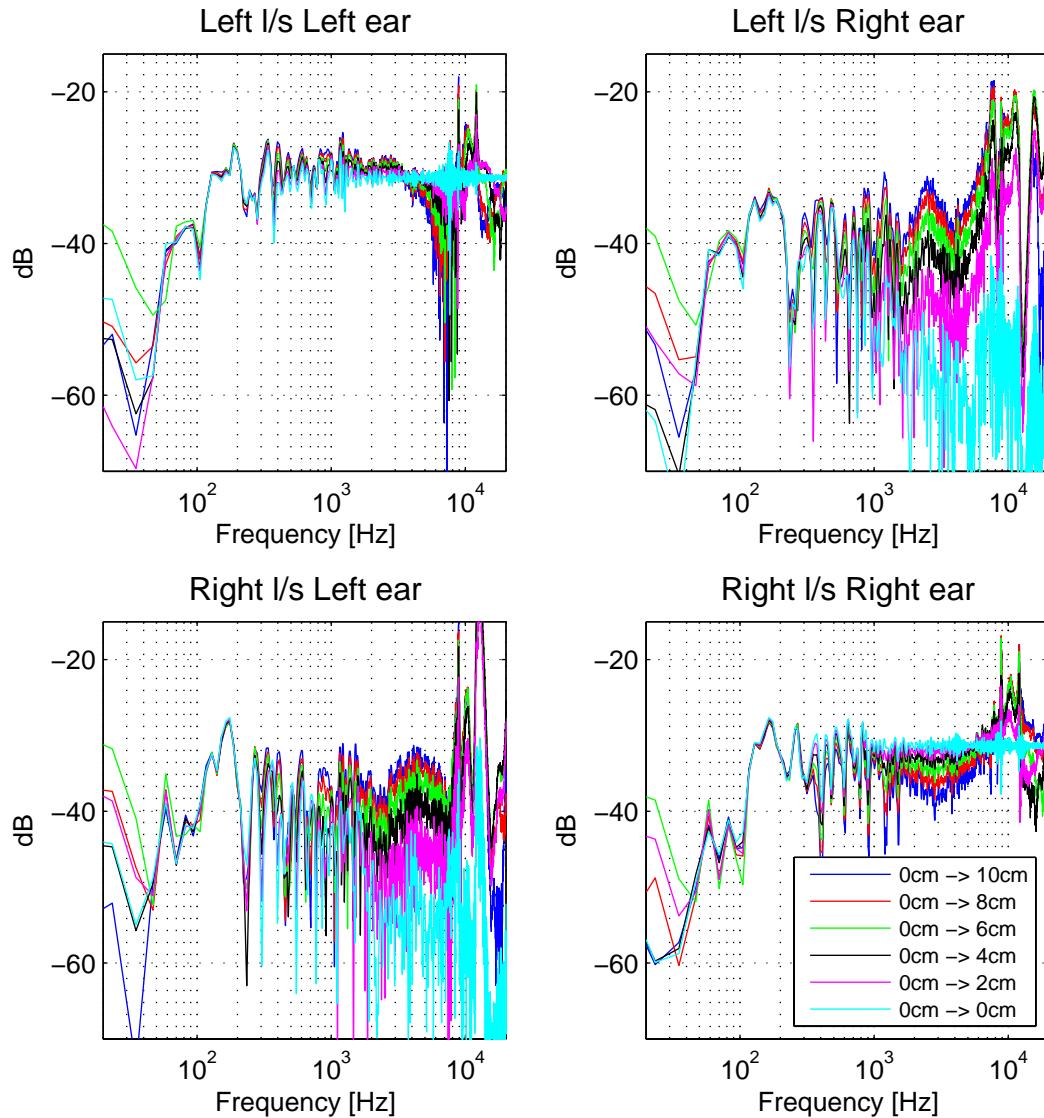


**Figure 5-16:** Simulated inversion results for listener misplacement up to 20mm to the left in 4mm increments. The inverse is determined using the  $\mathbf{C}_{1740}$  model of the plant for the 0mm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 coefficients.

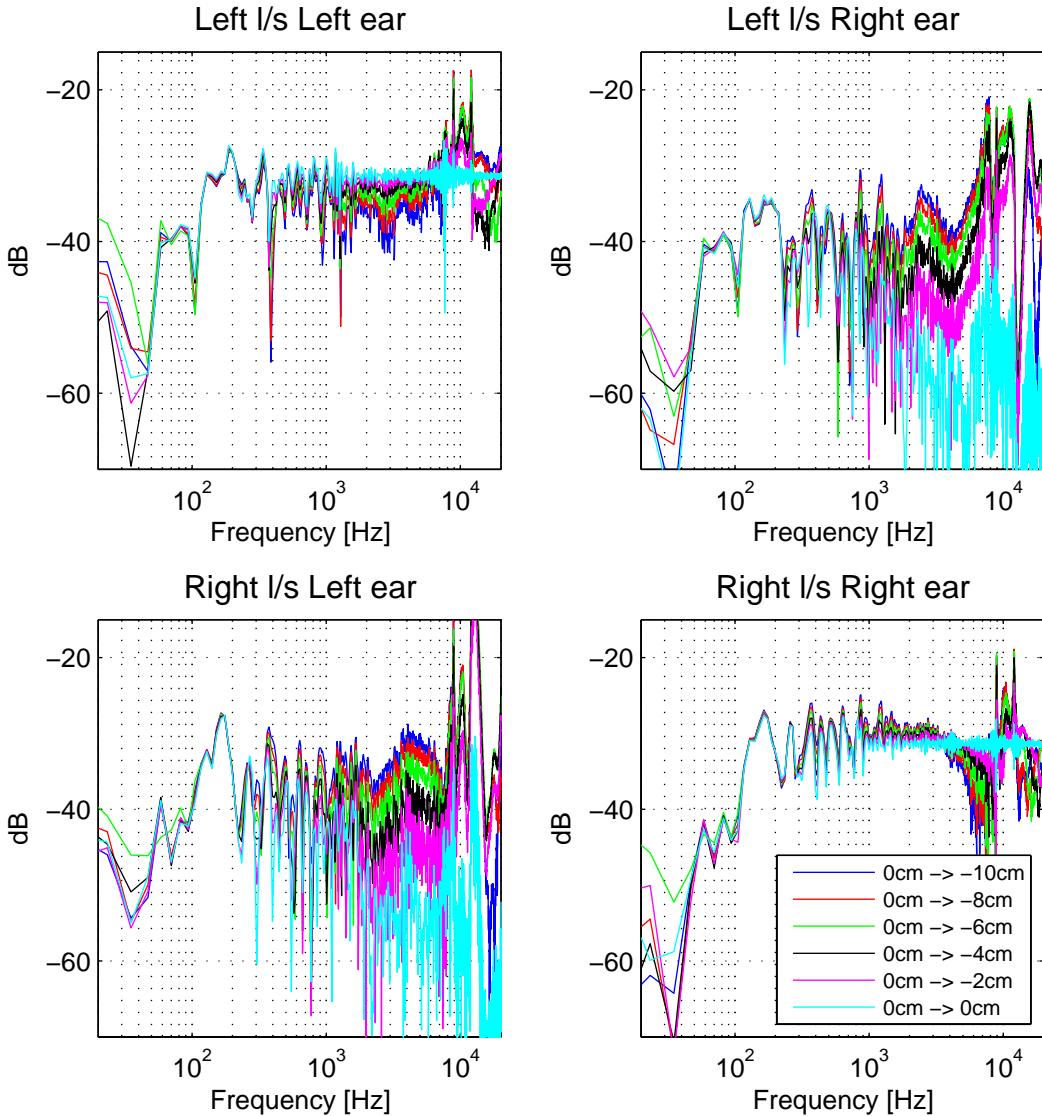
We now return to the original question posed in the beginning of this section, namely that of the comparison of the results obtained in the presence of misalignment error with those obtained under ideal conditions in the previous chapter. A further set of results is presented in figure 5-16. These depict the case where the inverse is determined on the basis of the  $\mathbf{C}_{1740}$  model of the plant. Following the conclusions

drawn from §4.2, the regularisation is set to  $\beta=10^{-4}$  and the inverse length to 10000 coefficients. As is seen in the figure, the advantage of the ipsilateral equalisation with this  $\mathbf{H}_{1740}$  inverse model compared to that of the  $\mathbf{H}_{190}$  model of figure 5-10, is retained up to approximately 7kHz even for the 20mm displacement case. Conversely, the cross-talk cancellation efficiency drops very fast (even from the 4mm displacement case) to approximately the same level obtained by the  $\mathbf{H}_{190}$  model. It thus becomes clear that the use of a very accurate inverse model such as  $\mathbf{H}_{1740}$  can be of practical interest only when the positioning accuracy of the reproduction arrangement is guaranteed to be in the order of no more than a few millimetres. No results are presented for higher regularisation values, as such an arrangement would primarily influence the low frequency region which as is seen is behaving very well for reasonably small positioning errors.

A final set of results is presented in figures 5-17 and 5-18 where we plot the inversion results for a listener misalignment of up to a maximum of 10cm to the left and right in 2cm increments. The asymmetry in the increase of the introduced delays in the misaligned plant is reflected in a consequent asymmetry in the cross-talk cancellation degradation. As can be seen in figure 5-17, the cross-talk from the right loudspeaker to the left ear remains substantially lower than the left side ipsilateral transmission even for the 10cm misalignment. Conversely, the cross-talk from the left loudspeaker to the right ear reaches the level of the right side ipsilateral transmission already for an 8cm misalignment. The situation is, naturally, reversed in the results of figure 5-18 where the misalignment of the listener to the right is considered.



**Figure 5-17:** Simulated inversion results for listener misplacement up to 10cm to the left in 2cm increments. The inverse is determined using the strictly anechoic model of the plant  $\mathbf{C}_{190}$  for the 0cm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 4000 coefficients.



**Figure 5-18:** Simulated inversion results for listener misplacement up to 10cm to the right in 2cm increments. The inverse is determined using the strictly anechoic model of the plant  $\mathbf{C}_{190}$  for the 0cm position. The regularisation is set to  $\beta=10^{-4}$  and the inverse length to 4000 coefficients.

Overall, the results of figures 5-17 and 5-18 show that a small amount of cross talk cancellation is still achieved in the mid-frequency region for up to 6-8cm of left-right misalignment. Note that in the Stereo Dipole geometry used here there is practically no natural head-shadowing (see figure 4-22). So even such a slight amount of cross-talk cancellation can achieve some widening of the soundstage presented by the loudspeakers compared to their physical positioning. This, however, comes with a severe colouration of the ipsilateral paths and cross-talk spill-over at higher frequencies. Even though the use of a non-individualised plant model should be

expected to make the influence of misalignment less clear-cut, both these effects have been verified in informal listening tests of the author with material presented with non-individualised inverse filters over the Stereo Dipole arrangement in anechoic conditions.

In conclusion then, the results of this section make apparent that a misalignment of 2cm renders the inversion of the frequency region above 10kHz pointless. Such misalignment will most probably exist in implementations without tracking of the listener's position in real-time. Hence, in such applications, a band-limited implementation of the inverse could reduce the complexity of the problem without any perceptible degradation of the performance. The inversion of the model of the plant containing the full decay of the measured responses (in this case  $\mathbf{C}_{1740}$ ) does indeed achieve an improvement in the ipsilateral equalisation (comparison of figures 5-10 and 5-16) that is retained in the mid-frequency frequency region up to about 8kHz for a reasonably large amount of listener misalignment. The advantage in terms of the cross-talk cancellation is on the other hand lost when the misalignment exceeds a limit of approximately 4mm. Such a fine spatial resolution is challenging to achieve even when real-time head-tracking is implemented. On the other hand, the analysis of the error pattern in the misaligned results proves that if such accurate knowledge of the listener's position is available, a simple geometrical calculation of the path length differences together with the inclusion of the head shadowing parameter (which can be deduced from the non-misaligned plant model) can very efficiently rectify the error due to the misalignment. Such an implementation would only incur the scaling of the inverse responses by a constant and the introduction of appropriate fractional delay corrections (Laakso et al. 1996).

## ***5.4 Non-individualised plant model***

A third source of error in the inversion, that of the modelling of the plant with non-individualised responses, is considered in this section. The influence of this element of error has been previously investigated Gardner (1998), Köring and Schmitz (1993)

and Takeuchi (2001)<sup>69</sup>. In (Gardner 1998), the restriction of the inverse filtering stage to the frequency range below 6kHz is seen to significantly reduce the influence of the use of a non-individualised plant model for the determination of the inverse. An analysis based on subjective experiments in (Takeuchi 2001) shows that the subjects' set can be divided in a class of “good localisers” and a class of “bad localisers”. The use of KEMAR HRTFs for the determination of the inverse is shown to retain most of the presented material's directional information in the first group of subjects and the contrary is shown for the second group. These findings are shown to be correlated with the degree of deviation of the individual subject's HRTFs from those of KEMAR. The analysis is however restricted to the range up to 8kHz due to restrictions of individual HRTF measurement apparatus. Based again on subjective experiments, the results of Köring and Schmitz (1993) show that the reproduction is perceived as virtually perfect when individualised HRTFs are used in both the binaural synthesis and inversion parts of the design. On the other hand, a considerable degradation of the results is reported when this is not the case. No quantitative inversion results are presented in that study.

As was the case with the previous sections, the objective here is to give a quantitative measure of the error introduced in the inversion achieved by the design under investigation, due to this element of error. To achieve this, one would ideally want to have knowledge of the full-length model of the plant resulting from a large number of different listeners. The convolution of this realistic model of the plant with an inverse computed on the basis of a plant derived from the HRTFs of a manikin (e.g. KEMAR) would isolate the actual degradation imposed by the use of the non-individual inverse model.

Note that a process whereby an inverse  $\mathbf{H}_1$  computed on the basis of one listener's HRTFs is convolved with this same listener's HRTF plant matrix  $\mathbf{C}_1$  and the result  $\mathbf{C}_1 * \mathbf{H}_1$  is compared to the convolution of  $\mathbf{C}_1 * \mathbf{H}_2$  where  $\mathbf{H}_2$  stands for the inverse computed on the basis of some other listener's HRTF matrix is not very useful. In such a process, the result of  $\mathbf{C}_1 * \mathbf{H}_1$  would be virtually perfect whereas the result of  $\mathbf{C}_1 * \mathbf{H}_2$  would drop to a much more modest amount of inversion accuracy. This

---

<sup>69</sup> It should be noted that a significant amount of literature exists regarding the associated issue of the use of individualised or non-individualised HRTFs in the binaural synthesis part of the problem (see also fn. 6). This is however outside the scope of this thesis.

deterioration however would by no means reflect the actual distance between the inversion quality in the two cases. As was extensively shown in chapter 4, in real-world situations the inversion accuracy will be significantly lower than the one estimated by  $\mathbf{C}_1^* \mathbf{H}_1$  even when the individual set of HRTFs is known. This was specifically made apparent with the results of figure 3-8 where it was shown that the results of the convolution of the strictly anechoic part of the plant with the inverse are excessively optimistic compared to the actually achieved inversion results. Furthermore, as was shown with the results of figure 4-25, an additional element of error is imposed on the inversion when the inverse is computed on the basis of HRTFs rather than on the basis of a model of the plant that also incorporates the responses of the loudspeakers.

It thus becomes apparent that the difference between the actual inversion accuracy when the inverse is computed on the base of individualised and non-individualised HRTFs should be much smaller than estimated by a process such as that described above. The exact objective therefore is to estimate the distance between the quality achieved when the individual HRTF-based inverse  $\mathbf{H}_1$  is used to invert the actual plant and the one achieved when a non-individual inverse  $\mathbf{H}_2$  (based on, say, KEMAR's HRTFs) is used for the same purpose.

Strictly speaking, to achieve such an objective one would need to have measurements of the full-length model of a set of individualised plants. Such a measurement has not been feasible during the course of the current research project. Hence a different approach has been taken that utilises an existing dataset available in the public domain. This is the CIPIC HRTF database (Algazi et al. 2001) that contains measured HRTFs of 43 human subjects and also two sets of KEMAR HRTFs fitted with the small and large pinna models. Thus, for the results presented, the inverse is determined of each one of these 45 HRTF-based plants<sup>70</sup>. Our measured 3000-coefficients full-length model of the plant (downsampled to 44.1kHz to meet the sampling rate of the CIPIC database) is then convolved with each of the 45 inverse models. In simple terms this inversion process represents the case where KEMAR is presented with binaural material using our ISVR anechoic chamber reproduction arrangement with the inverse processing network designed on the basis of some other

---

<sup>70</sup> In the results presented in this section the length of the inverse is always set to a high enough value to avoid the presence of truncation effects, typically a few thousand coefficients.

listener's HRTFs. The additional amount of error introduced by this inversion process compared to the inversion designed using KEMARs own HRTFs is a good estimate of the additional amount of error that would be introduced in real conditions of reproduction to a listener due to the design of the inverse using KEMAR HRTFs.

This is explained in equations (5-11) and (5-12). In these equations, with  $\mathbf{C}_{ind}$  and  $\mathbf{C}_{KEMAR}$  we denote the actual plants corresponding to an individual listener and KEMAR respectively. These contain the responses of the loudspeakers used for the reproduction and the late decay tail. Hence  $\mathbf{C}_{KEMAR}$  is well modelled with our full-length 3000-coefficients measured model of the plant. With  $\mathbf{H}_{ind}^{HRTF}$  and  $\mathbf{H}_{KEMAR}^{HRTF}$  we denote the models of the inverse determined using the HRTF-based plant matrix corresponding to an individual listener and KEMAR respectively. Hence the error matrix quantity  $\mathbf{E}(\omega)$  in equation (5-11) corresponds to the amount of error superimposed on the inversion by the use of the non-individualised inverse  $\mathbf{H}_{KEMAR}^{HRTF}$  instead of the individualised inverse  $\mathbf{H}_{ind}^{HRTF}$ . This is the error quantity we would ideally want to measure but are not able to measure directly because we lack a set of measurements of  $\mathbf{C}_{ind}$  for different listeners.

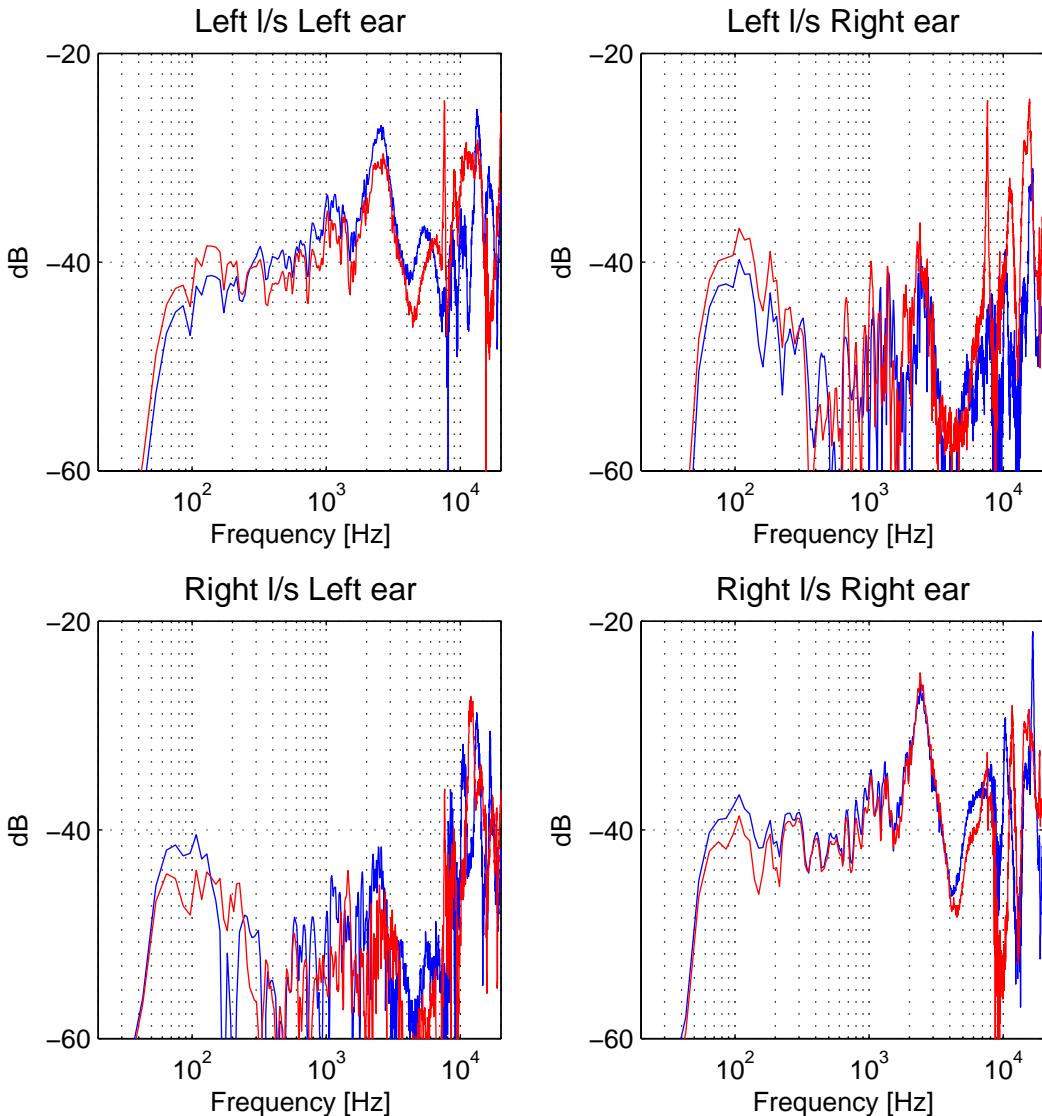
$$\begin{aligned}\mathbf{E}(\omega) &= \mathbf{C}_{ind}(\omega) \cdot \mathbf{H}_{KEMAR}^{HRTF}(\omega) - \mathbf{C}_{ind}(\omega) \cdot \mathbf{H}_{ind}^{HRTF}(\omega) \\ &= \mathbf{C}_{ind}(\omega) \cdot [\mathbf{H}_{KEMAR}^{HRTF}(\omega) - \mathbf{H}_{ind}^{HRTF}(\omega)]\end{aligned}\quad (5-11)$$

Instead, the error quantity  $\tilde{\mathbf{E}}(\omega)$  of equation (5-12) is that estimated by the simulation process described above.

$$\begin{aligned}\tilde{\mathbf{E}}(\omega) &= \mathbf{C}_{KEMAR}(\omega) \cdot \mathbf{H}_{ind}^{HRTF}(\omega) - \mathbf{C}_{KEMAR}(\omega) \cdot \mathbf{H}_{KEMAR}^{HRTF}(\omega) \\ &= \mathbf{C}_{KEMAR}(\omega) \cdot [\mathbf{H}_{ind}^{HRTF}(\omega) - \mathbf{H}_{KEMAR}^{HRTF}(\omega)]\end{aligned}\quad (5-12)$$

We see that the two error quantities in equations (5-11) and (5-12), albeit not identical, are essentially of the same form. More specifically, we see that in both cases the error quantity is equal to the difference between  $\mathbf{H}_{ind}^{HRTF}$  and  $\mathbf{H}_{KEMAR}^{HRTF}$  left-multiplied (or equally convolved in the time-domain) with the full-length model matrix of a given plant. Even though in equation (5-11) this is the plant matrix corresponding to the individual listener whereas in equation (5-12) it corresponds to KEMAR, the responses contained in these two matrices are roughly of the same form.

We thus expect the amount of error estimated by each one of these two expressions to be of approximately the same magnitude.



**Figure 5-19:** Simulated inversion results. The inverse is determined using a plant model containing KEMAR HRTFs from the CIPIC database with KEMAR fitted with the small pinna shape (red line) and the large pinna model (blue line). The regularisation is set to  $\beta=10^{-4}$ . The inversion process is simulated by use of our measured full-length model of the plant.

The quantity  $\mathbf{C}_{KEMAR} \mathbf{H}_{KEMAR}^{HRTF}$  described above is depicted in figure 5-19. In this figure we plot the results obtained from the above process using the two KEMAR plant models (small and large pinna shapes) from the CIPIC database. Simply speaking, these results correspond to the case where a listener (in this case KEMAR) has had

his/her HRTF database measured in an anechoic facility and then uses this database to process binaural material which is to be reproduced in another anechoic environment by a different audio setup. Such a scenario, which thus represents the realistic optimum when the individual HRTFs of the listener are known but the plant corresponding to the specific reproduction arrangement is not measured *in situ*, incorporates a number of error sources. These include (in a roughly ascending order of importance) the sampling rate conversion, the fact that the electroacoustic responses of the reproduction arrangement are not corrected (see the discussion in pp. 115-120), the presence of measurement errors in the HRTF measurement and the probability of positioning discrepancies between the HRTF measurement geometry and the reproduction geometry.

The last of these sources of error is probably the most important. For the case considered here the angles of the reproduction arrangement were identical with those of the HRTF measurement arrangement. On the other hand the CIPIC database HRTFs were measured with the measurement loudspeaker positioned 1m way from the centre of the listener's head whereas in our reproduction arrangement this distance was slightly more than 1.5m. At this distance range, the plane wave local approximation to the spherical spreading can be appreciably erroneous and hence an increase of the distance can introduce error. This assumption is supported by the difference in the results corresponding to the small and the large pinna shapes at low frequencies. At a few hundred Hz the wavelength of the sound-field is in the range metres, so a variation of the plant by an object whose difference in shape is in the order of millimetres should not affect this frequency region.

Moving then to the results of figure 5-19, we see that a significant amount of cross-talk cancellation ranging from 5dB to more than 20dB is retained in the region from 100Hz to 10kHz. The ipsilateral equalisation in this region displays an error of approximately  $\pm 5$ dB which should also be considered tolerable. Note that the broad peak in the 2-3kHz region is due to the fact the CIPIC HRTFs were measured at the entrance of a blocked ear canal whereas our plant model was measured at the position of the eardrum. The ear canal resonance, typically appearing as a broad peak in this region, is thus left uncorrected and it would not be present both measurements were

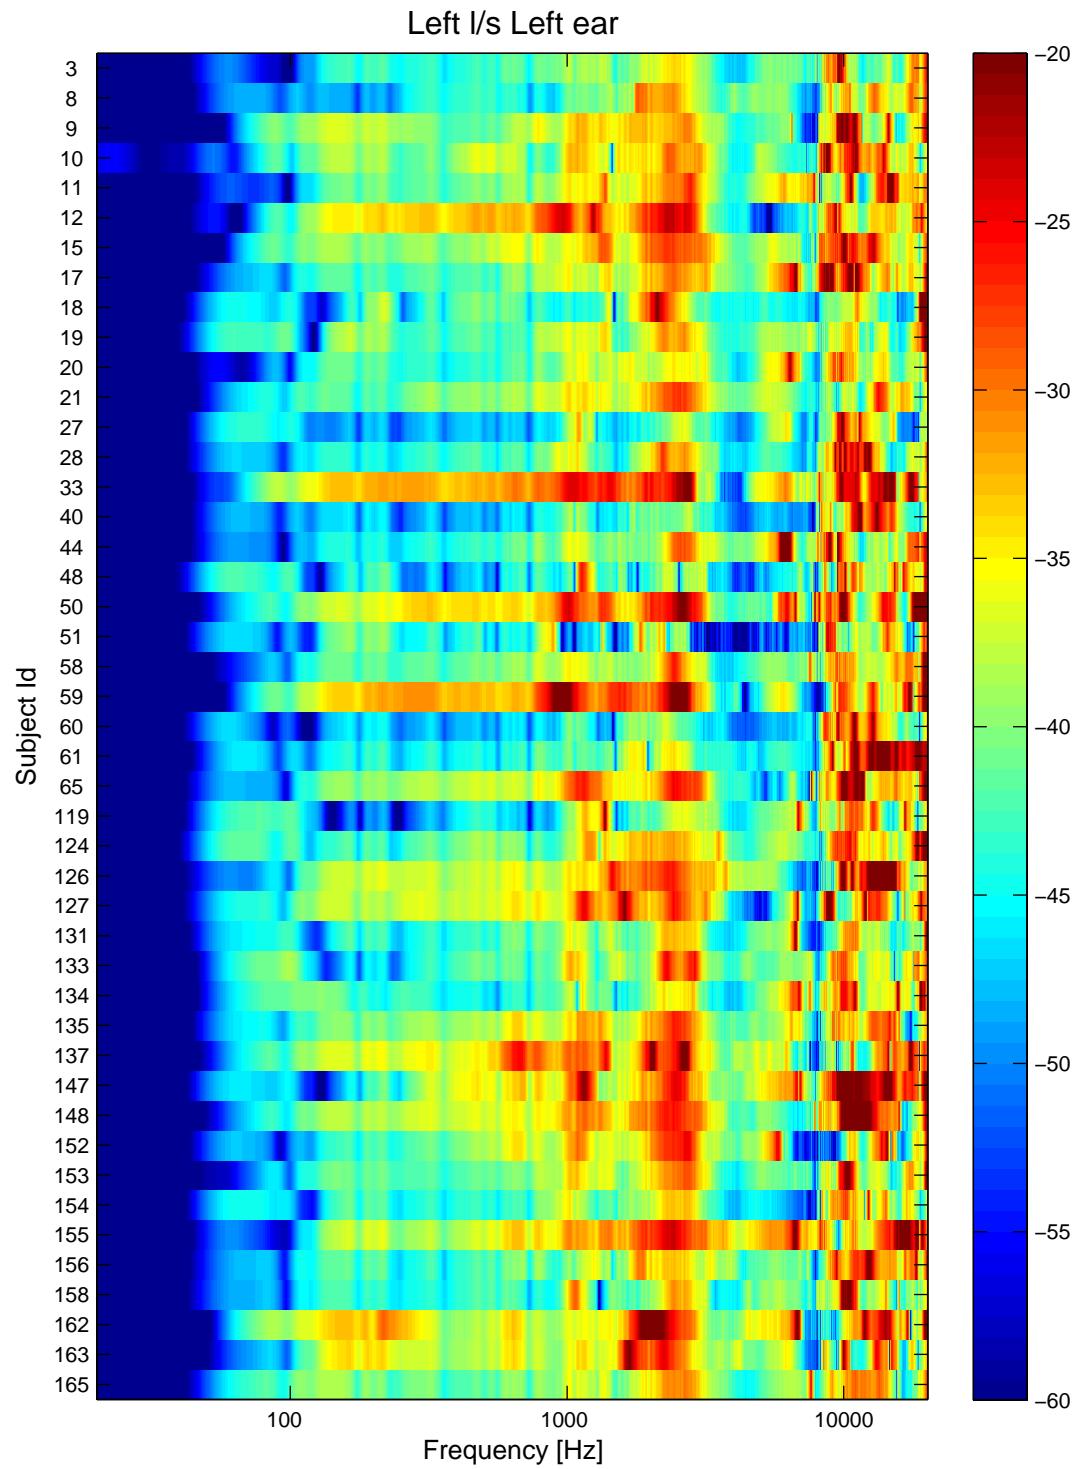
made at the same point along the ear canal<sup>71</sup>. For frequencies above 10kHz both cross-talk cancellation and ipsilateral equalisation can be seen to fail.

The collective comparison of the results of figure 5-19 with the results obtained with non-individualised inverse models are plotted in figures 5-20 and 5-21. In figure 5-20 we plot the frequency response of the (1,1) element of the matrix  $\mathbf{C}_{KEMAR} \cdot \mathbf{H}_{ind}^{HRTF}$  with  $\mathbf{H}_{ind}^{HRTF}$  derived from the set of 45 subjects' HRTFs of the CIPIC database. KEMAR fitted with the small pinna shape is subject 165, so the last line of the figure is the same with the results (red line) of figure 5-19. Lines with minimal colour variation show subject cases where the left channel ipsilateral equalisation has succeeded despite the use of a non-individualised inverse model. Such a case is shown, for example, by the subject number 3 (first line in figure 5-20) which would correspond to the “good localisers” class described by Takeuchi (2001). A few more subjects (e.g. 9, 15, 58, 124, 126 and 154) can be seen for which the use of a non-individualised inverse model seems to also work fairly well. On the other hand, in a few subjects (e.g. 33, 51, 137 158) the equalisation seems to fail altogether. Those would correspond to the group of “poor localisers” of (Takeuchi 2001). Note that the measurement artefact due to the non-correction of the ear canal resonance is present in nearly all subject cases.

