Causal MSE-Optimal Filters for Personal Audio Subject to Constrained Contrast

Simon Widmark

***Abstract*—A novel design method that generates causal pre-compensation filters is formulated. The resulting filters are constrained with respect to the amount of acoustic contrast they generate and are intended to be used for personal audio. The proposed method provides a more direct method for trading bright zone behavior against acoustic contrast as compared to other causal methods available. It also provides improved control over the temporal properties of the resulting filters as compared to the pre-existing non-causal methods. The resulting filters are analyzed by means of simulations, based on measured impulse responses of the sound-system–room interactions. The results of the simulations are compared to simulations of a frequency-domain optimal method with comparable objective, as proposed by Cai *et al.* and the results of the comparison are explained using the design equations. It is shown that the proposed method is viable, but that unattainable contrasts have a detrimental impact on the spectral bright zone behavior. A few different strategies for dealing with this problem are also proposed. It is demonstrated that the detrimental effect of increasingly strict causality constraints mainly concerns the lower frequency bright zone behavior in the system under investigation, but that the very highest attainable contrast levels may also be reduced somewhat.**

***Index Terms*—Signal processing, acoustic signal processing, filters, IIR filters.**

I. Introduction

THE topic of the current paper is filter design for personal audio applications. The purpose of personal audio is to divide a given volume of space into two or more acoustical zones with respect to a sound generating system. In at least one of these zones, the sound system should ideally reproduce a desired sound field, while at least one zone is left completely quiet. The quiet zone(s) is referred to as the dark zone(s) in the literature, while the non-quiet zone(s) is termed the bright zone(s) [1]. Due to the complexity of the sound-field, constraints on the number of loudspeakers and their positioning, and requirements on the physical extent of the zones, it is normally unrealistic to expect perfect cancellation in the dark zone and simultaneous, perfect, sound field reconstruction in the bright zone. Compromise is therefore necessary in practice.

Manuscript received June 29, 2018; revised January 24, 2019; accepted February 25, 2019. Date of publication March 26, 2019; date of current version April 4, 2019. This work was supported in part by the Swedish Research Council under Contract 621-2014-5871 and in part by the Knut and Alice Wallenberg Foundation. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Huseyin Hacihabiboglu.

The author is with the Signals and Systems Group, Uppsala University, 752 36 Uppsala, Sweden, and also with Dirac Research AB, 753 20 Uppsala, Sweden (e-mail: simon.widmark@hotmail.com).

Digital Object Identifier 10.1109/TASLP.2019.2904839

A reasonable such compromise is to substitute 'absence of noticeable disturbance' for complete silence in the dark zone, where the word disturbance is used to capture all ill effects in the dark zone that result from the reproduction of a desired sound field in the bright zone. The most widely used measure with relation to disturbance is the acoustic contrast, which is defined as the ratio of sound power in the bright zone to the sound power in the dark zone at a single frequency [1]. The acoustic contrast is thus a measure of the physical part of the psychoacoustical property of disturbance.

Several studies have tried to quantify the acoustic contrast needed to produce reliably low disturbance levels in a variety of situations, see e.g. [2]–[4]. These results are however hard to extrapolate to a wider range of real world situations due to the inherent difficulties in quantifying psychoacoustic effects.

We will here utilize the average contrasts within a set of frequency bands as a design constraint, thus exclusively targeting the physically quantifiable part of the disturbance reduction problem. By the flexibility of this measure, it can be expected to remain relevant in a wide variety of situations. For a given type of sound (speech, music, etc.) and an assumed background masking noise level, a set of contrast levels can for instance be defined that would provide an adequately low level of disturbance. Further, provided a method that produces a given contrast in a system, psychoacoustically motivated pre-processing of the input signal can be applied if the raw, physical sound level difference between the zones is not adequate. The proposed formulation also allows for different levels of acoustic contrast to be demanded in different frequency bands or between different control point constellations. This may be beneficial, e.g., if a certain frequency–contrast profile is beneficial from a psychoacoustical point of view.

Just as the degree of disturbance in the dark zone can be subject to compromise, so can the sound quality in the bright zone. The methods that produce the greatest acoustic contrast [1], [5] do so by completely ignoring the bright zone sound quality. Normally, however, the most desirable solution lies somewhere between no consideration of the bright zone and no consideration of the dark zone. A large number of papers exploring methods to find a perfect compromise for any given situation have been published, see, e.g., [6]–[11]. None of these methods do, however, consider the constraint on causality of the resulting practical implementation in the optimization stage.

Introducing causality of the generated filters as a pre-requisite already in the filter optimization can be advantageous to the final implementation. While any Finite Impulse Response (FIR) filter can be time-shifted into the causal domain, doing so removes any

2329-9290 © 2019 IEEE. Translations and content mining are permitted for academic research only. Personal use is also permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

explicit control over the temporal properties of the compensated system from the designer. As a result, the generated filter may cause non-zero outputs before the start of the desired impulse response of the compensated system, this is referred to as 'pre-ringing'. Pre-ringing often has a negative effect on the perceived sound quality and using design methods that consider causality as a constraint allows a designer control over the temporal extent of the generated pre-ringing.

In this paper, we explore a causally constrained[1] personal audio design that minimizes the weighted sum of the sound field synthesis error in the bright zone and the filter power gain, subject to a constraint specifying the minimum allowed contrast between the bright and dark zone(s). The proposed method thus produces the best possible sound field reproduction in the bright zone (for a given filter power gain penalty) provided that a certain acoustic contrast (here parametrized by the sound power difference) is generated.

A. Relation to Other Causal Methods

There are a few other design methods for personal audio that consider the causality of the generated filters in the optimization phase. The earliest of these is the causal acoustic contrast maximization method with constrained filter power gain [12]. This approach generates very poor performance in the bright zone, and several modifications to the method have been published since, [13]–[16]. These methods attempt a trade between maximizing contrast and generating a flat spectrum in the bright zone by limiting the (phase and-) magnitude deviation from a reference frequency, or a set of reference frequencies of the compensated response in the bright zone. The phase target is thus implicitly specified and the group delay follows from this specification with limited control over phase, delay and pre-ringing in the bright zone as a result. The herein proposed method instead aims to attain a pre-specified contrast level and beyond that optimizes the adherence to a pre-specified target of the bright zone sound field. This target entails both a magnitude target and a phase target with explicitly specified group delay.

In the works [17], [18], the full desired bright zone behaviour is explicitly specified and the adherence to the target sound field in the bright zone is traded against the goal of generating no sound in the dark zone. This type of formulation differs from the proposed as the generated contrast is not explicitly considered but becomes a by-product of the conflicting goal of generating sound in one zone and silence in another. Additionally, it is also a weighted approach where a designer tunes weights until a satisfactory design is found whereas the proposed method is derived subject to a constraint. A similar formulation is also found in [19], but with the addition of a shaping filter, that may be used to penalize pre-ringing in the generated *filter*.[2]

[1]Here taken to mean methods that consider the causality of the generated filters at the optimization stage. The term 'causal' is also sometimes used.

[2]Note that a filter with pre-ringing tendencies may or may not produce pre-ringing in the system, while a filter with no such tendencies can never produce pre-ringing in the compensated system. The approach proposed herein where the pre-ringing of the system output can be treated is thus less conservative. On the other hand, the approach proposed in [19] allows for a soft transition between no penalty and a heavy penalty on pre-ringing over time while the approach discussed herein instead prohibits all output before a specified time but accepts all output after this time.

In addition to the differences in objective between the proposed and previously investigated methods, all but [17] are based on Toeplitz-type system descriptions. Such filter design methods generally encounter computational challenges in that large matrices are often encountered when a moderate to large number of loudspeakers or a large number[3] of filter taps are required.

In the publications [20], [21], the goal of the bright zone sound field manipulation is to obtain a certain geometrical distribution rather than to generate a given acoustical behaviour. The method further requires geometrical models of the acoustical environment under consideration which may become a concern in more complex, e.g. non-anechoic, situations.

B. Contribution

The proposed method provides improved control over the trade-off between contrast and bright zone behaviour compared to the previously published causal methods.

The methods presented herein also differ from the previously presented methods (except [17]) in that they generate Infinite Impulse Response (IIR) controller filters. This can be advantageous when filters with long decay times are to be implemented in a real system. Further, the filter parametrization is specified by the method, generating a filter that is optimally parametrized for the provided system description. This eliminates the need to iterate the solution with different filter lengths in order to find a satisfactory trade-off between performance and filter length.

IIR filters generally provide a more succinct implementation with respect to the number of multiplications and additions needed for each sample and a substantial amount of research on the topic of IIR filter design can be found in the literature. Causal and constrained IIR filters do however seem relatively under-explored for audio applications. A general framework for designing casual IIR filters subject to quadratic constraints is presented in [22].[4]

An important difference between the constrained, causal filters investigated herein and the constrained (frequency-domain optimal) filters investigated previously in the literature, that do not take causality into account, is that the latter are constrained per-frequency. There, a constraint is adhered to for every design frequency, under a user-defined DFT resolution. The constraints of the designs presented here are specified as averages over the frequency ranges covered by the corresponding constraint functions. As a constraint on each frequency bin is more restrictive than one on the average over several frequencies, the causal controllers may, in some instances, achieve a better objective function value than their frequency-domain optimal counterparts. One drawback of the constraint on average power is that it allows the constraint function to assume a high value for parts of the frequency range by balancing this with a correspondingly low value in other parts. This behaviour can be countered by utilizing several constraints that are focused on different frequency

[3]On the order of a few thousand.

[4]Note however that the presently investigated formulation, where acoustic contrast is considered in the constraint function is not admissible in the framework presented in [22]. This is due to a negative semi-definite term that we will introduce in the mathematical formulation of the herein investigated constraint function.

bands, effectively reducing the number of possible trades available to the controller.

The main objective of this paper is to present a casual method with good control over the bright zone properties that is constrained with respect to the lowest allowed contrast produced by the resulting filter. The resulting filter is compared to an alternative, previously known frequency-domain optimal filter [9] designed with the same objective.

C. Outline of the Paper

A mathematical structure, serving as the foundation of the work described in the present paper is introduced in Section II. The causal and stable optimizing filter is presented in Section III, together with a previously known frequency-domain optimal design with a similar objective. The causally constrained design method is investigated and analysed, both in absolute terms and in relation to the frequency-domain optimal design in Section IV. The paper is summarized in Section V and the relevant proofs and motivations are found in Appendix A.

II. Preliminaries

A. Notation

Throughout this paper, scalar quantities are denoted by regular typeface symbols, a or A. Vectors and matrices are denoted by bold lower case, $\boldsymbol{a}$, and capital, $\mathbf{A}$, symbols respectively. Positive definiteness and semi-definiteness of a matrix $\mathbf{A}$ is denoted by $\mathbf{A} \succ 0$ and $\mathbf{A} \succeq 0$ respectively. The number of rows, R and columns, C of a matrix is expressed $R|C$. Polynomial matrices are matrices of polynomials and rational matrices are matrices of rational functions. Both the polynomial functions and the rational functions are functions in the discrete-time shift operator q^{-1}, such that $q^{-1}x(t) = x(t-1)$ and $qx(t) = x(t+1)$. Polynomial matrices are denoted by italic bold face capital symbols, $\boldsymbol{A}$, and rational matrices are denoted by bold, script font, capital letters $\boldsymbol{\mathcal{A}}$. Rational (polynomial) matrices in the time shift operator represent transfer operator matrices with IIR (FIR) filters as elements. The product of a set of rational matrices in the shift operator corresponds to the convolution of all impulse responses represented by them. Allowing the resulting matrix of polynomials in the delay operator to operate on an input signal then yields the input-output relation of a system with all delayed versions of the input signal correctly represented. In the frequency domain, we substitute the complex variables z and z^{-1} for the shift operators q and q^{-1} respectively. The frequency domain variable, ω denotes the normalized angular frequency and $j = \sqrt{-1}$. Matrices can be transposed $\mathbf{A}^T$, complex conjugate transposed $\mathbf{A}^H$ and polynomial matrices can be complex conjugate transposed and delay conjugated so that $\boldsymbol{A}_*(q)$ is found by substituting q for q^{-1} in $\boldsymbol{A}^H(q^{-1})$. Operator dependence will be omitted in the text where this can be done without impacting readability. The degree of a polynomial matrix $\boldsymbol{M}(q, q^{-1})$ is the highest absolute exponent value $|x|$ of any polynomial element $\sum_{-Y}^{Z} q^x m_x$, for which $m_x \neq 0$ and is denoted $\deg(\boldsymbol{M})$.

B. Model Structure and Parametrization

A linear and time-invariant dynamic system is assumed to be correctly and completely described by a transfer operator matrix, represented by the stable rational matrix $\boldsymbol{\mathcal{H}}(q^{-1})$, containing transfer operators with real valued coefficients. The matrix of $M_B|N$ rational functions describing the transfer functions between N loudspeakers and M_B control points in one or more bright zones (in which we want to reproduce a specified sound field) is denoted $\boldsymbol{\mathcal{H}}_B(q^{-1})$. Similarly, the dark zone(s) (in which we ideally want to reproduce no sound) are described by the $M_D|N$ rational matrix $\boldsymbol{\mathcal{H}}_D(q^{-1})$.