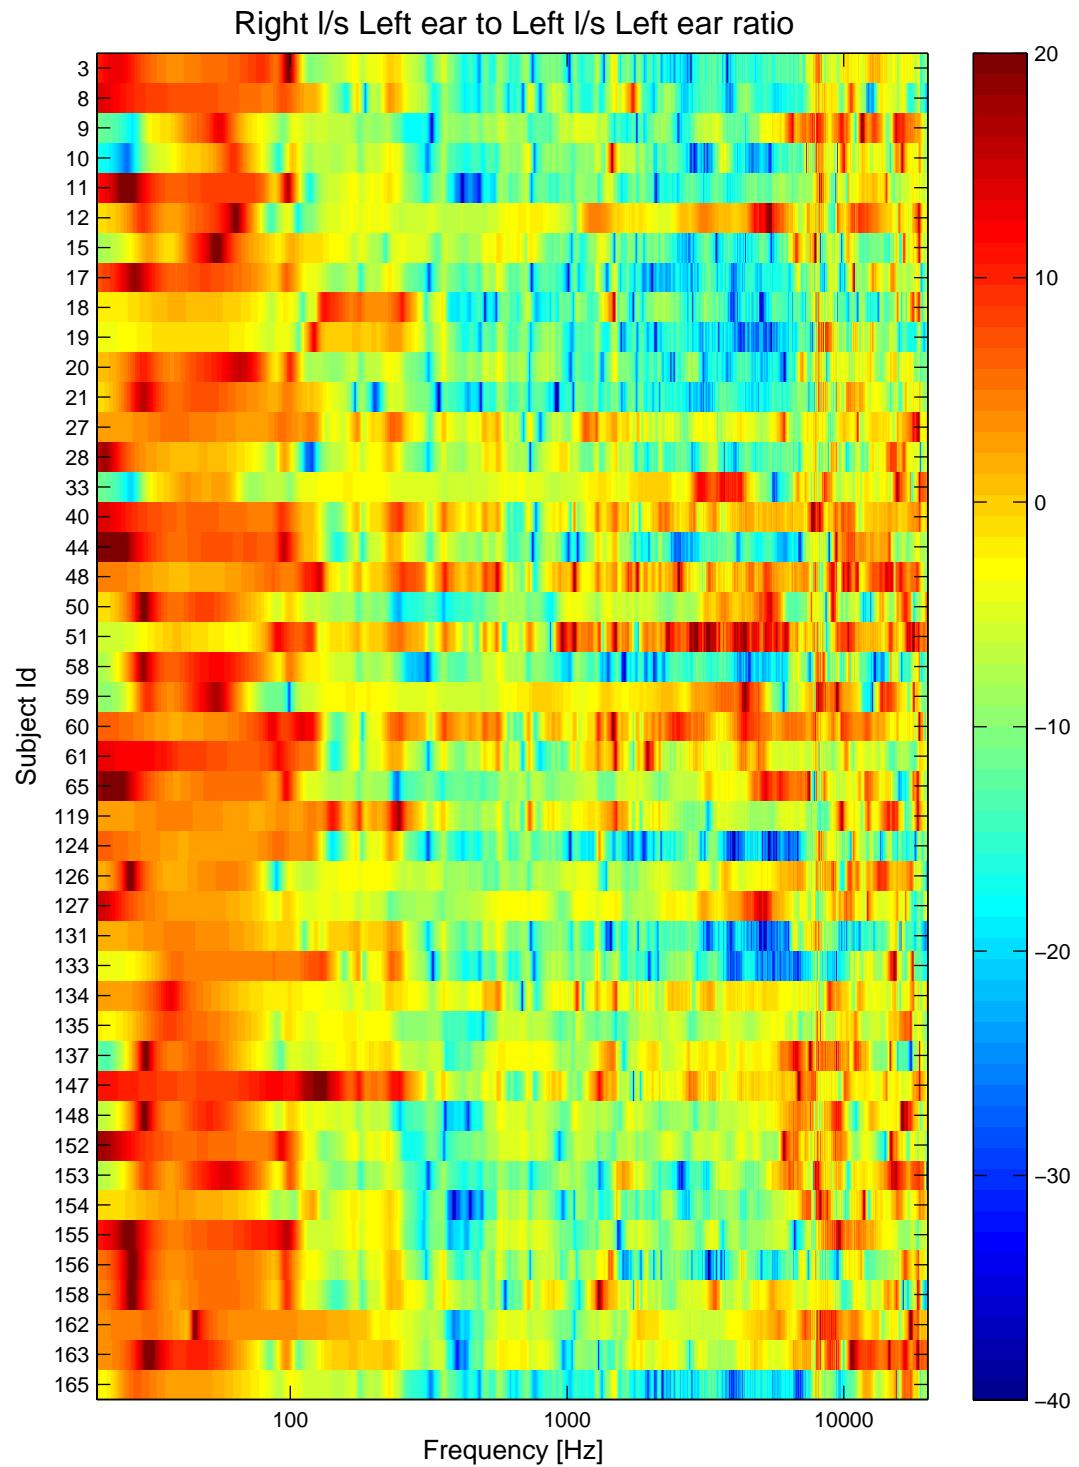
In figure 5-21 we plot the ratio of the (1,2) over the (1,1) element of the matrix  $\mathbf{C}_{KEMAR} \cdot \mathbf{H}_{ind}^{HRTF}$ . This corresponds to the level difference (in dB) between the signal reaching the left ear from the left loudspeaker and the signal reaching the same ear from the right loudspeaker. The closer any line of the figure is to blue, the better the cross-talk cancellation is for that subject. Normally, the best result is seen for KEMAR fitted with the small pinna shape (subject 165) but a good level of cross-talk cancellation is retained for a number of other subjects too. Those largely coincide with those for which the ipsilateral equalisation was also effective. In all cases however, as was to be expected from the results of figure 5-19, the inversion can be seen to be completely ineffective in the high frequency region above 8kHz or so.

---

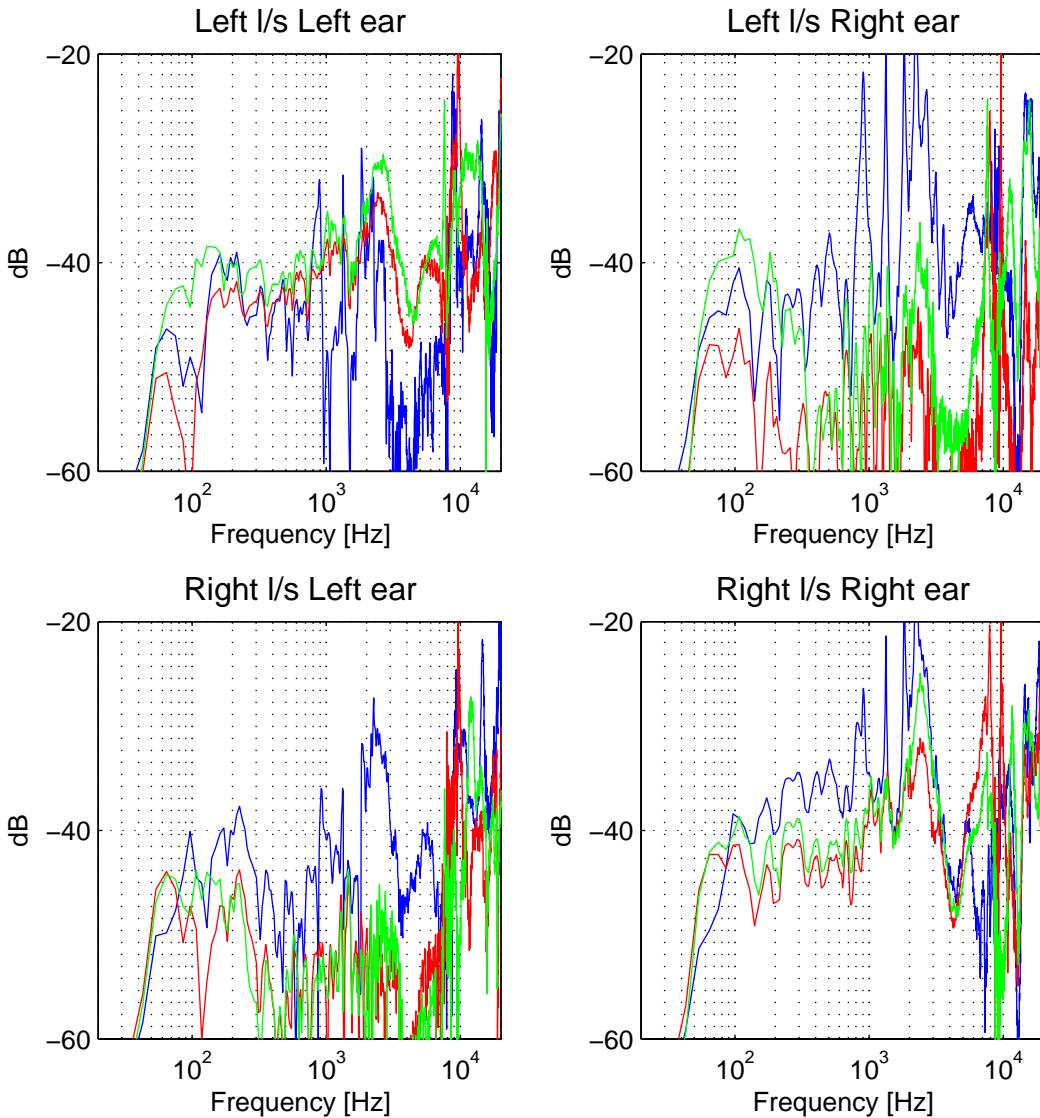
<sup>71</sup> For a detailed discussion on the issue of the optimal HRTF measurement position see (Köring and Schmitz 1993), (Møller 1992) and (Møller et al. 1995b).



**Figure 5-20:** Results of left channel ipsilateral equalisation. The inverse is determined using KEMAR HRTFs (subject 165) and with the regularisation set to  $\beta=10^{-4}$ .



**Figure 5-21:** Results of right channel to left ear cross-talk suppression. The inverse is determined using KEMAR HRTFs (subject 165) and with the regularisation set to  $\beta=10^{-4}$ .



**Figure 5-22:** Simulated inversion results for subjects 165 (green line), 51 (blue line) and 3 (red line). The inverse is determined using KEMAR's HRTFs (small pinna shape, subject 165). The regularisation is set to  $\beta=10^{-4}$ .

A clearer depiction of two extreme instances of the results presented above is given in figure 5-22 where we plot the results obtained for subjects 3, 51 and 165 (KEMAR fitted with small pinna shape). As is seen in the figure, the use of non-individualised inverse seems to leave the inversion virtually uninfluenced for subject 3. This is hardly the case for subject 51 who should be expected to receive only a slight cross-talk cancellation in the region of 200Hz to 1kHz and a large amount of colouration throughout the spectrum of the presented material. Assuming, however, that the analysis of equations (5-11)-(5-12) is correct, one would have to presume that even this level of inversion quality is enough to retain a small portion of the localisation

information in a binaural recording or a binaurally synthesised pair of signals. This is because, after an extensive record of demonstrations of such material to literally hundreds of visitors in ISVR, the author has never come across a single listener who failed to perceive virtual images from the “easy-to-achieve” directions, namely those in the front half of the horizontal plane.

In conclusion then we see that in all case where the inverse is determined using a non-individual HRTF model of the plant, the region above 8kHz is practically pointless to invert. This would even be true in the case where an individual HRTF model is known but not measured *in situ*. Further to that, the results show that a significantly large group of the CIPIC subjects can indeed be accommodated with the use of KEMAR based inverse. It is possible that the remaining cases could be classified in one or more different sets and that an equally satisfactory inverse representative could be found for each one of them. In such an event a set of, maybe as few as three or four, subject-representative HRTF database measurements could cover the entire population in terms of designing a satisfactory approximation to the ideal individualised inverse matrix. Finally we note that the results of figure 5-22 strongly suggest that the ample use of regularisation should be in order for the case of listeners like subject 51. Such a heavily regularised inverse should not at all degrade the already poor quality of the inversion and at the same time result in the associated benefits in terms of the design complexity discussed at length in chapter 4. On the other hand, such a choice could be damaging for cases like subject 3 in which the inversion is successful despite the use of a non-individualised inverse. This could be addressed in real-world implementations by the incorporation of a user-adjustable amount of regularisation that is set to the optimum value by means of trial-and-error.

## 5.5 Conclusions

In this chapter we expanded the evaluation of the inverse filtering design of §2.3 to cover a number of errors and approximations that are typically present in everyday reproduction situations.

First we considered the case where the reproduction arrangement contains strong early reflections. It was shown that when a model of the plant is available that

contains the effect of these reflections, then their presence can be efficiently treated by the inverse and the inversion results can be of equivalent accuracy to that of the anechoic inversion. This, however, incurs an increased requirement in the inverse filters' length. Conversely, if such early and strong reflections are ignored in the determination of the inverse, the quality of the inversion is dramatically influenced. The localisation cues that are mostly influenced are the monaural spectral cues the efficient delivery of which is determined by the accuracy of the ipsilateral equalisation. The failure to deliver these cues has as a result the ineffectiveness of the system in creating virtual images in the median plane and at the rear of the listener. In such an event, it was demonstrated that the best level of achievable inversion accuracy can be obtained with a very compact, heavily regularised model of the inverse and that any higher specification for the inverse is pointless. As a final point regarding the presence of reflections in the plant, we illustrated the fact that the presence of such early and strong reflections enhances the non-minimum phase character of the measured plant responses.

Subsequently we considered the case where positioning error is present in the reproduction arrangement. The measured results verified an expected effect, namely that the high frequency part of the inversion begins to fail for positional errors of the order of a few millimetres. The error in the inversion exhibited a systematic pattern which was explained with the use of an analytical model. This analytical model was obtained by applying a slight modification to an already existing analysis of the inversion process. The excellent agreement between the measured results and the proposed analytical model shows that (if an accurate knowledge of the position of the listener is available) the error introduced by the misalignment of a few tens of millimetres can be efficiently alleviated by the introduction of fractional-sample delays and constant scaling in the inverse filters without any need to recalculate the inverse filters. The need for high spatial-resolution head-tracking for such an implementation was highlighted. When such high-resolution head-tracking is not available and consequently positional errors in the order of a few millimetres are inevitable, the results verify that the inversion of the region above 10kHz is not justified. The case of positional error of the order of a few centimetres was also examined. The measured results again verified a previously observed fact, namely that the sweet-spot size of the Stereo Dipole geometry for the creation of virtual

images in the frontal half of the horizontal plane extends to approximately  $\pm 8\text{cm}$  to the left and right of the listener.

Finally, using data from the CIPIC database, we examined the effectiveness of the use of a non-individual HRTF model of the plant for the determination of the inverse. It was shown that a number of listeners are indeed well represented by the KEMAR-based HRTF model while others are not. It is thus possible that a collective database of a few dummy-head measurements of various sizes and shapes could cover all individuals. On the other hand, as was also the case with the positioning misalignment type of error, the results showed that, even when individual HRTFs are used for transaural reproduction in anechoic conditions, the inversion of the frequency range above 10kHz is not justified. Hence there is considerable room for complexity reduction and computational savings if the inverse filtering design is implemented at lower sampling rates.

## ***6 Recursive inverse models –***

### ***Forward-backward time filtering***

#### ***6.1 Introduction***

In this chapter we present a method for the implementation of the inverse filter matrix  $\mathbf{H}(z)$  with recursive filters. The fundamental difficulty in such a design is the fact that the determinant polynomial of the plant matrix is typically non-minimum phase, a fact long recognised in the literature (see also the discussion in §2.3.1). Hence, the recursive implementation that follows directly from the rational form of  $\mathbf{H}(z)$  of equation (2-14) is non-realisable. Recursive implementations exist in the literature that overcome this difficulty by various approximations or reductions of the complexity of the problem. The aim of the work presented here, however, is the investigation of the possibility of a recursive design that follows the FIR inverse models of the previous chapter, copying thus the local sound field reconstruction results that were presented there.

Some initial work on the possibility of such a recursive design of the filters comprising the inverse matrix  $\mathbf{H}(z)$  was undertaken in a previous project by the author (Papadopoulos 2000). From that work it became apparent that the standard zero-pole form of a recursive filter fails to effectively model the rising anti-causal part of the inverse filters considered in our case. Hence, a different approach is taken in the method introduced in this chapter. More specifically, the transfer functions comprising the matrix  $\mathbf{H}(z)$  are decomposed into two parts, one that is realisable in forward time and a second that is realisable in backward time. With this decomposition all the responses involved in the realisation of  $\mathbf{H}(z)$  correspond to decaying impulse responses that can more naturally be modelled with recursive filters. As is shown in the following, this can indeed lead to more efficient implementations of  $\mathbf{H}(z)$ .

The technique of backward-time filtering appears in the literature mainly in association with the problems of zero-phase filtering (Czarnach 1982), (Gilbert and Morrison 1996), (Gustafsson 1996), (Powell and Chau 1990a), (Powell and Chau 1990b), (Powell and Chau 1991) and subband signal coding and reconstruction in image processing (Chen and Vaidyanathan 1992), (Ramstad 1988). The details of its implementation are not covered however in any of the standard signal processing textbooks and we thus begin the presentation in this chapter by giving in §6.2 the related terminology and the basic formulae involved. In §§6.3–6.4 we describe how the technique is applied to the specific inverse filtering problem at hand and present a method for the numerical calculation of the transfer functions involved. In §6.5 we discuss a number of shortcomings associated with the basic formulation of the method and propose methods of improvement. Finally, in §6.6 we present and discuss the simulation results obtained when the method is used for the inversion of the plant matrix corresponding to the Stereo Dipole audio reproduction geometry. We begin the presentation in a single-channel context and subsequently we describe its application to the two-channel case. The formulation in this chapter assumes an off-line implementation of the method. A block processing algorithm the can be used for the on-line implementation of the method is presented in the next chapter.

## 6.2 *Convolution in backward time*

The basic idea behind the proposed inverse filtering technique is that, given a system with impulse response  $h(n)$  and an input signal  $x(n)$  the filtering of  $x(n)$  through  $h(n)$  can be implemented either in forward or in backward time. That is, one can either compute the convolution of the input time-series with the system's impulse response to obtain the output  $y(n)=h(n)*x(n)$  or alternatively, assuming that the whole input is known beforehand, reverse both the input series and the system's impulse response in time to get  $\hat{x}(n)$  and  $\hat{h}(n)$ <sup>72</sup>, and after computing the convolution  $u(n)=\hat{x}(n)*\hat{h}(n)$  reverse the outcome to acquire the same output as before:  $\hat{u}(n)=y(n)$ . As this process forms the basis of all the inverse filtering implementations described in the following, we give here its exact analytical formulation.

---

<sup>72</sup> In what follows we use the  $\hat{\cdot}$  symbol to denote the reversal in time of a given sequence.

We introduce the following terminology, to which we adhere throughout the following discussion. For any given sequence  $x(n)$  of finite length<sup>73</sup>  $N=N_2-N_1+1$  that is non-zero for the interval  $N_1 \leq n \leq N_2$  we associate the  $(N \times 1)$  vector  $\underline{\mathbf{x}}$  and the  $(N+M-1 \times M)$  Toeplitz matrix  $\mathcal{X}^M$  with  $M$  a given positive number, as in equation (6-1)<sup>74</sup>. We also use the symbols  $(\cdot)^T$ ,  $(\cdot)^R$  and  $(\cdot)^C$  to denote the transposition, reversal of rows and reversal of columns of the matrix or vector in the brackets.

$$\underline{\mathbf{x}} = \begin{bmatrix} x(N_1) \\ x(N_1+1) \\ \vdots \\ x(N_2-1) \\ x(N_2) \end{bmatrix} \quad \mathcal{X}^M = \begin{bmatrix} x(N_1) & 0 & \cdots & 0 \\ x(N_1+1) & x(N_1) & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ x(N_2) & x(N_2-1) & \cdots & x(N_2-M+1) \\ 0 & x(N_2) & \cdots & x(N_2-M+2) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & x(N_2) \end{bmatrix} \quad (6-1)$$

We can then see that, for a finite length signal  $x(n)$  which is non-zero in the interval  $N_1 \leq n \leq N_2$  and a finite length impulse response  $h(n)$  which is non-zero in the interval  $M_1 \leq n \leq M_2$ , the convolution sum  $y(n)=h(n)*x(n)$  which is equal to

$$y(n) = \sum_{k=-\infty}^{\infty} x(k)h(n-k) = \sum_{k=N_1}^{N_2} x(k)h(n-k) \quad (6-2)$$

can be written in the matrix-vector notation introduced as

$$\underline{\mathbf{y}} = \underline{\mathcal{H}}^M \underline{\mathbf{x}} =$$

$$= \begin{bmatrix} h(M_1) & 0 & \cdots & 0 \\ h(M_1+1) & h(M_1) & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ h(M_2) & h(M_2-1) & \cdots & h(M_2-N_2+N_1) \\ 0 & h(M_2) & \cdots & h(M_2-N_2+N_1+1) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h(M_2) \end{bmatrix} \begin{bmatrix} x(N_1) \\ x(N_1+1) \\ \vdots \\ x(N_2) \end{bmatrix} = \begin{bmatrix} y(M_1+N_1) \\ \vdots \\ y(M_2+N_2) \end{bmatrix} \quad (6-3)$$

<sup>73</sup> The sequence  $x(n)$  can be either a signal or a system's impulse response.

<sup>74</sup> The corresponding symbols for the sequences  $y(n)$  and  $h(n)$  are  $\underline{\mathbf{y}}$ ,  $\underline{\mathcal{Y}}^M$  and  $\underline{\mathbf{h}}$ ,  $\underline{\mathcal{H}}^M$  respectively. Similarly, to the time reversed sequence  $\hat{y}(n)$  we assign the symbols  $\hat{\mathbf{y}}$  and  $\hat{\mathcal{H}}^M$  and so on for the remaining symbols.

Now it follows directly from the above that the vector  $\underline{\hat{x}}$  corresponding to the  $N$ -samples-long time-reversed signal  $\hat{x}(n)$  can be obtained by reversing the rows of the vector  $\underline{x}$  corresponding to the original  $N$ -samples-long signal  $x(n)$

$$(\underline{x})^R = \underline{\hat{x}} \quad (6-4)$$

Similarly, the matrix  $\hat{\mathcal{H}}^N$  corresponding to the  $M$ -samples-long time-reversed impulse response  $\hat{h}(n)$  can be obtained by reversing both the rows and the columns of the matrix  $\mathcal{H}^N$

$$\left( (\mathcal{H}^N)^C \right)^R = \hat{\mathcal{H}}^N, \quad \text{for any given positive integer } N \quad (6-5)$$

Finally, we make use of the algebraic identity

$$\left( (\mathbf{A})^C \right)^R (\mathbf{v})^R = (\mathbf{A}\mathbf{v})^R \quad (6-6)$$

which holds for any matrix  $\mathbf{A}$  and vector  $\mathbf{v}$  of appropriate dimensions<sup>75</sup>.

Using equations (6-3), (6-4), (6-5) and (6-6) we obtain

$$\begin{aligned} \underline{\hat{y}} &= (\underline{y})^R \\ &= (\mathcal{H}^N \underline{x})^R \\ &= \left( (\mathcal{H}^N)^C \right)^R (\underline{x})^R \\ &= \hat{\mathcal{H}}^N \underline{\hat{x}} \end{aligned} \quad (6-7)$$

Equation (6-7) proves what was stated in the beginning of this section. That is, if  $y(n)$  is the convolution of  $h(n)$  with  $x(n)$  (i.e.  $\underline{y} = \mathcal{H}^N \underline{x}$ ) then convolving a time reversed version of  $x(n)$  with a time reversed version of  $h(n)$  results in a time reversed version of  $y(n)$  (i.e.  $\underline{\hat{y}} = \hat{\mathcal{H}}^N \underline{\hat{x}}$ ).

A very important note that has to be made here, however, is that the analysis above assumes the computation of the full-length convolution of the input  $x(n)$  with the

---

<sup>75</sup> A proof of the identity in equation (6-6) is given in the appendix (see §A2)

impulse response  $h(n)$ . Obviously this computation is not realisable when  $h(n)$  is of infinite duration, i.e. when it corresponds to a rational transfer function as is the case with the inverse filtering implementations discussed below. A formulation in which zero-padded finite length filtering is used to approximate the infinite length convolution is presented in the appendix (see §A1).

## 6.3 Decomposition of the inverse into minimum and maximum phase parts

Having seen how the convolution of a given input with a given impulse response can be implemented either in forward or in backward time, we describe in this section how a given transfer function can be decomposed into two parts, one realisable in forward time and one in backward time. We show that this decomposition can be effected in two ways, which we term the *parallel* and the *cascade* connection after a similar formulation presented in the context of linear phase filtering by Czarnach (1982). For notational simplicity we first describe the application of these two decompositions on the single-channel transfer function  $1/C(z)$ . Again we assume that the polynomial  $C(z)$  contains no pure delay and is factorised in first order terms as in equation (6-8).

$$C(z) = c_0 + c_1 z^{-1} + \dots + c_{N-1} z^{-N+1} = c_0 \prod_{i=1}^{N-1} (1 - p_i z^{-1}) \quad (6-8)$$

### 6.3.1 Parallel connection

Starting then with the rational expression  $1/C(z)$  we can take the partial fraction expansion<sup>76</sup> of equation (6-9) in a similar manner as was described above (see §2.3.1)

---

<sup>76</sup> As we also did in the analysis of §2.3.1, we assume in this section that the transfer function  $C(z)$  has only single order zeros and that none of its zeros is of magnitude equal to one. Even though the first assumption is not necessarily true we will see that the numerical method we propose for the decomposition of (6-9) (see §6.4) is formulated without such an assumption. As far as the second condition is concerned, the presence of a zero with magnitude equal to one is equivalent to a plant that completely removes the input's content at a certain frequency point. This cannot be true in cases where the plant is modelled with measured (and typically averaged) impulse responses that always exist on

$$\frac{1}{C(z)} = \frac{1}{c_0} \sum_{i=1}^{N-1} \frac{A_i}{1 - p_i z^{-1}} \quad (6-9)$$

We can then separate the, say  $L$ , terms for which  $|p_i| > 1$  and the, say  $K$ , terms for which  $|p_i| < 1$

$$\frac{1}{c_0} \sum_{i=1}^{N-1} \frac{A_i}{1 - p_i z^{-1}} = \frac{1}{c_0} \sum_{k=1}^K \frac{A_k}{1 - p_k z^{-1}} + \sum_{l=1}^L \frac{A_l}{1 - p_l z^{-1}}, \quad K + L = N - 1 \quad (6-10)$$

And finally regrouping these two groups of terms we obtain

$$\frac{1}{C(z)} = \frac{B_{min}(z)}{C_{min}(z)} + \frac{B_{max}(z)}{C_{max}(z)} \quad (6-11)$$

where the orders of the polynomials  $B_{min}(z)$ ,  $B_{max}(z)$ ,  $C_{min}(z)$  and  $C_{max}(z)$  are defined in equation (6-12). The  $K$  roots of  $C_{min}(z)$  are all of magnitude less than unity, while the  $L$  roots of  $C_{max}(z)$  are all of magnitude higher than unity. Hence  $C_{min}(z)$  is of strictly minimum phase and  $C_{max}(z)$  of strictly maximum phase. Note that this is not the case for  $B_{min}(z)$  and  $B_{max}(z)$  in which the subscripts *min* and *max* are used for notational convenience.

$$\begin{aligned} B_{min}(z) &= b_{min,0} + b_{min,1}z^{-1} + \dots + b_{min,K-1}z^{-K+1} \\ C_{min}(z) &= c_{min,0} + c_{min,1}z^{-1} + \dots + c_{min,K}z^{-K} \\ B_{max}(z) &= b_{max,0} + b_{max,1}z^{-1} + \dots + b_{max,L-1}z^{-L+1} \\ C_{max}(z) &= c_{max,0} + c_{max,1}z^{-1} + \dots + c_{max,L}z^{-L} \end{aligned} \quad (6-12)$$

Following the discussion in §2.3.1 we see that the term  $B_{min}(z)/C_{min}(z)$  in equation (6-9) corresponds to a causal and stable impulse response whereas the impulse response corresponding to the term  $B_{max}(z)/C_{max}(z)$  will either be unstable in forward time or stable in backward time. We denote with  $h_{left}(n)$  the left-sided<sup>77</sup>, anti-causal and stable impulse response given by

---

top of a noise floor. Furthermore, in a case where such a nullification of a certain frequency was effected by the plant, the inverse transfer function would have a pole at this frequency and its frequency response magnitude at this frequency would be infinite. The introduction of regularisation would then certainly constrain the magnitude of the inverse at this point by replacing the ill-conditioned pole with a pair of poles and a zero.

<sup>77</sup> For the terms right-sided sequence and left-sided sequence see (Oppenheim and Schafer 1975).

$$h_{left}(n) = Z^{-1} \left\{ \frac{B_{max}(z)}{C_{max}(z)} \right\}, \quad |z| < \min \{ |p_i| : C_{max}(p_i) = 0 \} \quad (6-13)$$

We now define the “reversed” polynomials  $\hat{B}_{max}(z)$  and  $\hat{C}_{max}(z)$  as

$$\begin{aligned} \hat{B}_{max}(z) &= b_{max,L-1} + b_{max,L-2}z^{-1} + \dots + b_{max,0}z^{-L+1} \\ \hat{C}_{max}(z) &= c_{max,L} + c_{max,L-1}z^{-1} + \dots + c_{max,0}z^{-L} \end{aligned} \quad (6-14)$$

and we note that these will satisfy the relation

$$\frac{B_{max}(z^{-1})}{C_{max}(z^{-1})} = \frac{z^{-1}\hat{B}_{max}(z)}{\hat{C}_{max}(z)} \quad (6-15)$$

i.e. the expression resulting by substituting  $z^{-1}$  in the place of the  $z$  variable in the ratio  $B_{max}(z)/C_{max}(z)$ , is equal to the expression  $z^{-1}\hat{B}_{max}(z)/\hat{C}_{max}(z)$ .

We then make use of the time-reversal property of the  $z$ -transform (Proakis and Manolakis 1992) according to which, if  $x(n)$  and  $X(z)$  are a  $z$ -transform pair then  $\hat{x}(n)$  and  $X(z^{-1})$  are also a transform pair. It thus follows that the right-handed, stable and causal impulse response  $h_{right}(n)$  of equation (6-16) will satisfy equation (6-17).

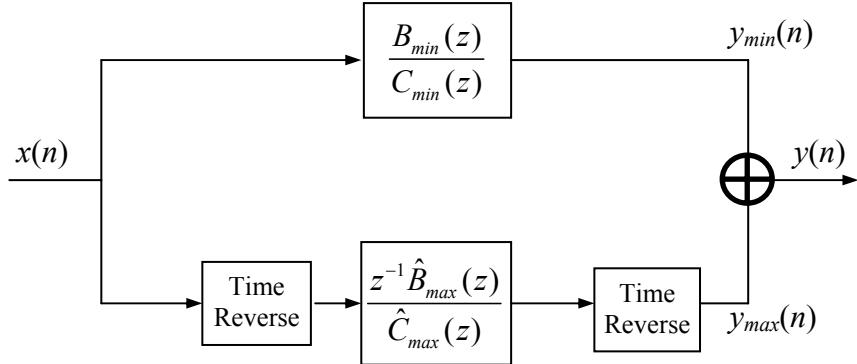
$$h_{right}(n) = Z^{-1} \left\{ \frac{z^{-1}\hat{B}_{max}}{\hat{C}_{max}} \right\}, \quad |z| > \max \{ |p_i| : \hat{C}_{max}(p_i) = 0 \} \quad (6-16)$$

$$h_{right}(n) = h_{left}(-n) \quad \text{for all } n \quad (6-17)$$

In other words the transfer function  $z^{-1}\hat{B}_{max}(z)/\hat{C}_{max}(z)$  will correspond to a realisable impulse response that is the exactly symmetrical around the time-origin to the impulse response corresponding to  $B_{max}(z)/C_{max}(z)$ .

Hence, upon determination of the polynomials  $B_{min}(z)$ ,  $\hat{B}_{max}(z)$ ,  $C_{min}(z)$  and  $\hat{C}_{max}(z)$  of equations (6-12) and (6-14) and with reference to what was said in §6.2, the filtering through the inverse transfer function  $1/C(z)$  can be implemented as is described in

figure 6-1<sup>78</sup> from which the reason for the term “parallel connection” becomes evident.



**Figure 6-1:** Inverse filtering (parallel connection). The thick-line addition symbol in the figure denotes the shift-add process.

### 6.3.2 Cascade connection

A second way of decomposing the inverse transfer function  $1/C(z)$  into two parts, one realisable in forward time and the other in backward time, is as follows. Starting again from the factorisation of  $C(z)$  into first order terms as in equation (6-8) we see that the reciprocal transfer function  $1/C(z)$  can be written as the product of two terms as in equation (6-18). The polynomials  $C_{min}(z)$  and  $C_{max}(z)$  are the same as in equation (6-12) so again the term  $1/C_{min}(z)$  will correspond to a causal and stable impulse response and the term  $1/C_{max}(z)$  to an anti-causal and stable impulse response.

$$\frac{1}{C(z)} = \frac{1}{c_0 \prod_{i=1}^{N-1} (1 - p_i z^{-1})} = \frac{1}{c_0 \prod_{k=1}^K (1 - p_k z^{-1})} \frac{1}{\prod_{l=1}^L (1 - p_l z^{-1})} = \frac{1}{C_{min}(z)} \frac{1}{C_{max}(z)} \quad (6-18)$$

In the same manner as in the previous section we see that equation (6-19) will hold.

$$\frac{1}{C_{max}(z^{-1})} = \frac{z^{-L}}{\hat{C}_{max}(z)} \quad (6-19)$$

<sup>78</sup> The thick-line addition symbol in the figure denotes the shift-add process (explained in detail in the appendix, see §A1) in which samples of  $y_{min}(n)$  and  $y_{max}(n)$  with the same time-index are added to produce the (longer) output  $y(n)$ .

Thus the left-handed sequence

$$h_{left}(n) = Z^{-1} \left\{ \frac{1}{C_{max}(z)} \right\}, \quad z < \min \{ |p_i| : C_{max}(p_i) = 0 \} \quad (6-20)$$

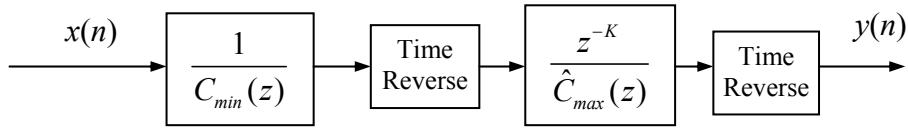
and the right-handed sequence

$$h_{right}(n) = Z^{-1} \left\{ \frac{z^{-L}}{\hat{C}_{max}} \right\}, \quad z > \max \{ |p_i| : \hat{C}_{max}(p_i) = 0 \} \quad (6-21)$$

will be a mirrored version of each other.

$$h_{right}(n) = h_{left}(-n) \quad \text{for all } n \quad (6-22)$$

The filtering through the inverse transfer function of equation (6-18) can thus be implemented by means of the “cascade connection” of figure 6-2.



**Figure 6-2:** Inverse filtering (cascade connection)

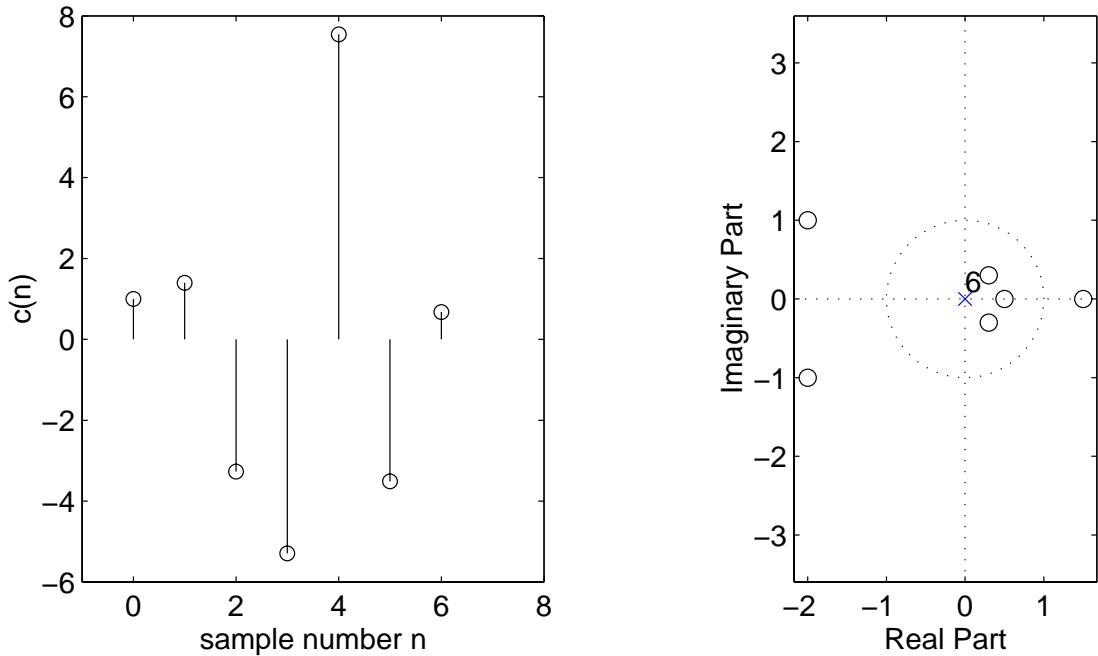
### 6.3.3 Computation of the decomposition determining the poles and residues of the inverse transfer function

As an example of the parallel and the cascade connection decompositions described above, we consider here the 6<sup>th</sup> order mixed-phase plant of equation (6-23) the impulse response  $c(n)$  together with the zero locations of which are shown in figure 6-3.

$$C(z) = 1 + 1.4z^{-1} - 3.27z^{-2} - 5.29z^{-3} + 7.545z^{-4} - 3.51z^{-5} + 0.675z^{-6} \quad (6-23)$$

Using then the `residuez` function included in Matlab’s Signal Processing Toolbox (MathWorks 2000) we can compute the roots  $p_i$  and residues  $A_i$  of equation (6-9)

listed in table 6-1 which, in turn, are used for the determination of the polynomials  $B_{min}(z)$ ,  $B_{max}(z)$ ,  $\hat{B}_{max}(z)$ ,  $C_{min}(z)$ ,  $C_{max}(z)$  and  $\hat{C}_{max}(z)$  of equations (6-12) and (6-14).