The rational matrices $\boldsymbol{\mathcal{H}}_\bullet$ can be represented in polynomial matrix form by a right matrix fraction description (MFD) so that $\boldsymbol{\mathcal{H}}_\bullet(q^{-1}) = \boldsymbol{B}_\bullet(q^{-1})\boldsymbol{A}^{-1}(q^{-1})$ [23, pp. 367–370]. Here, $\bullet$ is used for a placeholder and can be substituted for either B or D. We will in the following assume that the polynomial matrix $\boldsymbol{A}$ is diagonal. We will further use the fact that the modes of a room are common for all control points, see, e.g., [24]–[26], to avoid spatial dependence of the $N|N$ denominator matrix $\boldsymbol{A}(q^{-1})$. The numerator matrix $\boldsymbol{B}_\bullet$ assumes the dimension $M_\bullet|N$.

The system is here assumed to be driven by L sound sources, modelled as white, zero mean noises and described by the $L|1$ vector $\boldsymbol{r}(t)$ with covariance $E\left\{\boldsymbol{r}(t)\boldsymbol{r}^T(t)\right\} = \mathbf{I}$. All expected values $E\{\dots\}$ are here taken with respect to the driving stochastic signal $\boldsymbol{r}(t)$.[5]

The desired sound field in the bright zone is described by FIR transfer functions in the form of the $M_B|L$ polynomial target matrix $\boldsymbol{D}(q^{-1})$. Note that this means that both desired spectral and temporal properties, such as e.g. delays are modelled in $\boldsymbol{D}(q^{-1})$. Assuming that the shortest propagation path between any loudspeaker–control point combination is zero, it is sometimes convenient to partition the target matrix by $\boldsymbol{D}(q^{-1}) = q^{-d_0}\boldsymbol{D}_0(q^{-1})$. The variable d_0 here represents the common desired delay of the elements of the matrix $\boldsymbol{D}$, so the smallest delay of any element of $\boldsymbol{D}_0$ is zero. The common delay d_0 is also referred to as the modelling delay.

A stable, causal and linear feed-forward controller, $\boldsymbol{\mathcal{R}}(q^{-1})$ of dimension $N|L$ is included in the signal path, manipulating the driving signal $\boldsymbol{r}(t)$ before feeding it to the DAC-amplifier-loudspeaker system as an input signal $\bar{\boldsymbol{u}}(t)$. A weighted compensated input signal can thus be expressed

$$\boldsymbol{u}(t) = \boldsymbol{W}(q^{-1})\bar{\boldsymbol{u}}(t) = \boldsymbol{W}(q^{-1})\boldsymbol{\mathcal{R}}(q^{-1})\boldsymbol{r}(t). \tag{1}$$

An $N|N$ polynomial matrix weight $\boldsymbol{W}(q^{-1})$ is here introduced and can be used to penalize large filter gains differently for different loudspeakers and frequency bands.

We can now describe the sound field synthesis error vector $\bar{\boldsymbol{\varepsilon}}(t)$ by the difference between an input filtered through the target matrix and the same input filtered through the pre-compensated

[5]The methods may, however, be extended to also account for uncertainties in the transfer functions along the lines of [27], and to take into account the performance in-between positions where acoustic transfer functions have been measured, as described in [24] and [28]. If this is done, the expected values will also represent expectations with respect to all stochastic variables that are introduced to describe the transfer function uncertainties.

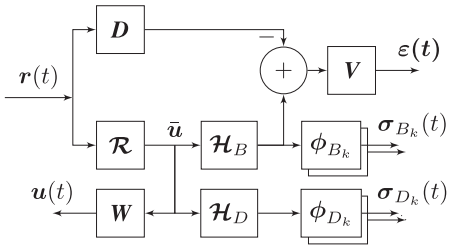

Fig. 1. Block diagram representation of the system model and relationships of the relevant signals, the weighted bright zone control error $\boldsymbol{\varepsilon}(t)$ is given by (2), the weighted vector of control signals $\boldsymbol{u}(t)$ by (1) and the weighted bright and dark zone sound pressure vectors $\boldsymbol{\sigma}_{B_k}(t),\ \boldsymbol{\sigma}_{D_k}(t)$ by (3).

system bright zone. A weighted error vector is then defined by

$$\begin{aligned}\boldsymbol{\varepsilon}(t) &= \boldsymbol{V}(q^{-1})\bar{\boldsymbol{\varepsilon}}(t)\\ &= \boldsymbol{V}(q^{-1})\left(\boldsymbol{\mathcal{H}}_B(q^{-1})\boldsymbol{\mathcal{R}}(q^{-1}) - \boldsymbol{D}(q^{-1})\right)\boldsymbol{r}(t). \quad (2)\end{aligned}$$

The $M|M$ polynomial weighting matrix $\boldsymbol{V}(q^{-1})$ above can be used to assign different weight to different aspects of the sound field synthesis part of the optimization that we will formalize below.

Note that the sound field synthesis error vector contains information about both the magnitude, phase and temporal properties of the deviations from the desired impulse responses.

The control variables and their relation to the system to be controlled are illustrated by Fig. 1.

We can also model the sound pressure at the control points by constructing a set of K dark- and bright zone sound pressure vectors, indexed by k, $1 \le k \le K$ as

$$\boldsymbol{\sigma}_{D_k}(t) = \boldsymbol{\phi}_{D_k}(q^{-1})\boldsymbol{B}_D(q^{-1})\boldsymbol{A}^{-1}(q^{-1})\boldsymbol{\mathcal{R}}(q^{-1})\boldsymbol{r}(t) \quad (3a)$$

$$\boldsymbol{\sigma}_{B_k}(t) = \boldsymbol{\phi}_{B_k}(q^{-1})\boldsymbol{B}_B(q^{-1})\boldsymbol{A}^{-1}(q^{-1})\boldsymbol{\mathcal{R}}(q^{-1})\boldsymbol{r}(t). \quad (3b)$$

Additional polynomial weighting matrices $\boldsymbol{\phi}_{\bullet_k}(q^{-1})$ are here introduced for assigning different weight to the sound pressure in the dark or bright zone at different frequency bands and/or at different control points. The dark and bright zone sound pressure vectors will later be used in the constraint function of the constrained optimization problem.

III. Optimal Filter Design

A. The Causal Contrast Constrained Pre-Compensator

The generated acoustic contrast is an important feature of any filter for personal audio. It is, however, rarely the only property of interest. Given a contrast beyond a level where the inter-zone-separation is 'good enough', we can expect 'good' bright zone behaviour to be more important than additional contrast generation. A filter that produces the critical amount of contrast and beyond that focuses on optimizing the bright zone properties would in such situations be advantageous.

Further, as the filters are to be implemented in a causal system, preferably without truncation or windowing of the filter, without introducing unreasonable delays in the audio system, and without causing excessive pre-ringing, a causal filter formulation is beneficial.

We therefore propose a design of a causal stable pre-compensator $\boldsymbol{\mathcal{R}}$ with K constraints on the minimum allowed contrast in different frequency bands and/or spatial regions. The objective is then to minimize a criterion function with respect to $\boldsymbol{\mathcal{R}}(q^{-1})$ so that the K constraints are satisfied:

$$\begin{aligned}\arg\min_{\boldsymbol{\mathcal{R}}} \quad & J = E\left\{\boldsymbol{\varepsilon}^T(t)\boldsymbol{\varepsilon}(t) + \boldsymbol{u}^T(t)\boldsymbol{u}(t)\right\}\\ \text{s.t.} \quad & C_k = E\left\{\alpha_k\boldsymbol{\sigma}_{D_k}^T(t)\boldsymbol{\sigma}_{D_k}(t) - \boldsymbol{\sigma}_{B_k}^T(t)\boldsymbol{\sigma}_{B_k}(t)\right\} \le 0,\\ & k = 1,\dots,K. \quad (4)\end{aligned}$$

Here, the weighted sum of a quadratic synthesis error for a target sound field and a quadratic control effort term, is minimized subject to K constraints on the minimum allowed acoustic contrast, averaged over the statistics of $\boldsymbol{r}(t)$ and the control points of the zones. The scalars α_k are interpreted as the desired contrasts (see explanation below). In this way, we specify a set of minimum allowed contrast levels and obtain the optimal (with respect to the criterion function J in (4)) causal filter that generates at least the specified contrast levels.

As we have assumed that the driving noise, $\boldsymbol{r}(t)$, is white, the squared term $E\left\{\boldsymbol{u}^T(t)\boldsymbol{u}(t)\right\}$ defined above can be regarded as a weighted filter power gain and will henceforth be referred to as the *filter power gain*. We refer to the quantity $E\left\{\boldsymbol{\varepsilon}^T(t)\boldsymbol{\varepsilon}(t)\right\}$ as the Sound Field Synthesis (SFS) error.

The constraints above can be understood as a manipulation of the acoustic contrast measure, which can be seen by rearranging (4), (assuming $E\left\{\boldsymbol{\sigma}_{Dk}^T(t)\boldsymbol{\sigma}_{Dk}(t)\right\} \neq 0$ and $\alpha_k > 0$)

$$\begin{aligned}\arg\min_{\boldsymbol{\mathcal{R}}} \quad & J = E\left\{\boldsymbol{\varepsilon}^T(t)\boldsymbol{\varepsilon}(t) + \boldsymbol{u}^T(t)\boldsymbol{u}(t)\right\}\\ \text{s.t.} \quad & \frac{E\left\{\boldsymbol{\sigma}_{Bk}^T(t)\boldsymbol{\sigma}_{Bk}(t)\right\}}{E\left\{\boldsymbol{\sigma}_{Dk}^T(t)\boldsymbol{\sigma}_{Dk}(t)\right\}} \ge \alpha_k, k = 1,\dots,K. \quad (5)\end{aligned}$$

The re-formulation of (5) as (4) was originally suggested in [5] for an unconstrained per-frequency design. Note that by posing the problem as in (4) or (5), the constraint concerns only the relative pressure levels between the bright and the dark zones and not the absolute pressure levels in either.

For signal vectors defined in (1)–(3), where the penalty matrix $\boldsymbol{W}(z^{-1})$ is assumed to be of full rank on $z^{-1} = e^{-j\omega}$, the filter $\boldsymbol{\mathcal{R}}(q^{-1})$ that solves (4) is found as the unique stable and causal pre-compensator minimizing the Lagrange function defined by

$$\begin{aligned}\mathcal{L}(\boldsymbol{\mathcal{R}}, \lambda_1, \cdots, \lambda_K) &= E\Bigg\{\boldsymbol{\varepsilon}^T(t)\boldsymbol{\varepsilon}(t) + \boldsymbol{u}^T(t)\boldsymbol{u}(t)\\ &\quad + \sum_{k=1}^{K}\lambda_k\left(\alpha_k\boldsymbol{\sigma}_{D_k}^T(t)\boldsymbol{\sigma}_{D_k}(t) - \boldsymbol{\sigma}_{B_k}^T(t)\boldsymbol{\sigma}_{B_k}(t)\right)\Bigg\}\\ &= J(\boldsymbol{\mathcal{R}}) + \sum_{k=1}^{K}\lambda_k C_k(\boldsymbol{\mathcal{R}}). \quad (6)\end{aligned}$$

As over-satisfying the constraint functions is not beneficial with respect to the objective function value, the smallest non-negative real valued scalar multipliers λ_k for which the minimization of $\mathcal{L}$ results in a controller $\boldsymbol{\mathcal{R}}$ that satisfies the constraints $C_k \le 0$ for all $k = 1,\dots,K$, are optimal. This property is further discussed in [22].

It is shown in Section VI-C of the Appendix that the stable and causal rational filter solving the optimization problem (4)

can be expressed by

$$\boldsymbol{\mathcal{R}}(q^{-1}) = \boldsymbol{A}(q^{-1})\boldsymbol{\beta}^{-1}(q^{-1})\boldsymbol{Q}(q^{-1}). \tag{7}$$

Above, $\boldsymbol{\beta}(q^{-1})$ is a stably and causally invertible $N|N$ polynomial matrix that satisfies the spectral factorization

$$\begin{aligned}\boldsymbol{\beta}_*(q)\boldsymbol{\beta}(q^{-1}) &= \boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A} \\ &+ \sum_{k=1}^{K}\lambda_k\left(\alpha_k\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D_k*}\boldsymbol{\phi}_{D_k}\boldsymbol{B}_D - \boldsymbol{B}_{B*}\boldsymbol{\phi}_{B_k*}\boldsymbol{\phi}_{B_k}\boldsymbol{B}_B\right)\end{aligned} \tag{8}$$

and attains a degree of

$$\begin{aligned}&\deg(\boldsymbol{\beta}) \le \max\left(\deg(\boldsymbol{B}_B) + \deg(\boldsymbol{V}), \deg(\boldsymbol{A}) + \deg(\boldsymbol{W}),\right. \\ &\deg(\boldsymbol{B}_D) + \deg\left(\boldsymbol{\phi}_{D_1}\right), \ldots, \deg(\boldsymbol{B}_D) + \deg\left(\boldsymbol{\phi}_{D_K}\right), \\ &\left.\deg(\boldsymbol{B}_B) + \deg\left(\boldsymbol{\phi}_{B_1}\right), \ldots, \deg(\boldsymbol{B}_B) + \deg\left(\boldsymbol{\phi}_{B_K}\right)\right).\end{aligned} \tag{9}$$

We note that by (8), the optimal spectral factor matrix, $\boldsymbol{\beta}(q^{-1})$, is obtained from a sum with three principal terms. The first is related to the system whose behaviour we wish to replace with that of $\boldsymbol{D}(q^{-1})$. The second term is influenced by the pre-compensator penalty term, $\boldsymbol{W}(q^{-1})$, which also acts as a regularization of the inverse $\boldsymbol{\beta}^{-1}(q^{-1})$ in (7). The third term represents the constraint functions, i.e. the power level differences between the bright and the dark zones, which must be satisfied (here ≤ 0) for the resulting solution to be acceptable.

The polynomial matrix spectral factorization equation (8) above can only be solved if the right hand side represents a valid spectral density matrix when substituting $z = e^{j\omega}$ for q. The right hand side of (8) is therefore required to be positive definite on the unit circle $z = e^{j\omega}$. Due to the negative sign preceding the last term, an important and non-trivial aspect of a proof of optimality for this problem, which differs from the formulation of [22], is to assure solvability of (8), for all relevant values of λ_k. This is shown in Section VI-B of the Appendix, with the logic summarized in Section III-A1 below.

Given a solution $\boldsymbol{\beta}(q^{-1})$ of (8), the polynomial matrix $\boldsymbol{Q}(q^{-1})$ of degree $\deg(\boldsymbol{V}) + \deg(\boldsymbol{D})$ in (7), together with a polynomial matrix $\boldsymbol{L}_*(q)$ of degree $\max\left(\deg(\boldsymbol{V}) + \deg(\boldsymbol{B}_B), \deg(\boldsymbol{\beta})\right) - 1$ are then obtained as the $N|L$ polynomial matrices that constitute the unique solution to the polynomial matrix Diophantine equation

$$\boldsymbol{\beta}_*\boldsymbol{Q} - \boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{D} = q\boldsymbol{L}_*. \tag{10}$$

A proof of this is given in Section VI-C of the Appendix.

Note that the desired system behaviour is introduced into the pre-compensator (7) exclusively via the causal factor $\boldsymbol{Q}(q^{-1})$. The Diophantine equation (10) separates the causal part, $\boldsymbol{Q}(q^{-1})$, of the compensation from the anti-causal part $\boldsymbol{L}_*(q)$, which is not used in the causal solution. It can be shown that letting $d_0 \to \infty$ produces the (time-shifted) non-causal Wiener solution, in which $\boldsymbol{L}_*(q) = 0$.

The solution procedure thus consists of solving (8) and (10) to obtain (7), taking a step in the search space of $\{\lambda_1, \ldots, \lambda_K\}$ and iterate until all constraints are satisfied.

Note that the filter parametrization of the IIR pre-compensator $\boldsymbol{\mathcal{R}}(q^{-1})$ is not explicitly specified by the user but is given by the design method itself. The resulting solution is automatically given an optimal parametrization with respect to the system model. Another aspect of the solution that is worth emphasizing is that the optimal pre-compensator (7) is given as a set of IIR filters. This is beneficial where long controller lengths are required, which is often the case in compensation of systems with long decay times.

The constraint functions $C_k(\boldsymbol{\mathcal{R}})$ are in the following assumed to be approximately independent of each other, so that $dC_i(\boldsymbol{\mathcal{R}})/d\lambda_p \approx 0$, where $i \ne p$. This assumption holds if the different constraint functions act on different frequency bands and these are specified with negligible overlap. This will allow us to simplify a few steps and utilize simple search methods.

1) Summary of the Proof of Appendix A: We can solve the optimization problem (4) if we can find a global minimizing $\boldsymbol{\mathcal{R}}(q^{-1})$, for a set of λ_k that provide constraint satisfaction, to the Lagrange function (6). It is shown in Appendix A that this can be done if the matrix (where the complex variable z is substituted for q and z^{-1} for q^{-1} in all polynomials)

$$\begin{aligned}\boldsymbol{M}(z, z^{-1}) = \boldsymbol{A}_*^{-1}\Bigg(&\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A} \\ &\sum_{k=1}^{K}\lambda_k + \left(\alpha_k\boldsymbol{B}_{D*}\boldsymbol{\phi}_{Dk*}\boldsymbol{\phi}_{Dk}\boldsymbol{B}_D - \boldsymbol{B}_{B*}\boldsymbol{\phi}_{Bk*}\boldsymbol{\phi}_{Bk}\boldsymbol{B}_B\right)\Bigg)\boldsymbol{A}^{-1},\end{aligned} \tag{11}$$

is positive definite on the unit circle, $z = e^{j\omega}$, since this will guarantee solvability of (8). It is also shown that we can limit the search space for λ_k by $\left\{\lambda_k | \boldsymbol{M}(z, z^{-1}) \succ 0\right\}$ without excluding the minimizing solution $\boldsymbol{\mathcal{R}}$ that provides $\lambda_k C_k(\boldsymbol{\mathcal{R}}) = 0$, and $C_k(\boldsymbol{\mathcal{R}}) \le 0$, which guarantees optimality.

The solution space which provides a positive definite matrix $\boldsymbol{M}(z, z^{-1})$ by choice of λ_k, $\left\{\lambda_k | \boldsymbol{M}(z, z^{-1}) \succ 0\right\}$, is defined by

$$0 \le \lambda_k < \lambda_k^{max} \qquad k = 1, \ldots, K. \tag{12}$$

The limiting multiplier λ_k^{max} can be found, when each constraint function considers a frequency band that has negligible overlap with the frequency band of any other constraint function, as the inverse of the largest positive eigenvalue, $\gamma(z)$, at any $z = e^{j\omega}$, $-\pi < \omega < \pi$ of the eigenvalue problem

$$\begin{aligned}\gamma(z)\Lambda(z) = &-\left(\left(\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A}\right)\boldsymbol{A}^{-1}\right)^{-1} \\ &\times\left(\alpha_k\boldsymbol{B}_{D*}\boldsymbol{\phi}_{Dk*}\boldsymbol{\phi}_{Dk}\boldsymbol{B}_D - \boldsymbol{B}_{B*}\boldsymbol{\phi}_{Bk*}\boldsymbol{\phi}_{Bk}\boldsymbol{B}_B\right)\boldsymbol{A}^{-1}\Lambda(z).\end{aligned} \tag{13}$$

A frequency domain search for $\gamma(z)$ of (13) can be utilized, if care is taken to perform the evaluation densely enough in the frequency domain, so that the true maximal value of $\gamma\left(e^{j\omega}\right)$ for all ω is not overlooked.

The procedure for finding the filter that solves (4) (to a specified numerical precision μ) for a single constraint is described in Algorithm 1, with variable dimensions summarized in Table I.

TABLE I
SUMMARY OF THE DIMENSIONS OF THE DIFFERENT RELEVANT VARIABLES FOR A SYSTEM WITH N LOUDSPEAKERS, M_B CONTROL POINTS IN THE BRIGHT ZONE, M_D CONTROL POINTS IN THE DARK ZONE, L INPUT CHANNELS AND K CONSTRAINTS

Variable	$\boldsymbol{B}_B$	$\boldsymbol{B}_D$	$\boldsymbol{A}$	$\boldsymbol{D}$	$\boldsymbol{W}$
Dimension	$M_B\|N$	$M_D\|N$	$N\|N$	$M_B\|L$	$N\|N$
Variable	$\boldsymbol{V}$	$\boldsymbol{\phi}_B$	$\boldsymbol{\phi}_D$	$\boldsymbol{\mathcal{R}}$	$C_k\,(\boldsymbol{\mathcal{R}})$
Dimension	$M_B\|M_B$	$M_B\|M_B$	$M_D\|M_D$	$N\|L$	$1\|1$
Variable	α_k	λ_k	μ		
Dimension	$1\|1$	$1\|1$	$1\|1$		

Algorithm 1: Compute the Optimal Causal Filter Under a Single Constraint on Minimum Allowed Contrast Using a Bisection Search.

Data: System models $\boldsymbol{B}_B(q^{-1})$, $\boldsymbol{B}_D(q^{-1})$, $\boldsymbol{A}(q^{-1})$. Weight matrices $\boldsymbol{W}(q^{-1})$, $\boldsymbol{V}(q^{-1})$, $\boldsymbol{\phi}_B(q^{-1})$, $\boldsymbol{\phi}_D(q^{-1})$. Target matrix $\boldsymbol{D}(q^{-1})$, minimal allowed contrast α, and computational precision μ
Result: The stable, causal IIR filter $\boldsymbol{\mathcal{R}}(q^{-1})$ that minimizes the weighted sum of target error and filter power gain for a specified contrast α.
Find the maximum admissible multiplier λ^{max} by solving (13) over a densely sampled frequency domain representation;
Set $\lambda_{it} \leftarrow 0$;
Form the matrix $\boldsymbol{\beta}_*\boldsymbol{\beta}$ by (8);
Find $\boldsymbol{\beta}_*(q)$ and $\boldsymbol{\beta}(q^{-1})$ by performing a spectral factorization on (8);
Find $\boldsymbol{Q}(q^{-1})$ by solving the Diophantine equation (10);
$\boldsymbol{\mathcal{R}}(q^{-1})$ is then given by (7);
Compute C by (4);
if $C \leq 0$ **then**
  return $\boldsymbol{\mathcal{R}}(q^{-1})$;
else
  Set $\lambda^{min} \leftarrow 0$;
  while $|C| > \mu$ **do**
    Set $\lambda_{it} \leftarrow \left(\lambda^{max} - \lambda^{min}\right)/2$;
    Find $\boldsymbol{\beta}_*(q)$ and $\boldsymbol{\beta}(q^{-1})$ by performing a spectral factorization on (8);
    Find $\boldsymbol{Q}(q^{-1})$ by solving the Diophantine equation (10);
    $\boldsymbol{\mathcal{R}}(q^{-1})$ is then given by (7);
    Compute C by (4);
    if $C < 0$ **then**
      Set $\lambda^{max} \leftarrow \lambda_{it}$;
    else
      Set $\lambda^{min} \leftarrow \lambda_{it}$;

B. Computational Numerics

As part of finding the optimal causal pre-compensator proposed herein, a spectral factorization has to be solved. A mathematical implication of this is that the polynomial matrix variable $\boldsymbol{\beta}_*\boldsymbol{\beta}$ must be positive definite on the unit circle $|z| = 1$ in order to be successfully computable. This can be guaranteed to hold under mild assumptions[6] but the accuracy and possibly speed of convergence of the numerical methods employed may be compromised if the matrix $\boldsymbol{\beta}_*\boldsymbol{\beta}$ is close to being indefinite on $|z| = 1$. The proposed method is particularly sensitive when a multiplier λ_k is required that is close to what is maximally allowed, a situation that occurs more frequently when a large contrast is demanded.

Additionally, we make the assumption that $\frac{\delta C_i}{\delta \lambda_j} \approx 0$. If this is not entirely true, the numerical search for λ_j^{opt} may cause the problem to become unsolvable in the frequency band pertaining to the constraint function C_i, unless precautions are taken.

C. Feasibility

For the following discussion, it is important to note that there is a difference between attaining a certain *contrast* and satisfying a certain *constraint*. There are two fundamental ways in which satisfying a constraint as specified in (4) may differ from attaining the level of contrast specified with the constraint.

First, the herein utilized constraint function in (4) is formulated as a difference while contrast is formulated as a ratio, similar to (5). This means that (as the pre-compensator is a common factor in all terms) the problems investigated are always feasible, the constraints are always satisfied by $\boldsymbol{\mathcal{R}}(q^{-1}) = 0$. This is fundamentally different from a formulation where the original contrast measure is utilized in the constraint functions, which then becomes infeasible if the contrast specified in the constraints cannot be attained.

Second, each constraint function herein covers an *interval* of frequencies. The generally utilized contrast measure is defined per frequency, or in practice, per frequency bin of a discrete Fourier transform.

Satisfaction of constraint k implies (roughly) that the difference between the dark zone acoustic power over the specified frequency interval multiplied with α_k and the bright zone acoustic power over the frequency interval equals zero. In frequency bands where the highest attainable contrast of the system is smaller than α_k, this difference gives a positive contribution to the total constraint value. A negative contribution can be given if instead a larger level of contrast is generated than α_k. By scaling the filter gain in these frequencies, the magnitude of the contribution can be controlled by the pre-compensator. Scaling the filter gain does however also affect the SFS error. This introduces a balancing effect, where the filter gain is reduced in frequency bands where poor contrast (less than α_k) can be produced and increased in frequency band where good contrast can be produced in such a way that:

1) The *constraint* is satisfied.
2) The SFS error is kept at a minimum.

This has a few implications:

1) Any conceivable constraint $\alpha \geq 0$ has a feasible solution.[7]
2) Low filter gain (zero if the tolerance $\mu = 0$) in a frequency band indicates that the desired *contrast* cannot be attained in this frequency band.

[6] E.g. that the weighting matrix $\boldsymbol{W}_*\boldsymbol{W}$ is of full rank on $|z| = 1$ and that $\lambda_k < \lambda_k^{max}$.

[7] This does not necessarily mean that the desired contrast is attained at all (or, indeed *any*) frequencies within the interval.

3) Specifying a constraint that tries to attain a contrast that is unattainable in parts of the targeted frequency band and attainable in other parts leads to a filter with relatively high gain at frequencies where the *contrast* can be attained and low gain in frequencies where the *contrast* cannot be attained. The pre-compensator is balanced to produce the smallest possible SFS error while satisfying the *constraint*.

Failure to attain the contrast specified by a constraint indicates that the system is not capable of producing such contrast levels.[8]

If the specified constraint level is a 'locally hard' constraint, in the sense that a contrast of α or greater must to be satisfied at every covered frequency, then the acoustical properties of the environment must be redesigned.