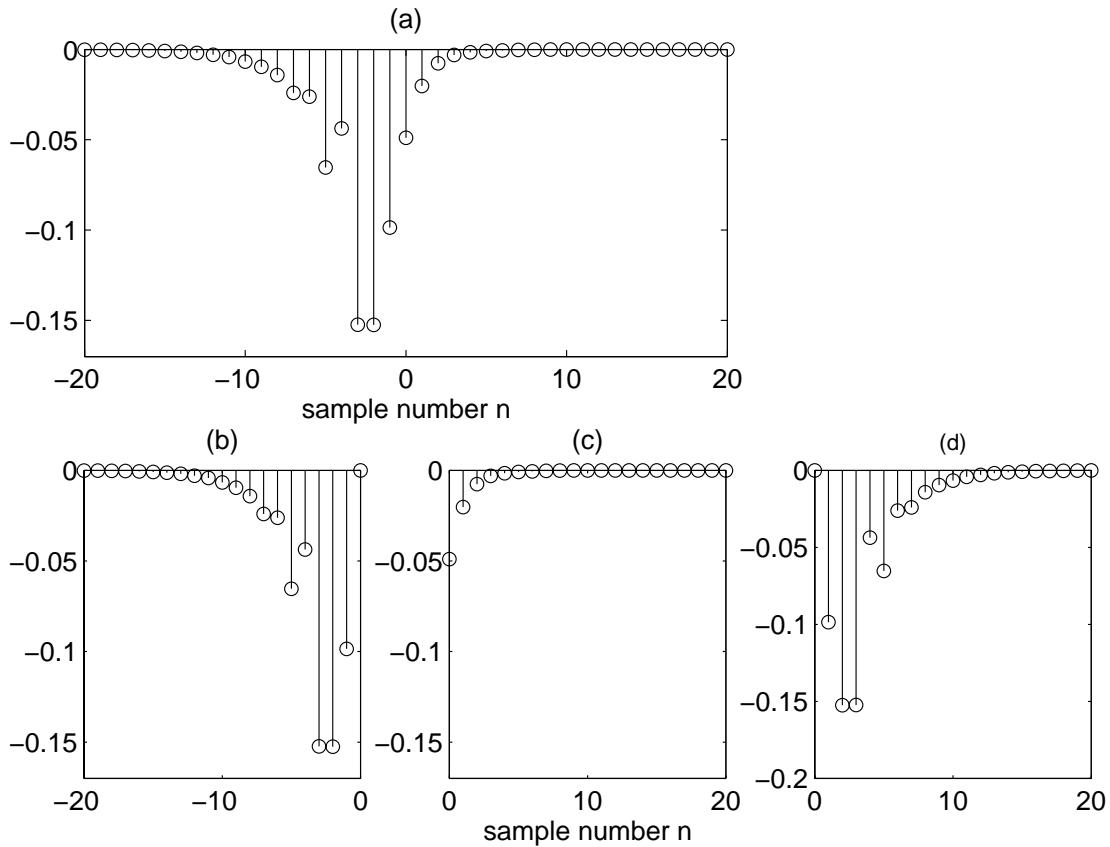


**Figure 6-3:** Impulse response and zero-map of the plant of equation (6-23)

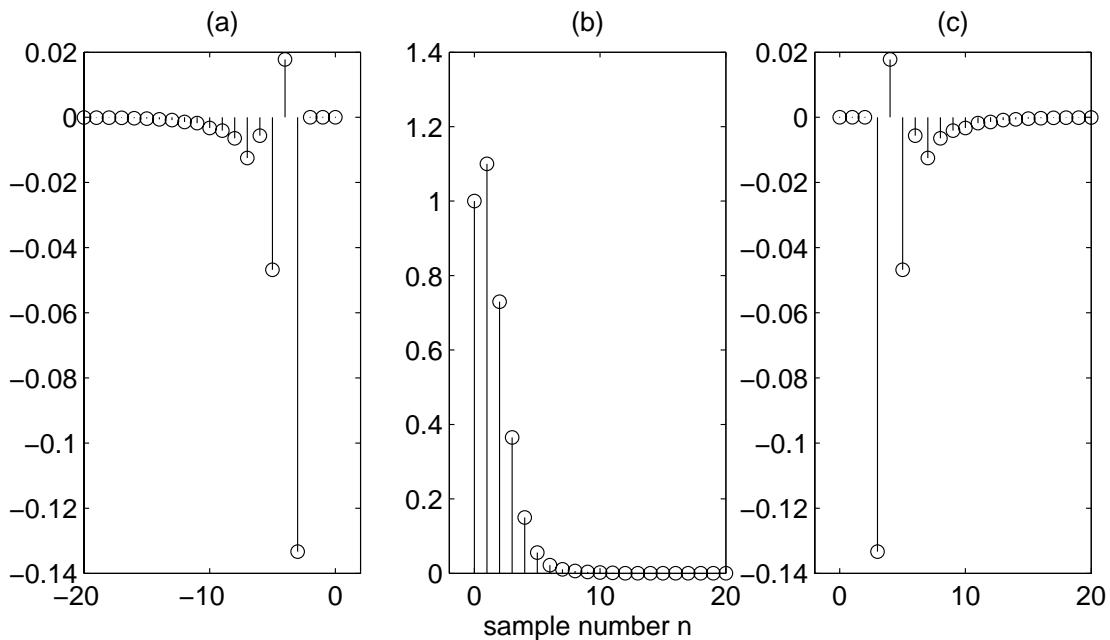
$i$	$p_i$	$A_i$
1	$-2+i$	$0.3372+0.2965i$
2	$-2-i$	$0.3372-0.2965i$
3	$1.5$	$0.3746$
4	$0.5$	$-0.0332$
5	$0.3+0.3i$	$-0.0079-0.0018i$
6	$0.3-0.3i$	$-0.0079+0.0018i$

**Table 6-1:** Roots and residues of the inverse transfer function of equation (6-23)

The 41-samples-long FIR approximation to  $1/C(z)$  with the modelling delay set to 20 samples will then be as in part (a) of figure 6-4, while the 20 first samples of the impulse responses corresponding to the rational transfer functions  $B_{min}(z)/C_{min}(z)$ ,  $B_{max}(z)/C_{max}(z)$  and  $z^{-1}\hat{B}_{max}(z)/\hat{C}_{max}(z)$  of the parallel connection implementation will be as in parts (b), (c) and (d) of the same figure. Similarly, in parts (a), (b) and (c) of figure 6-5 we plot the first 20 samples of the impulse responses corresponding to the rational transfer functions  $1/C_{min}(z)$ ,  $1/C_{max}(z)$  and  $z^{-3}\hat{C}_{max}(z)$  of the cascade connection implementation.



**Figure 6-4:** (a) Impulse response of the FIR approximation to the inverse, (b), (c) and (d) First 20 samples of the impulse responses of  $B_{max}(z)/C_{max}(z)$ ,  $B_{min}(z)/C_{min}(z)$  and  $z^{-1}B_{max}(z)/C_{max}(z)$



**Figure 6-5:** (a)-(c) First 20 samples of the impulse responses of  $1/C_{max}(z)$ ,  $1/C_{min}(z)$  and  $z^{-3}C_{max}(z)$

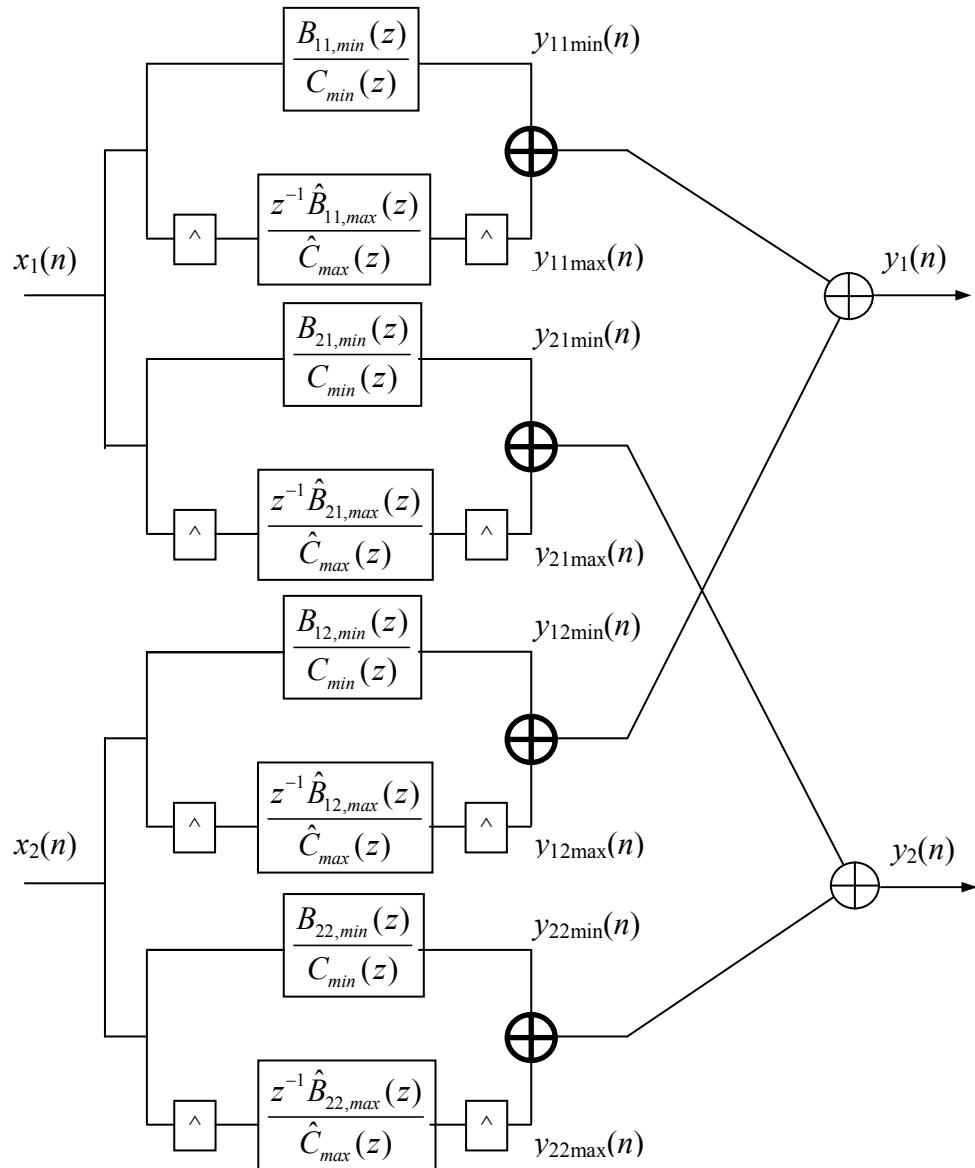
### 6.3.4 Two-channel case

Concluding the basic formulation of the proposed inverse filtering method we give here the decomposition formulae corresponding to the two-channel case where the plant is modelled by the  $2 \times 2$  matrix  $\mathbf{C}(z)$  containing four transfer functions  $C_{ij}(z)$  and the ideal  $2 \times 2$  inverse matrix  $\mathbf{H}(z)$  is as in equation (2-14). Combining the direct and common denominator arrangements of §2.4.2 and the parallel and the cascade connection arrangements of §6.3.1 and §6.3.2 we see that the filtering through the inverse matrix  $\mathbf{H}(z)$  can be implemented in four ways depicted in the block diagrams of the figures 6-6 – 6-8. It is easy to see that that the formulae for the determination of the polynomials appearing in these block diagrams will be as in equations (6-24) for the direct arrangement (parallel and cascade connection) cases and equations (6-25) for the common denominator arrangement (parallel and cascade connection) cases with the minimum and maximum phase polynomials  $C_{min}(z)$  and  $C_{max}(z)$  being as in equation (6-26).

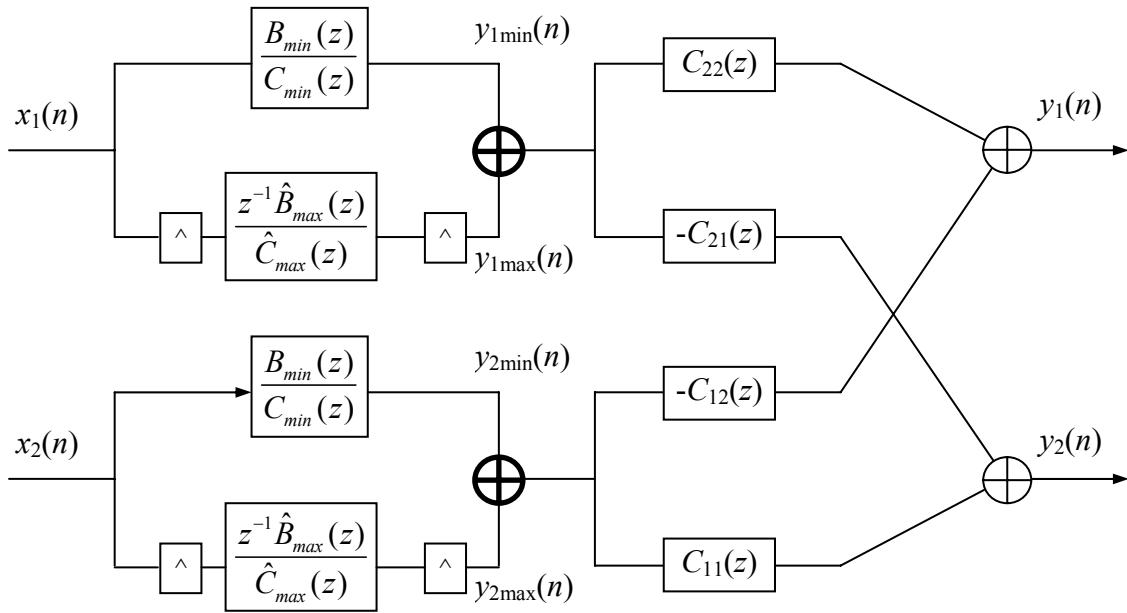
$$\begin{aligned} H_{11}(z) &= \frac{C_{22}(z)}{\det[\mathbf{C}(z)]} = \frac{B_{11,min}(z)}{C_{min}(z)} + \frac{B_{11,max}(z)}{C_{max}(z)} = \frac{C_{22}(z)}{C_{min}(z) C_{max}(z)} \frac{1}{C_{min}(z) C_{max}(z)} \\ H_{12}(z) &= \frac{-C_{12}(z)}{\det[\mathbf{C}(z)]} = \frac{B_{12,min}(z)}{C_{min}(z)} + \frac{B_{12,max}(z)}{C_{max}(z)} = \frac{-C_{12}(z)}{C_{min}(z) C_{max}(z)} \frac{1}{C_{min}(z) C_{max}(z)} \\ H_{21}(z) &= \frac{-C_{21}(z)}{\det[\mathbf{C}(z)]} = \frac{B_{21,min}(z)}{C_{min}(z)} + \frac{B_{21,max}(z)}{C_{max}(z)} = \frac{-C_{12}(z)}{C_{min}(z) C_{max}(z)} \frac{1}{C_{min}(z) C_{max}(z)} \\ H_{22}(z) &= \frac{C_{11}(z)}{\det[\mathbf{C}(z)]} = \frac{B_{22,min}(z)}{C_{min}(z)} + \frac{B_{22,max}(z)}{C_{max}(z)} = \frac{C_{11}(z)}{C_{min}(z) C_{max}(z)} \frac{1}{C_{min}(z) C_{max}(z)} \end{aligned} \quad (6-24)$$

$$\frac{1}{\det[\mathbf{C}(z)]} = \frac{B_{min}(z)}{C_{min}(z)} + \frac{B_{max}(z)}{C_{max}(z)} = \frac{1}{C_{min}(z)} \frac{1}{C_{max}(z)} \quad (6-25)$$

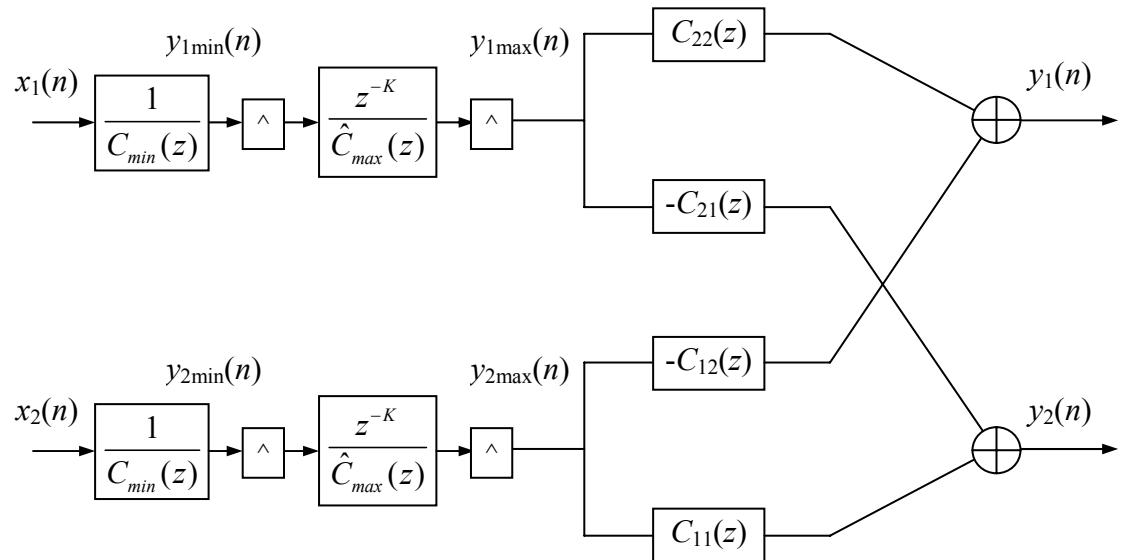
$$\det[\mathbf{C}(z)] = C_{min}(z) C_{max}(z) \quad (6-26)$$



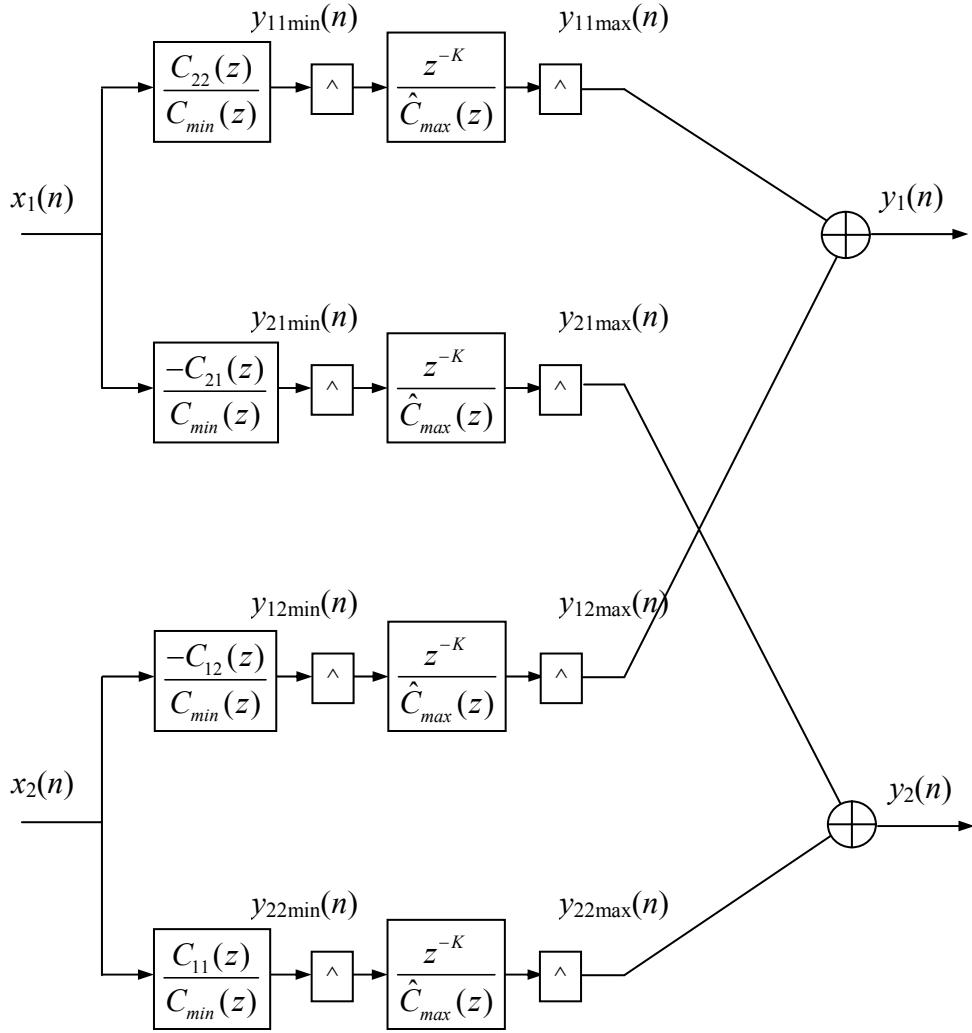
**Figure 6-6:** Parallel connection direct filtering arrangement



**Figure 6-7:** Parallel connection common denominator arrangement



**Figure 6-8:** Cascade connection common denominator arrangement



**Figure 6-9:** Cascade connection direct filtering arrangement

We denote the computational costs in the four cases of the parallel connection direct arrangement, parallel connection common denominator arrangement, cascade connection direct arrangement and cascade connection common denominator arrangement with  $\mathcal{O}_{PDI}$ ,  $\mathcal{O}_{PCD}$ ,  $\mathcal{O}_{CDI}$  and  $\mathcal{O}_{CCD}$  respectively. These denote the number of multiply-add operations needed for the computation of each output sample (madds/output sample) and will be as in equations (6-27)-(6-30). In these equations, with  $N_C$  we denote the order of the chosen FIR model of the plant transfer functions  $C_{ij}(z)$  and we also make use the equations  $N_{Cmin} + N_{Cmax} = 2N_C$ , and  $N_{Bij,min} + N_{Bij,max} = N_{Bmin} + N_{Bmax} = 2N_C - 2$  (obvious from the formulation of §§6.3.1-6.3.2) which relate the orders of the polynomials  $C_{min}(z)$ ,  $C_{max}(z)$ ,  $B_{min}(z)$ ,  $B_{max}(z)$ ,  $B_{ij,min}(z)$  and  $B_{ij,max}(z)$  to the order  $N_C$  of each transfer function in the plant matrix.

$$\mathcal{O}_{PDI} = 4(2N_c + 2N_c - 2) \approx 16N_c \quad \text{madds/output sample} \quad (6-27)$$

$$\mathcal{O}_{PCD} = 2(2N_c + 2N_c - 2) + 4N_c \approx 12N_c \quad \text{madds/output sample} \quad (6-28)$$

$$\mathcal{O}_{CDI} = 4(2N_c + N_c) = 12N_c \quad \text{madds/output sample} \quad (6-29)$$

$$\mathcal{O}_{CCD} = 2 \cdot 2N_c + 4N_c = 8N_c \quad \text{madds/output sample} \quad (6-30)$$

## 6.4 Numerical method for the decomposition of the inverse transfer function

In the previous section we showed how a given rational transfer function  $B(z)/C(z)$  can be decomposed into a sum or a product of two rational transfer functions one of which can be realised in forward and the other in backward time and we gave an example of this decomposition based on the use of Matlab's `residuez` function. However, a computation like that described in §6.3.3 cannot be relied upon for the decomposition of the inverse transfer functions resulting when the plant's model comprises a few hundreds of coefficients as it presupposes a highly sensitive numerical computation for the determination of the roots of the polynomial  $C(z)$ <sup>79</sup>.

In this section we present an alternative method for the determination of the polynomials of equations (6-24)-(6-26). First, in §6.4.1, we show how the minimum and maximum phase polynomials  $C_{min}(z)$  and  $C_{max}(z)$  of equation (6-26) can be determined by means of computing the cepstrum of the polynomial  $\det[\mathbf{C}(z)]$ . As is evident from equations (6-24) and (6-25), the accuracy of that computation solely determines the accuracy of the implementation of the cascade connection (both for the direct and the common denominator arrangements) while for the implementation of the parallel connection one has to also determine the polynomials  $B_{ij,min}(z)$  and

---

<sup>79</sup> The `residuez` function makes use of the `roots` function for the determination of the roots of  $C(z)$ . For the details of the specific implementation see (Mathworks 2000). Methods that promise better accuracy of the numerical solution of the root-computation problem for high-order polynomials have been published recently (see (Sitton et al. 2003) and references therein). These methods however are tailored to the case of polynomials of much higher order than those encountered here and are thus not considered in this work.

$B_{ij,max}(z)$  of equation (6-24) (or the polynomials  $B_{min}(z)$  and  $B_{max}(z)$  of equation (6-25) in the case where the common denominator filtering arrangement is used). We show in §6.4.2 how this can be achieved by means of solving the equivalent Diophantine equation resulting in each case.

#### 6.4.1 Determination of $C_{min}(z)$ and $C_{max}(z)$

A way to avoid computing the roots of  $C(z)$  in the determination of the polynomials  $C_{min}(z)$  and  $C_{max}(z)$  is to solve equation (6-26) in the cepstral domain. Denoting with  $c_{cep}(n)$  the cepstrum of a given time-sequence  $c(n)$  as in equation (6-31)

$$c_{cep}(m) = Z^{-1} \{ \log [C(z)] \}, \quad -\infty < m < \infty \quad (6-31)$$

the minimum and maximum phase polynomials  $C_{min}(z)$  and  $C_{max}(z)$  satisfying the equation  $C(z) = C_{min}(z)C_{max}(z)$  will be given by the following<sup>80</sup>.

$$\begin{aligned} C_{max}(z) &= \exp \left[ Z \{ c_{cep}(m) \} \right], \quad m \leq 0 \\ C_{min}(z) &= \exp \left[ Z \{ c_{cep}(m) \} \right], \quad m \geq 0 \end{aligned} \quad (6-32)$$

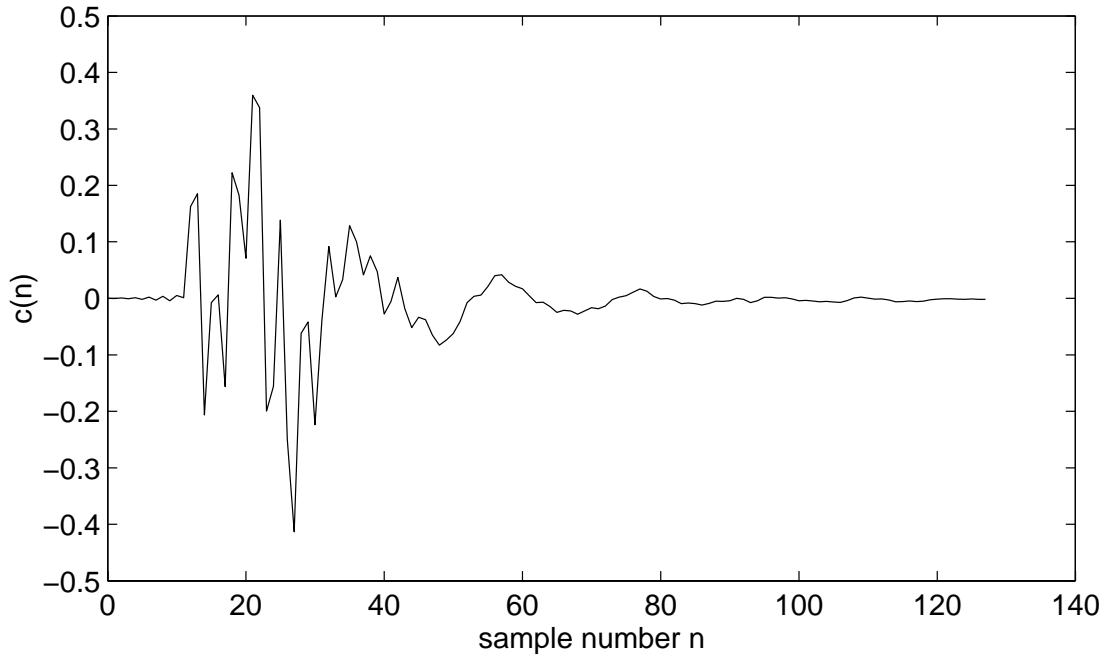
The computation of the complex cepstrum of a given time series and of the time series corresponding to the inverse of a given complex cepstrum are implemented by `cceps` and `icceps` Matlab's functions<sup>81</sup> which we used for the results presented in this section. The implementation of these Matlab functions directly follows the relevant discussion in (Oppenheim and Schafer 1975) where a method for the calculation of  $c_{cep}(n)$  based on the computation of the DFT  $C(k)$  of  $c(n)$  is presented. As explained in (Oppenheim and Schafer 1975), the critical point in the computation of the cepstrum sequence  $c_{cep}(n)$  is the proper unwrapping of the phase of  $C(k)$  in a continuous function, an issue that can be successfully addressed if the length (and hence the frequency resolution) of the corresponding DFT is made large enough. Additionally, an adequately long DFT guarantees that the time-aliasing in the computation of the inverse DFT  $c_{cep}(n) = \text{IDFT}\{\log[C(k)]\}$  becomes insignificant.

---

<sup>80</sup> For a detailed discussion of the related theory of homomorphic signal processing see (Oppenheim and Schafer 1975) and (Oppenheim and Tribollet 1976).

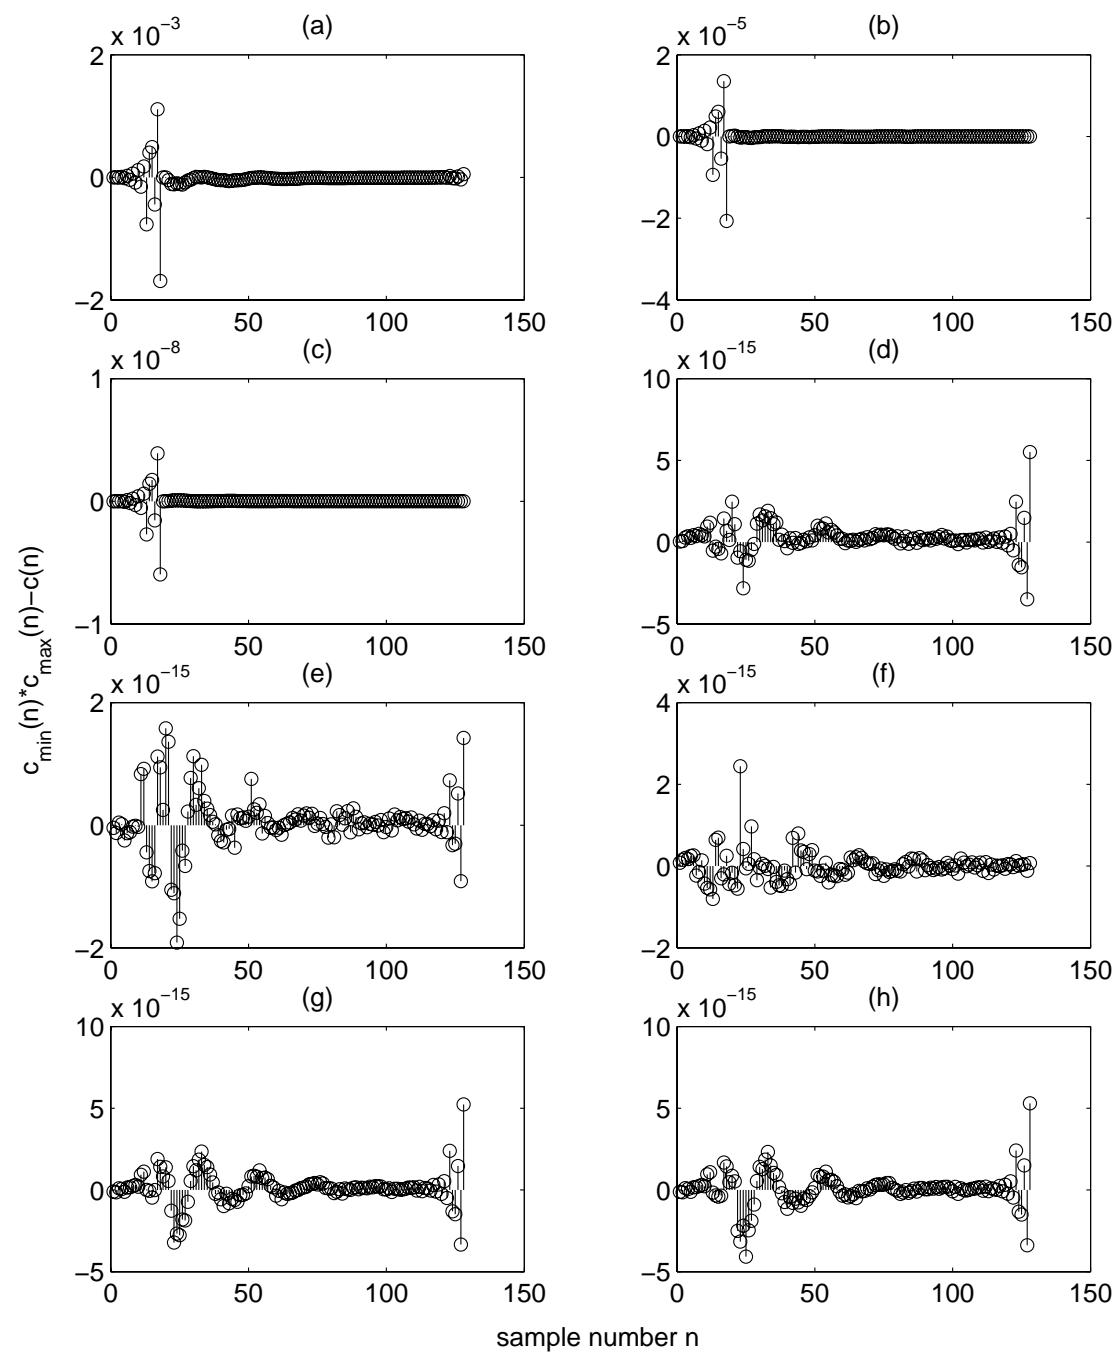
<sup>81</sup> For the details concerning the implementation of `cceps` and `icceps` see (Mathworks 2000)

As an example of the cepstrum-based determination of the minimum and maximum phase parts of a given impulse response, we apply the method to the decomposition of a 128-samples-long HRIR  $c(n)$  from the MIT-database (Gardner and Martin 1994). The HRIR used corresponds to a source directly in front of the listener and is plotted in figure 6-10.

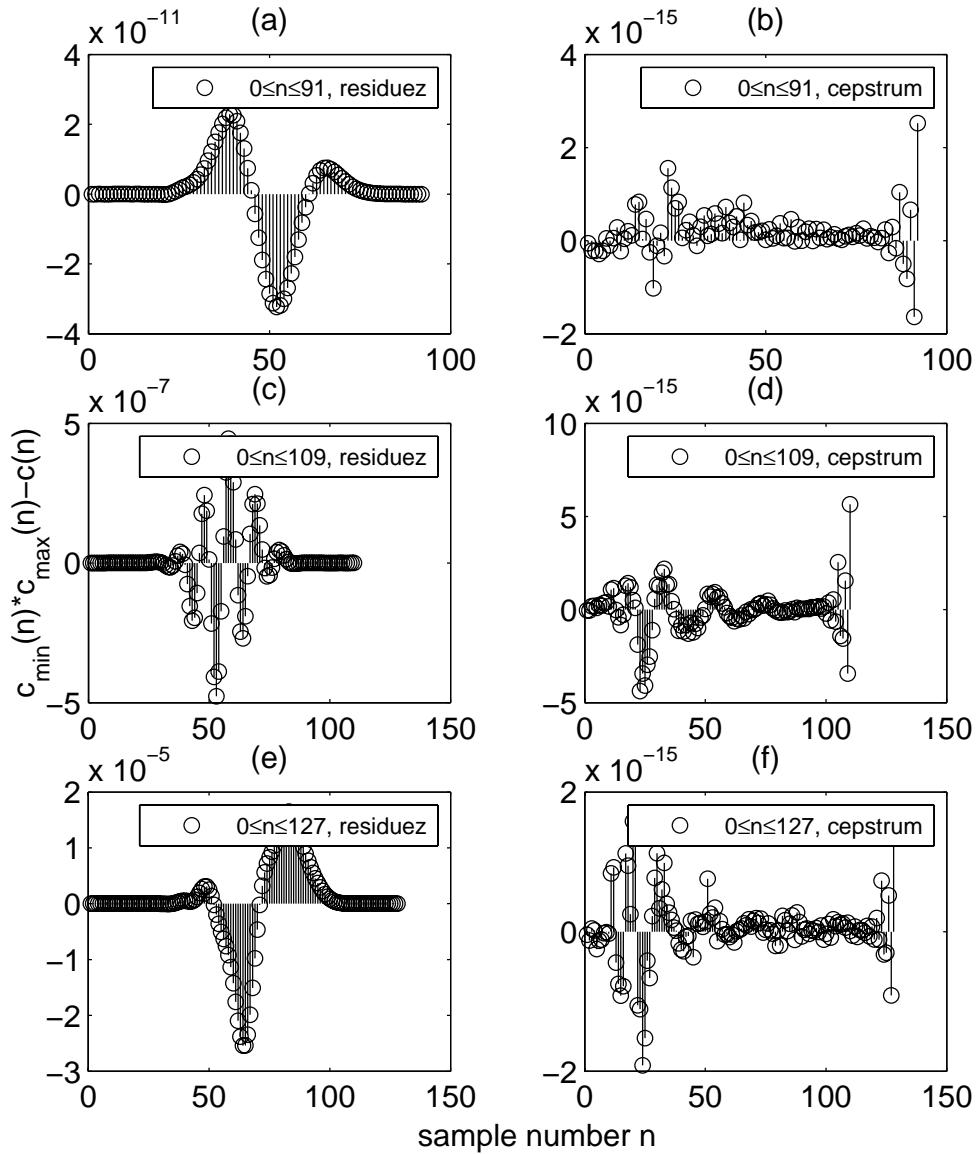


**Figure 6-10:** MIT-database HRIR corresponding to a source directly in front of the listener

The effect of the DFT length on the accuracy of the cepstrum-based computation is demonstrated in figure 6-11 where we have plotted the computation error  $c(n) - c_{min}(n)*c_{max}(n)$  achieved with DFT lengths ranging from  $2^{10}$  points in part (a) up to  $2^{17}$  points in part (h). As can be seen in the figure the accuracy of the solution increases as the length of the computed DFT goes up from  $2^{10}$  points to  $2^{13}$  points. As we go further up in the DFT length the error stays roughly constant so, evidently, both the phase unwrapping and time-aliasing issues discussed above are being properly addressed when a DFT length of  $2^{13}$  points or higher is used.



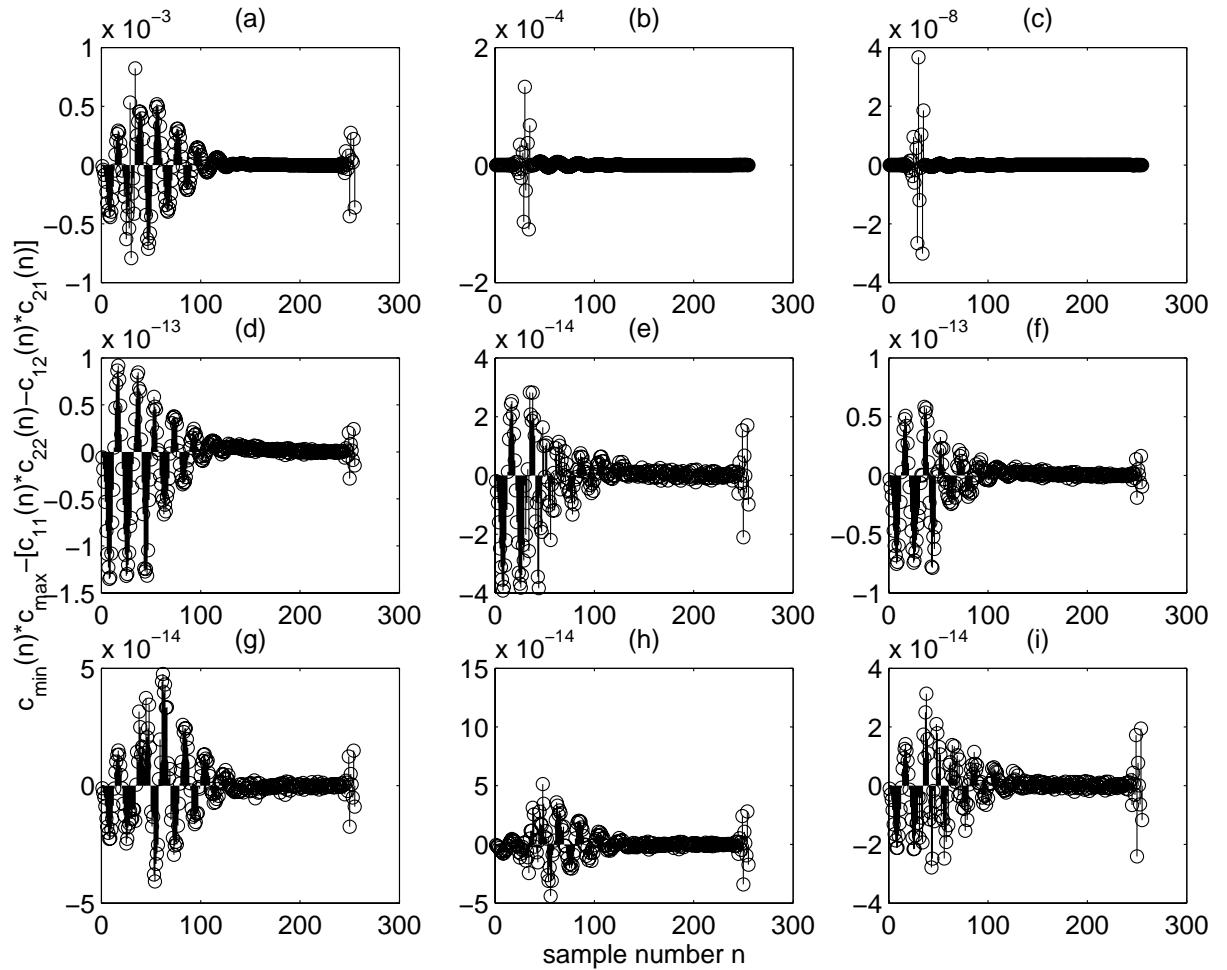
**Figure 6-11:** (a)-(h) Precision of the cepstrum-based decomposition for DFT lengths ranging from  $2^{10}$  points to  $2^{17}$  points



**Figure 6-12:** (a), (c), (e) precision of the roots-based decomposition, (b), (d), (f) precision of the cepstrum-based decomposition with the DFT length set to  $2^{14}$  points.

When, on the other hand, the `roots` Matlab function is used for the determination of  $c_{min}(n)$  and  $c_{max}(n)$  the results are far inferior. This is demonstrated in figure 6-12 where we compare the accuracy of the cepstrum-based and the roots-based solutions when the decomposition is applied to the full impulse response  $c(0 \leq n \leq 127)$  and two shorter versions  $c(0 \leq n \leq 91)$  and  $c(0 \leq n \leq 109)$ . As can be seen in the figure the roots-based method works very well for lower order plants but as the order increases the performance drops dramatically. The accuracy in the determination of  $c_{min}(n)$  and

$c_{max}(n)$  is of course of extreme importance in the implementation of the method. This is because a small error in the coefficients of the polynomials  $C_{min}(z)$  and  $C_{max}(z)$  that are placed in the taps of the recursive loop of the resulting filters will inevitably result in increased error in the output of the filters. Even more important, unlike the cepstrum-based method where the computed impulse responses  $c_{min}(n)$  and  $c_{max}(n)$  are guaranteed to be of strictly minimum and maximum phase respectively as long as the cepstrum computation length is chosen to be sufficiently high, the roots-based solution can easily diverge to give mixed phase impulse responses. Such an error is obviously detrimental as the reciprocal of such mixed phase responses is unrealisable in both forward and backward time.



**Figure 6-13:** (a)-(i) Accuracy of the cepstrum-based decomposition of  $\det[\mathbf{C}(z)]$  for DFT lengths ranging from  $2^{10}$  points to  $2^{18}$  points.

Similar results as above are obtained in the two-channel case where, as was discussed in §6.3.4, the decomposition into minimum and maximum phase parts is applied to

the transfer function  $\det[\mathbf{C}(z)]^{82}$ . The accuracy of the decomposition is depicted in figure 6-13 where it can be seen that, as expected, like for like accuracy with the single-channel case of figure 6-11 is achieved here for higher values of the DFT length with the highest accuracy being achieved for a DFT length of  $2^{14}$  points or more.

### 6.4.2 Determination of $B_{min}(z)$ and $B_{max}(z)$

Having computed the polynomials  $C_{min}(z)$  and  $C_{max}(z)$  as described in the previous section, the calculation of the transfer functions related to the parallel connection decomposition reduces to the solution of the four equations (6-24) in the unknown polynomials  $B_{ij,min}(z)$  and  $B_{ij,max}(z)$  or equation (6-25) in the unknowns  $B_{min}(z)$  and  $B_{max}(z)$  in the case where the common denominator is implemented separately<sup>83</sup>.

It is now easy to see that equation (6-25) can be written as

$$B_{min}(z)C_{max}(z) + B_{max}(z)C_{min}(z) = \frac{C_{min}(z)C_{max}(z)}{\det[\mathbf{C}(z)]} \quad (6-33)$$

which using equation (6-26) takes the form of the Diophantine equation

$$B_{min}(z)C_{max}(z) + B_{max}(z)C_{min}(z) = 1 \quad (6-34)$$

Similarly, equations (6-24) can be transformed to the Diophantine form

$$\begin{aligned} B_{11,min}(z)C_{max}(z) + B_{11,max}(z)C_{min}(z) &= C_{22}(z) \\ B_{12,min}(z)C_{max}(z) + B_{12,max}(z)C_{min}(z) &= -C_{12}(z) \\ B_{21,min}(z)C_{max}(z) + B_{21,max}(z)C_{min}(z) &= -C_{21}(z) \\ B_{22,min}(z)C_{max}(z) + B_{22,max}(z)C_{min}(z) &= C_{11}(z) \end{aligned} \quad (6-35)$$

In general then both equations (6-24) and (6-25) can be written in the general Diophantine equation form

---

<sup>82</sup> For the results shown here we use the matrix  $\mathbf{C}(z)$  containing the MIT-database HRTFs corresponding to  $\pm 5^0$  azimuth and  $0^0$  degrees elevation (Stereo Dipole format). The impulse response and the root distribution corresponding to the polynomial  $\det[\mathbf{C}(z)]$  are those already plotted in figures 2-2 and 2-3 respectively.

<sup>83</sup> As was already noted above (see p. 188), only the determination of the  $C_{min}(z)$  and  $C_{max}(z)$  polynomials is necessary for the implementation of the cascade connection arrangement.

$$P(z)X(z) + Q(z)Y(z) = R(z) \quad (6-36)$$

where the known polynomials  $P(z)$  and  $Q(z)$  are equal to  $C_{min}(z)$  and  $C_{max}(z)$ , the polynomials  $X(z)$  and  $Y(z)$  are the unknowns  $B_{ij,min}(z)$  and  $B_{ij,max}(z)$  for the direct filtering arrangement or the unknowns  $B_{min}(z)$  and  $B_{max}(z)$  for the common denominator arrangement and the polynomial  $R(z)$  in the right-hand side is equal to one of  $\pm C_{ij}(z)$  for the direct filtering arrangement or equal to 1 for the common denominator arrangement.

A detailed exposition of the properties of the general form the Diophantine equation (6-36) can be found in (Kučera 1979) where it is shown that a solution in the unknowns  $X(z)$  and  $Y(z)$  exists if and only if the greatest common divisor  $D(z)$  of the known polynomials  $P(z)$  and  $Q(z)$  is a divisor of the right-hand side polynomial  $R(z)$ . Furthermore, with  $P'(z)$  and  $Q'(z)$  defined as in equation (6-37) and  $S(z)$  any polynomial, it is shown in (Kučera 1979) that if  $X_0(z)$  and  $Y_0(z)$  is a solution then the pair  $X(z)=X_0(z)+Q'(z)S(z)$  and  $Y(z)=Y_0(z)+P'(z)S(z)$  is also a solution of (6-36). In the specific case with which we are concerned here, a solution is then guaranteed to exist given that the polynomials  $C_{min}(z)$  and  $C_{max}(z)$  share no common roots<sup>84</sup> and thus their greatest common divisor is equal to one.

$$\begin{aligned} P'(z) &= \frac{P(z)}{D(z)} \\ Q'(z) &= \frac{Q(z)}{D(z)} \end{aligned} \quad (6-37)$$

The available numerical methods for the solution of Diophantine equations are reviewed in (Kučera et al. 1991) where they are categorised in the three families of matrix, polynomial and mixed methods. A numerical analysis is also presented in (Kučera et al. 1991) of the “best-known representatives” of the three categories, respectively the *indeterminate coefficients*, the *polynomial reduction* and the *state-space realisation* methods. The conclusion is that the method of indeterminate coefficients is the most reliable in the ill-conditioned case where the polynomials  $P(z)$  and  $Q(z)$  have roots close to each other but not close to a root of  $R(z)$ . Hence we used this method for the computation of the results presented below.

---

<sup>84</sup> This is because one of them is of strictly minimum phase and the other of strictly maximum phase.

In simple terms, the method of indeterminate coefficients amounts to arranging the coefficients of the polynomials  $P(z)$  and  $Q(z)$  in a matrix  $M$ <sup>85</sup>, the coefficients of  $R(z)$  in a vector  $N$  and the coefficients the unknown polynomials  $X(z)$  and  $Y(z)$  in a vector  $K$  so that (6-36) becomes equivalent to the matrix-vector equation  $MK=N$  and solving for the unknown vector  $K$  (for a more detailed presentation of the method's formulation see (Kučera 1979)).

Again, in order to compare the accuracy of the method with the accuracy achieved when the partial fraction expansion is used for the computation of the decomposition, we apply it to the single-channel transfer function  $C(z)$  of figure 6-10. That is, we use the minimum and maximum phase polynomials  $C_{min}(z)$  and  $C_{max}(z)$  computed as was explained in §6.4.1<sup>86</sup> to solve equation (6-38) that is equivalent to the decomposition  $B_{min}(z)/C_{min}(z)+B_{max}(z)/C_{max}(z)=1/C(z)$ <sup>87</sup>.

$$C_{max}(z)B_{min}(z) + C_{min}(z)B_{max}(z) = 1 \quad (6-38)$$

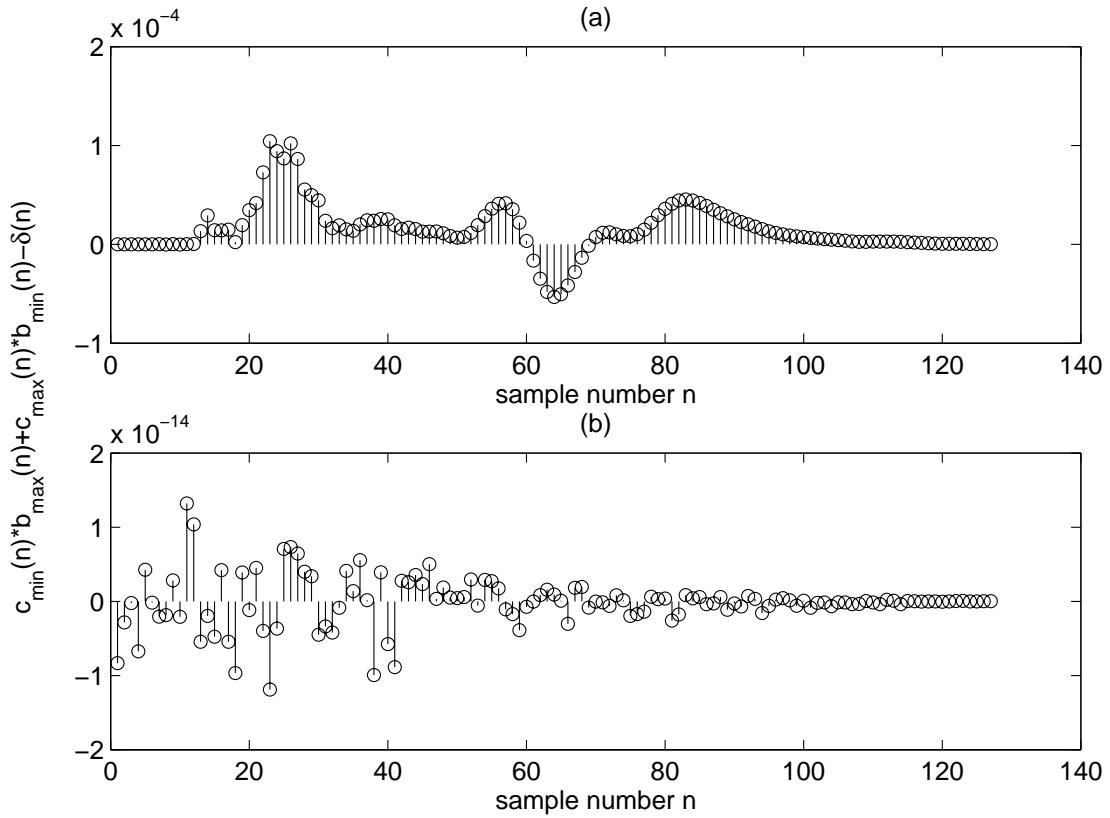
The corresponding error quantity  $c_{min}(n)*b_{max}(n)+c_{max}(n)*b_{min}(n)-\delta(n)$  is plotted in part (b) of figure 6-14 where it can be seen to be greatly superior compared to the same error quantity -plotted in part (a) of figure 6-14- achieved when the decomposition is effected by computing the partial fraction expansion of  $1/C(z)$  (using Matlab's `residuez` function).

---

<sup>85</sup> This matrix will then have the specific form a Sylvester matrix (Åström and Wittenmark 1997)

<sup>86</sup> The solution for  $C_{min}(z)$  and  $C_{max}(z)$  is the one obtained with the DFT length used for the transition from the cepstral domain and back set to  $2^{14}$  points.

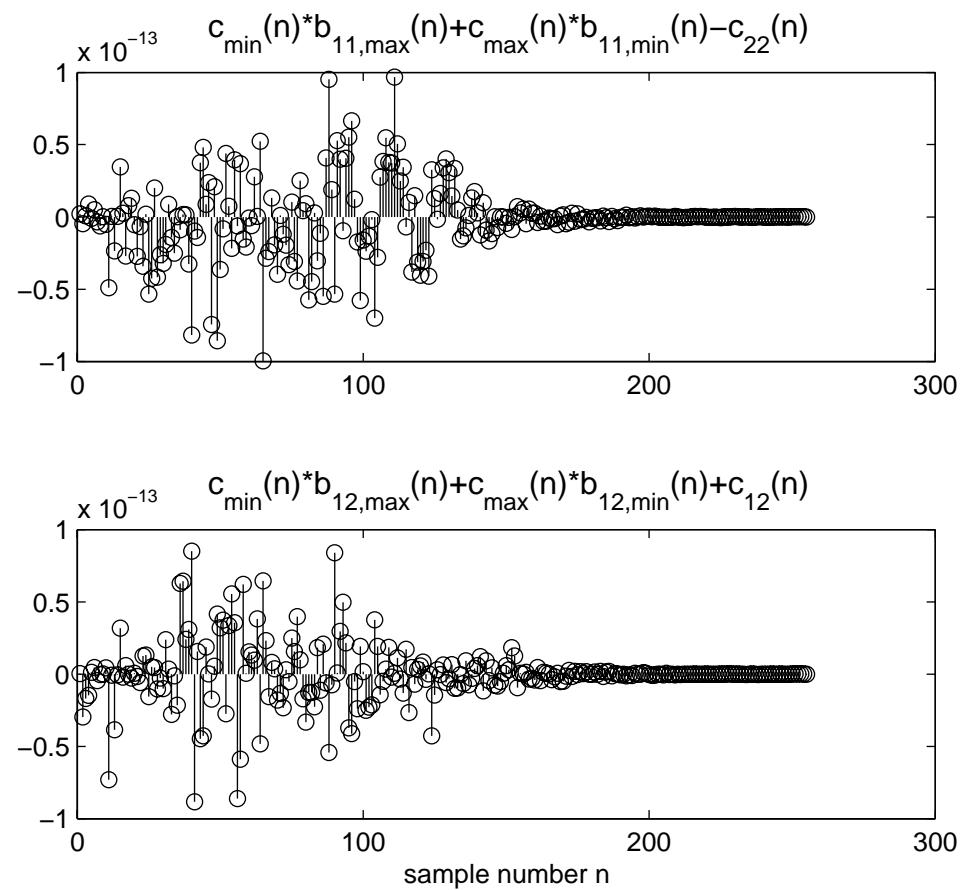
<sup>87</sup> It is interesting to note the similarity in form between this formulation and the MINT formulation presented by Miyoshi and Kaneda (1986). Unlike the MINT, where the element corresponding to the non-minimum phase part of the plant is inverted through an additional electroacoustic channel, in this case this part is implemented in negative time and the two parts of the inverse are combined in the same electroacoustic channel.



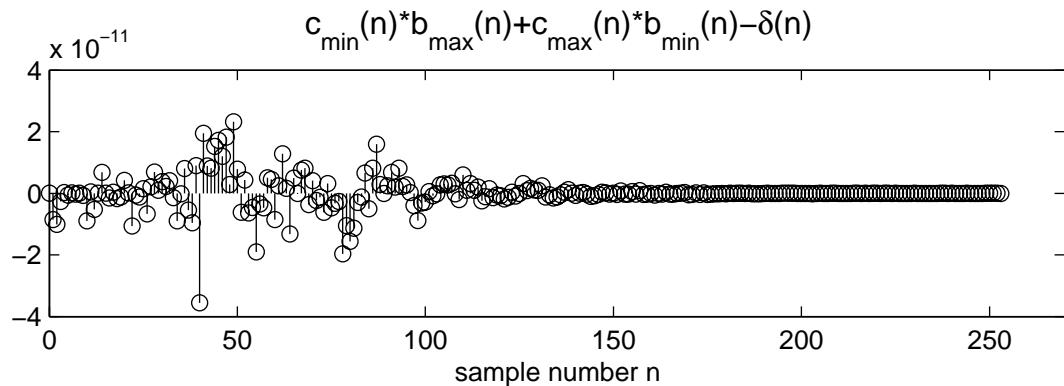
**Figure 6-14:** Accuracy of the parallel connection decomposition of the transfer function  $1/C(z)$ , (a) using the partial fraction expansion and (b) solving the equivalent Diophantine equation using the method of indeterminate coefficients.

Finally we move to the two-channel case. The error quantity  $c_{min}(n)*b_{ij,max}(n)+c_{max}(n)*b_{ij,min}(n)\pm c_{ij}(n)$  corresponding to the solution of equations (6-35) when the MIT-database HRTFs are used for the modelling of the Stereo Dipole geometry are plotted in figure 6-15<sup>88</sup>. The error quantity  $c_{min}(n)*b_{max}(n)+c_{max}(n)*b_{min}(n)-\delta(n)$  demonstrating the precision of the solution of equation (6-34) corresponding to the common denominator arrangement is plotted in figure 6-16.

<sup>88</sup> Given the symmetry of the plant when it is modelled using the MIT-database HRTFs, the first and fourth as well as the second and third lines in equation (6-35) will be identical so we only plot the results corresponding to the first two lines of equation (6-35) in figure 6-15.



**Figure 6-15:** Precision of the solution of equation (6-35)



**Figure 6-16:** Precision of the solution of equation (6-34)

## ***6.5 Issues regarding the practical implementation of the method – Reduced order inverse models***

In the previous sections we showed how the exact (non-regularised) inverse matrix  $\mathbf{H}(z)$  of equation (2-14) can be realised by implementing its elements with rational transfer functions in forward and backward time. As is demonstrated with the simulation results of the next section, when the decomposition in forward and backward time parts is computed in Matlab's double-precision arithmetic by use of the numerical method of §6.4, a highly accurate model of  $\mathbf{H}(z)$  can be obtained.

On the other hand as was shown with the results of chapter 3, the actual accuracy of the inversion, is significantly limited in the real-world implementation of the design, even when this implementation takes place in ideal conditions of reproduction. Hence, the high accuracy achieved by the exact recursive modelling of  $\mathbf{H}(z)$  described above, is of limited practical interest in this specific audio signal processing application.

Its practical interest would be further reduced by the fact that the order of the resulting recursive filters (and hence the number of multiplications/additions needed for the computation of their output) is very high. From equations (6-27)-(6-30) it can be seen that, for the typical 4ms-long anechoic plant model described in §3.4, the total order of the four elements in  $\mathbf{H}(z)$  turns out to range from about 1500 coefficients (for the common denominator cascade arrangement) to about 3000 coefficients (for the direct parallel arrangement). Compared to the non-regularised FIR inverse models considered in chapter 4 which needed a number of coefficients in the order of 8000 to 16000 (see the results in figures 4-2 and 4-4 respectively) in order to avoid the presence of truncation end-effects, we see that the exact recursive models do indeed offer a reduction in the number of coefficients but this reduction is not dramatic.

Further to that, the realisation of recursive filters of such order would pose a serious design challenge in lower-precision arithmetic platforms that are typically used in audio applications. Even in the double precision arithmetic of Matlab, the numerical decomposition becomes less robust as the length of the plant model increases. A final

drawback of the forward backward time filtering method as described above is that it does not accommodate the use of regularisation in the inverse<sup>89</sup>.

In order to address these issues, two alternative approaches for the modelling of the inverse are examined: In the first, the recursive model of the inverse is determined in the frequency-domain. The frequency response of the elements  $H_{ij}(\omega)$  of the (regularised or non-regularised) inverse matrix  $\mathbf{H}$  of equation (2-20) are computed over a chosen grid of  $N$  points  $\omega_n$  on the frequency axis.

$$\mathbf{H}(e^{j\omega}) = \left[ \mathbf{C}^H(e^{j\omega}) \mathbf{C}(e^{j\omega}) + \beta(\omega) \mathbf{I} \right]^{-1} \mathbf{C}^H(e^{j\omega}) \quad (6-39)$$

with  $0 \leq \omega = \omega_0, \omega_1, \dots, \omega_{N-1} \leq \pi$

Subsequently, rational transfer functions  $B_{ij}(z)/A_{ij}(z)$  of chosen order are fitted to these responses.

$$\frac{B_{ij}(z)}{A_{ij}(z)} \approx H_{ij}(z) \quad (6-40)$$

This is done by minimising the quantity of equation (6-41) over the coefficients of the rational transfer function  $B_{ij}(z)/A_{ij}(z)$ .

$$\min \left\{ \sum_n \left| B_{ij}(e^{j\omega_n}) - A_{ij}(e^{j\omega_n}) H_{ij}(e^{j\omega_n}) \right| \right\} \quad (6-41)$$

This practically amounts matching the frequency response of  $B_{ij}(z)/A_{ij}(z)$  to  $H_{ij}(\omega)$  at the chosen set of points  $\omega_n$  on the frequency axis. In the results presented here this was done using the *invfreqz* function of Matlab (Mathworks 2000).

This optimisation procedure retains both the magnitude and the phase of the desired responses  $H_{ij}(\omega)$  and hence the resulting rational expression  $B_{ij}(z)/A_{ij}(z)$  will have poles outside the unit circle. A realisable design is then obtained by applying the decomposition to forward and backward time realisable recursive filters using the method of §§6.2-6.4. As will be seen with the results of the next section, a reduction

---

<sup>89</sup> The above list of issues does not include the fact that, as presented up to this point, the method is restricted to off-line implementation. This issue is addressed in the next chapter.

in the order of the recursive models is indeed achieved in comparison with the exact model but this reduction is only minor.

In the second approach, the recursive filters that model the inverse are determined in the time-domain. First, the method described in §2.3 is employed to compute a (regularised or non-regularised) FIR model  $\mathbf{h}(n)$  of the inverse matrix  $\mathbf{H}$  of chosen length  $N$ . As is explained in §2.3, a chosen amount of coefficients, say  $\Delta$ , out of the total of  $N$  coefficients in each  $h_{ij}(n)$  will correspond to the anti-causal part of the inverse. The method then proceeds by fitting a recursive model  $B_{min,ij}(z)/A_{min,ij}(z)$  to the causal part of each  $h_{ij}(n)$

$$Z^{-1} \left\{ \frac{B_{min,ij}(z)}{A_{min,ij}(z)} \right\} \simeq h_{ij}(n), \quad n = \Delta + 1, \Delta + 2, \dots, N - 1 \quad (6-42)$$

and another recursive model  $B_{max,ij}(z)/A_{max,ij}(z)$  to the time-reversed version of the anti-causal part of each  $h_{ij}(n)$ .

$$Z^{-1} \left\{ \frac{B_{max,ij}(z)}{A_{max,ij}(z)} \right\} \simeq h_{ij}(-n), \quad n = 0, 1, \dots, \Delta - 1 \quad (6-43)$$

The method for the determination of the recursive models that was found to give the best results was the Steiglitz-McBride algorithm as implemented by the *stmcb* function in Matlab (Mathworks 2000), (McBride et al. 1966), (Steiglitz and McBride 1965). The recursive models of equations (6-42) and (6-43) are then directly applicable to the direct arrangement parallel connection form of §6.3 (see figure 6-6). As will be seen with the results of the following section, with both the responses in the right-hand sides of equations (6-42) and (6-43) decaying in time, the modelling turns out to be more efficient than the frequency-domain method which was described above.

In the results presented in the next section we examine the application of these two methods to the direct arrangement parallel connection form of the inverse matrix  $\mathbf{H}$ . Similar formulations could be devised for the remaining three cases of figures 6-7 – 6-9. However the issue brought up in §2.4.2 should be expected to render the two common denominator arrangements impractical. The conclusions drawn here for the direct arrangement parallel connection are expected to be valid also for the direct

arrangement cascade connection case, however the actual implementation and evaluation of that case is left as further work.

## 6.6 Inversion results

We present here the results obtained when the recursive inverse models described above are used for the inversion of the plant measured as described in §§3.3-3.4. The results are presented in the same manner as the results of chapter 4. That is we plot the logarithm of the squared impulse responses from the input to the inverse filter matrix to the output at the listener's ears and/or the logarithm of the squared magnitude of the corresponding frequency responses. In each case presented, a comparison with the corresponding FIR inversion results is made. The FIR inversion results are obtained with the simulation process described in §4.2, i.e. by convolving the 3000-samples-long model of the plant matrix with the FIR model of the inverse matrix. For the inversion results of the recursive models we first compute the two impulse responses (for the part realised in forward time and of the part realised in backward time) up to a chosen length<sup>90</sup> (set here to 5000 points). We then concatenate the reversed-in-time version of the backward-time realisable part with the forward-time realisable part to obtain the 10000 points FIR equivalent of the recursive model. This FIR model is then used to obtain the results with the same simulated convolution process as in the FIR case<sup>91</sup>.

First, in figure 6-17, we compare the inversion results for the 3000-tap non-regularised FIR  $\mathbf{H}_{190}$ <sup>92</sup> with the results obtained with the exact recursive model of the inverse that was described in §§6.2-6.4. The squared impulse responses plotted in the

---

<sup>90</sup> A detailed exposition of the effective length of the corresponding filter impulse responses taken into account when the rational transfer functions considered here are implemented in forward and backward time is given in §A1. As is shown there, this length is determined by the length of the input signal. If this is large enough then full length of the associated inverse impulse responses is effectively taken into account in the convolution process. On the other hand, as if further elaborated with the results of chapter 7, in the block-processing on-line implementation of the method, zero padding has to be applied to the (short) blocks of input and consequently the choice of this length needed to be taken into account significantly influences the computational cost of the algorithm. For the presentation of this chapter where the algorithm is implemented off-line we took this length to be higher and follow the decay of the determined recursive filters further down in positive and negative time.

<sup>91</sup> We use here the symbols  $\mathbf{H}_{190}$  (introduced in §3.4) to refer to the inverse computed on the basis of the  $\mathbf{C}_{190}$  measured model of the plant.