If the specified constraint level instead represents a 'locally soft' constraint, specifying an ideal but not strictly necessary contrast at all frequencies covered by the constraint function, then we have three options:

1) Accept poor performance in the bright zone (low gain in frequency bands with lower contrast). The sound power level in the dark zone is then reduced by the filter gain reduction at frequencies where less contrast than is specified in the constraint can be produced.
2) Use the actually generated contrast as a 'blueprint' for specifying new desired constraint values and constraint function frequency intervals. This works by ensuring that only attainable contrasts are specified by the constraints, negating the need for gain trades between frequencies. Repeat the filter optimization for the updated set of constraints.
3) Apply a post-correction filter to the generated contrast constrained filter. This filter must be applied to all loudspeaker channels and adjusts the gain and phase properties equally for all channels to not affect the generated contrast.

The first of these three options will generally produce a far poorer bright zone target adherence, but will produce a low dark zone gain. The latter options conversely produce better bright zone properties at the expense of increased by-frequency fluctuations in, and overall higher levels of, the dark zone gain. The two last options have different advantages and drawbacks.

The filter produced using option 2 will satisfy the optimality conditions specified in (4) but large fluctuations in the bright zone may still be present if the desired contrast cannot be attained in narrow frequency bands. Alternatively, the desired contrast level must be chosen lower than or equal to the lowest contrast generated in the corresponding frequency band, if nulls in the bright zone are to be avoided. A drawback of this approach is also that it requires an additional iterative optimization, which, depending on the problem, may be more or less time consuming.

Conversely, the filter produced using option 3 will not satisfy the original constraint functions. It will still, however, generate the same per-frequency contrast as the optimal filter. It is recommendable to utilize minimum phase filters for this correction so that the advantages of generating causal filters are not lost. This however imposes a certain 'smoothness' in the spectral domain, meaning that we cannot expect to correct very narrow frequency band imperfections. Note also that it is important to not introduce unreasonable filter gains when the post-correction filter is applied.

To summarize the practical behaviour of the proposed contrast constrained pre-compensator: There is a difference between satisfying the *constraints* and attaining the level of *contrast* postulated by said constraints.

The optimal solution will be the one that obtains the average constraint satisfaction with the smallest possible criterion value J in (4). If the attainable contrast of the system in the *entire* frequency band addressed by a constraint is lower than the specified constraint level, a filter gain of zero in this frequency band is the only feasible solution, assuming $\mu = 0$. If the search tolerance is $\mu \neq 0$, then a strictly speaking non-feasible solution (to the tolerance μ), and thus a non-zero pre-compensator, will instead be found.

It is also worth noting that these differences between the investigated constraint function and the utilized contrast measure exist also in the frequency-domain optimal constraint based formulations such as [9]. The phenomenon of gain balancing between frequencies will however be much less noticeable as far narrower constraint function frequency bands are then normally utilized.

D. The Frequency-Domain Optimal Contrast Constrained Pre-Compensator

A causally unconstrained filter design method with a very similar design objective as the proposed method was proposed and investigated by Cai *et al.* in [9]. It is also, being computed per-frequency, more mathematically tractable as compared to the Toeplitz matrix based causally constrained methods. It will therefore serve as our base-line in a practical, comparative investigation of the proposed causal method.

In the formulation of the frequency-domain optimal pre-compensator used in the present paper we also introduce a weighted filter gain penalty, $\mathbf{W}(e^{-j\omega})$, so that we can more readily compare the resulting filters of the two methods. The filter gain penalty plays an important role in producing realizable filters, as a loudspeaker can normally only operate within a limited frequency band without generating non-linear distortion effects.

We specify, for each bin of a discrete frequency domain representation of the system described in Section II-B, the optimization problem

$$\begin{aligned} \underset{\mathbf{R}}{\arg\min} \quad & \|\mathbf{V}\left(\mathbf{H}_B\mathbf{R}-\mathbf{D}\right)\|_2^2 + \|\mathbf{W}\mathbf{R}\|_2^2 \\ \text{s.t.} \quad & \mathbf{R}^H\left(\alpha\mathbf{H}_D^H\boldsymbol{\phi}_D^H\boldsymbol{\phi}_D\mathbf{H}_D - \mathbf{H}_B^H\boldsymbol{\phi}_B^H\boldsymbol{\phi}_B\mathbf{H}_B\right)\mathbf{R} \leq 0. \end{aligned} \quad (14)$$

where $\mathbf{H}_B = \boldsymbol{\mathcal{H}}_B\left(e^{-j\omega}\right)$ and so forth, for each center frequency ω defined by the chosen DFT resolution. Also the filter matrix $\mathbf{R}$ is here a constant matrix that optimizes (14) for a frequency bin of the spectral system representation.

[8]It may also indicate that filter power gain reduction is prioritized over SFS error reduction in the criterion function J, see Section IV-C.

We can reformulate this as a problem in real variables only, with the parameters defined by

$$\mathbf{R}_r = \begin{bmatrix} \text{Real}\{\mathbf{R}\} \\ \text{Imag}\{\mathbf{R}\} \end{bmatrix}, \quad \mathbf{X}_r = \begin{bmatrix} \text{Real}\left\{\mathbf{B}_B^H \mathbf{D}\right\} \\ \text{Imag}\left\{\mathbf{B}_B^H \mathbf{D}\right\} \end{bmatrix}$$

$$\mathbf{H}_1 = \begin{bmatrix} \text{Real}\left\{\mathbf{B}_B^H \mathbf{B}_B + \mathbf{W}^H \mathbf{W}\right\} & -\text{Imag}\left\{\mathbf{B}_B^H \mathbf{B}_B + \mathbf{W}^H \mathbf{W}\right\} \\ \text{Imag}\left\{\mathbf{B}_B^H \mathbf{B}_B + \mathbf{W}^H \mathbf{W}\right\} & \text{Real}\left\{\mathbf{B}_B^H \mathbf{B}_B + \mathbf{W}^H \mathbf{W}\right\} \end{bmatrix}$$

$$\mathbf{S} = \alpha \mathbf{B}_D^H \boldsymbol{\phi}_D^H \boldsymbol{\phi}_D \mathbf{B}_D - \mathbf{B}_B^H \boldsymbol{\phi}_B^H \boldsymbol{\phi}_B \mathbf{B}_B$$

$$\mathbf{H}_2 = \begin{bmatrix} \text{Real}\{\mathbf{S}\} & -\text{Imag}\{\mathbf{S}\} \\ \text{Imag}\{\mathbf{S}\} & \text{Real}\{\mathbf{S}\} \end{bmatrix}, \tag{15}$$

as

$$\begin{aligned} \min_{\mathbf{R}_r} \quad & \mathbf{R}_r^T \mathbf{H}_1 \mathbf{R}_r - 2\mathbf{X}_r^T \mathbf{R}_r \\ \text{s.t.} \quad & \mathbf{R}_r^T \mathbf{H}_2 \mathbf{R}_r \leq 0. \end{aligned} \tag{16}$$

The term $\left\|\mathbf{D}^H \mathbf{D}\right\|_2^2$ was excluded from the optimization above as it is independent of the optimization variable $\mathbf{R}_r$, and will thus not influence the resulting pre-compensator.

A minimizing controller to the Lagrange function associated with the problem above can be found (it will be bounded from below), if $\mathbf{H}_1 + \lambda \mathbf{H}_2 \succeq 0$ and if the target is attainable, i.e. $\mathbf{X}_r = (\mathbf{H}_1 + \lambda \mathbf{H}_2)\, \mathbf{R}_r$. The minimizing solution will then be $\mathbf{R}_r = (\mathbf{H}_1 + \lambda \mathbf{H}_2)^{-1} \mathbf{X}_r$. The Lagrangian dual problem can be formulated

$$\begin{aligned} \arg\max_{\lambda} . \quad & -\mathbf{X}_r^T \left(\mathbf{H}_1 + \lambda \mathbf{H}_2\right)^{-1} \mathbf{X}_r \\ \text{s.t.} \quad & \lambda \geq 0 \\ & \mathbf{H}_1 + \lambda \mathbf{H}_2 \succeq 0 \\ & \mathbf{X}_r = (\mathbf{H}_1 + \lambda \mathbf{H}_2)\, \mathbf{R}_r . \end{aligned} \tag{17}$$

We can formulate this, using the Schur complement in order to avoid the above matrix inverse, as

$$\begin{aligned} \arg\max_{\lambda} \quad & \xi \\ \text{s.t.} \quad & \lambda \geq 0 \\ & \begin{bmatrix} \mathbf{H}_1 + \lambda \mathbf{H}_2 & -\mathbf{X}_r \\ -\mathbf{X}_r^T & \xi \end{bmatrix} \succeq 0, \end{aligned} \tag{18}$$

where ξ is a real valued scalar.

The above problem can be solved for each frequency bin with widely available software, such as e.g. the CVX toolbox for MATLAB [29]. The solutions then represent the optimal filter at each design frequency and the temporal filter can be reconstructed by collating the solutions and inverse transforming the result.

1) A Note on Implementation: As with the proposed causal design method, the desired bright zone spectrum matrix $\mathbf{D}$ lends us control over the desired phase and thus delay in the design of the frequency-domain optimal controller. However, since the design methodology does not account for causality, the generated filters often require a substantial number of non-zero coefficients for negative group delays, preceding the start of the desired impulse response. If the modelling delay in the desired impulse response does not admit all required negative delays, the remaining filter taps will, in the IFFT be wrapped around to the end of the generated filter. This may generate noticeable post-echoes and poor performance between design frequencies. In the experiments described in the following, a base delay (a common delay to all design positions) is therefore introduced that corresponds to half the chosen FIR filter length (which in the experiments herein is 0.6 s). This approximately, on average, centres the main peaks of the produced filters.

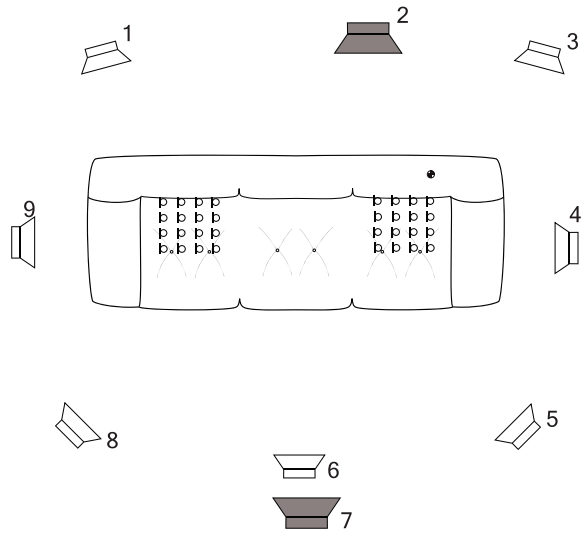

Fig. 2. The experiment setup. White loudspeakers indicate full band speakers, while dark speakers are woofers. The rings hovering roughly at ear height above the sofa indicate the microphone measurement positions.

IV. Simulations

We shall here investigate the proposed method and illuminate various aspects of it and its relationship to the frequency-domain optimal method via a set of simulations.[9] The design parameters of the investigated filters are primarily chosen in order to illustrate these properties and certain aspects of the generated filters may therefore be less than optimal for some applications. Care is however taken to ensure that the filters are implementable in the physical system on which the filter designs and simulations are based without producing clipping or other non-linear artefacts.

A. Experimental Conditions

The simulations are based on mathematical transfer function models of the sound generating system with parameters obtained by measurement-based estimation in a room. A sofa is placed roughly in the middle of a room of dimensions 2.6 by 4.5 by 5.8 meters, with a reverberation time T_{60} slightly less than 0.4 seconds below 400 Hz. Surrounding the sofa, nine loudspeakers are placed, see Fig. 2, the transfer functions of which are estimated in two grids of 4 by 4 equidistant measurement

[9]The word 'simulation' should here be interpreted as the predicted performance based on measured transfer functions.

TABLE II
INVESTIGATED SYSTEM PARAMETER DIMENSIONS

Variable	M_B	M_D	N	L	K
Value	16	16	9	1	1 and 11

positions each. The smallest distance between two of those positions is 0.1 meters and the smallest distance between the two grids is 0.7 meters. One of the grids defines the bright zone whereas the other grid defines the dark zone for our simulations. The various relevant parameter dimensions of the experimental system are summarized in Table II.

The highest frequency completely captured by the chosen measurement position spacing is, using Kirkeby's rule of thumb [30], approximately 1200 Hz. All transfer functions are measured with a sampling rate of 44.1 kHz.

Care has been taken at the design stage to produce filters that behave in a reasonable way with respect to gains and other practical implementation issues. The loudspeaker penalty matrix $\boldsymbol{W}(q^{-1})$ in (1) is therefore chosen so that all filters generate output magnitudes in the linear regions of the loudspeakers. This is in turn achieved by constructing $\boldsymbol{W}(q^{-1})$ as a diagonal matrix with band stop FIR filters on the diagonal, where a high penalty (30 dB above the target gain) is applied outside of each loudspeaker's frequency band and a lower, yet non-negligible penalty (10 dB below the target gain) is applied at the frequencies the loudspeaker is intended to reproduce. Thus, $\boldsymbol{W}(z^{-1})$ will have full rank on $|z| = 1$. In the system on which the simulations are based, loudspeaker 7 is active between 60 Hz and 200 Hz, loudspeaker 2 is active between 30 Hz and 400 Hz, while the remaining loudspeakers are active between 60 Hz and 10 kHz. The bandpass filters are constructed using an FIR impulse response approximation of a set of sixth order IIR Butterworth filters.