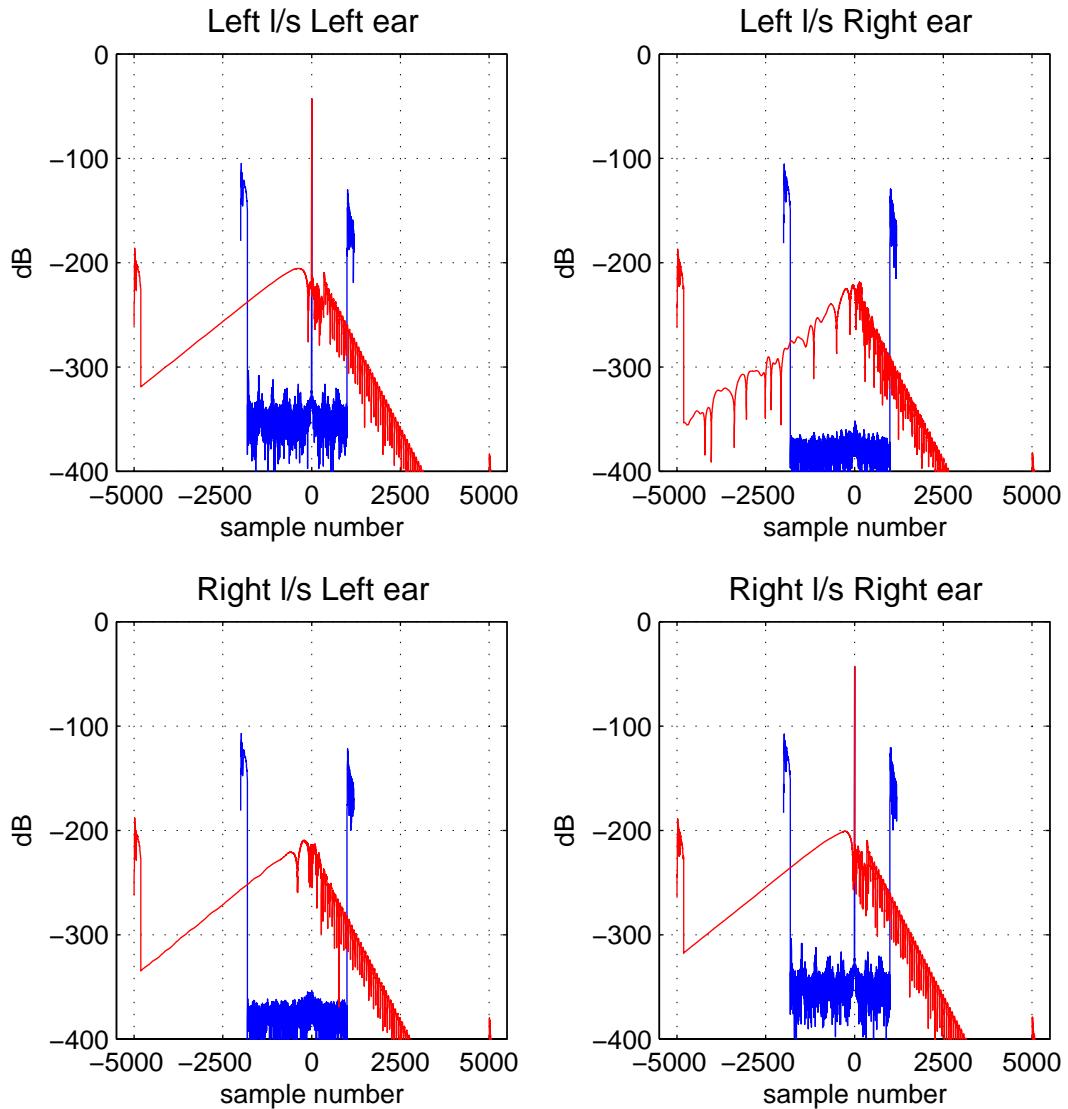
<sup>92</sup> We set the anticausal length of the inverse equal to 2000 taps and the causal length equal to 1000 taps.

figure are obtained by convolving  $\mathbf{C}_{190}$  with the inverse matrix. As was discussed in §3.5, the responses obtained with this simulation are excessively optimistic compared to the actually obtained inversion accuracy. They are included here, however, as an indication of the adequacy of the proposed numerical solution for the determination of the polynomials of equation (6-24).

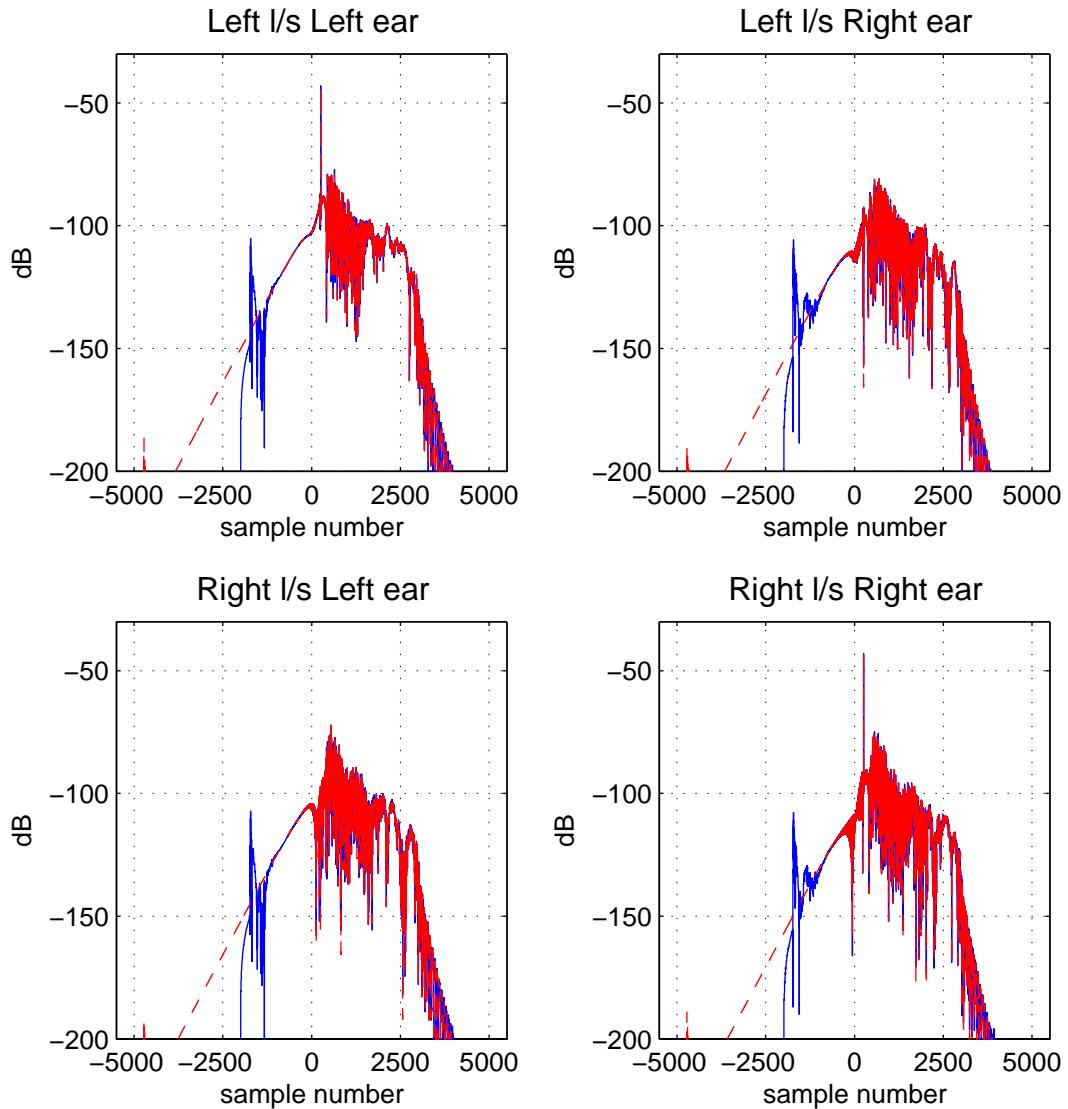
As can be seen in the figure, the accuracy of the inverse model determined with that method succeeds in suppressing the error down to approximately -200dBFS. This is of course seen to be a modest result compared to the accuracy of the FIR model determined by the frequency-domain method of § 2.3 (which in this case was implemented with a DFT length of  $2^{13}$  points). Nevertheless it is quite sufficient with reference to the actual accuracy obtained in the best practical conditions in which the error rises much above this level as was seen in the results of chapter 4. Furthermore, even though the use of a DFT of this length is enough to suppress the time-aliasing effect down to more than -300dBFS as can be seen in the figure, the 2000 coefficients used for the FIR realisation of the anti-causal part of the inverse are just enough to keep the truncation error down to -100dBFS. In the recursive model on the other hand, the length of the filter's impulse response that is effectively used in the filtering depends on the length of the input. Hence, as long as the input is long enough or appropriately zero-padded, the end-effects can be suppressed at will (see also §A1). In the decomposition of the polynomial corresponding to the determinant of the plant matrix  $\det[\mathbf{C}_{190}(z)]$ , the order of the minimum phase part is 344 samples and that of the maximum phase part is 38 samples<sup>93</sup>. Hence, as should be expected, the total order of the transfer functions  $C_{ij}(z)/\det[C(z)]$  is divided into a high-order forward-time realisable recursive structure and a much lower-order backward-time realisable recursive structure.

---

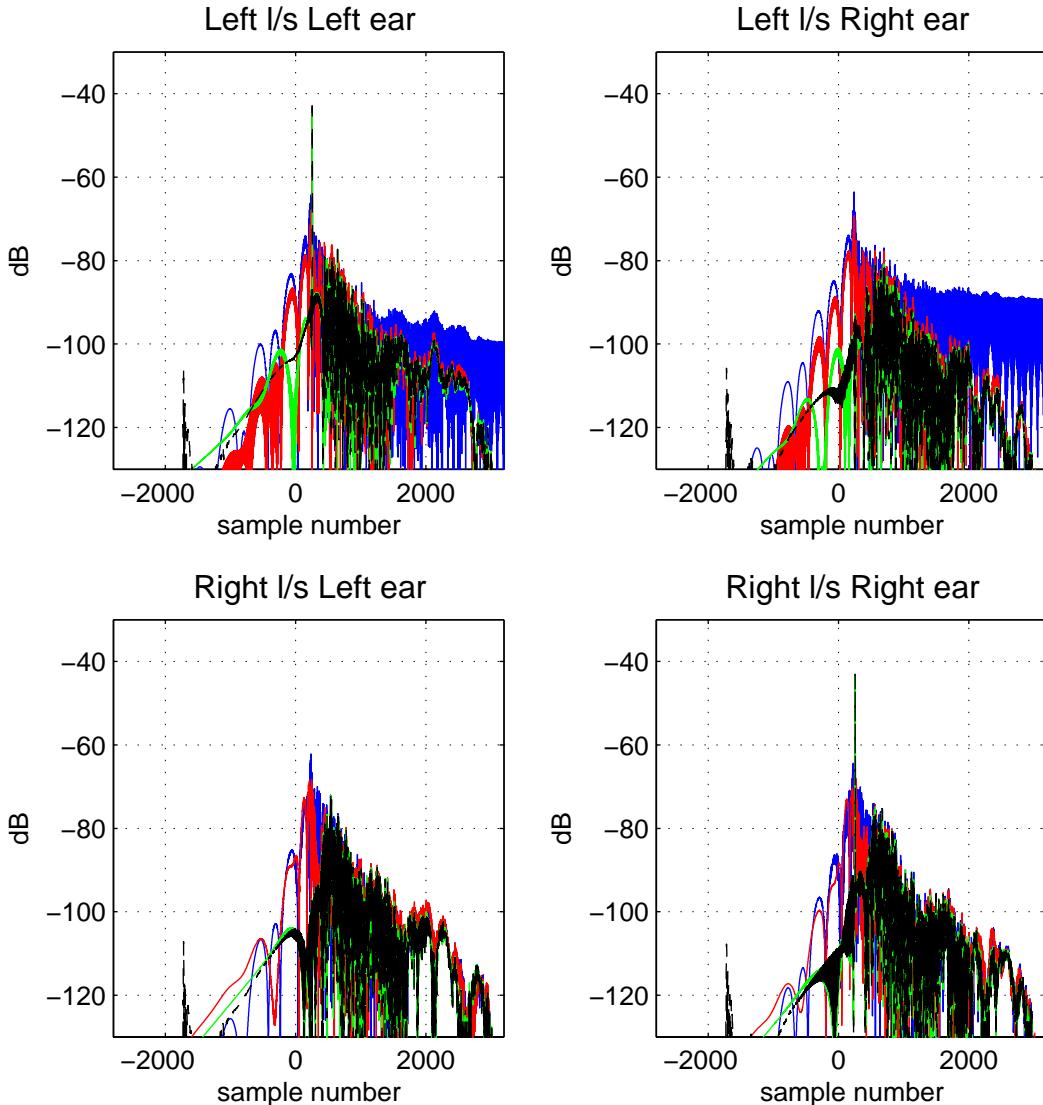
<sup>93</sup> The length of each response in  $\mathbf{C}_{190}(n)$  is 191 samples, hence the determinant polynomial  $\det[\mathbf{C}_{190}(z)]$  has 381 coefficients and this is in agreement with these orders of the minimum and maximum phase parts as is described in equation (6-26).



**Figure 6-17:** Convolution of  $\mathbf{C}_{190}$  with the 3000-tap non-regularised ( $\beta=0$ ) FIR inverse (blue line) and with 10000-points impulse response matrix of the exact recursive model of eq. (6-24) and figure 6-6 (red dashed line).



**Figure 6-18:** Actual performance of the 3000-taps non-regularised ( $\beta=0$ ) FIR inverse (blue line) and of the exact recursive model of equation (6-24) and figure 6-6 (red dashed line).



**Figure 6-19:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 3000-tap (2000 anti-causal and 1000 causal), (black dashed line), (ii) a rational transfer function of total order (100,120) determined using *invfreqz* and decomposed as described in §6.4, (blue line), (iii) a rational transfer function of total order (120,150) as above, (red line) and (iv) a rational transfer function of total order (190,200) as above, (green line). In all cases the regularisation is set to  $\beta=0$ .

What was said above regarding the actual accuracy of the inversion is illustrated in figure 6-18 in which we plot the results obtained with the same inverse models as in figure 6-17 but this time simulating the convolution by use of the full-length 3000-samples model of the plant. As is seen in the figure, the actual performance of the

exact recursive model of the inverse is identical to the FIR model<sup>94</sup> except for the truncation effects. We see then that the order of the exact recursive model (approximately equal to  $4 \times 760 \approx 3000$ , see equation (6-27)) is indeed lower than the order of the FIR inverse (equal to  $4 \times 3000 = 12000$ ) that gives the same results. This reduction in the order is nevertheless rather minor and also not satisfactory due to the reasons exposed in §6.5. Still, the fact that the exact recursive model achieves an excess of accuracy, even if it proves to be redundant in the actual implementation of the inversion, suggests that a reduction in the required order of the recursive model should be possible. In the following we present the results obtained when this reduction in the order of the inverse is attempted with the two methods proposed in §6.5.

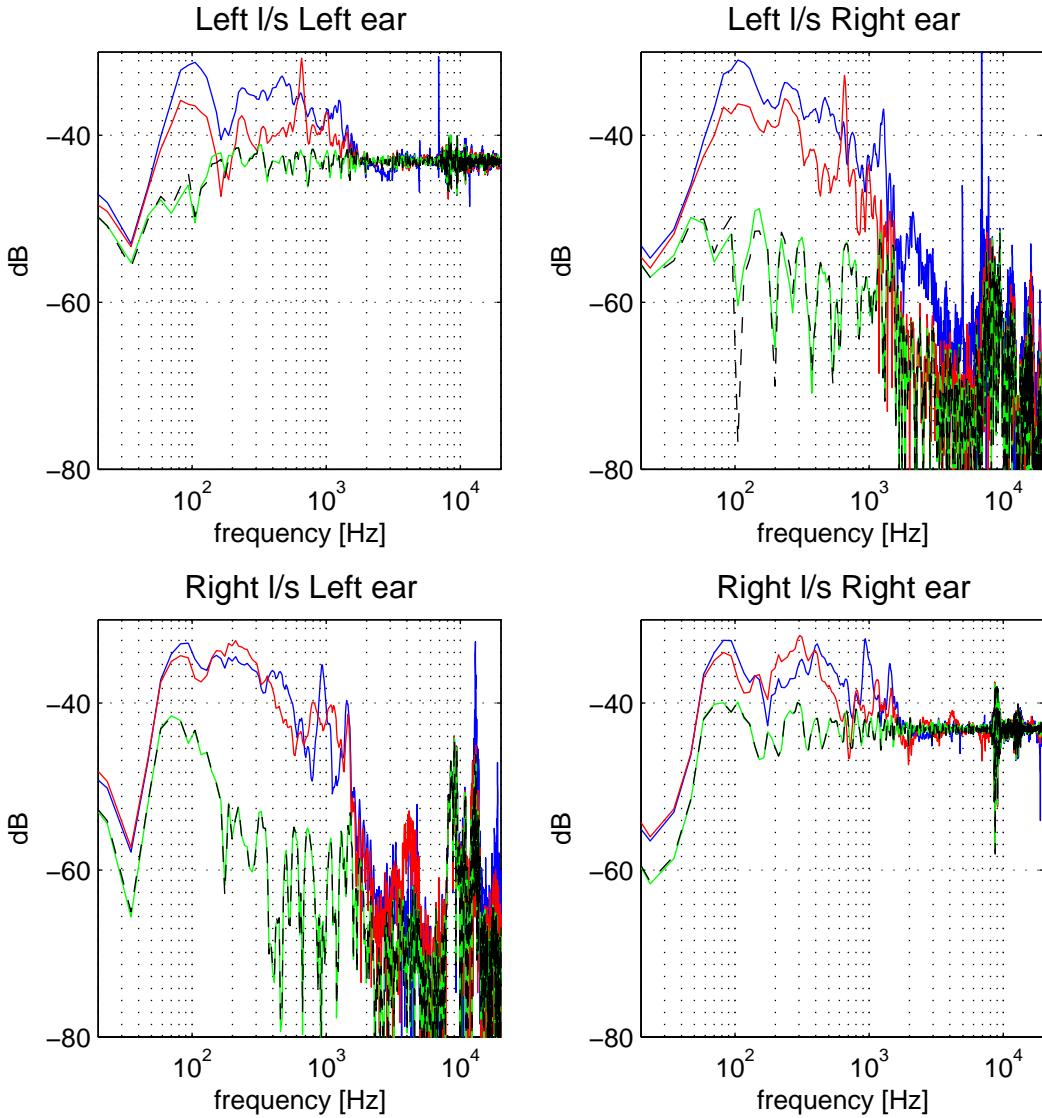
We start by showing in figure 6-19 the performance of the recursive models determined by the frequency-domain method described in §6.5. The results of figure 6-19 correspond to the non-regularised case ( $\beta=0$ ) and the order of the three recursive models examined is (100,120), (120,150) and (190,200)<sup>95</sup>. As is seen in the figure, the first two cases give poor results. In the (100,120) model, the effect of one or more minimum phase poles of either  $H_{11}(z)$  or  $H_{21}(z)$  that have shifted out of position in relation to the zeros of the model, can be seen in the form of a long ringing in the equalised plant. Finally, the model of order (190,200) can be seen to achieve the optimum inversion results, i.e. the same accuracy as the 3000-tap FIR. Obviously, in all cases of the inversion with the recursive models there are no truncation end-effects.

Figure 6-20 shows the results of the same inverse models in the frequency-domain. For the models of order (100,120) and (120,150) the inversion can be seen to be effective only above 1.5kHz with deviations up to 10dB in the ipsilateral equalisation and no cross-talk cancellation below this frequency. The pole-zero misplacement in the (100,120) model is now manifested as a sharp high resonance at 7kHz and a lower one at 5kHz at the left-loudspeaker-left-ear and left-loudspeaker-right-ear responses.

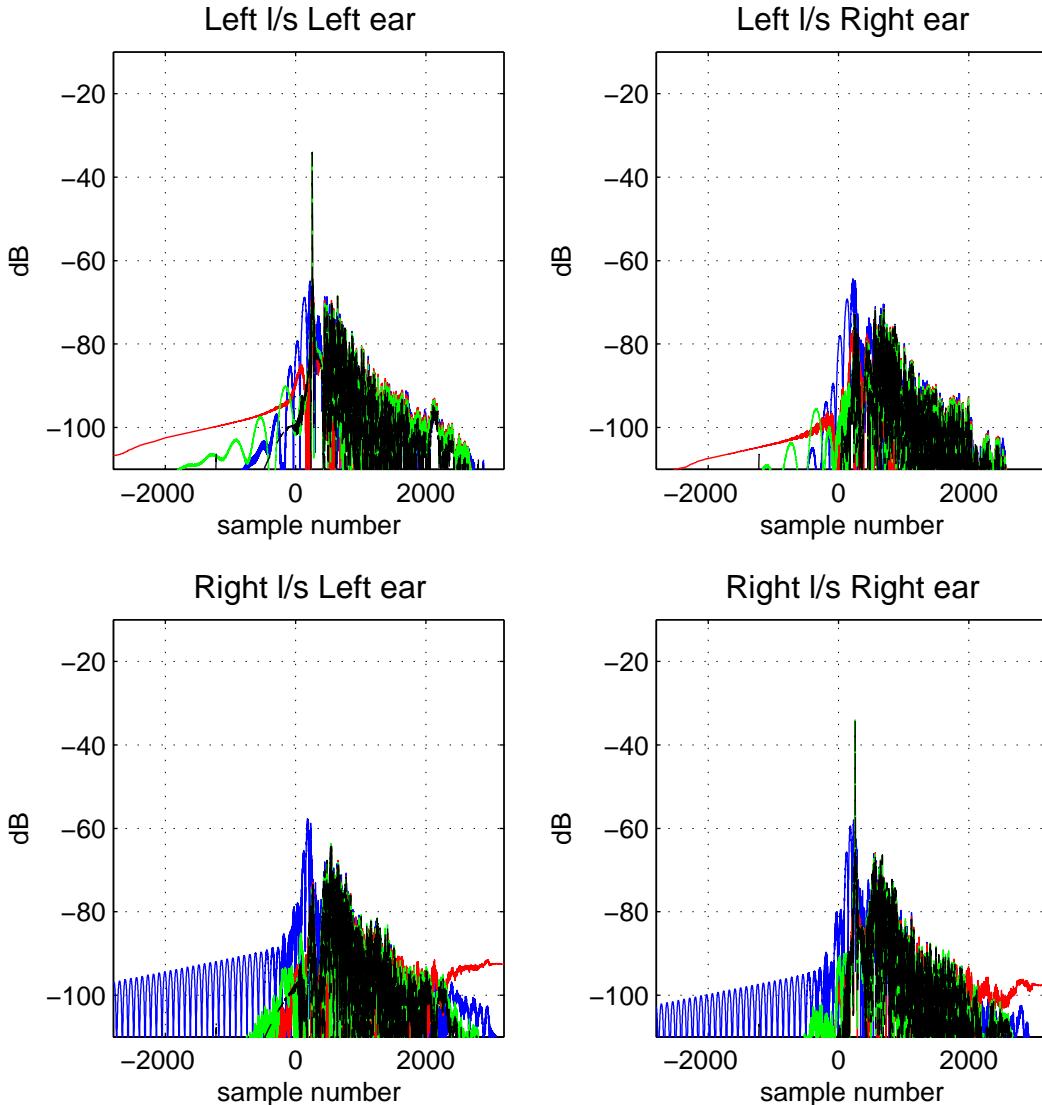
---

<sup>94</sup> This result can be seen to be the same above the measurement noise floor as the directly measured result of figure 4-3.

<sup>95</sup> The first number in the brackets is the total number of zeros and the second the total number of poles used for the modelling of the each one of the rational transfer functions  $C_{ij}(z)/\det[C(z)]$  of the inverse matrix  $\mathbf{H}$ . The order of the forward and backward time realisable parts resulting in each case varies.



**Figure 6-20:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 3000-tap (2000 anti-causal and 1000 causal), black dashed line, (ii) a rational transfer function of total order (100,120) determined using *invfreqz* and decomposed as described in §6.4, blue line, (iii) a rational transfer function of total order (120,150) as above, red line and (iv) a rational transfer function of total order (190,200) as above, green line. In all cases the regularisation is set to  $\beta=0$ .

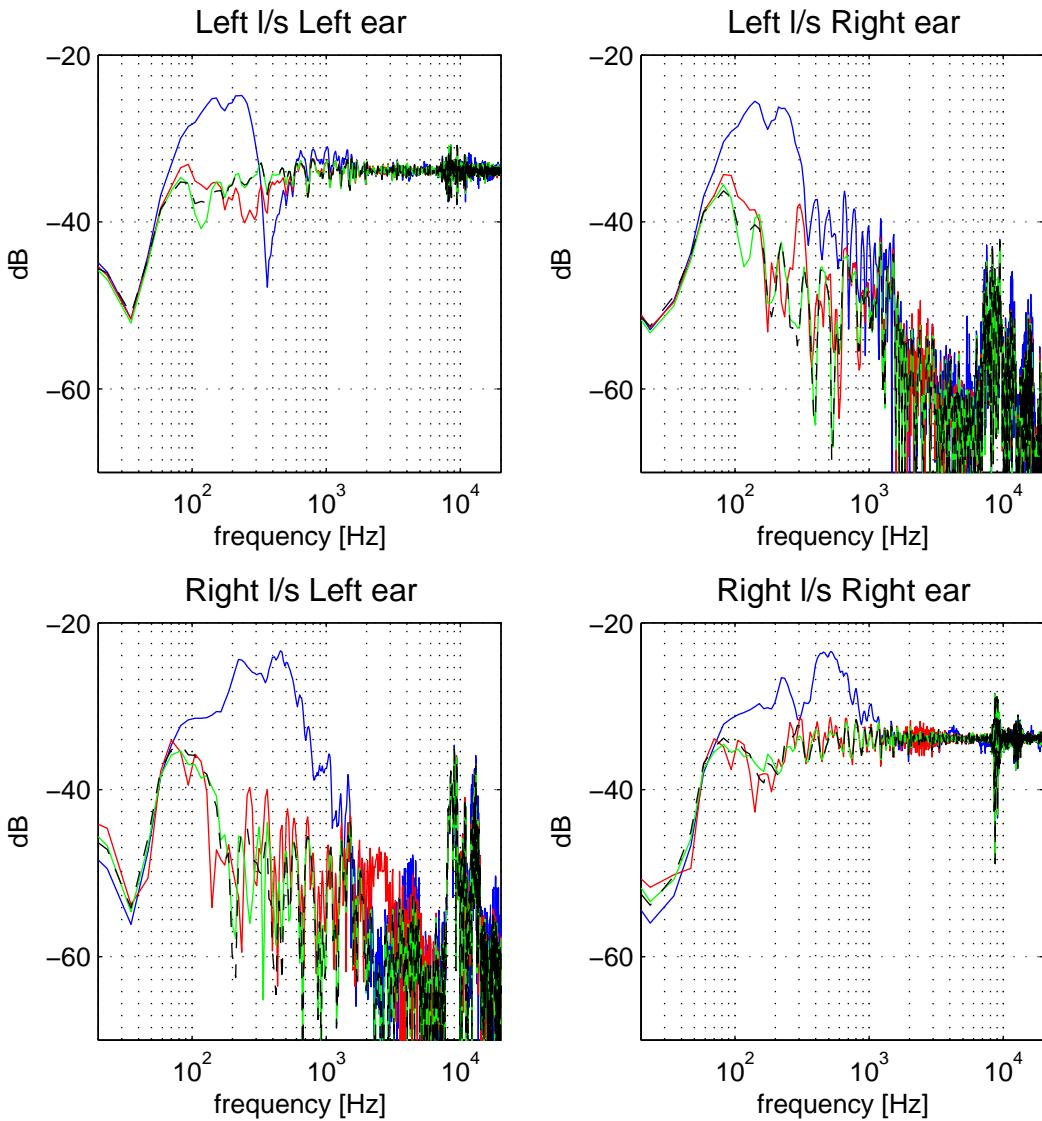


**Figure 6-21:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 2500-tap (1500 anti-causal and 1000 causal), black dashed line, (ii) a rational transfer function of total order (120,150) determined using *invfreqz* and decomposed as described in §6.4, blue line, (iii) a rational transfer function of total order (150,180) as above, red line and (iv) a rational transfer function of total order (200,250) as above, green line. In all cases the regularisation is set to  $\beta=10^{-4}$ .

Next, in figures 6-21 and 6-22 we plot the results obtained for the regularised inversion ( $\beta=10^{-4}$ ). In this case the recursive models are of orders (120,150), (150,180) and (200,250). The FIR inversion case included for comparison is now of slightly lower order than before with the filters implemented with 2500 taps (1500 acausal and 1000 causal)<sup>96</sup>. Surprisingly, the orders of the recursive models needed to

<sup>96</sup> In fact, as was shown in §4.4 (see pp. 106-108) this length of the inverse is higher than required in this case.

match the FIR performance are now higher than in the previous (non-regularised) case. As is seen in both the time and the frequency domain plots, the (120,150) model fails altogether. The (150,180) model is only marginally worse than the FIR inverse but in order to exactly match the performance of the FIR inverse one would need to use a (200,250) recursive model.

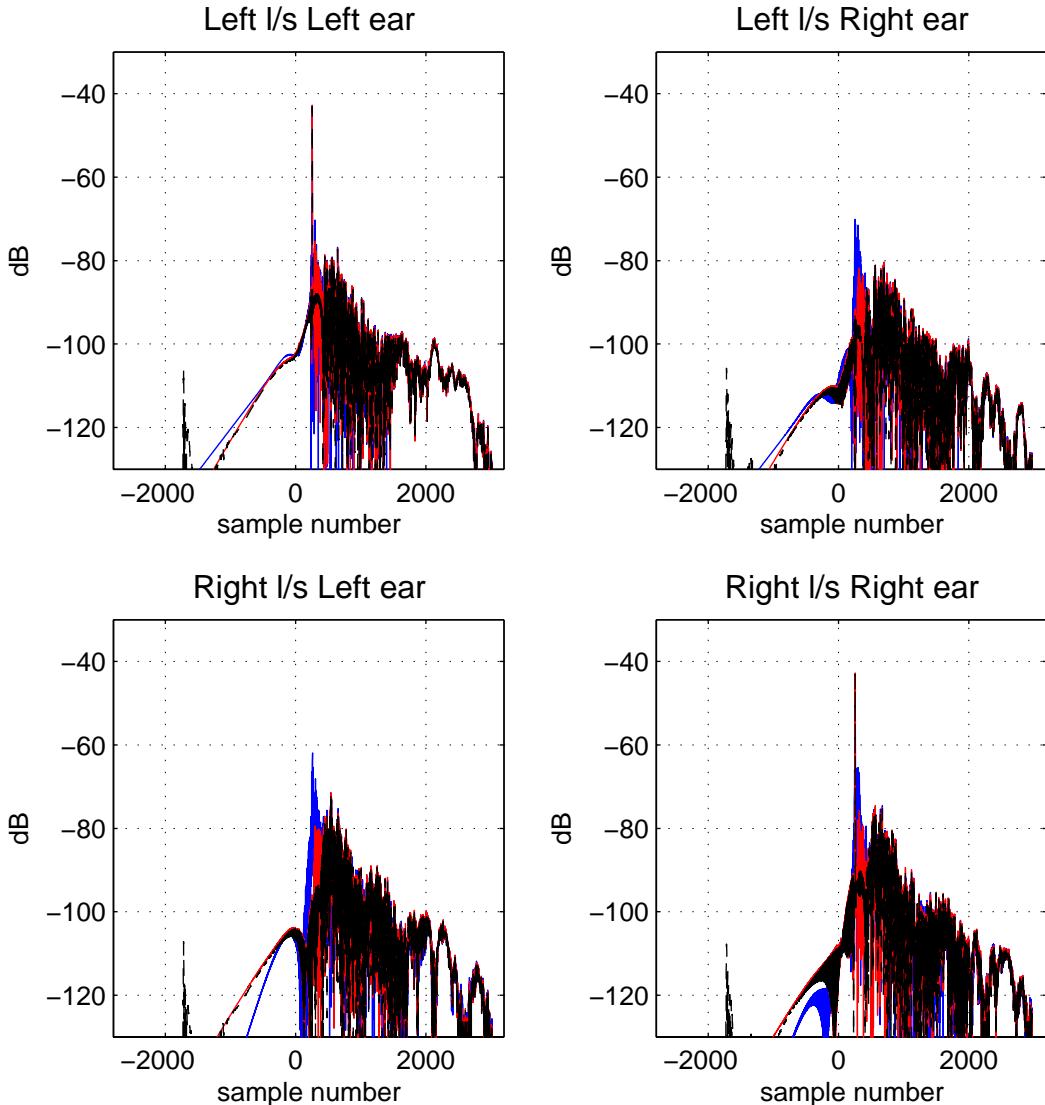


**Figure 6-22:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 2500-tap (1500 anti-causal and 1000 causal), black dashed line, (ii) a rational transfer function of total order (120,150) determined using *invfreqz* and decomposed as described in §6.4, blue line, (iii) a rational transfer function of total order (150,180) as above, red line and (iv) a rational transfer function of total order (200,250) as above, green line. In all cases the regularisation is set to  $\beta=10^{-4}$ .