A desired system response was specified by $\boldsymbol{D}(q^{-1})$ as a perfect impulse with propagation delays over the grid of control points corresponding to a sound wave emanating from loudspeaker 6 of Fig. 2. The longest common delay, or modelling delay, d_0 is arbitrarily chosen to 10 ms.[10]

Causally constrained designs with acoustic contrast greater than or equal to -10, $+20$ and $+50$ dB over the entire audible spectra are computed based on the system described above. The latter contrast is chosen in order to push the filters beyond what they can be expected to deliver for the investigated system. The first is included as an illustration of what can be attained when no contrast constraints are active.

A multiply constrained (computed subject to multiple constraints) filter is also computed with the intent of investigating option 2 in Section III-C. This filter is based on the actual attained contrast distribution, as obtained with the 20 dB contrast constrained filter described above. The constraints of this filter are -10 dB in the frequency bands 0–100 Hz and 1000–$f_s/2$ Hz, 22 dB in the frequency band 100–200 Hz, 20 dB in the frequency band 200–300 Hz, 15 dB in the frequency bands 300–400 Hz and 400–500 Hz, 10 dB in the frequency band 500–600 Hz, and 5 dB in the frequency bands 600–700 Hz, 700–800 Hz, 800–900 Hz and 900–1000 Hz.

The system itself is modelled as a set of FIR transfer functions ($\boldsymbol{A} = \mathbf{I}$ in Section II-B) of 20 000 taps, covering 0.45 s, as this is deemed sufficient to capture the essential parts of the room acoustics. The choice to model the entire system by a polynomial matrix description, i.e. utilizing only FIR filters simplifies the model-acquisition somewhat, as the model order then only contains one parameter - the FIR length. The drawback of this approach is instead that the denominator, $\boldsymbol{\beta}(q^{-1})$, of the resulting pre-compensator from (7) attains a degree of at least 19 999, see (9). This means that finding the inverse of $\boldsymbol{\beta}(q^{-1})$ may become mathematically problematic and time consuming. If a polynomial matrix (FIR) pre-compensator is acceptable, the FIR approximation to the inverse of $\boldsymbol{\beta}(q^{-1})$ may instead be computed, with a tractable computational load. Note, however, that this approach necessitates an additional design choice in selecting the length of the FIR approximation of the IIR filter.

If an IIR pre-compensator is desired, more effort should be spent on producing a well-parametrized IIR system model. For the purpose of investigating the behavioural properties of the herein proposed design strategy, however, we can tolerate an FIR pre-compensator of a rather high order, and so we can simplify the model acquisition phase by choosing the above described FIR parametrization.

Computing the FIR inverse approximation using a frequency domain method requires, in the particular case investigated here, on the scale of a few (<10) seconds of computational time on an 'HP EliteBook 8540w' laptop purchased in 2011. This simplification also allows direct comparison to the behaviour of the frequency-domain optimal FIR filter. All investigated filters herein consist of 28146 taps, which also corresponds to the DFT length. This filter length ensures that any effects of the truncation have negligible impact on the resulting compensated system.

Frequency-domain optimal designs were generated with the frequency domain representations of the weighting matrices as used in the causal designs. While the proposed causal design method utilizes a constraint on the average contrast over a range of frequencies, the frequency-domain optimal method is defined subject to a constraint at each design frequency. Choosing contrast levels to use as constraints in the causally unconstrained case therefore presents a challenge in terms of comparability to the causal design method. In order to produce a relatively fair comparison, the frequency-domain optimal contrast constraint was set to match the contrast that was attained by the causal filter at each design frequency. A regularization $\mathbf{W}$ was used, with values at each frequency corresponding to those of the matrix $\boldsymbol{W}$ utilized in the causally constrained filter generation, evaluated at that frequency.

Note that while the entire audible spectrum is included in the constraint function, we cannot make any assertions regarding the compensated system outside the design positions for frequencies above the spatial Nyquist frequency. All frequencies covered by the utilized loudspeakers are however shown in most plots in the present section, as they help elucidate the behaviour of the filters produced by the investigated methods.

[10]In practice, also the shortest propagation delay from loudspeaker 6 to any bright zone design position is taken into account so that the assumption of zero propagation delay (see Section II-B) is not violated.

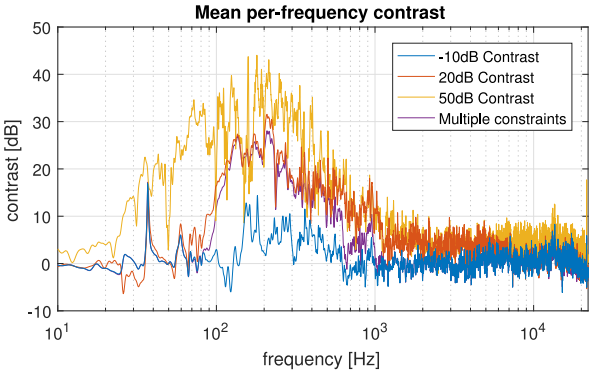

Fig. 3. Contrast per frequency for three different contrast constraint levels, −10 dB (bottom, blue), 20 dB (middle, red) and 50 dB (top, yellow). Note how greater contrasts are primarily achieved in the band between 60 Hz and 1000 Hz, and that the demanded contrast of 50 dB is not attained at any frequency by the associated filter.

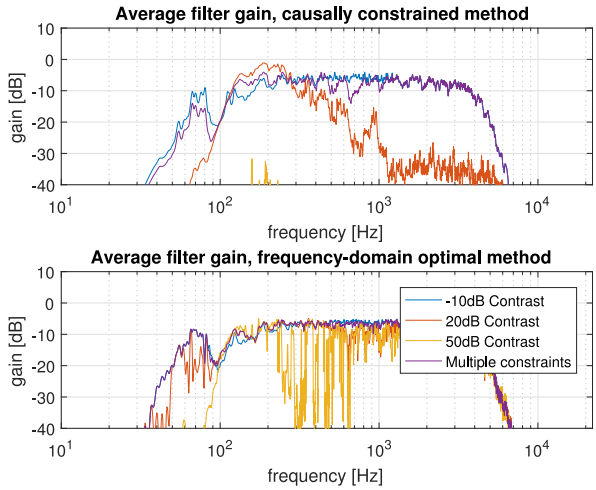

Fig. 4. Filter gain per frequency of the four experimentally investigated contrast constrained filters. A trend of reduced filter bandwidth with increasing contrast can be discerned for the three singly constrained, causal filters (top).

B. Attained Contrast

The spectral distribution of the actual attained contrast of the investigated constrained causal filters (which is used as target in the frequency-domain optimal design for the experiments herein) is shown in Fig. 3. Note that the utilized contrast measure is evaluated with respect to the average over the bright zone control points and dark zone control points. Examining the contrast in any sub-set may produce different results.

High contrast levels are, for all except the 50 dB constrained filter, predominantly concentrated to the frequency band between 80 Hz and 1000 Hz. The lower limit on this band can be explained by the majority of the loudspeakers being inactive at frequencies below 80 Hz due to the design of $\boldsymbol{W}(q^{-1})$. The decline of attained contrast at high frequencies can be attributed to increasing sound field complexity. This is manifested by a rather abrupt increase in contrast over frequency (due to more loudspeakers becoming active), as compared to a far more gradual decrease over frequency, due to the increasing sound field complexity.[11] The demanded contrast of 50 dB is not attained at any frequency by the associated filter. As explained in Section III-C, this filter can still constitute a valid solution to the optimization problem (4) by reducing the controller gain until the constraint is satisfied to the desired tolerance.

One aspect of the results illustrated in Fig. 3 that may seem contradictory to the discussion in Section III-C is that the system clearly is able to attain contrasts in excess of 20 dB below 100 Hz, but the 20 dB constrained filter does not attain this contrast level. The explanation for this behaviour lies in the filter gain penalty matrix $\boldsymbol{W}(q^{-1})$ and the interaction between the two objectives in the objective function J of (4). The magnitude of the penalty matrix $\boldsymbol{W}(q^{-1})$ is in this particular case approximately equal to that of the system model matrix in the bright zone $\boldsymbol{B}_B(q^{-1})$ and the target matrix $\boldsymbol{D}(q^{-1})$ somewhere in the neighbourhood of 90 Hz (due to the finite slope of the filters in $\boldsymbol{W}(q^{-1})$). For frequencies below this, it is more beneficial from an optimization point of view to reduce the filter gain than it is to reduce the SFS error. Reducing the filter gain also, per the discussion in Section III-C, reduces the constraint function magnitude independently of the attained contrast. This effect can also be observed for higher frequencies, where the gain of the produced pre-compensators (Fig. 4) declines with increasing frequency beyond 4 000 Hz even though the loudspeakers output significant power up to 10 000 Hz. A steeper slope of the penalty matrix filters would have reduced this effect.

[11]The likelihood of the nine loudspeakers being able to produce significant contrast between the two groups of 16 control points each decreases with increasing frequency.

C. Filter Gains and Bright Zone Gains

In this section, we investigate one particular aspect of the quality of the produced pre-compensators: the magnitude response of the compensated system. The best possible performance in this aspect would be perfect adherence to the desired magnitude response. We will also investigate the gains of the produced pre-compensators and relate this back to the gain-trading behaviour described in Section III-C. Let us first establish a base-line to which we can relate the other results.

The filter gain and the generated bright zone gain of the unconstrained pre-compensator, is illustrated by the blue line in Figs. 4 and 5 respectively. As the contrast constraint is here set so that it is never active, the only mechanism for limiting the filter gain, except from attaining the desired target gain, is the penalty matrix $\boldsymbol{W}(q^{-1})$. Bright zone gain here is measured as average gain from signal input to microphone output of the system and is proportional to the sound pressure level in the bright zone.

We note, by comparing the bottom and top figures of Fig. 5, that constraining the pre-compensator to be causal introduces a reduction in attained reproduction quality. This reduction further appears to be more prominent in lower frequencies. In terms of average filter gain, Fig. 4, the differences between the causally constrained method and the frequency-domain optimal method are near negligible, although the filter gain is generally somewhat higher with the latter in low frequencies.

Comparing the bright zone gain of the multiply constrained pre-compensator (purple line) to the base-line established above,

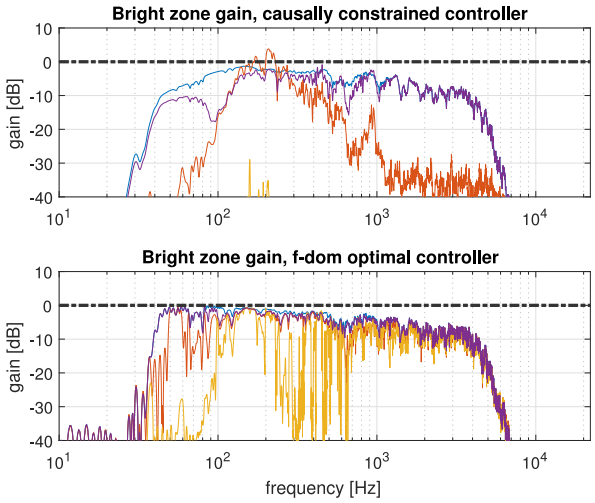

Fig. 5. Gain in the bright zone of the compensated system for three different constraint levels, −10 dB (top, blue), 20 dB (middle, red) and 50 dB (lower, yellow). A multiply constrained filter, as specified in Section IV-A is also included (purple). The dashed horizontal line indicates the desired bright zone gain specified in the target matrix $\boldsymbol{D}$.

it is clear that the SFS performance reduction caused by contrast generation can indeed be kept quite low. It is clear, by Fig. 3, that the proposed method is capable of producing a significant improvement in generated contrast at the cost of a small deterioration in bright zone performance.

The 50 dB constrained pre-compensator (yellow line) represents the extreme point to the base-line pre-compensator, as the desired contrast can here not be attained at any frequency. By Fig. 4 (top), the filter gain reduction effect due to unattainable contrast constraints, as discussed in Section III-C, is clear (the yellow line largely falls outside of the figure). The negligible filter gain also translates to negligible bright zone gain, Fig. 5 (top).

The more nuanced effects of the gain trading can be observed in the 20 dB constrained pre-compensator (red). High filter- and bright zone gain is here produced where the desired contrast can be generated, Figs. 3–5. Where the desired contrast levels cannot be reproduced, the filter gain is instead reduced but a higher gain is retained also in these frequencies, as compared to the 50 dB constrained filter. This is a direct effect of the gain trading of the filter optimization, which compensates the gain in frequencies with less contrast than required with the gain in the frequency bands where more than the desired contrast is attained.

The difference (all quantities in dB) between the bright zone power and the dark zone power at any single frequency is characterized by the acoustic contrast. We can thus get a sense of the magnitude gain profile of the dark zone by comparing Figs. 3 and 5.

Such a comparison indicates that the 20 dB constrained pre-compensator produces low dark zone gains across all frequencies. In contrast, the multiply constrained pre-compensator produces low dark zone gains in the frequencies where significant contrast is attained but higher such gains where little contrast is generated.

As the per-frequency contrast produced by the causally constrained filters is used to set the constraints of the frequency-domain optimal filters, the desired contrasts can in these cases always be attained (albeit with some apparent difficulty at the very highest contrasts, evident by the erratic fluctuation of the average filter gain of the 50 dB constrained controller), see Fig. 4 (bottom). This means that a higher filter gain can be expected from the frequency-domain optimal filters, which is also evident in the figure. Further, the bandwidth of each contrast constraint as defined in the frequency-domain optimal design method is far smaller than the band width of any constraint function explored for the causally constrained method. This also means that 'gain trades' over a frequency band, as may occur in the causally constrained method are generally not sought by the frequency-domain optimal design method.