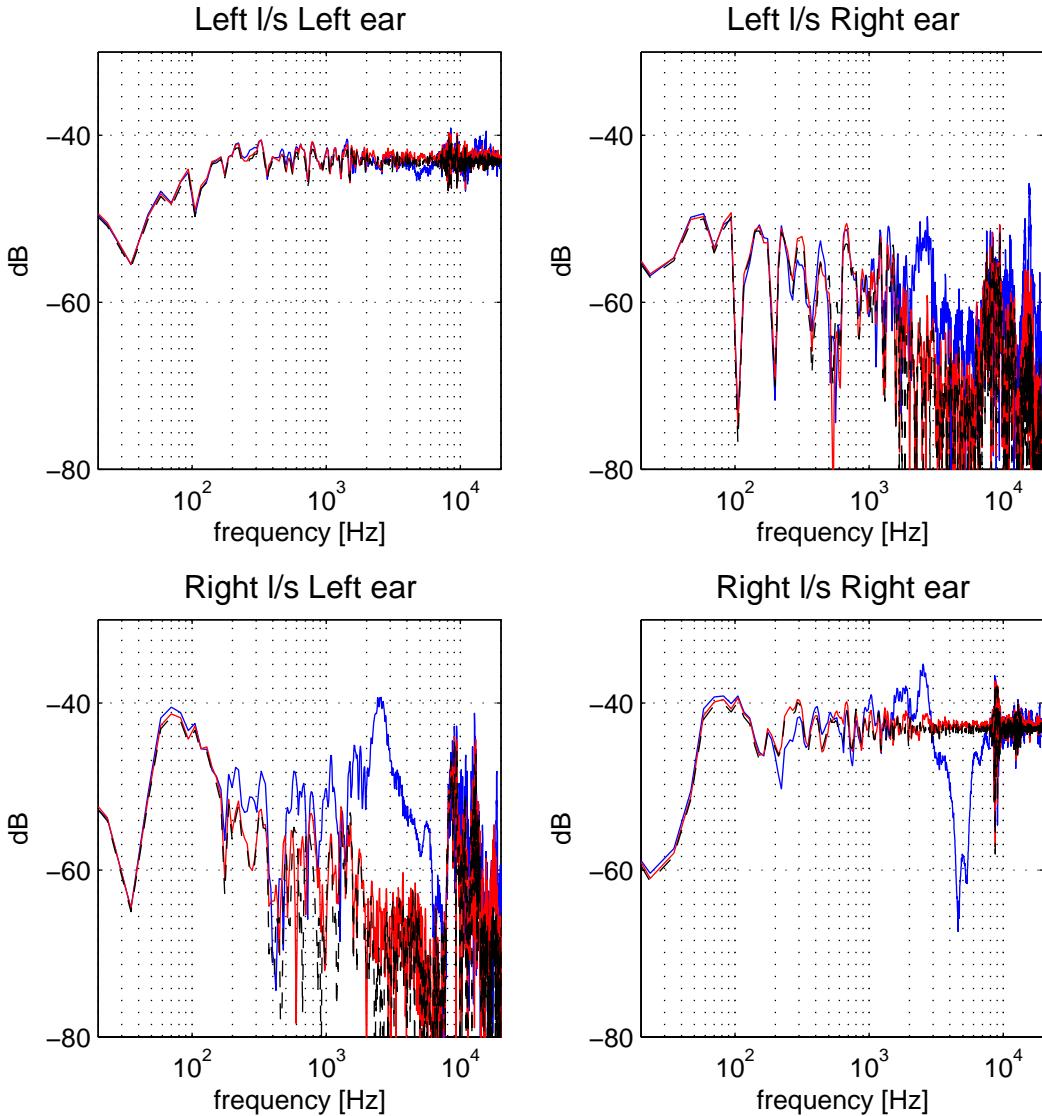
Overall then, of the two modified methods proposed in §6.5 for the recursive modelling of the inverse, that computed in the frequency-domain is seen to yield poor performance. It does achieve a reduction by half of the total order of the recursive models which from  $4 \cdot N_C = 4 \cdot 190 \approx 800$  per element of  $\mathbf{H}$  (see equation (6-24)) is reduced to  $200 + 190 \approx 400$ . But it does worse in the regularised case and is certainly not attractive compared to the approximately 2000 taps needed for the FIR modelling of each element of  $\mathbf{H}$  in that case.

We then move on to examine the same inversion cases for the second modified method proposed in §6.5. This is that of fitting one recursive model to each of the causal and anti-causal parts of the (pre-calculated) time-domain responses of the four elements  $H_{ij}$  of the inverse matrix. The results obtained are plotted in figures 6-23, 6-24, 6-25 and 6-26. Figures 6-23 and 6-24 depict the time and frequency domain results for the non-regularised inversion case and figures 6-25 and 6-26 the results obtained with the regularisation parameter set to  $\beta = 10^{-4}$ . In all cases, the recursive filters are set to model the first 1000 coefficients of the causal and anti-causal parts of the responses  $H_{ij}$ . In both the regularised and non-regularised results two cases are considered. In the first, the anti-causal parts are modelled with filters of order (20,10) and the causal parts with filters of order (30,15). These results are plotted in the blue line. In the second case, the anti-causal parts are modelled with filters of order (25,20) and the causal parts with filters of order (40,25). These results are plotted in the red line. The results of the FIR inversion (identical with the ones presented above) are again plotted in black dashed line.

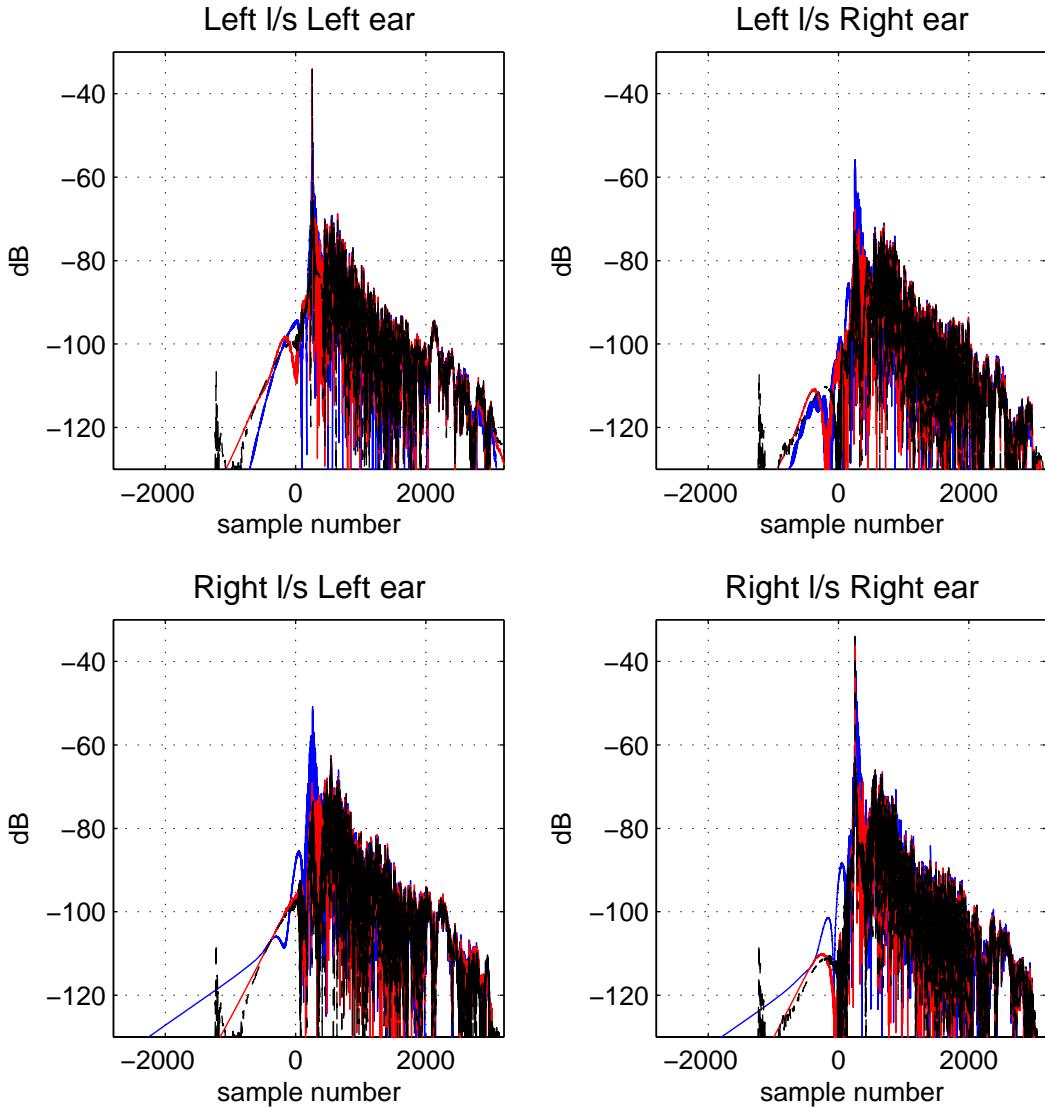
As is seen in these results, the larger order recursive models succeed in achieving virtually identical results with the FIR inverse. These models are of a total order of  $(40,25) + (25,20) = (65,45)$  which constitutes a significant reduction with the order of the frequency-domain method presented above. A further order reduction will probably not be possible as in all cases, the lower order models can be seen to fail.



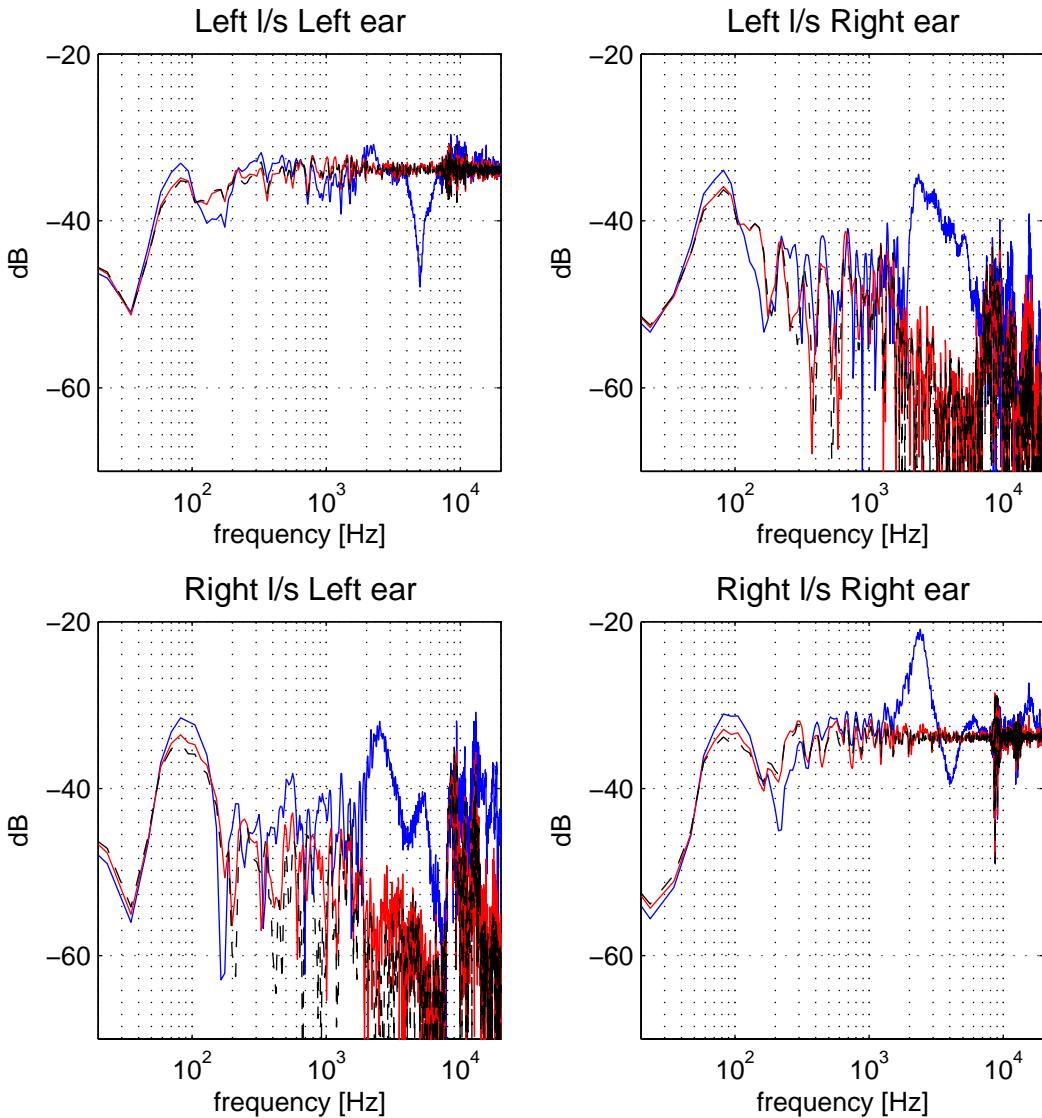
**Figure 6-23:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 3000-tap (2000 anti-causal and 1000 causal), black dashed line, (ii) recursive filters of total order (50,25) determined using the time-domain method of §6.5, blue line and (iii) recursive filters of total order (65,45) determined using the time-domain method of §6.5, red line. In all cases the regularisation is set to  $\beta=0$ .



**Figure 6-24:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 3000-tap (2000 anti-causal and 1000 causal), black dashed line, (ii) recursive filters of total order (50,25) determined using the time-domain method of §6.5, blue line and (iii) recursive filters of total order (65,45) determined using the time-domain method of §6.5, red line. In all cases the regularisation is set to  $\beta=0$ .



**Figure 6-25:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 2500-tap (1500 anti-causal and 1000 causal), black dashed line, (ii) recursive filters of total order (50,25) determined using the time-domain method of §6.5, blue line and (iii) recursive filters of total order (65,45) determined using the time-domain method of §6.5, red line. In all cases the regularisation is set to  $\beta=10^{-4}$ .



**Figure 6-26:** Inversion results. Each one of the transfer functions  $H_{ij}$  is realised with (i) an FIR filter of 2500-tap (1500 anti-causal and 1000 causal), black dashed line, (ii) recursive filters of total order (50,25) determined using the time-domain method of §6.5, blue line and (iii) recursive filters of total order (65,45) determined using the time-domain method of §6.5, red line. In all cases the regularisation is set to  $\beta=10^{-4}$ .

## 6.7 Conclusions

A technique was presented in this chapter for the decomposition of the elements  $H_{ij}$  of the inverse matrix  $\mathbf{H}$  in two parts, one realisable in forward and one in backward time. The technique was first presented in a highly analytical form in §§6.2-6.4. In that form, an overwhelmingly accurate recursive representation of the inverse can be obtained. It was argued however that in such a form the technique can be of very

limited practical interest, mainly due to the excessively high order of the resulting recursive structures. Despite the fact that this order is slightly lower than the FIR order needed to model the inverse with equivalent accuracy, such a level of accuracy would be lost in any real-world implementation of the system. Even in cases when such an accurate implementation of the inverse is of importance, the FIR equivalent could be implemented with one of the available frequency-domain block-processing algorithms (Press et al. 2002), (Gardner 1995) at a fraction of the computational cost.

Two modifications were proposed in the method for the determination of the recursive filters that model the forward and backward time realisable parts of the inverse. One of them was found to give very promising results in terms of reducing the order of the involved recursive filters without reducing the accuracy of the inversion below the level that one would realistically expect to achieve in real-world implementations. Due to time restrictions the presented investigation of the application of the proposed inverse structures was confined to a small “sample” set among the inversion cases previously examined in chapters 4 and 5. It is quite probable that this example case is not the one most suitable to show the full potential of the proposed method. This argument is expanded as follows.

As was shown in the analysis of chapters 4 and 5, depending on a number of factors, the required length of the inverse filters FIR model ranges from several thousands of coefficients (in the case where an accurate model of the plant is known that contains the responses of the loudspeakers) to no more than a few hundreds of coefficients. This last case would correspond to most everyday virtual acoustic imaging applications where the plant contains non-individual HRTFs and some degree of listener misalignment and reverberation is inevitable. This wide range of possible requirements for the order of the inverse filters translates to a corresponding computational cost requirement in a somewhat non-linear fashion. When filters of many thousands of coefficients need be implemented, the existence of frequency-domain block-processing algorithms guarantees an extremely efficient implementation. The relatively recent development of such algorithms in zero (or very low) input-output latency formulations, e.g. by Gardner (1995), makes them practically impossible to beat when such high filter orders are considered. This is more clearly shown in chapter 7 where we give the computational cost of the block-

processing on-line version of the proposed method. At the opposite end of the spectrum, when filters of no more than one to two hundred coefficients are to be implemented, the direct FIR sample-by-sample filtering implementation is usually advantageous as it is more robust to lower-precision arithmetic platforms and easier to design. The results of this chapter, however, show that the recursive forward-backward-time modelling of inverse impulse responses of the order of several hundred to a few thousand coefficients can indeed be meaningful. The closer investigation of the actual merit of the proposed method in such a type of implementation is proposed as further work.

# ***7 On-line implementation of the forward-backward time inverse filtering***

## ***7.1 Introduction***

In this chapter, in §7.2, we describe a method for the implementation of the forward-backward time filtering of the input signal in segments. This forms the basis for the on-line implementation of the inverse recursive models that were presented in the previous chapter<sup>97</sup>. In §7.3 we discuss and give the formulae describing the effect of the segmenting in the total processing cost of the filtering and the resulting input-output latency. An example of the computational cost and input-output latency in the application of the algorithm is given based on the inverse recursive model used for the results of figures 6-25 and 6-26. This is compared with the corresponding cost and input-output latency of the available block-processing frequency-domain algorithms.

## ***7.2 On-line implementation of the inverse filtering***

### ***algorithm – Filtering of the input in segments***

Let us assume here that the input signal  $x(n)$  corresponding to the vector<sup>98</sup>  $\underline{x}=[x(0) \ x(1) \ \dots \ x(N-1)]^T$  is partitioned in segments  $x_k(n)$  corresponding to the vectors  $\underline{x}_k=[x_k(0) \ x_k(1) \ \dots \ x_k(L-1)]^T=[x(kL) \ x(kL+1) \ \dots \ x(kL+L-1)]^T$  of length  $L$ <sup>99</sup>, each one available at the time index  $kL+L-1$ ,  $k=0, \dots, N/L-1$  and that we are to compute the output  $y(n)$  partitioned in the corresponding segments  $\underline{y}_k$ . In general, when we filter an input

<sup>97</sup> A similar formulation is presented in (Powell and Chau 1990a) and (Powell and Chau 1991).

<sup>98</sup> We use here the notation introduced in §6.2.

<sup>99</sup> For notation simplicity we assume here that the signal's length  $N$  is an exact multiple of the segment's length  $L$  and, hence, that the integer index  $k$  takes the values  $0, 1, \dots, N/L-1$ . Of course this does not restrict the generality of the presentation as, if this is not the case,  $k$  will take the values  $0, 1, \dots, [N/L]-1, [N/L]$  but the last segment's length will be equal to  $N \bmod L$

signal in segments, we have to address the issue of the initial conditions with which implement the filtering of each segment. As an example of that let us consider the case where the  $k$ -th input segment  $x_k(n)$  is to be convolved (in the normal forward-time sense) with an  $M$ -samples-long, strictly causal impulse response  $h(n)$ <sup>100</sup>. The first output sample of the convolution  $h(n)*x_k(n)$  will be equal to

$$y_k(0) = \sum_{l=0}^{M-1} h(l)x_k(-l) \quad (7-1)$$

Now given that the  $l^{\text{th}}$  sample of the  $k^{\text{th}}$  segment of the input will be the  $(kL+l)^{\text{th}}$  sample of the input:  $x_k(l)=x(kL+l)$  we get

$$y_k(0) = \sum_{l=0}^{M-1} h(l)x(kL-l) \quad (7-2)$$

or equivalently

$$y_k(0) = h(0)x(kL) + h(1)x(kL-1) + \dots + h(M-1)x(kL-M+1) \quad (7-3)$$

and replacing back the segmented time indices we can write

$$y_k(0) = h(0)x_k(0) + h(1)x_{k-1}(L-1) + \dots + h(M-1)x_{k-1}(L-M+1) \quad (7-4)$$

In other words, as is described in equation (7-4), the first sample of the  $k$ -th output segment  $y_k(n)$  is a function of input samples contained not only in the corresponding  $k$ -th input segment  $x_k(n)$  but also in the previous input segment  $x_{k-1}(n)$  (and possibly, depending on the relative lengths  $M$  and  $L$  of the impulse response and the input segments, of samples contained in the even earlier input segments  $x_{k-2}(n)$ ,  $x_{k-3}(n)$  etc.).

Now in the forward-time filtering case, this issue can be trivially addressed by storing the state of the filter's internal buffer at the end of each input segment and using it as initial conditions for the filtering of the next segment. If we want to compute the full-length  $(N+M-1)$ -samples-long convolution  $y(n)=x(n)*h(n)$ , then the full-length  $(L+M-1)$ -samples-long of the last segment  $x_{N/L-1}(n)*h(n)$  has to be computed (for example by zero-padded sample-by-sample filtering as is explained in §A1).

---

<sup>100</sup> For the purposes of the current discussion  $h(n)$  can be considered to be the approximation up to the  $M^{\text{th}}$  sample of an infinitely long impulse response corresponding to a recursive filter.

However, this approach is clearly not feasible in the backward-time filtering case since, due to the reversal in time, the first sample of the input segment  $x_k(n)$  to “enter” the filtering process is  $x_k(L-1)$  followed by  $x_k(L-2)$  and so on, the last sample to be processed being  $x_k(0)$ . With this arrangement, the sample “preceding”  $x_k(L-1)=x(kL+L-1)$  in this reversed sequence would be  $x(kL+L)=x_{k+1}(0)$  i.e. the first sample of the next input segment. This means that for the computation of the output segment  $y_{k,max}(n)$  we need knowledge of a number of samples of the next input sample  $x_{k+1}(n)$  that is yet not available.

A simple illustration of this process is given in the next example in which we assume that the causal input  $x(n)$  corresponding to the vector  $\underline{x}=[x(0) \ x(1) \ \dots \ x(7)]^T$  is to be convolved in backward-time with the strictly anti-causal impulse response  $h(n)$  corresponding to the vector  $\underline{h}=[h(-1) \ h(-2) \ h(-3)]^T$ . The 10-samples-long output  $y(n)$  calculated by means of zero-padded sample-by-sample filtering (see the discussion in §A1) will then be as in equation (7-5):

$$\begin{bmatrix} h(-1) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ h(-2) & h(-1) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ h(-3) & h(-2) & h(-1) & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & h(-3) & h(-2) & h(-1) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & h(-3) & h(-2) & h(-1) & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & h(-3) & h(-2) & h(-1) & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & h(-3) & h(-2) & h(-1) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & h(-3) & h(-2) & h(-1) & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & h(-3) & h(-2) & h(-1) & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & h(-3) & h(-2) & h(-1) \end{bmatrix} \begin{bmatrix} x(7) \\ x(6) \\ x(5) \\ x(4) \\ x(3) \\ x(2) \\ x(1) \\ x(0) \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} y(6) \\ y(5) \\ y(4) \\ y(3) \\ y(2) \\ y(1) \\ y(0) \\ y(-1) \\ y(-2) \\ y(-3) \end{bmatrix} \quad (7-5)$$

If now this computation is carried by segmenting the input into two segments  $\underline{x}_0=[x(0) \ x(1) \ x(2) \ x(3)]^T$  and  $\underline{x}_1=[x(4) \ x(5) \ x(6) \ x(7)]^T$  then the two corresponding output segments  $\underline{y}_0$  and  $\underline{y}_1$  will be as in equation (7-6) which, by comparison with equation (7-5), readily shows that the output  $y(n)$  will be related to the segments  $\underline{y}_0$  and  $\underline{y}_1$  as in equation (7-7), i.e. it will be equal to an overlap-addition of the two (reversed-in-time) output segments  $\underline{y}_0$  and  $\underline{y}_1$ .

$$\begin{aligned}
 \mathbf{y}_0 &= \begin{bmatrix} h(-1) & 0 & 0 & 0 & 0 & 0 \\ h(-2) & h(-1) & 0 & 0 & 0 & 0 \\ h(-3) & h(-2) & h(-1) & 0 & 0 & 0 \\ 0 & h(-3) & h(-2) & h(-1) & 0 & 0 \\ 0 & 0 & h(-3) & h(-2) & h(-1) & 0 \\ 0 & 0 & 0 & h(-3) & h(-2) & h(-1) \end{bmatrix} \begin{bmatrix} x(3) \\ x(2) \\ x(1) \\ x(0) \end{bmatrix} = \begin{bmatrix} y_0(2) \\ y_0(1) \\ y_0(0) \\ y_0(-1) \\ y_0(-2) \\ y_0(-3) \end{bmatrix} \\
 \mathbf{y}_1 &= \begin{bmatrix} h(-1) & 0 & 0 & 0 & 0 & 0 \\ h(-2) & h(-1) & 0 & 0 & 0 & 0 \\ h(-3) & h(-2) & h(-1) & 0 & 0 & 0 \\ 0 & h(-3) & h(-2) & h(-1) & 0 & 0 \\ 0 & 0 & h(-3) & h(-2) & h(-1) & 0 \\ 0 & 0 & 0 & h(-3) & h(-2) & h(-1) \end{bmatrix} \begin{bmatrix} x(7) \\ x(6) \\ x(5) \\ x(4) \end{bmatrix} = \begin{bmatrix} y_1(2) \\ y_1(1) \\ y_1(0) \\ y_1(-1) \\ y_1(-2) \\ y_1(-3) \end{bmatrix}
 \end{aligned} \tag{7-6}$$

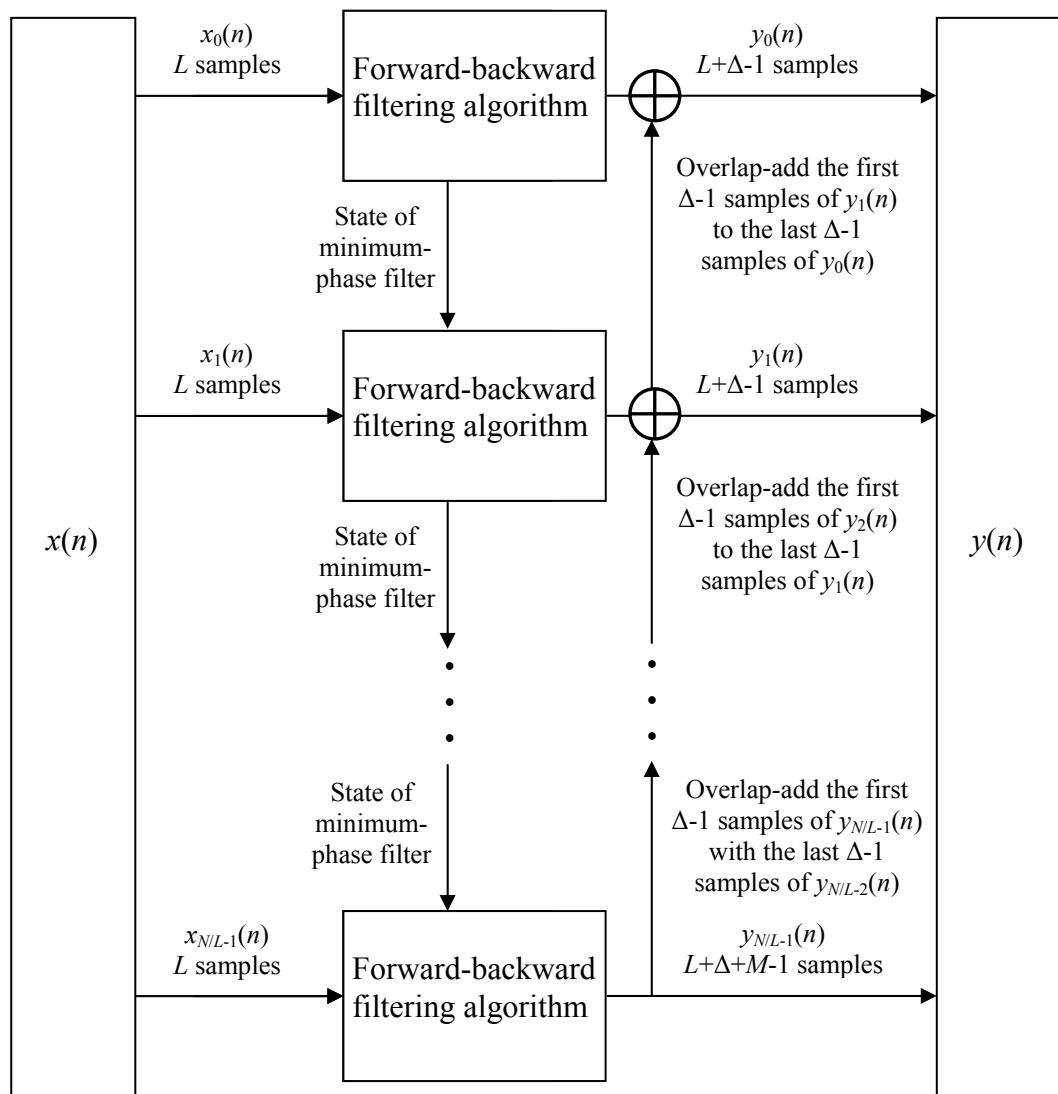
$$\begin{bmatrix} y(-3) \\ y(-2) \\ y(-1) \\ y(0) \\ y(1) \\ y(2) \\ y(3) \\ y(4) \\ y(5) \\ y(6) \end{bmatrix} = \begin{bmatrix} y_0(-3) \\ y_0(-2) \\ y_0(-1) \\ y_0(0) \\ y_0(1)+y_1(-3) \\ y_0(2)+y_1(-2) \\ y_1(-1) \\ y_1(0) \\ y_1(1) \\ y_1(2) \end{bmatrix} \tag{7-7}$$

In the general case then, where the  $N$ -samples-long input  $x(n)$  partitioned in segments  $\underline{x}_k$  each of length  $L$  is to be convolved with the  $\Delta$ -samples-long anti-causal impulse response  $h(n)$ <sup>101</sup> we see that the  $(N+\Delta-1)$ -samples-long output  $y(n)$  will be related to the  $(L+\Delta-1)$ -samples-long segments  $\underline{y}_k$  by means of the overlap-addition of equation (7-8). Note that in this equation we have implicitly assumed that the segment length  $L$  is larger than the (effective) length of the impulse response  $\Delta$ . If this is not the case then more than two consecutive output segments have to be overlapped for the creation of the output  $y(n)$ .

<sup>101</sup> Again, in the case where  $h(n)$  is of infinite length,  $M$  stands for the effective length we want to take into account in the calculation of the convolution (see the discussion in §A1)

$$\begin{bmatrix} y(-\Delta) \\ \vdots \\ y(-\Delta+L-1) \\ y(-\Delta+L) \\ \vdots \\ y(L-2) \\ \vdots \\ y(N-2) \end{bmatrix} = \begin{bmatrix} y_0(-\Delta) \\ \vdots \\ \frac{y_0(-\Delta+L-1)}{y_0(-\Delta+L)} \\ \vdots \\ y_0(L-2) \\ \vdots \\ \frac{y_1(-\Delta+L-1)}{y_1(-\Delta+L)} \\ \vdots \\ y_0(L-2) \end{bmatrix} + \dots + \begin{bmatrix} y_1(-\Delta) \\ \vdots \\ \frac{y_1(-\Delta+L-1)}{y_1(-\Delta+L)} \\ \vdots \\ y_0(L-2) \\ \vdots \\ \frac{y_{N/L-1}(-\Delta+L-1)}{y_{N/L-1}(-\Delta+L)} \\ \vdots \\ y_{N/L-1}(L-2) \end{bmatrix} \quad (7-8)$$

Combining then the forward and backward time cases presented above we see that the (infinitely long) full-length convolution of an  $N$ -samples-long input  $x(n)$  with a doubly infinite non-causal impulse response can be approximated up to the index  $\Delta$  in backward time and  $M-1$  in forward time by the  $(N+M+\Delta-1)$ -samples-long output  $y(n)$  computed in segments by means of the algorithm explained in the block diagram of figure 7-1. For the direct arrangement case considered here, each one of the blocks labelled “forward-backward time filtering algorithm” in figure 7-1 corresponds to the algorithm described in the block diagram of figure 7-2. In this block diagram, for the processing of the last input segment  $x_{N/L-1}(n)$  through the forward time branch we have to zero-pad the segment  $x_{N/L-1}(n)$  with  $M-1$  zeros at its end in order to compute the  $(M-1)$ -samples-long tail of the convolution output (hence the zero-padding of the dashed block is only applied for the filtering of the last input segment). As is further discussed below, with the basic assumption that the total computing time needed for the processing of each input segment is less than the time-duration of the segment, the algorithm of figure 7-1 can be implemented on-line.



**Figure 7-1:** Forward-backward time convolution in segments

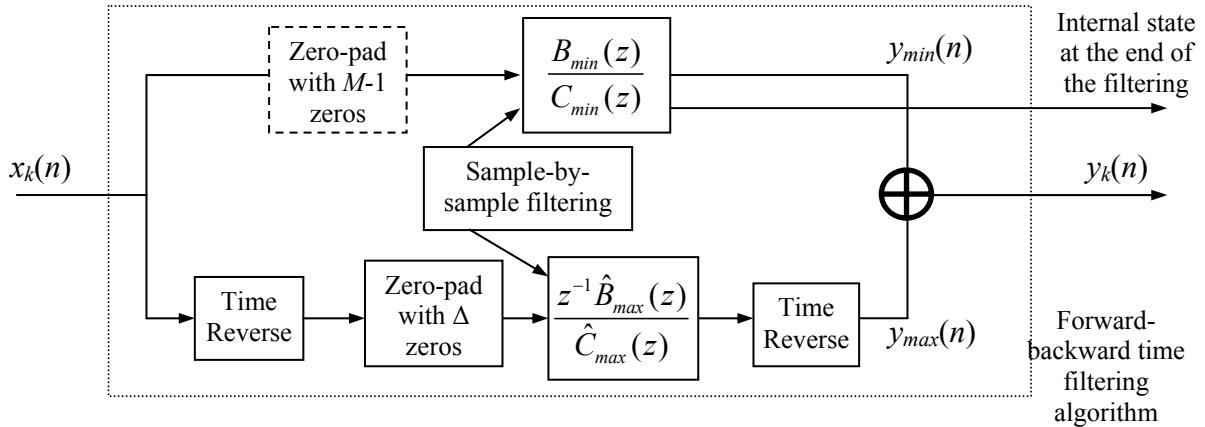


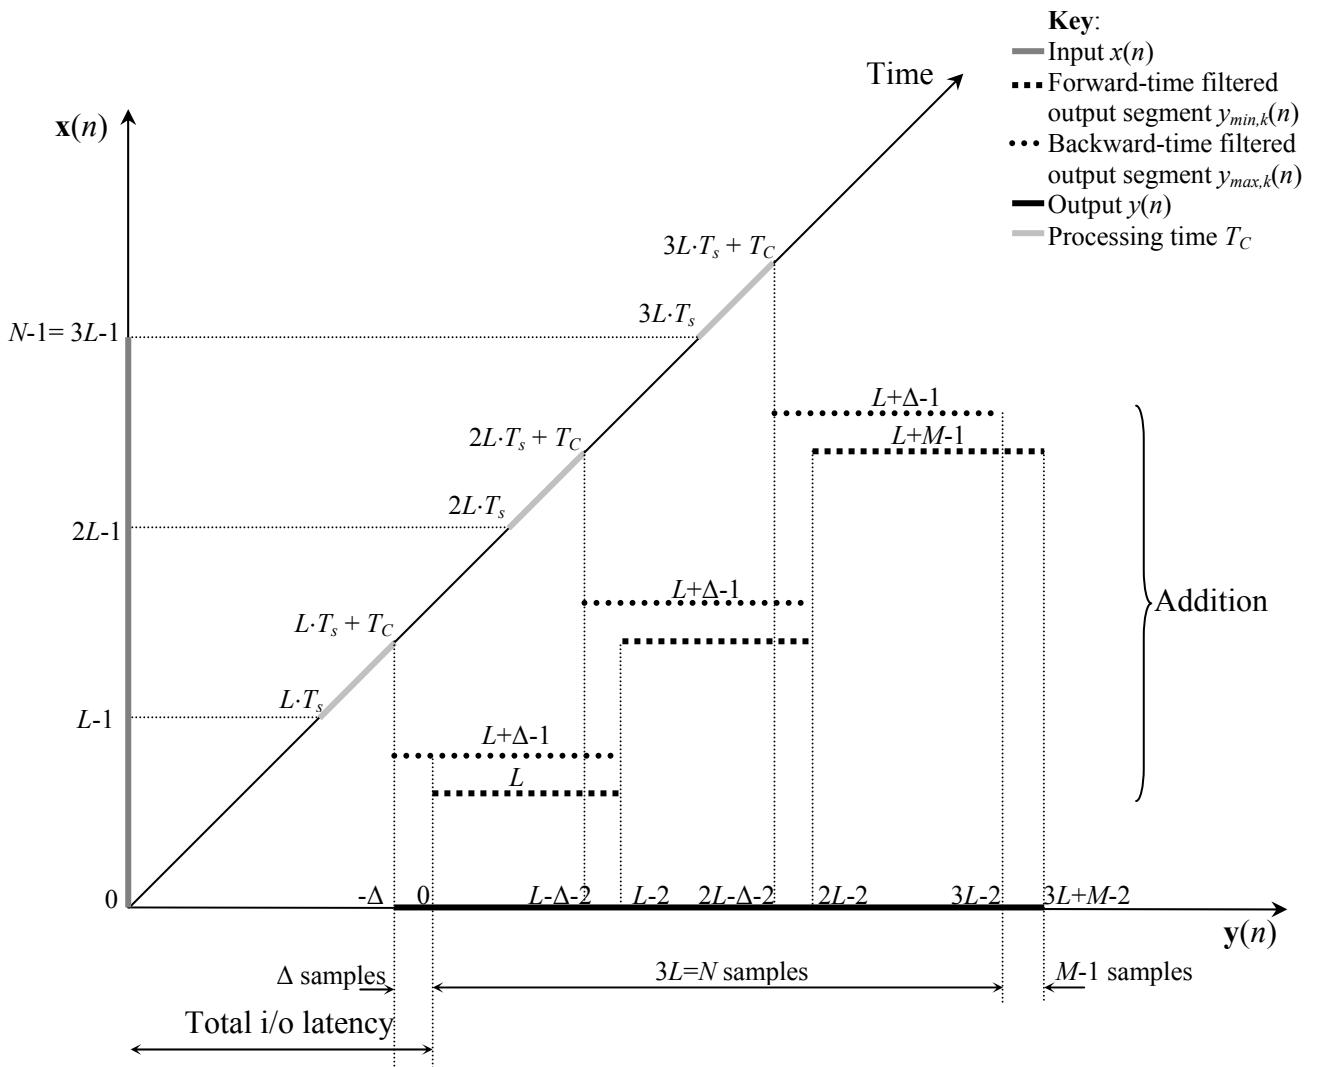
Figure 7-2: Forward-backward time convolution of each input segment

### 7.3 Computational cost and input-output latency of the on-line algorithm

We give here the computation cost for the implementation of the algorithm and the related delay between input and output. The on-line application of the forward-backward time filtering for each inverse response  $H_{ij}$  is described in the schematic of figure 7-3. In the figure, for notational simplicity, we have assumed that the input is segmented in 3 segments. In the vertical and horizontal axes of this figure we line up the input and output sequences while the third axis corresponds to the time progression and thus input and output samples projecting at the same point on this axis occur at the same point in time. In the same figure, with  $T_s$  we denote the sampling period and with  $T_C$  the processing time needed for the computation of each output segment  $y_k(n)$ . With this notation the real-time requirement is equivalent to  $T_C < L \cdot T_s$ , in other words the time  $T_C$  needed for the processing of the  $L$ -samples-long input segment has to be less than the time-duration  $L \cdot T_s$  of the segment.