The frequency-domain optimal method has a potential advantage in that a constraint level is set for each individual frequency bin. The price for this is instead increased pre-ringing and reduced control over the absolute delays of the filter, as will be illustrated in the following section.

D. Temporal Properties

The main motivation for choosing the causal problem description in the first place was to gain control over the temporal properties of the compensated system. The (normalized) impulse response simulated at a microphone in the middle of the bright zone is shown in Fig. 6. The top figure shows compensated impulse responses using the causal and frequency-domain optimal filters, with relative magnitudes and absolute delays. The lower figures show a version where the main peak amplitudes are individually normalized in order to illuminate the relative amplitudes of the pre- and post ringing of the compensated systems. In addition, the causal filter is delayed to match the larger delay of the frequency-domain optimal pre-compensator. Only the behaviour generated by the multiply constrained filter is shown in Fig. 6, as it is found to be representative also for all other investigated cases. An exception to this is the causal 50 dB constrained filter, which produces negligible output and therefore also is of little practical use.

The delay difference between the two filters is apparent in Fig. 6 (top). The magnitude of the main impulses of the compensated systems differ due to the different ways in which the filter gains are traded, as discussed in Section IV-C. It is also clear from Fig. 6 (middle and bottom) that the causally unconstrained controller causes pre-ringing in the compensated system with a far longer duration than the causally constrained controller does. The filter found using the proposed algorithm causes pre-ringing for a duration of roughly 10 ms. The pre-ringing amplitude of the filter produced using the frequency-domain optimal algorithm at more than 100 ms before the main peak are still larger than those at 15 ms before the main peak of the filter generated with the proposed algorithm. The amount of generated pre-ringing will vary between different systems and control point layouts. The perceptibility of a given pre-ringing envelope is further a complicated topic, but pre-ringings may be audible from as early as 20 ms down to as late as 5 ms before the main impulse (see

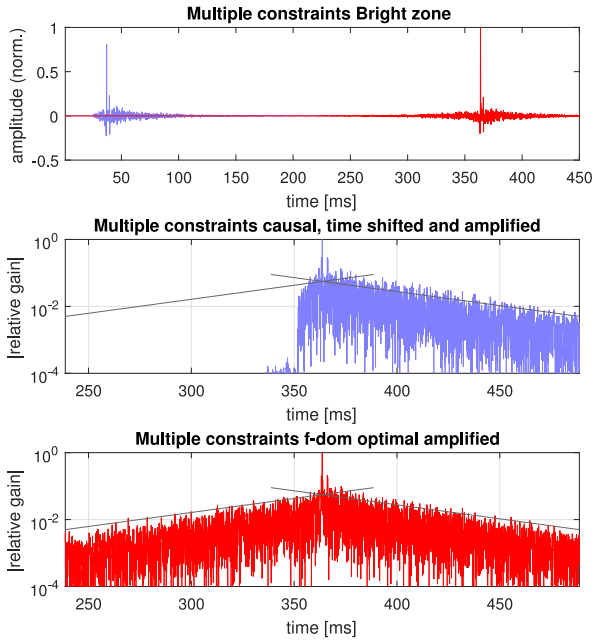

Fig. 6. Impulse response of the compensated system simulated in the middle of the bright zone, for the multiply constrained controller. Lighter (blue) impulse illustrates the causal controller while the darker (red) impulse illustrates the frequency-domain optimal controller. Middle figure depicts the causal impulse response magnitude relative to the largest magnitude of the impulse response, time-shifted to match the frequency-domain optimal impulse response and plotted in logarithmic scale. The bottom figure depicts the frequency-domain optimal filter relative to its largest peak and plotted on a logarithmic scale. The sloped lines in the middle and lower plots are reference lines introduced in order to allow comparison between the two figures.

e.g. [31], [32]). Perceptibility notwithstanding, the above results demonstrate that the proposed method allows a far greater level of control over any delay or pre-ringing related issues that may arise than does the causally unconstrained method.

E. The Effects of Increased Modelling Delay

By increasing the modelling delay, d_0, the behaviour of the causal design becomes increasingly similar to a time-shifted version of the non-causal Wiener filter formulation. The constraint is however still specified with respect to the average contrast of the produced filter, and so we do not expect it to approximate the investigated frequency-domain optimal method, unless we also increase the number of constraints over increasingly narrow sub-bands.

The effect of increasing the modelling delay on the attained contrast and on the bright zone spectral properties is shown in Fig. 7 for the multiply constrained filter. It is worth emphasizing that, since the spatial sampling frequency of the experiment set-up corresponds to a temporal frequency of about 1200 Hz, any results beyond this limit are valid only at the exact design positions. For this reason, the frequency axis in Fig. 7 does not extend beyond 1200 Hz.

It is apparent from Fig. 7 that increasing the modelling delay will have some positive effects on the attained contrast in the system under investigation. This effect appears to be somewhat more prominent at lower frequencies. It also appears that the

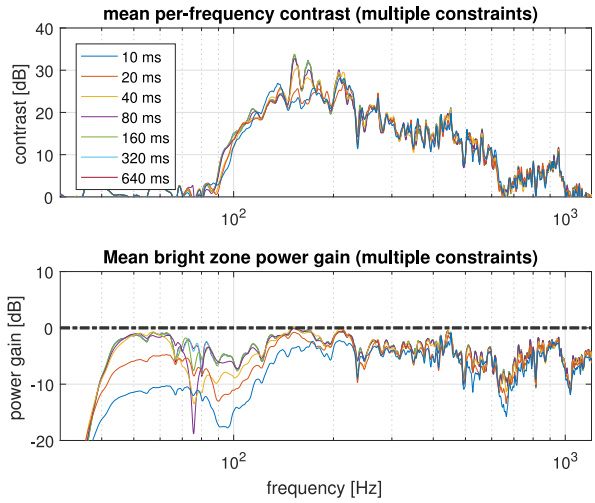

Fig. 7. Experiments with increasing modelling delay, d_0 for the multiply constrained filter. The different investigated modelling delays are, by increasing line brightness 10 ms (blue), 20 ms (red), 40 ms (yellow), 80 ms (purple), 160 ms (green), 320 ms (teal) and 640 ms (maroon). The latter two are often obscured by the 160 ms curve. The dashed line visible in the lower figure indicates the desired bright zone gain, so controllers that follow the dashed line more closely perform better in terms of (spectral magnitude) bright zone target following.

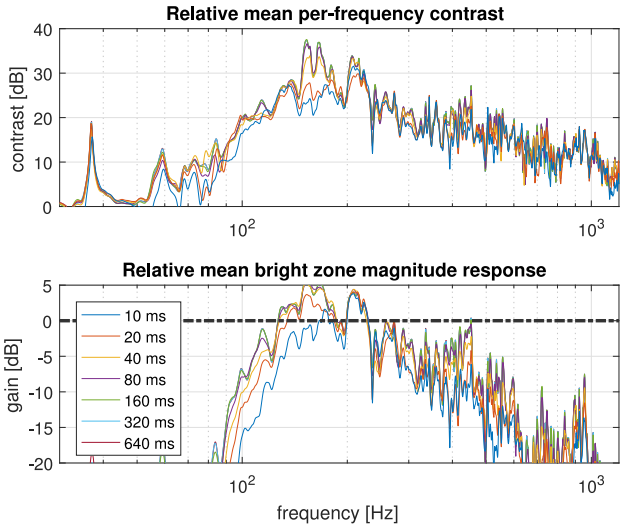

Fig. 8. Experiments with increasing modelling delay, d_0 for a filter requiring an acoustic contrast of at least 20 dB. The different investigated modelling delays are, by increasing line brightness 10 ms (blue), 20 ms (red), 40 ms (yellow), 80 ms (purple), 160 ms (green), 320 ms (teal) and 640 ms (maroon). The latter two are often obscured by the 160 ms curve. The dashed line visible in the lower figure indicates the desired bright zone gain, so controllers that follow the dashed line more closely perform better in terms of (spectral magnitude) bright zone target following.

highest contrast levels can (here) only be attained with modelling delays in excess of 20 ms.

The improvement in terms of bright zone behaviour is clear (Fig. 7 (bottom)). As was the case for the -10 dB constrained pre-compensator (Fig. 5), improvements due to reduced strictness of the causality constraint is also here more prominent in lower frequencies. This behaviour is also observed in, e.g., [33].

Looking instead at the behaviour of the 20 dB singly constrained pre-compensator, the improved contrast in the

frequency band between 150–200 Hz means that a more aggressive gain-trading can be performed, Fig. 8.

The result is that the bright zone target adherence is improved with an increasing modelling delay at most frequencies, except those where good contrast is attained, where the target gain is increasingly overshot with increasing modelling delay.

The price of increasing the modelling delay is increased delay and increased pre-ringing duration of the compensated system.

V. Concluding Remarks

A novel design method, producing a causal, stable IIR pre-compensator that is constrained with respect to the produced average contrast of the compensated system is presented and compared to a known frequency-domain optimal pre-compensator. A significant discrepancy between the two formulations in terms of the constraint being formulated as an average over a wider frequency band rather than point wise in the frequency domain, is discussed. The causal method can also be formulated with several constraints, each defined for a different frequency band. This version of the filter is shown to be able to achieve results that are closer to those of the frequency-domain optimal filter, while retaining the causal property of the generated filters.

The delay of the main impulse of the compensated system is freely controllable using the causal controller. Contrastingly, the delay of the filter produced with the frequency-domain optimal method is affected by implementation issues arising from the non-causal properties of the design method. This in turn allows the filter derived using the causally constrained method to generate pre-ringing of shorter duration in the compensated system than the filter produced by the frequency-domain optimal method.

The term 'constraints' is used loosely in the design process described herein. The filters are indeed computed subject to a set of user-specified constraints and if minimum allowed contrasts in a set of pre-specified frequency bands are known, the optimal filter adhering to these constraints will be produced. In cases where the minimum required contrast is loosely specified, however, the constraint levels can be viewed as a tuning parameter controlling the achieved contrast which is traded against the sound field synthesis error.

The presented framework can be generalized in several ways. Interesting research that does not fit in the scope of the present paper, but is left for future efforts includes expansion of the method to be robust to modelling errors and the inclusion of an explicit pre-ringing constraint.

Appendix A
Proofs of Section II

A. Outline

As the Lagrange multiplier method will yield an optimal solution to the primal problem when a global minimizer can be found, we will focus on finding this minimizing controller. The problem at hand offers a particular complication in that the spectral factorization that we shall see is part of our solution cannot always be solved. A solution only exists if the negative term(s), $E\left\{\boldsymbol{\sigma}_{Bk}^T\boldsymbol{\sigma}_{Bk}\right\}$ in (6) do not dominate the other quadratic terms for any choice of $\boldsymbol{\mathcal{R}}$. We shall therefore also provide proof that the optimal solution is contained in the set of multipliers λ_k that guarantees this situation will not arise. In order to facilitate the numerical search for the solution, an upper limit on λ_k, below which the spectral factorization will not fail is also derived.

The proofs of the claims made in Section III-A are divided into two parts. In Section VI-B, we will show that the sought optimal solution is included in the solution space defined by a Lagrange function that is bounded from below. We will also show *how* the multiplier λ_k can be limited in such a way that the Lagrange function is bounded from below, while not excluding the sought optimal value. Finally, in Section VI-C, we derive the expressions (7)–(10) describing the minimizing controller for any $\lambda_1, \ldots, \lambda_K$ that satisfies (12).

B. Lower Boundedness of the Lagrange Function $\mathcal{L}(\boldsymbol{\mathcal{R}}, \lambda)$

A global minimizer to the Lagrange function can be found in the present case if the Lagrange function is bounded from below, as the function surface is then a quadratic 'bowl'. We shall here establish a condition that guarantees the Lagrange function to be bounded from below and show that the optimal constrained solution is not excluded by this condition.

First, we will here establish, for $K = 1$, that by assuring that the expression

$$\boldsymbol{M}(z, z^{-1}) = \boldsymbol{A}_*^{-1}\left(\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A} + \lambda\left(\alpha\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D*}\boldsymbol{\phi}_D\boldsymbol{B}_D - \boldsymbol{B}_{B*}\boldsymbol{\phi}_{B*}\boldsymbol{\phi}_B\boldsymbol{B}_B\right)\right)\boldsymbol{A}^{-1}, \quad (19)$$

is positive definite on the unit circle, $|z| = 1$, and by choosing $\boldsymbol{W}(q^{-1})$ so that it is of full rank, the Lagrange function $\mathcal{L}(\boldsymbol{\mathcal{R}}, \lambda)$ will be bounded from below. When this is shown, we will show that the optimal solution (by the Karush-Kuhn-Tucker conditions, $\lambda C\left(\boldsymbol{\mathcal{R}}\right) = 0$, where $C(\boldsymbol{\mathcal{R}}) \leq 0$) is contained in this new solution set. We will then show how to choose the interval for λ that defines our new solution set.

1) The Lagrange Function is Bounded From Below: Choosing $\boldsymbol{M}(z, z^{-1}) \succ 0$ on $|z| = 1$ implies that the integrand below is a valid spectral density matrix and that

$$f\left(\boldsymbol{\mathcal{R}}, \lambda\right) = \frac{1}{2\pi j}\oint_{|z|=1} \operatorname{tr}\left[\boldsymbol{\mathcal{R}}_*\boldsymbol{A}_*^{-1}\left(\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A} + \lambda\alpha\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D*}\boldsymbol{\phi}_D\boldsymbol{B}_D - \lambda\boldsymbol{B}_{B*}\boldsymbol{\phi}_{B*}\boldsymbol{\phi}_B\boldsymbol{B}_B\right)\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\right]\frac{dz}{z} > 0. \quad (20)$$

Using $E\left\{\boldsymbol{r}^T(t)\boldsymbol{r}(t)\right\} = \mathbf{I}$, the function (20) can then be interpreted as the following sum of correlation matrices in the time

domain by using (1), (2), (3) and (6)

$$\begin{aligned} f(\boldsymbol{\mathcal{R}},\lambda) &= E\Big\{\left(\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad + \left(\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad + \lambda\alpha\left(\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad -\lambda\left(\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\Big\} \\ &= \mathcal{L}(\boldsymbol{\mathcal{R}},\lambda) - g(\boldsymbol{\mathcal{R}}) > 0, \forall\boldsymbol{\mathcal{R}}(q^{-1}), \end{aligned} \tag{21}$$

where

$$\begin{aligned} g(\boldsymbol{\mathcal{R}}) &= E\Big\{\left(\boldsymbol{V}\boldsymbol{D}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{D}\boldsymbol{r}(t) \\ &\quad - \left(\boldsymbol{V}\boldsymbol{D}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad - \left(\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{D}\boldsymbol{r}(t)\Big\}. \end{aligned} \tag{22}$$

Note that all individual terms of (21) are quadratic in $\boldsymbol{\mathcal{R}}(q^{-1})$ and, since $g(\boldsymbol{\mathcal{R}})$ contains only linear terms in $\boldsymbol{\mathcal{R}}(q^{-1})$, $f(\boldsymbol{\mathcal{R}},\lambda)$ will dominate $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$ for large values of $E\left\{\left(\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right\}$. As $f(\boldsymbol{\mathcal{R}},\lambda) > 0$ for all $\boldsymbol{\mathcal{R}} \neq 0$, the only possible way that $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$ can be unbounded from below is if $g(\boldsymbol{\mathcal{R}})$ could be influenced by $\boldsymbol{\mathcal{R}}$ independently of $f(\boldsymbol{\mathcal{R}},\lambda)$. This in turn is not possible since we assume the polynomial matrix $\boldsymbol{W}$ to be of full rank at all frequencies, meaning that there is no possible way in which we can increase $E\left\{\left(\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right\}$ that does not cause a quadratic increase in $f(\boldsymbol{\mathcal{R}},\lambda)$.

In conclusion, since $f(\boldsymbol{\mathcal{R}},\lambda) > 0$ for all $\boldsymbol{\mathcal{R}}$, no $\boldsymbol{\mathcal{R}}$ can cause $f(\boldsymbol{\mathcal{R}},\lambda) \to -\infty$. Since $\boldsymbol{W}$ is of full rank for all frequencies, choosing $\boldsymbol{\mathcal{R}}$ so that $g(\boldsymbol{\mathcal{R}}) \to -\infty$ linearly, also causes $f(\boldsymbol{\mathcal{R}},\lambda) \to \infty$ quadratically, and so $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda) = f(\boldsymbol{\mathcal{R}},\lambda) + g(\boldsymbol{\mathcal{R}})$ is bounded from below.

2) The Sought Optimal Solution Is a Part of Our Search Space: We will now show that the conditions $\boldsymbol{M}(z,z^{-1}) \succ 0$ on $|z|=1$ and $\boldsymbol{W}(q^{-1})$ of full rank do not exclude our sought solution, where $\lambda C\left(\boldsymbol{\mathcal{R}}(q^{-1})\right) = 0$.

For a minimizing solution $\boldsymbol{\mathcal{R}}$ to $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$ for a given λ, to produce a non-positive value when inserted into (21) (i.e. for the minimization problem to be unbounded from below), it would have to be such that

$$\begin{aligned} &E\Big\{\left(\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad + \left(\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad + \lambda\alpha\left(\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\Big\} \\ &\leq E\left\{\lambda\left(\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right\}, \end{aligned} \tag{23}$$

since all terms individually are always non-negative by construction. As $\lambda \geq 0$ by definition, and $\lambda = 0 \Rightarrow \boldsymbol{M}(z,z^{-1}) \succeq 0$, we must have $\lambda > 0$. Since we assume that $\boldsymbol{W}$ is of full rank, we also know that $E\left\{\left(\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right\} > 0$, and would therefore need to assume that

$$\begin{aligned} &E\Big\{\left(\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &\quad + \lambda\alpha\left(\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\Big\} \\ &< E\left\{\lambda\left(\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right\}. \end{aligned} \tag{24}$$

Inserting the above condition into the constraint function yields

$$\begin{aligned} C(\boldsymbol{\mathcal{R}}) = E\Big\{&\alpha\left(\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_D\boldsymbol{B}_D\boldsymbol{\mathcal{R}}\boldsymbol{r}(t) \\ &- \left(\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\right)^T\boldsymbol{\phi}_B\boldsymbol{B}_B\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\Big\} < 0. \end{aligned} \tag{25}$$

We therefore conclude that $\lambda C(\boldsymbol{\mathcal{R}}) < 0$ in all cases that would make $\boldsymbol{M}(z,z^{-1})$ non positive definite on $|z|=1$. By the Karush-Kuhn-Tucker conditions, we are interested in the solution $\boldsymbol{\mathcal{R}}(q^{-1})$ where $\lambda C(\boldsymbol{\mathcal{R}}(q^{-1})) = 0, C(\boldsymbol{\mathcal{R}}) \leq 0$. If such a solution exists, it must therefore do so for $\boldsymbol{M} \succ 0$, when $\boldsymbol{W}$ is of full rank.

The condition that $\boldsymbol{W}$ must have full rank is in practice not very restrictive, such a matrix, serving our purpose, is e.g. readily designed with band stop filters along the diagonal, that have non-zero gains for all $z = e^{j\omega}$.

3) Finding the Limiting Multiplier λ: Assume that $\boldsymbol{\mathcal{R}}_a$ is the minimizer to $\mathcal{L}(\boldsymbol{\mathcal{R}},a\eta)$ and that $\boldsymbol{\mathcal{R}}_{a+1}$ minimizes $\mathcal{L}(\boldsymbol{\mathcal{R}},(a+1)\eta)$, $a \in \mathbb{N}, \eta > 0$. Then

$$\begin{aligned} \mathcal{L}(\boldsymbol{\mathcal{R}}_a,(a+1)\eta) &= J(\boldsymbol{\mathcal{R}}_a) + a\eta C(\boldsymbol{\mathcal{R}}_a) + \eta C(\boldsymbol{\mathcal{R}}_a) \\ &= \mathcal{L}(\boldsymbol{\mathcal{R}}_a,a\eta) + \eta C(\boldsymbol{\mathcal{R}}_a) \geq \mathcal{L}(\boldsymbol{\mathcal{R}}_{a+1},(a+1)\eta) \\ &= \mathcal{L}(\boldsymbol{\mathcal{R}}_{a+1},a\eta) + \eta C(\boldsymbol{\mathcal{R}}_{a+1}). \end{aligned} \tag{26}$$

letting $\eta \to 0$, we have that increasing λ cannot increase the value of the associated constraint function as $\mathcal{L}(\boldsymbol{\mathcal{R}}_a,a\eta) \leq \mathcal{L}(\boldsymbol{\mathcal{R}}_{a+1},a\eta)$. This has previously been shown in e.g. [22]. The λ that provide $\boldsymbol{M}(z,z^{-1}) \succ 0$ on $|z|=1$ must therefore be smaller than the smallest λ that provides an indefinite $\boldsymbol{M}(z,z^{-1})$ anywhere on $|z|=1$. This can in turn be found as the smallest λ for which at least one eigenvalue of $\boldsymbol{M}(z,z^{-1})$ is zero anywhere on the unit circle $|z|=1$,

$$\begin{aligned} &\boldsymbol{A}_*^{-1}\left(\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B + \boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A}\right. \\ &\quad \left.+\gamma^{-1}\left(\alpha\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D*}\boldsymbol{\phi}_D\boldsymbol{B}_D - \boldsymbol{B}_{B*}\boldsymbol{\phi}_{B*}\boldsymbol{\phi}_B\boldsymbol{B}_B\right)\right)\boldsymbol{A}^{-1}\Lambda = 0\Lambda. \end{aligned} \tag{27}$$

Here, γ^{-1} has been substituted for λ. We can now rearrange (27) into an eigenvalue problem as in (13), with $K=1$.

The derivation above considers only a single constraint function. When several constraint functions are considered the

equation instead becomes

$$\boldsymbol{A}_*^{-1}\Bigg(\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B+\boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A}+\sum_{k=1}^{K}\gamma_k^{-1}\left(\alpha_k\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D_k*}\boldsymbol{\phi}_{D_k}\boldsymbol{B}_D-\boldsymbol{B}_{B*}\boldsymbol{\phi}_{B_k*}\boldsymbol{\phi}_{B_k}\boldsymbol{B}_B\right)\Bigg)\boldsymbol{A}^{-1}\Lambda=0\Lambda. \tag{28}$$

We can simplify (28) into (13) if the constraints describe different, well separated frequency bands with negligible overlap. This is possible if all polynomial weight matrices $\boldsymbol{\phi}_{\bullet_k}(z^{-1})$ are approximately zero at all frequencies where any other such filter matrix is non-zero.[12]

In conclusion, we will search for each optimal multiplier λ_k in the direction of increasingly satisfied constraint functions. The practical reason for doing so is that the way we find the minimizing controller $\boldsymbol{\mathcal{R}}(q^{-1})$ relies on $\boldsymbol{M}(q,q^{-1})$ being a valid spectral factorization problem, as will be clear in the following section. This also means that the search will be performed in the direction of generally increasing λ_k, where the first tested hypothesis is always $\lambda_k=0$. By performing the search in this direction and keeping within the boundaries of $0\leq\lambda_k<\lambda_k^{max}$, we always have a solvable problem with the optimal solution contained within our search space.

C. Minimizing $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$ w.r.t $\boldsymbol{\mathcal{R}}(q^{-1})$

We will here show that the filter $\boldsymbol{\mathcal{R}}(q^{-1})$ that minimizes $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$ for a given scalar λ is found as the solution to

$$\boldsymbol{\mathcal{R}}(q^{-1})=\boldsymbol{A}\boldsymbol{\beta}^{-1}\boldsymbol{Q}, \tag{29}$$

where,

$$\boldsymbol{\beta}_*(q)\boldsymbol{\beta}(q^{-1})=\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B+\boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A}+\sum_{k=1}^{K}\lambda_k\left(\alpha_k\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D_k*}\boldsymbol{\phi}_{D_k}\boldsymbol{B}_D-\boldsymbol{B}_{B*}\boldsymbol{\phi}_{B_k*}\boldsymbol{\phi}_{B_k}\boldsymbol{B}_B\right) \tag{30}$$

and

$$\boldsymbol{\beta}_*\boldsymbol{Q}-\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{D}=q\boldsymbol{L}_*. \tag{31}$$

The derivations will be performed for the case $K=1$ in order to keep the mathematics succinct, but all steps below can be easily performed by substituting a sum of constraint functions for the single constraint.

We will show that (29) is optimal by adding a stable and causal variation filter, $\boldsymbol{\mathcal{T}}$ to the filter $\boldsymbol{\mathcal{R}}$ and showing that if the filter $\boldsymbol{\mathcal{R}}$ is chosen according to (29)–(31), then no $\boldsymbol{\mathcal{T}}$ can possibly achieve a smaller value of $\mathcal{L}(\boldsymbol{\mathcal{R}},\lambda)$. This derivation echoes that of e.g. [34].