As is shown in the figure, upon the acquisition of the first  $L$ -samples-long input segment  $x_0(n)$  at time  $L \cdot T_s$ , we start computing the first output segment  $y_0(n)$  which we start reproducing at time  $L \cdot T_s + T_C$ . As is explained in §7.2 the first  $L$  samples of  $y_0(n)$  correspond to the desired output values but to the last  $\Delta-1$  samples we have to add the

first  $\Delta-1$  samples of  $y_1(n)$ . As can be then seen in figure 7-3, the  $\Delta^{\text{th}}$  sample of  $y_0(n)$  occurs at time  $2L \cdot T_s + T_C$ , i.e. (if the real-time assumption  $T_C < L \cdot T_s$  holds) at the time when we have finished the processing of  $y_1(n)$  and we are thus in position to keep reproducing the properly overlap-added output. The same holds for the case where the segmenting length  $L$  becomes smaller than the anti-causal length  $\Delta$  but in this case it is easy to infer from figure 7-3 that the overlap-addition process will involve more than two consecutive output segments.



**Figure 7-3:** On-line implementation of the computation of the output  $y(n)$  of each of the four elements  $H_{ij}$  of the inverse matrix  $\mathbf{H}$ .

As is then illustrated in the figure, the total input-output latency of this on-line implementation will be described by equation (7-9) where the first term is due to the

non-causal approximation to the inverse and is the same in the FIR inversion. The remaining two terms correspond to the extra delay due to the segmenting and the time needed for the block processing.

$$\text{Latency} = \Delta \cdot T_s + L \cdot T_s + T_C \quad (7-9)$$

Assuming that the real-time-implementation condition holds, the input-output latency of equation (7-9) is readily seen to be bounded by the expression of equation (7-10). In other words, the total input-output latency will be no more than the sum of the modelling delay plus twice the block size.

$$\text{Latency} < (\Delta + 2 \cdot L) \cdot T_s \quad (7-10)$$

The computation cost  $\mathcal{O}_{sample}$  per sample-pair of the two-channel output  $\mathbf{y}(n)$  will be equal to the cost for the computation of each output segment  $\mathcal{O}_{seg}$  divided by the length  $L$  of the input segment. This segment cost  $\mathcal{O}_{seg}$  is given in equation (7-11) as a sum of three terms. The term  $4 \cdot N_{tot} \cdot L$  corresponds to the filtering of the  $L$ -samples-long segment through the 4 recursive structured modelling the elements of the inverse matrix  $\mathbf{H}$ . Hence in this term, the factor  $N_{tot}$  stands for the total number of coefficients needed to model each of the elements  $H_{ij}$ . The second term  $4 \cdot N_{max} \cdot \Delta$  corresponds to the extra zero-padding length  $\Delta$  that has to be filtered through the backward-time realisable part of the 4 recursive models. Hence in this term, the factor  $N_{max}$  stands for the total number of coefficients needed to model the anti-causal part of the response of each of the elements  $H_{ij}$ . The third term corresponds to the overlap-additions of the consecutive output segments. This last term is neglected since in the equation as it can be seen to always be very small compared to the remaining two terms<sup>102</sup>.

$$\begin{aligned} \mathcal{O}_{sample} &= \frac{\mathcal{O}_{seg}}{L} \\ &= \frac{(4 \cdot N_{tot} \cdot L + 4 \cdot N_{max} \cdot \Delta + \mathcal{O}_{overlap-add})}{L} \\ &\approx 4 \cdot N_{tot} + \frac{8 \cdot \Delta}{L} \cdot N_{max} \quad \text{madds/pair of output samples} \end{aligned} \quad (7-11)$$

<sup>102</sup> In the case where the segment length  $L$  is of comparable value with  $\Delta$  this term is negligible compared to the  $4 \cdot N_{tot}$  term while in the case  $L \ll \Delta$  where the factor  $8 \cdot \Delta / L$  becomes increasingly large, this term has again a negligible contribution to the total cost compared to the second term.

The computational cost of equation (7-10) for different input segment lengths is listed in table 7-1. This is done for the example-case of figures 6-21 and 6-22 (red line). In that case the total order for the modelling of each of the four elements  $H_{ij}$  of the inverse was  $65+45=110$  taps and the total order for the modelling of the anti-causal part of each of them was  $25+20=45$  taps. Note that this cost reflects strictly the number of multiplications/additions needed and does not take into account any extra arithmetic-specific load for the implementation of the recursive filtering (internal scalings to avoid overflows etc.). In the same table we list the total-input output delay of the implementation as estimated by equation (7-10) for a sampling rate of 48kHz. The same results of the computational cost  $\mathcal{O}_{sample}$  versus the input segment length are plotted (with the blue line) in figure 7-4. The total-input output delay corresponding to each input segment length is given in the top  $x$ -axis of the figure.

The performance of the proposed algorithm is compared in table 7-1 and figure 7-4 with that of three other algorithms for the computation of the filter output. These are;

- (a) The direct sample-by-sample computation of the output as described by the convolution sum of equation (6-2). This computation requires a number of multiplications and additions equal to the length  $N_H$  of the filter for the computation of each output sample and has an input-output equal to one sampling period.

$$\mathcal{O}_{direct} = N_H \quad \text{madds/output sample} \quad (7-12)$$

- (b) The well documented overlap-save algorithm (Press et al. 2002). In that case the convolution sum of equation (6-2) is computed by means of circular convolutions implemented in the frequency-domain. These circular convolutions are between properly overlapped segments of the input with the filter's impulse response. For the implementation of the algorithm, a segment of the input of length equal to the length of the impulse response has to be obtained first. Hence, the input-output latency in the implementation of this algorithm is typically taken to be twice the length of the impulse response of the filter. A detailed analysis of the computational cost for the

implementation of the overlap-save algorithm is presented by Gardner (1995) where it is demonstrated that this cost for a reasonably optimised realisation is equal to<sup>103</sup>

$$\mathcal{O}_{overlap-save} = \begin{cases} 3N_H \log_2 N_H + 6N_H & \text{multiplications/output sample} \\ \frac{9N_H}{2} \log_2 N_H + 8N_H & \text{additions/output sample} \end{cases} \quad (7-13)$$

(c) The zero-delay modification of the overlap-save algorithm introduced by Gardner (1995). In that algorithm, the impulse response of the filter is partitioned in segments of increasing length. The convolution with the first (and smallest) of these segments is implemented with direct sample-by-sample filtering. Hence the need to obtain a segment of the input beforehand is alleviated and the algorithm does not incur any extra input-output delay compared to the direct filtering case. The convolution with the remaining segments of the impulse response is implemented using overlap-save convolutions and thus reaps the computational efficiency of the FFT. The computational cost of this algorithm is approximated in (Gardner 1995) by<sup>104</sup>

$$\mathcal{O}_{zero-delay} = 34 \log_2 N_H - 151 \quad \text{multiplications/output sample} \quad (7-14)$$

It should be noted here, with reference to this algorithm, that a variant is described in (Gardner 1995) which trades a small amount of input-output delay for a further reduction of the computation cost.

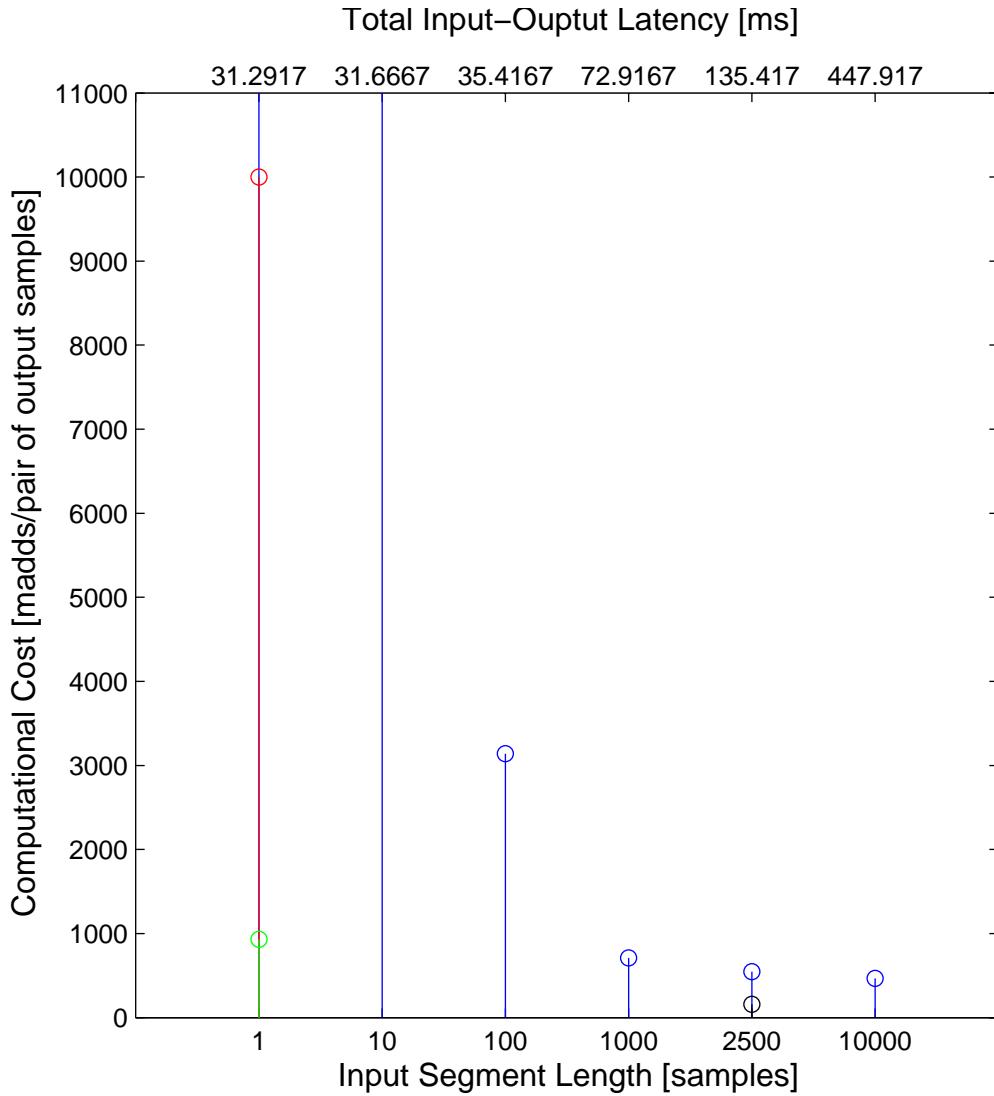
The performance estimates of the these three algorithms for the same inversion case as above are listed in table 7-1 and plotted in figure 7-4 in red, black and green line respectively. In this inversion case (dashed black line results of figures 6-21 and 6-22) the length of the FIR inverse filters was set to 2500 coefficients and the modelling delay to 1500 samples. The cost estimate for the computation of each pair of output samples will be four times the cost of equations (7-12), (7-13) and (7-14) (one for each element of the inverse filter matrix).

<sup>103</sup> As this algorithm is mainly about computing FFTs, the computational cost is not directly linked to the multiply/add process and hence the number of multiplications is different from that of additions. As is the usual practice, we estimate the cost on the basis of the number of multiplications.

<sup>104</sup> Again, we use the number of multiplications needed for the computation for each output sample to estimate the cost of the algorithm. It should also be noted that in the derivation by Gardner (1995) of both formulae (7-13) and (7-14) additional assumptions are made regarding the length  $N_H$  being equal to a power of 2 or equal to a sum of powers of 2. If this is not true, as for instance in the cases examined in Table 7-1, the associated costs should be expected to be slightly higher. Nevertheless we still use these formulae for the results of Table 7-1.

Algorithm	Input segment length [samples]	Computational cost [madds/pair of output samples]	Total input-output latency [ms]
<b>Forward-backward time filtering</b>	1	270440	31.3
<b>Forward-backward time filtering</b>	10	27440	31.7
<b>Forward-backward time filtering</b>	100	3140	35.4
<b>Forward-backward time filtering</b>	1000	710	72.9
<b>Forward-backward time filtering</b>	2500	548	135.4
<b>Forward-backward time filtering</b>	10000	467	447.9
<b>Direct sample-by-sample filtering</b>	1	10000	31.3
<b>Zero-delay frequency-domain method (Gardner 1995)</b>	1	931	31.3
<b>Overlap-save frequency-domain filtering</b>	2500	159	135.4

**Table 7-1:** Computational cost and total input-output latency of the introduced block-processing algorithm for various segment lengths. Computational cost and total input-output latency for the sample-by-sample FIR filtering, the zero-delay algorithm of (Gardner 1995) and the overlap-save algorithm.



**Figure 7-4:** Computational cost of the introduced block-processing algorithm versus segmenting length (blue line). Cost of direct sample-by-sample FIR filtering (red line). Cost of zero delay algorithm of (Gardner 1995) (green line). Cost of overlap-save algorithm (black line). The corresponding input-output delay is noted on the top axis.

Returning then to our forward-backward time filtering algorithm, as is displayed in the results of table 7-1 and figure 7-4, when the input segment length is set to be very small, the input-output delay approaches the hard limit of the implementation which is equal to the modelling delay. In that case, however, the cost of the algorithm overshoots. Evidently this is due to the influence of the second term in equation (7-11). In other words, the filtering of the extra zero-padded length has to be implemented at every input sample.

As the segment length is increased so is the total input-output latency. However, a segment length of 100 samples can be seen to result in a very modest increase from 31ms to 35ms and at a significant computational saving compared to the direct sample-by-sample FIR filtering. For input segment length of above 1000 samples the influence of the second term in equation (7-11) has completely abated and the cost of the proposed algorithm reaches its optimum. For this input segment length value the total input-output latency has only doubled compared to the lowest physically obtainable value.

As far as the frequency-domain filtering algorithms are concerned, the results show that, when the overlap-save algorithm is used, the computational cost saving is tremendous at the expense of a significantly increased input-output latency. At the other end, the zero-delay algorithm yields a quite significant reduction compared to the direct sample-by-sample FIR implementation without any increase of the input-output latency. A formulation achieving a trade-off between the input-output latency and the computational cost is given in (Gardner 1995). On the basis of that we can estimate that using the same input-output latency that our algorithm needs, the frequency-domain algorithm of (Gardner 1995) should achieve better results. Nevertheless, the results of figure 7-4 make clear that the computational cost of our algorithm is broadly on a par with that of the frequency-domain based algorithms. This in itself is far from trivial in the broader context of the comparison between recursive and frequency-domain filtering. Furthermore, it is quite probable that a more exhaustive examination of its performance in the wide range of cases covered throughout chapters 4 and 5 together with possible optimisations would prove the proposed method to be a preferable candidate in certain cases of implementation (see also the relevant discussion in §6.7).

## 7.4 Conclusions

In this chapter we presented a block-processing algorithm for the on-line implementation of the forward-backward time recursive filter structures introduced in chapter 6. Formulae for the overall input-output latency of the algorithm and the associated computational cost were given. As in any block-processing formulation,

the basic characteristic of the on-line implementation is the trade-off that was shown to exist between the input-output latency and the computational cost. It was demonstrated that our algorithm is impractical when the input-output latency needs to be set close to the hard limit physically imposed by the modelling delay. In those cases, the segment length has to be very small and, as was made apparent, the computational cost overshoots. Nonetheless, when the latency requirement is modestly relaxed to no more than double the minimum physically obtainable value, the cost of the algorithm drops dramatically. Hence the algorithm showed genuine promise for implementations where such a minor amount of extra input-output delay is tolerable.

Compared to the frequency-domain based algorithms, two possible advantages of the proposed algorithm already at this first stage of its development are (i) the significantly lower order of the filter structures involved and hence the lower memory required for its implementation and (ii) the fact that its design and realisation does not depend on the length of the inverse responses to be modelled, unlike the zero-delay frequency-domain filtering algorithm where the optimal partitioning of the inverse filters' impulse response is not trivial for any filter length.

## ***8 Conclusions – Proposed further research***

The objective of the research presented in this thesis was the investigation of a specific signal processing design that has previously been proposed for the determination of the control signals needed for the reproduction of a pair of binaural signals at the ears of a single static listener. The further development of the design in order to replace the used FIR filters with recursive filters was undertaken. A detailed examination of the properties, the capacity and the limitations of this design was carried that was based on objective experiments. Based on this examination, a significant amount of conclusions related to the real-world implementation of the design were drawn that were not identifiable from previous studies based on subjective experiments or computer simulations. A method was introduced for the implementation of the inverse using forward-backward time recursive filtering. The application of the method was exemplified using the results of the objective experiments and it was shown to achieve a significant reduction of the required order of the inverse models. The method was modified to accommodate its on-line implementation and its properties were compared to those of the already available highly efficient frequency-domain filtering methods.

The key findings of the presented research can thus be summarised as follows.

- In the first set of results of chapter 4 we established and quantified the optimal inversion accuracy that can be realistically achieved with the inverse design under consideration. We showed that under ideal conditions of implementation highly accurate inversion results are possible. In such cases it is possible to obtain ipsilateral equalisation confined to  $\pm 1\text{dB}$  in the region above 200Hz and cross-talk cancellation of 15-30dB in that region. On the other hand, it was demonstrated in chapter 4 that these results are restricted to the case where the plant corresponding to the individual listener and the specific reproduction system is measured *in situ*. In cases where such an *in situ* measurement is not possible (and hence the late decay part of the plant's model is not available) a perceptually significant amount of error is introduced even when knowledge of the strictly anechoic part of the plant corresponding to the

specific listener and reproduction system is available. Furthermore it was shown that the aforementioned optimal inversion accuracy is only possible if a small degree of sub-optimality is allowed in advance by the introduction of a small amount of regularisation.

- The next stage of approximation, namely the cases where (i) the inverse filters are determined based on plant models that contain HRTFs and hence do not contain the transfer functions of the specific reproduction chain and (ii) do not describe the individual listener's HRTF, were considered in §4.5 and §5.4 respectively. In the former case it was shown that a significant amount of error is introduced due to the fact that the responses of the reproduction loudspeakers are not corrected which in practice relaxes the requirement of accuracy in the inversion. Thus higher regularisation can be applied and, consequently, shorter inverse filters can be used. In the latter case, using data from the CIPIC database, we demonstrated that a number of listeners are indeed well represented by the KEMAR-based HRTF model while others are not. It is thus possible that a collective database of a few dummy-head measurements of various sizes and shapes could cover all individuals. On the other hand the results made apparent that, even when individual HRTFs are used for transaural reproduction in anechoic conditions, the inversion of the frequency range above 10kHz is not justified. Hence there is considerable room for complexity reduction and computational savings if the inverse filtering design is implemented at lower sampling rates
- The issue of the loss of dynamic range associated with the application of the inverse filtering was also highlighted with the aid of the results of chapter 4. The role of regularisation as a means to address this issue was examined in detail and it was shown that in certain cases its use can considerably enhance the overall dynamic range of the inversion without significantly, if at all, degrading its accuracy. The use of a frequency-varying regularisation penalty was also examined. We showed that, in cases where the very low frequency content can be reproduced by a separate transducer circuit, this frequency-varying regularisation can achieve a further increase of the dynamic range by acting as linear-phase frequency separation stage. The overall evaluation of the actual advantages of such an approach (especially in view of

the increased requirement for the inverse filters' length that it entails) was left outside the scope of the thesis.

- The justification of the choice to invert the frequency region below 200Hz as well as that of correcting the phase of the plant were also examined §4.5 where it was demonstrated that the region below 200Hz is practically uncontrollable with a system of the type considered here. As far as the second issue is concerned, it was shown that the equalisation of the phase response of the plant is indeed possible if an accurate (measured *in situ*) model of the plant is available, but that such an attempt to correct the phase is probably futile in any other case as in such cases the use of a minimum-phase approximation to the inverse yields nearly equivalent results.
- The case where the reproduction arrangement contains strong early reflections was considered in §5.2. It was shown that when a model of the plant is available that contains the effect of these reflections, their presence can be efficiently treated by the inverse and the inversion results can be of equivalent accuracy to that of the anechoic inversion. This, however, incurs an increased requirement for the inverse filters' length. Conversely, if such early and strong reflections are ignored in the determination of the inverse, the quality of the inversion is dramatically influenced. In such an event, it was demonstrated that the best level of achievable inversion accuracy can be obtained with a very compact, heavily regularised model of the inverse and that any higher specification for the inverse is pointless.
- Furthermore, the case where positioning error is present in the reproduction arrangement was considered in §5.3. The measured results verified an expected effect, namely that the high frequency part of the inversion begins to fail for positional errors of the order of a few millimetres. The error in the inversion exhibited a systematic pattern which was explained with the use of an analytical model. The excellent agreement between the measured results and the proposed analytical model shows that (if an accurate knowledge of the position of the listener is available) the error introduced by the misalignment up to a few tens of millimetres can be efficiently alleviated by the introduction of fractional-sample delays and constant scaling in the inverse filters without any need to recalculate the inverse filters. Conversely, when such highly accurate knowledge of the listener's position is not available, the results verify that the inversion of the region above 10kHz is not justified. The case of

positional error of the order of a few centimetres was also examined. The measured results verified a fact previously observed in subjective experiments, namely that the sweet-spot size of the Stereo Dipole geometry for the creation of virtual images in the frontal half of the horizontal plane extends to approximately  $\pm 8\text{cm}$  to the left and right of the listener.

- Gathering all the above cases we can conclude that the required filter length for the implementation of the FIR-based inversion method under examination ranges from several tens of thousands of coefficients (for the highly accurate inversion case where the inverse is determined on the base of a plant model measured *in situ* and frequency varying regularisation is introduced) to no more than a few hundreds of coefficients, when one or more of the aforementioned elements of approximation and error are present.
- In chapter 6 we introduced a technique for the decomposition of the elements  $H_{ij}$  of the inverse matrix  $\mathbf{H}$  in two parts, one realisable in forward time and one in backward time. We showed that with the use of this technique, an overwhelmingly accurate recursive representation of the inverse can be obtained. We noted however that in such a form the technique is of limited practical interest, mainly due to the excessively high order of the resulting recursive structures. Two modifications were proposed in the method for the determination of the recursive filters that model the forward and backward time realisable parts of the inverse. One of them, formulated in the time-domain, was found to give very promising results in terms of reducing the order of the involved recursive filters without reducing the accuracy of the inversion below the level that one would realistically expect to achieve in real-world implementations. The evaluation of the proposed recursive inverse modelling method was confined to a small “sample” set among the inversion cases previously examined in chapters 4 and 5 while a more comprehensive evaluation is left a further work.
- Finally, in chapter 7, we introduced a block-processing algorithm for the on-line implementation of the forward-backward time recursive filter structures of chapter 6. We examined the overall input-output latency and the associated computational cost of the algorithm and gave the related formulae. As in any block-processing formulation, the basic characteristic of the on-line implementation is the trade-off that was shown to exist between the input-output latency and the computational cost. It

was demonstrated that our algorithm is impractical when the input-output latency is set close to the hard limit physically imposed by the modelling delay. Nonetheless, when the latency requirement is modestly relaxed to no more than double the minimum physically obtainable value, the cost of the algorithm drops dramatically. Hence the algorithm showed genuine promise for implementations where such a minor amount of extra input-output delay is tolerable

A number of further research directions of interest were identified during the course of this project. These are the following:

- **Subjective verification of the results of chapters 3-5:** We propose the validation of the findings of our objective experiments with subjective tests. Even though a significant amount of published research exists in this area, little knowledge exists with respect to the actual requirements of the inverse filter design that would achieve the best realistically obtainable results in different implementation situations. Based on the results of chapters 3-5, we believe that properly designed subjective experiments that simulate the real-world implementation of the system would show that a significant downgrade of the inverse filter design specifications could be achieved. Again based on those results we believe that the, admittedly cumbersome, additional requirement for the *in situ* measurement of the plant matrix could be proven with properly designed subjective experiments to be a justifiable means for the improvement of the system's performance.
- **Separate implementation of the cross-talk cancellation and ipsilateral equalisation stages** (see also §2.4.2): Even though such an implementation has been considered in the past, we believe that it has not been adequately exploited in the context of the inverse filtering design examined in this thesis. Some experimental results obtained during the acquisition of the measurements presented in chapters 3-5 show that, when an *in situ* measurement of the plant is available, a separate cross-talk cancellation stage can be made to work very effectively without the limitations inherent in the inverse filtering. In other words, the perfect cross-talk cancellation of figure 2-7 (which in the discussion of §2.4.2 was deemed to be an artificial effect of the simulation process) can to some extent be replicated in the actual implementation of the inversion. The overall-equalisation inverse filtering stage can then be implemented separately with the introduction of a significant amount of regularisation

and hence a more modest inversion specification. In this way, a tolerable amount of colouration would be introduced but the dynamic range issues would be addressed much more efficiently. Minimum-phase approximations could also be more suited to such a design.

- **Application of the forward-backward time recursive modelling to a wider selection of implementation cases of chapter 3:** As was also mentioned in the main text, time limitations did not allow the extensive application of the proposed recursive modelling methods to the wide range of cases that were presented in chapter 3. Given the increased amount of error in the inversion in certain implementation cases (especially those of §5), we expect the undertaking of such a task to show that much more compact recursive models can be achieved in those cases.
- **Improvement of the recursive modelling methods of §6.5:** General purpose Matlab functions were used for the recursive modelling methods presented in §6.5. It is thus probable that the obtained models were not optimal. This is supported by the fact that using other time-domain model-fitting methods (Prony's method and Shank's method) we obtained quite varying results. A more detailed investigation of the issue could thus result in more efficient modelling.
- **Application of the reduced order recursive models to lower-precision arithmetics:** Assuming that the introduced recursive inverse models reach a stage of development where their use becomes clearly advantageous, the possibility of their implementation in platforms of lower precision arithmetic (than the double-precision used here) should be examined. This is fairly straightforward to do in modern computing packages such as Matlab.
- **Examination of the possibility for optimisation of the block-processing algorithm:** Finally, it is possible that some optimisation could be achieved in the simple block-processing algorithm presented here. To this end, the block-processing form of the cascade connection arrangement should also be considered. One issue to be addressed, should such a research direction was to be followed, would be the modification of the algorithm to achieve uniform computational load over the segment length.

# Appendix

## A1. Convolution and filtering in backward time

In this section we give the analytical expressions and formulae describing the notions of full-length convolution and linear filtering in the context of forward and backward filtering through recursive filters. We describe how full-length convolution with infinitely long impulse responses is implemented in the results presented in the thesis using the `filter` command, which is the standard way of implementing recursive filtering in Matlab.

The formula for the convolution of two finite length sequences  $x(n)$  and  $h(n)$  that are non-zero in the intervals  $N_1 \leq n \leq N_2$  and  $M_1 \leq n \leq M_2$  respectively<sup>105</sup>, is given in equation (A1-1) from which it is readily seen that the output signal  $y_c(n)$  will be, generally, non-zero for the interval  $N_1+M_1 \leq n \leq N_2+M_2$ . Clearly, for an infinitely long input impulse response  $h(n)$  the convolution output  $y_c(n)$  will also be infinitely long and can thus be only approximately realised.

$$y_c(n) = \sum_{k=-\infty}^{\infty} x(k)h(n-k) = \sum_{k=N_1}^{N_2} x(k)h(n-k) \quad (\text{A1-1})$$

On the other hand, in the case of sample-by-sample filtering of the input  $x(n)$  through the filter  $h(n)$  resulting in the output  $y_f(n)$  we typically take  $y_f(n)$  to be of the same length  $N=N_2-N_1+1$  as the input, irrespective of the length  $M=M_2-M_1+1$  of the filter's impulse response (this is for example the case with Matlab's `filter` function). This practically means that we consider the filter's output  $y_f(n)$  to cease as soon as the input  $x(n)$  stops whereas, due to the “memory” of the impulse response  $h(n)$ , the convolution sum of equation (A1-1) continues to produce output even after the input

---

<sup>105</sup> With  $N_i, M_i$  either positive or negative integers

has stopped. The corresponding formula for the sample-by-sample filtering will then be as in equation (A1-2).

$$y_f(n) = \begin{cases} \sum_{k=-\infty}^{\infty} x(k)h(n-k) = \sum_{k=N_1}^{N_2} x(k)h(n-k), & N_1 + M_1 \leq n \leq N_2 + M_1 \\ 0 & , \text{ otherwise} \end{cases} \quad (\text{A1-2})$$

Equivalently, in the terminology introduced in §6.2, the operation of sample-by-sample filtering will translate to the matrix-vector equation  $\underline{y}_f = \mathcal{H}_N \underline{x}$  where the  $(N \times N)$  Toeplitz matrix  $\mathcal{H}_N$  is defined as in equation (A1-3) and where, depending on the relative lengths of  $x(n)$  and  $h(n)$ , either a triangle of the lower-right part of the matrix will be equal to zero or some of the last samples of  $h(n)$  will not be present in the matrix.

$$\mathcal{H}_N = \begin{bmatrix} h(M_1) & 0 & \cdots & 0 \\ h(M_1+1) & h(M_1) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h(M_2 - M + N) & h(M_2 - M + N - 1) & \cdots & h(M_1) \end{bmatrix} \quad (\text{A1-3})$$

We then see that the statement  $((\mathcal{H}^N)^C)^R = \hat{\mathcal{H}}^N$  of equation (6-5) which holds for the full-length convolution formula (A1-1) does not hold for the filtering formula (A1-2) as the filtering matrix  $\mathcal{H}_N$  does not satisfy the corresponding algebraic identity:  $((\mathcal{H}_N)^C)^R \neq \hat{\mathcal{H}}_N$ . In other words, the output  $u_f(n)$  that we obtain by filtering  $\hat{x}(n)$  through  $\hat{h}(n)$  is not equal to the time reversed version of  $y_f(n)$ :  $\hat{u}_f(n) \neq y_f(n)$ .

However, it is easy to see that when the input signal  $x(n)$  is replaced by  $x_{M-1}(n)$ , a zero-padded version with  $M-1$  zeros at its end, the corresponding  $(N+M-1 \times N+M-1)$  filtering matrix  $\mathcal{H}_{N+M-1}$  will be as in equation (A1-4) and equation (A1-5) will hold<sup>106</sup>. In other words, the convolution of the input  $x(n)$  with the  $M$ -samples-long impulse response  $h(n)$  is equivalent to the filtering of  $x_{M-1}(n)$  through  $h(n)$ .

---

<sup>106</sup> The partitioning lines in the matrix  $\mathcal{H}_{N+M-1}$  in equation (A1-4) correspond to  $N$  and  $M-1$  number of rows and columns with  $N$  and  $M$  the lengths of the input and the impulse response respectively.

$$\underline{\mathbf{y}} = \begin{bmatrix} y(M_1 + N_1) \\ \vdots \\ y(M_2 + N_2) \end{bmatrix} = \mathcal{H}_{N+M-1} \underline{\mathbf{x}}_{M-1} =$$

$$\left[ \begin{array}{cccc|ccc} h(M_1) & 0 & \cdots & 0 & 0 & \cdots & 0 \\ h(M_1+1) & h(M_1) & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ h(M_2-M+N) & h(M_2-M+N-1) & \cdots & h(M_1) & 0 & \cdots & 0 \\ \hline h(M_2-M+N+1) & h(M_2-M+N) & \cdots & h(M_1+1) & h(M_1) & \cdots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h(M_2) & h(M_2-1) & \cdots & h(M_1) \end{array} \right] \begin{bmatrix} x(N_1) \\ \vdots \\ x(N_2) \\ \hline 0 \\ \vdots \\ 0 \end{bmatrix} \quad (\text{A1-4})$$

$$\mathcal{H}_{N+M-1} \underline{\mathbf{x}}_{M-1} = \underline{\mathbf{y}} = \mathcal{H}^N \underline{\mathbf{x}} \quad (\text{A1-5})$$

Furthermore, as is shown in equation (A1-6), in the backward time filtering case the time-reversed full-length convolution output  $\hat{y}(n)$  can be obtained when  $\hat{x}_{M-1}(n)$ , i.e. a zero-padded version of the time-reversed input  $\hat{x}(n)$  again with  $M-1$  zeros at its end, is filtered through the time-reversed impulse response  $\hat{h}(n)$  corresponding to the filtering matrix  $\hat{\mathcal{H}}_{N+M-1}$ .

$$\hat{\mathcal{H}}^N \hat{\underline{\mathbf{x}}} = \hat{\underline{\mathbf{y}}} = \begin{bmatrix} y(M_2 + N_2) \\ \vdots \\ y(M_1 + N_1) \end{bmatrix} = \hat{\mathcal{H}}_{N+M-1} \hat{\underline{\mathbf{x}}}_{M-1} =$$

$$\left[ \begin{array}{cccc|ccc} h(M_2) & 0 & \cdots & 0 & 0 & \cdots & 0 \\ h(M_2-1) & h(M_2) & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ h(M_2-N+1) & h(M_2-N+2) & \cdots & h(M_2) & 0 & \cdots & 0 \\ \hline h(M_2-N) & h(M_2-N+1) & \cdots & h(M_2-1) & h(M_2) & \cdots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h(M_1) & h(M_1+1) & \cdots & h(M_2) \end{array} \right] \begin{bmatrix} x(N_2) \\ \vdots \\ x(N_1) \\ \hline 0 \\ \vdots \\ 0 \end{bmatrix} \quad (\text{A1-6})$$

We now consider the specific case where a causal  $N$ -samples-long input  $x(n)$  is to be convolved with the infinitely long and strictly causal impulse response  $h_{IIR}(n)$  in which case the convolution matrix  $\mathcal{H}_{IIR}^N$  will have an infinite number of rows. If however  $h_{IIR}(n)$  is decaying in positive time we can assume that it will be insignificantly small after a given index  $M$ :  $h_{IIR}(n) \approx 0$  for  $n \geq M$ , with  $M$  a positive integer. In that case the infinite length convolution  $h_{IIR}(n)^*x(n)$  can be approximated

to an arbitrary degree of accuracy by the  $(M+N-1)$ -samples-long convolution  $h(n)*x(n)$  with  $h(n)$  the finite length sequence of equation (A1-7). In the case then where sample-by-sample filtering is used for the computation of the convolution we see that in order to approximate the infinite length convolution up the index  $(M+N-1)$  we need to zero pad the input  $x(n)$  with  $M-1$  zeros at its end.

$$h(n) = \begin{cases} h_{IIR}(n) & , \quad 0 \leq n < M \\ 0 & , \quad \text{otherwise} \end{cases} \quad (\text{A1-7})$$

Next, we consider the case where a causal  $N$ -samples-long input  $x(n)$  is to be convolved with the infinitely long and strictly anti-causal impulse response  $h_{IIR}(n)$  of which we have a time-reversed model  $\hat{h}_{IIR}(n)$ . As before we assume that  $h_{IIR}(n)$  becomes insignificantly small for indices smaller than  $-M$ :  $h_{IIR}(n) \approx 0$  for  $n < -M$ , with  $M$  a positive integer and similarly for its time-reversed model  $\hat{h}_{IIR}(n)$ . In that case the  $(M+N-1)$ -samples-long a-causal convolution output  $y(n)=h(n)*x(n)$  will be as in equation (A1-8), where depending on the relative lengths of  $h(n)$  and  $x(n)$  some of the terms appearing as  $h(-N)$ ,  $h(-N+1)$  etc. in the convolution matrix will be falling out of the range  $0 < -n \leq M$  and will thus be considered of insignificantly small value.

$$\underline{y} = \underline{\mathcal{H}}^N \underline{x} =$$

$$= \begin{bmatrix} h(-M) & 0 & \cdots & 0 \\ h(-M+1) & h(-M) & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ h(-1) & h(-2) & \cdots & h(-N) \\ 0 & h(-1) & \cdots & h(-N+1) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h(-1) \end{bmatrix} \begin{bmatrix} x(0) \\ \vdots \\ x(N-1) \end{bmatrix} = \begin{bmatrix} y(-M) \\ \vdots \\ y(0) \\ \vdots \\ y(N-2) \end{bmatrix} \quad (\text{A1-8})$$

We now take the case where this convolution is implemented by sample-by-sample filtering through the time-reversed finite-length impulse response  $\hat{h}(n)$ , an approximation to infinitely long time-reversed model  $\hat{h}_{IIR}(n)$  up to the sample  $n=M$ . In this case, as said above, we have to use a zero-padded version  $\hat{x}_{M-1}(n)$  of the time-reversed input in which case the time-reversed output will be as in equation (A1-9). It becomes thus evident from equation (A1-9) that the amount of zero-padding applied to the time-reversed input signal determines the (anti-causal) length of the infinite

impulse response  $h_{IIR}(n)$  taken into account in the computation of the output  $y(n)$  and is in effect equivalent to the modelling delay that has to be introduced in order to make the indices of the output correspond to positive time:  $\Delta=M$ .

$$\begin{aligned}
 & \begin{bmatrix} \hat{y}(0) \\ \vdots \\ \hat{y}(N-1) \\ \hline \hat{y}(N) \\ \vdots \\ \hat{y}(N+M-1) \end{bmatrix} = \hat{\mathbf{y}} = \hat{\mathbf{H}}^N \hat{\mathbf{x}} = \hat{\mathbf{H}}_{N+M-1} \hat{\mathbf{x}}_{M-1} = \\
 & \begin{bmatrix} \hat{h}(0) & 0 & \cdots & 0 & | & 0 & \cdots & 0 \\ \hat{h}(1) & \hat{h}(0) & \cdots & 0 & | & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & | & \vdots & \cdots & \vdots \\ \hat{h}(N) & \hat{h}(N-1) & \cdots & \hat{h}(0) & | & 0 & \cdots & 0 \\ \hline \hat{h}(N+1) & \hat{h}(N) & \cdots & \hat{h}(1) & | & \hat{h}(0) & \cdots & 0 \\ \vdots & \vdots & & \vdots & | & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \hat{h}(M) & | & \hat{h}(M-1) & \cdots & \hat{h}(0) \end{bmatrix} = \\
 & \begin{bmatrix} h(-1) & 0 & \cdots & 0 & | & 0 & \cdots & 0 \\ h(-2) & h(-1) & \cdots & 0 & | & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & | & \vdots & \cdots & \vdots \\ h(-N) & h(-N+1) & \cdots & h(-1) & | & 0 & \cdots & 0 \\ \hline h(-N-1) & h(-N) & \cdots & h(-2) & | & h(-1) & \cdots & 0 \\ \vdots & \vdots & & \vdots & | & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h(-M) & | & h(-M+1) & \cdots & h(-1) \end{bmatrix} = \\
 & \begin{bmatrix} x(N-1) \\ \vdots \\ x(0) \\ \hline 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} y(N-2) \\ \vdots \\ y(-1) \\ \hline y(-2) \\ \vdots \\ y(-M) \end{bmatrix} \quad (A1-9)
 \end{aligned}$$

Combining then the two cases described above, we see that in order to approximate the infinitely long convolution of the  $N$ -samples-long input  $x(0 \leq n < N)$  with the infinitely long impulse response  $h(n) = Z^{-1}\{1/C(z)\}$  in the interval  $-\Delta \leq n < N+M_2-1$  by means of the parallel connection (see figure 6-1 in the main text) using sample-by-sample filtering, we have to zero pad the input  $x(n)$  and its reversed version  $\hat{x}(n)$  with  $M_2-1$  and  $\Delta$  zeros respectively in order to compute the  $N+M_2-1$  and  $N+\Delta-1$  outputs  $y_{min}(n)$  and  $y_{max}(n)$ . These outputs have then to be shift-added with the  $(\Delta+1)^{th}$  sample of  $y_{max}(n)$  (i.e. the one with time index  $n=0$ ) added to the  $0^{th}$  sample of  $y_{min}(n)$  and so on, in order to acquire the  $(N+\Delta+M_2)$ -samples-long output  $y(n)$ , the first  $\Delta$  samples of which will correspond to negative time.

Finally, in the case of the cascade connection of figure 6-2 of the main text, we see that the decomposition is equivalent to the factorisation of the (infinite rows) convolution matrix  $\mathcal{H}^N$  corresponding to the non-causal impulse response  $h(n)=Z^{-1}\{1/C(z)\}$  as a product of the (infinite rows) convolution matrices  $\mathcal{H}_{caus}^N$  and  $\mathcal{H}_{acaus}^N$  corresponding to the strictly causal and anti-causal impulse responses  $h_{caus}(n)=Z^{-1}\{1/C_{min}(z)\}$  and  $h_{acaus}(n)=Z^{-1}\{1/C_{max}(z)\}$ . Assuming as before that  $h(n)\approx 0$  for  $n \leq -M_1$  and  $n \geq M_2$  (with  $M_1$  and  $M_2$  positive integers) the implementation of the convolution of an  $N$ -samples-long input  $x(n)$  with  $h(n)$  can be implemented by sample-by-sample filtering the signal  $x_{M_2-1}(n)$  through  $h_{caus}(n)$ , time-reversing it, zero-padding it with  $M_1-1$  zeros at its end, filtering it through  $\hat{h}_{acaus}(n)$  and time-reversing it again to obtain the  $(N+M_2+M_1-2)$ -samples-long output  $y(n)$ , the first  $M_1-1$  samples of which will again correspond to negative time indices.

## A2. Proof of the matrix identity of §6.2

In this section we give a proof of the matrix identity of equation (A2-1) that was used in the formulation of §6.2.

$$\left(\left(\mathbf{A}\right)^C\right)^R \left(\mathbf{v}\right)^R = \left(\mathbf{Av}\right)^R \quad (\text{A2-1})$$

In (A2-1),  $\mathbf{A}$  is a matrix with, say,  $M$  rows and  $N$  columns and  $\mathbf{v}$  is a vector with  $N$  elements. As was introduced in §6.2, the notation  $(\cdot)^R$  and  $(\cdot)^C$  is used to denote the reversal of rows and reversal of columns respectively of the matrix or vector inside the brackets. We also use the notation  $[\mathbf{A}]_{i,j}$  to denote the element in the  $i^{\text{th}}$  row and the  $j^{\text{th}}$  column of the matrix  $\mathbf{A}$  and similarly  $[\mathbf{v}]_i$  to denote the element in the  $i^{\text{th}}$  row of the vector  $\mathbf{v}$ . Hence, the algebraic definition of  $(\cdot)^R$  for the matrix  $\mathbf{A}$  will be

$$\left[\left(\mathbf{A}\right)^R\right]_{i,j} \triangleq [\mathbf{A}]_{M-i+1,j} \quad (\text{A2-2})$$

the algebraic definition of  $(\cdot)^C$  for the matrix  $\mathbf{A}$  will be

$$\left[\left(\mathbf{A}\right)^C\right]_{i,j} \triangleq [\mathbf{A}]_{i,N-j+1} \quad (\text{A2-3})$$

and the algebraic definition of  $(\cdot)^R$  for the vector  $\mathbf{v}$  will be

$$\left[\left(\mathbf{v}\right)^R\right]_i \triangleq [\mathbf{v}]_{N-i+1} \quad (\text{A2-4})$$

Finally, the definition of matrix-vector multiplication in the introduced notation will be

$$[\mathbf{Av}]_i = \sum_{k=1}^N [\mathbf{A}]_{i,k} [\mathbf{v}]_k \quad (\text{A2-5})$$

Hence, we see that the element in the  $i^{\text{th}}$  row of the  $M \times 1$  vector in the right-hand side of equation (A2-1) will be equal to

$$\begin{aligned}
 [(\mathbf{A}\mathbf{v})^R]_i &= [\mathbf{A}\mathbf{v}]_{M-i+1} \\
 &= \sum_{k=1}^N [\mathbf{A}]_{M-i+1,k} [\mathbf{v}]_k
 \end{aligned} \tag{A2-6}$$

whereas, the same element of the vector in the left-hand side of equation (A2-1) will be equal to

$$\begin{aligned}
 \left[ \left( (\mathbf{A})^C \right)^R (\mathbf{v})^R \right]_i &= \sum_{k=1}^N \left[ \left( (\mathbf{A})^C \right)^R \right]_{i,k} \left[ (\mathbf{v})^R \right]_k \\
 &= \sum_{k=1}^N \left[ (\mathbf{A})^C \right]_{M-i+1,k} [\mathbf{v}]_{N-k+1} \\
 &= \sum_{k=1}^N [\mathbf{A}]_{M-i+1,N-k+1} [\mathbf{v}]_{N-k+1}
 \end{aligned} \tag{A2-7}$$

and substituting the running index  $k$  that runs the interval from 1 to  $N$  with the index

$$\tilde{k} = N - k + 1 \tag{A2-8}$$

which runs the same interval in the opposite direction from  $N$  to 1 we can write equation (A2-7) as

$$\begin{aligned}
 \left[ \left( (\mathbf{A})^C \right)^R (\mathbf{v})^R \right]_i &= \sum_{\tilde{k}=N}^1 [\mathbf{A}]_{M-i+1,\tilde{k}} [\mathbf{v}]_{\tilde{k}} \\
 &= \sum_{\tilde{k}=1}^N [\mathbf{A}]_{M-i+1,\tilde{k}} [\mathbf{v}]_{\tilde{k}}
 \end{aligned} \tag{A2-9}$$

The equality of the right-hand sides of equations (A2-6) and (A2-9) proves the intended.

## ***References***

R. M. Aarts, A. W. M. Mathijssen, P. C. W. Sommen, and J. Garas (1998) "Efficient block frequency domain filtered-x applied to phantom sound source generation" presented at the AES 104th Convention, Amsterdam.

AES, Ed. (1986). *Stereophonic Techniques*, New York, Audio Engineering Society.

V. R. Algazi, R. O. Duda, D. M. Thompson, and C. Avendano (2001) "The CIPIC HRTF database" presented at the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics.

F. Asano, Y. Suzuki, and T. Sone (1990) "Role of Spectral Cues in Median Plane Localization" *Journal of the Acoustical Society of America*, vol. 88, pp. 159-168.

K. J. Åström and B. Wittenmark (1997) *Computer-controlled systems theory and design*, 3rd international ed. Upper Saddle River, N.J.: Prentice Hall.

B. S. Atal and M. R. Schroeder (1962) "Apparent sound source translator", *U.S. Patent*, 3,236,949.

J. Bauck and D. H. Cooper (1996) "Generalized transaural stereo and applications" *Journal of the Audio Engineering Society*, vol. 44, pp. 683-705.

B. B. Bauer (1961) "Stereophonic Earphones and Binaural Loudspeakers" *Journal of the Audio Engineering Society*, vol. 9, pp. 148-151.

D. R. Begault and E. M. Wenzel (1993) "Headphone Localization of Speech" *Human Factors*, vol. 35, pp. 361-376.

D. R. Begault (1994) *3-D sound for virtual reality and multimedia*. London: Academic Press Ltd..

J. S. Bendat and A. G. Piersol (1993) *Engineering applications of correlation and spectral analysis*, Wiley, New York.

A. J. Berkhouit, D. de Vries, and P. Vogel (1993) "Acoustic Control by Wave Field Synthesis" *Journal of the Acoustical Society of America*, vol. 93, pp. 2764-2778.

A. J. Berkhouit (1988) "A Holographic Approach to Acoustic Control" *Journal of the Audio Engineering Society*, vol. 36, pp. 977-995.

J. Blauert and P. Laws (1978) "Group delay distortions in electroacoustical systems" *The Journal of the Acoustical Society of America*, vol. 63, pp. 1478-1483.

J. Blauert (1996) *Spatial hearing : The psychophysics of human sound localization*. Cambridge, Mass.: MIT.

J. Borwick (1990) *Microphones, Technology and Technique*. London: Focal Press.

M. D. Burkhard and R. M. Sachs (1975) "Anthropometric manikin for acoustic research" *The Journal of the Acoustical Society of America*, vol. 58, pp. 214-222.

M. D. Burkhard, W. A. Yost, D. Preves, G. F. Kuhn, E. D. Burnett, and H. R. Silbiger (1985) "American National Standard: Specification for a Manikin for Simulated in-situ Airborn Acoustic Measurements" ANSI (ASA).

T. Chen and P. P. Vaidyanathan (1992) "General theory of time-reversed inversion for perfect reconstruction filter banks" presented at the Twenty-Sixth Asilomar Conference on Signals, Systems and Computers.

J. F. Claerbout and E. A. Robinson (1964) "The error in least-squares inverse filtering" *Geophysics*, vol. 29, pp. 118-120.

P. M. Clarkson, J. Mourjopoulos and J. K. Hammond (1985) "Spectral, phase, and transient equalization for audio systems" *Journal of the Audio Engineering Society*, vol. 33, pp. 127-32.

R. Czarnach (1982) "Recursive Processing by Non-Causal Digital Filters" *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. 30, pp. 363-370.

D. H. Cooper and J. L. Bauck (1989) "Prospects for Transaural Recording" *Journal of the Audio Engineering Society*, vol. 37, pp. 3-19.

P. Damaske (1971) "Head-Related Two-Channel Stereophony with Loudspeaker Reproduction" *The Journal of the Acoustical Society of America*, vol. 50, pp. 1109-1115.

S. A. Dyer (2003) "Partial-fraction expansion: Part 2" *Instrumentation & Measurement Magazine, IEEE*, vol. 6, pp. 48-50.

S. J. Elliott, I. M. Stothers, and P. A. Nelson (1987) "A multiple error LMS algorithm and its application to the active control of sound and vibration" *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP-35, pp. 1423-34.

M. J. Evans, J. A. S. Angus, and A. I. Tew (1998) "Analyzing head-related transfer function measurements using surface spherical harmonics" *Journal of the Acoustical Society of America*, vol. 104, pp. 2400-2411.

M. J. Evans, A. I. Tew, and J. A. S. Angus (2000) "Perceived performance of loudspeaker-spatialized speech for teleconferencing" *Journal of the Audio Engineering Society*, vol. 48, pp. 771-785.

A. Farina and E. Ugolotti (1997) "Use of digital inverse filtering techniques for improving car audio systems" presented at AES 103rd Convention, New York. An extended version of this work is presented in A. Farina and E. Ugolotti, "Spatial equalization of sound systems in cars" Informal report.

L. D. Fielder (2003) "Analysis of traditional and reverberation-reducing methods of room equalization" *Journal of the Audio Engineering Society*, vol. 51, pp. 3-26.

S. Flanagan, B. C. J. Moore and M. A. Stone (2005) "Discrimination of group delay in clicklike signals presented via headphones and loudspeakers" *Journal of the Audio Engineering Society*, vol. 53, pp. 593-611.

J. Garas and P. C. W. Sommen (1998) "Improving virtual sound source robustness using multiresolution spectral analysis and synthesis" presented at AES 105th Convention, San Francisco, California.

B. Gardner and K. Martin (1994) "HRTF Measurements of a KEMAR Dummy-Head Microphone" MIT, Report 280.

W. G. Gardner (1995) "Efficient Convolution without Input-Output Delay" *Journal of the Audio Engineering Society*, vol. 43, pp. 127-136.

W. G. Gardner (1998) *3-D audio using loudspeakers*. Boston: Kluwer Academic Publishers.

M. A. Gerzon (1973) "Periphony: With-Height Sound Reproduction" *Journal of the Audio Engineering Society*, vol. 21, pp. 2-10.

M. A. Gerzon (1975) "The Design of Precisely Coincident Microphone Arrays for Stereo and Surround Sound", presented at the AES 50th Convention, London.

D. M. Gilbert and I. F. Morrison (1996) "Investigation of real-time forward-backward filter structures" in *Isspa 96 - Fourth International Symposium on Signal Processing and Its Applications, Proceedings, Vols 1 and 2*. Brisbane: QUEENSLAND UNIVERSITY TECHNOLOGY pp. 419-420.

F. Gustafsson (1996) "Determining the initial states in forward-backward filtering" *IEEE Transactions on Signal Processing*, vol. 44, pp. 988-992.

D. Hammershoi and H. Moller (1996) "Sound transmission to and within the human ear canal" *The Journal of the Acoustical Society of America*, vol. 100, pp. 408-427.

W. M. Hartmann and A. Wittenberg (1996) "On the externalization of sound images" *The Journal of the Acoustical Society of America*, vol. 99, pp. 3678-3688.

E. Hulsebos, D. de Vries and E. Bourdillat (2002) "Improved, microphone array configurations for auralization of sound fields by wave-field synthesis" *Journal of the Audio Engineering Society*, vol. 50, pp. 779-790.

H. Irisawa, S. Shimada, H. Hokari and S. Hosoya (1998) "Study of a fast method to calculate inverse filters" *Journal of the Audio Engineering Society*, vol. 46, pp. 611-20.

J. M. Jot, V. Larcher, and O. Warusfel (1995) "Digital Signal Processing Issues in the Context of Binaural and Transaural Stereophony" presented at AES 98th Convention, Paris.

O. Kirkeby (1995) "Reproduction of acoustic fields" PhD Thesis, *Institute of Sound and Vibration Research*, University of Southampton.

O. Kirkeby, P. A. Nelson, H. Hamada and F. Orduna-Bustamante (1996) "Fast deconvolution of multi-channel systems using regularisation", Report 255, ISVR, University of Southampton.

O. Kirkeby, P. A. Nelson, and H. Hamada (1998a) "The "stereo dipole" - A virtual source imaging system using two closely spaced loudspeakers" *Journal of the Audio Engineering Society*, vol. 46, pp. 387-395.

O. Kirkeby, P. A. Nelson, and H. Hamada (1998b) "Local sound field reproduction using two closely spaced loudspeakers" *Journal of the Acoustical Society of America*, vol. 104, pp. 1973-1981.

O. Kirkeby, P. A. Nelson, H. Hamada, and F. Orduna Bustamante (1998c) "Fast deconvolution of multichannel systems using regularization" *IEEE Transactions on Speech and Audio Processing*, vol. 6, pp. 189-94.

O. Kirkeby and P. A. Nelson (1999) "Digital filter design for inversion problems in sound reproduction" *Journal of the Audio Engineering Society*, vol. 47, pp. 583-595.

O. Kirkeby, P. Rubak, and A. Farina (1999a) "Analysis of ill-conditioning of multi-channel deconvolution problems" presented at Workshop on Applications of Signal Processing to Audio and Acoustics, IEEE.

O. Kirkeby, P. Rubak, L. G. Johansen, and P. A. Nelson (1999b) "Implementation of cross-talk cancellation networks using warped FIR filters" presented at AES 16th International Conference.

D. J. Kistler and F. L. Wightman (1992) "A Model of Head-Related Transfer-Functions Based on Principal Components-Analysis and Minimum-Phase Reconstruction" *Journal of the Acoustical Society of America*, vol. 91, pp. 1637-1647.

J. Köring and A. Schmitz (1993) "Simplifying Cancellation of Cross-Talk for Playback of Head-Related Recordings in a 2-Speaker System" *Acustica*, vol. 79, pp. 221-232.

V. Kučera (1979) *Discrete linear control the polynomial equation approach*. Chichester: Wiley.

V. Kučera, J. Jezek, and M. Krupicka (1991) "Numerical-Analysis of Diophantine Equations" *Lecture Notes in Control and Information Sciences*, vol. 158, pp. 128-136.

A. Kulkarni and H. S. Colburn (1998) "Role of spectral detail in sound-source localization" *Nature*, vol. 396, pp. 747-749.

A. Kulkarni, S. K. Isabelle, and H. S. Colburn (1999) "Sensitivity of human subjects to head-related transfer-function phase spectra" *Journal of the Acoustical Society of America*, vol. 105, pp. 2821-2840.

C. Kyriakakis, P. Tsakalides, and T. Holman (1999) "Surrounded by sound" *Signal Processing Magazine, IEEE*, vol. 16, pp. 55-66.

T. I. Laakso, V. Välimäki, M. Karjalainen and U. K. Laine (1996) "Splitting the unit delay [FIR/all pass filters design]" *Signal Processing Magazine, IEEE*, vol. 13, pp. 30-60.

Lake (2003) *Huron Technical Manual - Software version 3.2*. Sydney Australia, Lake Technology Limited.

S. P. Lipshitz, M. Pocock and J. Vanderkooy (1982) "On the Audibility of Mid-Range Phase-Distortion in Audio Systems" *Journal of the Audio Engineering Society*, vol. 30, pp. 580-595.

E. A. Lopez-Poveda and R. Meddis (1996) "A physical model of sound diffraction and reflections in the human concha" *The Journal of the Acoustical Society of America*, vol. 100, pp. 3248-3259.

P. Mannerheim, M. Park, T. Papadopoulos and P. A. Nelson (2004) "The Measurement of a Database of Head Related Transfer Function", Samsung Contract Report 04/07, ISVR.

W. L. Martens (1997) "Unexamined Assumptions in the Commercialization of 3D Audio: Does KEMAR Sleep in a Procrustian Bed?" available at <http://www.u-aizu.ac.jp/~wlm/iwhit.html>.

R. J. Maxwell and M. D. Burkhard (1979) "Larger ear replica for KEMAR manikin" *The Journal of the Acoustical Society of America*, vol. 65, pp. 1055-1058.

L. E. McBride, H. W. Schaeffgen and K. Steiglitz (1966) "Time-domain approximation by iterative methods" *IEEE Transactions on Circuit Theory*, vol. 13, pp. 381-7.

MathWorks (2000) *Signal Processing Toolbox User's Guide Version 5* MathWorks.

S. Mehrgardt and V. Mellert (1977) "Transformation characteristics of the external human ear" *The Journal of the Acoustical Society of America*, vol. 61, pp. 1567-1576.

J. C. Middlebrooks, J. C. Makous and D. M. Green (1989) "Directional sensitivity of sound-pressure levels in the human ear canal" *The Journal of the Acoustical Society of America*, vol. 86, pp. 89-108.

M. Miyoshi and Y. Kaneda (1986) "Inverse control of room acoustics using multiple loudspeakers and/or microphones" presented at IEEE ICASSP, Tokyo.

M. Miyoshi and Y. Kaneda (1988) "Inverse filtering of room acoustics" *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 36, pp. 145-52.

H. Møller (1989) "Reproduction of Artificial-Head Recordings through Loudspeakers" *Journal of the Audio Engineering Society*, vol. 37, pp. 30-33.

H. Møller (1992) "Fundamentals of Binaural Technology" *Applied Acoustics*, vol. 36, pp. 171-218.

H. Møller, D. Hammershoi, C. B. Jensen and M. F. Sorensen (1995a) "Transfer Characteristics of Headphones Measured on Human Ears" *Journal of the Audio Engineering Society*, vol. 43, pp. 203-217.

H. Møller, M. F. Sorensen, D. Hammershoi, and C. B. Jensen (1995b) "Head-Related Transfer-Functions of Human-Subjects" *Journal of the Audio Engineering Society*, vol. 43, pp. 300-321.

M. Morimoto and Y. Ando (1980) "On the simulation of sound localization," *J. Acoust. Soc. Jpn. (E)*, vol.1, no. 3, pp. 167-174.

A. Mouchtaris, P. Reveliotis, and C. Kyriakakis (2000) "Inverse filter design for immersive audio rendering over loudspeakers" *Multimedia, IEEE Transactions on*, vol. 2, pp. 77-87.

J. Mourjopoulos, P. M. Clarkson and J. K. Hammond (1982) "A comparative study of least-squares and homomorphic techniques for the inversion of mixed phase signals", presented at the IEEE ICASSP, Paris.

J. Mourjopoulos (1994) "Digital equalization of room acoustics" *Journal of the Audio Engineering Society*, vol. 42, pp. 884-900.

J. N. Mourjopoulos (2003) "Comments on "analysis of traditional and reverberation-reducing methods of room equalization"" *Journal of the Audio Engineering Society*, vol. 51, pp. 1186-1188. [See also the correction by the author in *Journal of the Audio Engineering Society*, vol. 52, pp. 262-262].

S. Muller and P. Massarani (2001) "Transfer-function measurement with sweeps" *Journal of the Audio Engineering Society*, vol. 49, pp. 443-471. see also S. Muller and P. Massarani "Transfer-function measurement with sweeps - Director's cut including previously unreleased material." (manuscript available at [www.anselmgoertz.de/Page10383/Monkey\\_Forest\\_dt/Manual\\_dt/Aes-swp.pdf](http://www.anselmgoertz.de/Page10383/Monkey_Forest_dt/Manual_dt/Aes-swp.pdf))

S. T. Neely and J. B. Allen (1979) "Invertibility of a room impulse response" *Journal of the Acoustical Society of America*, vol. 66, pp. 165-9.

P. A. Nelson, H. Hamada, and S. J. Elliott (1992) "Adaptive inverse filters for stereophonic sound reproduction" *IEEE Transactions on Signal Processing*, vol. 40, pp. 1621-32.

P. A. Nelson (1994) "Active control of acoustic fields and the reproduction of sound" *Journal of Sound and Vibration*, vol. 177, pp. 447-77.

P. A. Nelson, F. Orduna-Bustamante, and H. Hamada (1995) "Inverse filter design and equalization zones in multichannel sound reproduction" *Speech and Audio Processing, IEEE Transactions on*, vol. 3, pp. 185-192.

P. A. Nelson, F. Orduna-Bustamante, D. Engler, and H. Hamada (1996) "Experiments on a system for the synthesis of virtual acoustic sources" *Journal of the Audio Engineering Society*, vol. 44, pp. 990-1007.

P. A. Nelson, O. Kirkeby, T. Takeuchi, and H. Hamada (1997) "Sound fields for the production of virtual acoustic images" *Journal of Sound and Vibration*, vol. 204, pp. 386-396.

P. A. Nelson and J. F. W. Rose (2005) "Errors in two-point sound reproduction" *Journal of the Acoustical Society of America*, vol. 118 pp. 193-204.

S. G. Norcross, G. A. Soulodre, and M. C. Lavoie (2004) "Subjective investigations of inverse filtering" *Journal of the Audio Engineering Society*, vol. 52, pp. 1003-1028.

A. V. Oppenheim and R. W. Schafer (1975) *Digital signal processing*. Englewood Cliffs, N.J.: Prentice-Hall.

A. K. Oppenheim, G.; Trbolet, J. (1976) "Signal analysis by homomorphic prediction" *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, pp. 327-332.

T. Papadopoulos (2000) "The use of IIR filtering in Virtual Acoustic Imaging Systems" MSc Thesis, *ISVR*, University of Southampton.

S. R. Powell and P. M. Chau (1990a) "Time reversed filtering in real-time" presented at IEEE International Symposium on Circuits and Systems.

S. R. Powell and P. M. Chau (1990b) "An improved noncausal filtering technique" presented at 24th Asilomar Conference on Signals, Systems and Computers.

S. R. Powell and P. M. Chau (1991) "A Technique for Realizing Linear-Phase IIR Filters" *IEEE Transactions on Signal Processing*, vol. 39, pp. 2425-2435.

D. Preis (1980) "Measures and perception of phase distortion in electroacoustical systems", presented at the IEEE International Conference on ICASSP.

W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery (2002) *Numerical recipes in C++ the art of scientific computing*, Cambridge University Press, Cambridge.

J. G. Proakis and D. G. Manolakis (1992) *Digital signal processing principles, algorithms, and applications*, 2nd ed. New York: Macmillan.

T. A. Ramstad (1988) "IIR filterbank for subband coding of images" presented at the IEEE International Symposium on Circuits and Systems.

J. F. W. Rose (2004) "Visually adaptive virtual acoustic imaging" PhD Thesis, *Institute of Sound and Vibration Research*, University of Southampton.

N. Sakamoto, T. Gotoh, T. Kogure and M. Shimbo (1982) "Controlling Sound-Image Localization in Stereophonic Reproduction: Part II" *Journal of the Audio Engineering Society*, vol. 30, pp. 719-722.

M. R. Schroeder (1975) "Models of hearing" *Proceedings of the IEEE*, vol. 63, pp. 1332-1350.

G. A. Sitton, C. S. Burrus, J. W. Fox, and S. Treitel (2003) "Factoring very-high-degree polynomials" *Signal Processing Magazine, IEEE*, vol. 20, pp. 27-42.

E. W. Start, D. de Vries and A. J. Berkhouit (1999) "Wave field synthesis operators for bent line arrays in a 3D space" *Acustica*, vol. 85, pp. 883-892.

K. Steiglitz and L. E. McBride (1965) "A technique for the identification of linear systems" *IEEE Transactions on Automatic Control*, vol., pp. 461-4.

B. Stuart and K. Howard (2005) "DSPerate times" *Hi-Fi News*, vol. 50, issue 6, pp. 90-93.

T. Takeuchi, P. A. Nelson, O. Kirkeby, and H. Hamada (1997) "Robustness of the Performance of the "Stereo Dipole" to Misalignment of Head Position" presented at AES 102nd Convention, Munich.

T. Takeuchi, P. A. Nelson, O. Kirkeby, and H. Hamada (1998) "Influence of Individual Head-Related Transfer Function on the Performance of Virtual Acoustic Imaging Systems" presented at AES 104th Convention, Amsterdam.

T. Takeuchi (2001) "Systems for virtual acoustic imaging using the binaural principle" PhD Thesis, *Institute of Sound and Vibration Research*, University of Southampton.

T. Takeuchi, P. A. Nelson, and H. Hamada (2001) "Robustness to head misalignment of virtual sound imaging systems" *Journal of the Acoustical Society of America*, vol. 109, pp. 958-971.

A. Vasileiadis and P. A. Nelson (2004) "Subjective evaluation of the importance of the ringing frequency in virtual acoustic imaging" presented at Akoustiki 2004, Hellenic Institute of Acoustics, Thessaloniki, Greece.

D. de Vries (1996) "Sound reinforcement by wavefield synthesis: Adaptation of the synthesis operator to the loudspeaker directivity characteristics" *Journal of the Audio Engineering Society*, vol. 44, pp. 1120-1131.

S. Weiss, G. W. Rice and R. W. Stewart (1999) "Multichannel equalization in subbands", presented at the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, New Paltz, New York.

E. M. Wenzel, M. Arruda, D. J. Kistler, and F. L. Wightman (1993) "Localization Using Nonindividualized Head-Related Transfer- Functions" *Journal of the Acoustical Society of America*, vol. 94, pp. 111-123.

F. L. Wightman and D. J. Kistler (1989a) "Headphone Simulation of Free-Field Listening .1. Stimulus Synthesis" *Journal of the Acoustical Society of America*, vol. 85, pp. 858-867.

F. L. Wightman and D. J. Kistler (1989b) "Headphone Simulation of Free-Field Listening .2. Psychophysical Validation" *Journal of the Acoustical Society of America*, vol. 85, pp. 868-878.