Adding a causal, stable variation, $\boldsymbol{\mathcal{T}}(q^{-1})$ to $\boldsymbol{\mathcal{R}}(q^{-1})$, the signals generated by $\boldsymbol{\mathcal{R}}$ and $\boldsymbol{\mathcal{T}}$ individually are

$$\boldsymbol{\varepsilon}(t)=\boldsymbol{V}\left(\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}-\boldsymbol{D}\right)\boldsymbol{r}(t),\quad \boldsymbol{\varepsilon}'(t)=\boldsymbol{V}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{T}}\boldsymbol{r}(t), \tag{32a}$$

$$\boldsymbol{u}(t)=\boldsymbol{W}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t),\quad \boldsymbol{u}'(t)=\boldsymbol{W}\boldsymbol{\mathcal{T}}\boldsymbol{r}(t), \tag{32b}$$

$$\boldsymbol{\sigma}_D(t)=\boldsymbol{\phi}\boldsymbol{B}_D\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\quad \boldsymbol{\sigma}'_D(t)=\boldsymbol{\phi}\boldsymbol{B}_D\boldsymbol{A}^{-1}\boldsymbol{\mathcal{T}}\boldsymbol{r}(t), \tag{32c}$$

$$\boldsymbol{\sigma}_B(t)=\boldsymbol{\phi}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}\boldsymbol{r}(t)\quad \boldsymbol{\sigma}'_B(t)=\boldsymbol{\phi}\boldsymbol{B}_B\boldsymbol{A}^{-1}\boldsymbol{\mathcal{T}}\boldsymbol{r}(t). \tag{32d}$$

Now, with the filter being $\boldsymbol{\mathcal{R}}+\boldsymbol{\mathcal{T}}$,

$$\begin{aligned}\mathcal{L}(\boldsymbol{\mathcal{R}},\boldsymbol{\mathcal{T}},\lambda)=E\Big\{&(\boldsymbol{\varepsilon}(t)+\boldsymbol{\varepsilon}'(t))^T(\boldsymbol{\varepsilon}(t)+\boldsymbol{\varepsilon}'(t))\\&+(\boldsymbol{u}(t)+\boldsymbol{u}'(t))^T(\boldsymbol{u}(t)+\boldsymbol{u}'(t))\\&+\lambda\alpha(\boldsymbol{\sigma}_D(t)+\boldsymbol{\sigma}'_D(t))^T(\boldsymbol{\sigma}_D(t)+\boldsymbol{\sigma}'_D(t))\\&-\lambda(\boldsymbol{\sigma}_B(t)+\boldsymbol{\sigma}'_B(t))^T(\boldsymbol{\sigma}_B(t)+\boldsymbol{\sigma}'_B(t))\Big\},\end{aligned} \tag{33}$$

which we can express as

$$\begin{aligned}\mathcal{L}(\boldsymbol{\mathcal{R}},\boldsymbol{\mathcal{T}},\lambda)=&\mathcal{L}_1(\boldsymbol{\mathcal{R}},\lambda)+\mathcal{L}_2(\boldsymbol{\mathcal{R}},\boldsymbol{\mathcal{T}},\lambda)+\mathcal{L}_3(\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{R}},\lambda)\\&+\mathcal{L}_4(\boldsymbol{\mathcal{T}},\lambda).\end{aligned} \tag{34}$$

Above,

$$\mathcal{L}_1(\boldsymbol{\mathcal{R}},\lambda)=E\left\{\boldsymbol{\varepsilon}^T\boldsymbol{\varepsilon}+\boldsymbol{u}^T\boldsymbol{u}+\lambda\left(\alpha\boldsymbol{\sigma}_D^T\boldsymbol{\sigma}_D-\boldsymbol{\sigma}_B^T\boldsymbol{\sigma}_B\right)\right\}, \tag{35a}$$

$$\mathcal{L}_2(\boldsymbol{\mathcal{R}},\boldsymbol{\mathcal{T}},\lambda)=E\left\{\boldsymbol{\varepsilon}^T\boldsymbol{\varepsilon}'+\boldsymbol{u}^T\boldsymbol{u}'+\lambda\left(\alpha\boldsymbol{\sigma}_D^T\boldsymbol{\sigma}'_D-\boldsymbol{\sigma}_B^T\boldsymbol{\sigma}'_B\right)\right\}, \tag{35b}$$

$$\mathcal{L}_3(\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{R}},\lambda)=E\left\{\boldsymbol{\varepsilon}'^T\boldsymbol{\varepsilon}+\boldsymbol{u}'^T\boldsymbol{u}+\lambda\left(\alpha\boldsymbol{\sigma}_D'^T\boldsymbol{\sigma}_D-\boldsymbol{\sigma}_B'^T\boldsymbol{\sigma}_B\right)\right\}, \tag{35c}$$

$$\mathcal{L}_4(\boldsymbol{\mathcal{T}},\lambda)=E\left\{\boldsymbol{\varepsilon}'^T\boldsymbol{\varepsilon}'+\boldsymbol{u}'^T\boldsymbol{u}'+\lambda\left(\alpha\boldsymbol{\sigma}_D'^T\boldsymbol{\sigma}'_D-\boldsymbol{\sigma}_B'^T\boldsymbol{\sigma}'_B\right)\right\}. \tag{35d}$$

Note that (35d) is always greater than or equal to zero so long as (21) holds, which it must do if we choose λ according to (12) and if $\boldsymbol{W}(q^{-1})$ is of full rank. If we find a filter $\boldsymbol{\mathcal{R}}(q^{-1})$ such that $\mathcal{L}_3=\mathcal{L}_2^T=0$, no addition $\boldsymbol{\mathcal{T}}(q^{-1})$ to that controller can result in a smaller value for $\mathcal{L}(\boldsymbol{\mathcal{R}},\boldsymbol{\mathcal{T}},\lambda)$.

Using the trace rotate property and Parseval's formula, we can rewrite $\mathcal{L}_3(\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{R}},\lambda)=0$ so that

$$\begin{aligned}\mathcal{L}_3(\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{R}},\lambda)=\frac{1}{2\pi j}\oint_{|z|=1}\mathrm{tr}\Big[\boldsymbol{\mathcal{T}}_*\boldsymbol{A}_*^{-1}\Big(\big(&\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{B}_B+\boldsymbol{A}_*\boldsymbol{W}_*\boldsymbol{W}\boldsymbol{A}+\lambda\alpha\boldsymbol{B}_{D*}\boldsymbol{\phi}_{D*}\boldsymbol{\phi}_D\boldsymbol{B}_D\\&-\lambda\boldsymbol{B}_{B*}\boldsymbol{\phi}_{B*}\boldsymbol{\phi}_B\boldsymbol{B}_B\big)\boldsymbol{A}^{-1}\boldsymbol{\mathcal{R}}-\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{D}\Big)\Big]\frac{dz}{z}=0.\end{aligned} \tag{36}$$

By the residue theorem, we now know that $\mathcal{L}_3(\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{R}},\lambda)=0$ if the integrand of (36) has no poles within $|z|\leq 1$. The factorization (30) is possible if the matrix on the right hand side

[12]Note that the spectral factorization will fail if the matrix $\boldsymbol{M}\left(e^{j\omega}\right)$ is not positive definite at *all* frequencies. By dividing the frequency band into several regions, where each constraint function only contributes to one region and each region is only affected by one constraint function, the condition becomes that every limiting multiplier λ_k^{max} must be small enough that the addressed constraint function has no negative eigenvalue in the frequency domain.

of the equation is a valid spectrum, which it is for all potentially optimal values of λ per Appendix VI-B2. The spectral factor $\boldsymbol{\beta}(q^{-1})$ of (8) thereby has all zeros (roots of $\det(\boldsymbol{\beta}(z^{-1}))=0$) within the unit circle and is stably invertible, so we can choose $\boldsymbol{\mathcal{R}}(q^{-1})$ according to (29). This yields

$$\mathcal{L}_3 = \frac{1}{2\pi j}\oint_{|z|=1} \operatorname{tr}\left[\boldsymbol{\mathcal{T}}_*\boldsymbol{A}_*^{-1}\left(\boldsymbol{\beta}_*\boldsymbol{Q}-\boldsymbol{B}_{B*}\boldsymbol{V}_*\boldsymbol{V}\boldsymbol{D}\right)\right]\frac{dz}{z}. \tag{37}$$

All conjugate matrices are stable polynomial or rational matrices, which have their poles mirrored in the unit circle, meaning they have all their poles outside of the unit circle and do not contribute any poles within $|z|\leq 1$. This leads us to the Diophantine equation (31), of which the solving matrix pair[13] $\boldsymbol{Q}(q^{-1})$ and $\boldsymbol{L}_*(q)$ guarantees that the integrand of (37) has no poles within the unit disc and therefore evaluates to zero. This in turn guarantees that no addition to the controller $\boldsymbol{\mathcal{R}}(q^{-1})$ can be beneficial and thus that the controller (29) is optimal.

ACKNOWLEDGMENT

The author wishes to express his gratitude towards Prof. M. Sternad for providing numerous and invaluable improvements on the manuscript.

REFERENCES

[1] J.-W. Choi and Y.-H. Kim, "Generation of an acoustically bright zone with an illuminated region using multiple sources," *J. Acoust. Soc. Amer.*, vol. 111, no. 4, pp. 1695–1700, Apr. 2002.

[2] W. F. Druyvesteyn and J. Garas, "Personal sound," *J. Audio Eng. Soc.*, vol. 45, no. 9, pp. 685–701, Sep. 1997.

[3] W. F. Druyvesteyn, R. M. Aarts, P. Gelat, and A. Ruxton, "Personal sound," *Proc. Inst. Acoust.*, vol. 16, no. 2, p. 3281, 1994.

[4] K. Baykaner, C. Hummersone, R. Mason, and S. Bech, "The prediction of the acceptability of auditory interference based on audibility," in *Proc. AES 52nd Int. Conf. Sound Field Control—Eng. Perception*, Sep. 2013, pp. 1–7.

[5] M. Shin *et al.*, "Maximization of acoustic energy difference between two spaces," *J. Acoust. Soc. Amer.*, vol. 128, no. 1, pp. 121–131, Jul. 2010.

[6] M. B. Møller, M. Olsen, and F. Jacobsen, "A hybrid method combining synthesis of a sound field and control of acoustic contrast," in *Proc. 132nd Audio Eng. Soc. Conv.*, Apr. 2012, Paper 8627.

[7] P. Coleman, P. J. B. Jackson, M. Olik, and J. Abildgaard Pedersen, "Personal audio with a planar bright zone," *J. Acoust. Soc. Amer.*, vol. 136, no. 4, pp. 1725–1735, Oct. 2014.

[8] T. Betlehem and P. D. Teal, "A constrained optimization approach for multi-zone surround sound," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2011, pp. 437–440.

[9] Y. Cai, M. Wu, and J. Yang, "Sound reproduction in personal audio systems using the least-squares approach with acoustic contrast control constraint," *J. Acoust. Soc. Amer.*, vol. 135, no. 2, pp. 734–741, Feb. 2014.

[10] M. R. Bai, J.-C. Wen, H. Hsu, and Y.-H. Hsieh, "Investigation on the reproduction performance versus acoustic contrast control in sound field synthesis," *J. Acoust. Soc. Amer.*, vol. 136, no. 4, pp. 1591–1600, Oct. 2014.

[11] J.-W. Choi, Y. Kim, S. Ko, and J.-H. Kim, "Super-directive loudspeaker array for the generation of personal sound zone," in *Proc. 125th Audio Eng. Soc. Conv.*, Oct. 2008, Paper 7620.

[13]Provided that a solution $\boldsymbol{\beta}(q^{-1})$ to the spectral factorization (8) exists, $\boldsymbol{Q}(q^{-1})$ will be unique up to a left orthogonal matrix factor $\boldsymbol{F}$, such that $\boldsymbol{F}^T\boldsymbol{F}=\mathbf{I}$, which cancels in the filter expression (7). The Diophantine equation (10), $\boldsymbol{\beta}_*(z)\boldsymbol{Q}(z^{-1})-\boldsymbol{B}_{B*}(z)\boldsymbol{V}_*(z)\boldsymbol{V}(z^{-1})\boldsymbol{D}(z^{-1})=z\mathbf{I}\boldsymbol{L}_*(z)$ is always solvable, since $\boldsymbol{\beta}^{-1}(q^{-1})$ is stable, implying that $\det\left(\boldsymbol{\beta}(z^{-1})\right)$ has all zeros in $|z|<1$. Therefore, $\det\left(\boldsymbol{\beta}_*(z)\right)$ has all zeros in $|z|>1$, while $z\mathbf{I}=0$ only for $z=0$, so $\boldsymbol{\beta}_*(z)$ and $z\mathbf{I}$ have no common factors. It therefore has a solution [35], that is unique, see [36].

[12] S. J. Elliott and J. Cheer, "Regularization and robustness of personal audio systems," 2011. [Online]. Available: http://eprints.soton.ac.uk/207989/1/Pub12685.pdf

[13] Y. Cai, M. Wu, and J. Yang, "Design of a time-domain acoustic contrast control for broadband input signals in personal audio systems," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, May 2013, pp. 341–345.

[14] Y. Cai, L. Liu, M. Wu, and J. Yang, "Robust time-domain acoustic contrast control design under uncertainties in the frequency response of the loudspeakers," in *Proc. INTER-NOISE NOISE-CON Congr. Conf. Proc.*, 2014, vol. 249, pp. 5775–5780.

[15] Y. Cai, M. Wu, L. Liu, and J. Yang, "Time-domain acoustic contrast control design with response differential constraint in personal audio systems," *J. Acoust. Soc. Amer.*, vol. 135, no. 6, pp. EL252–EL257, 2014.

[16] D. H. Schellekens, M. B. Møller, and M. Olsen, "Time domain acoustic contrast control implementation of sound zones for low-frequency input signals," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 365–369.

[17] S. Berthilsson, A. Barkefors, L.-J. Brännmark, and M. Sternad, "Acoustical zone reproduction for car interiors using a MIMO MSE framework," in *Proc. AES 48th Int. Conf. Autom. Audio*, Sep. 2012, pp. 1–9.

[18] M. F. Simón Gálvez, S. J. Elliott, and J. Cheer, "Time domain optimization of filters used in a loudspeaker array for personal audio," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 23, no. 11, pp. 1869–1878, Nov. 2015.

[19] M. B. Møller and M. Olsen, "Sound zones: On envelope shaping of FIR filters," in *Proc. 24th Int. Congr. Sound Vib.*, 2017, pp. 1–8.

[20] S. Zhao, X. Qiu, and I. Burnett, "Acoustic contrast control in an arc-shaped area using a linear loudspeaker array," *J. Acoust. Soc. Amer.*, vol. 137, no. 2, pp. 1036–1039, 2015.

[21] S. Zhao, X. Qiu, J. Lu, and J. Cheng, "Robust time domain acoustic contrast control in an arc-shaped area using a linear loudspeaker array," in *Proc. Int. Congr. Sound Vib.*, 2015, pp. 1–8.

[22] S. Widmark, "Causal IIR audio precompensator filters subject to quadratic constraints," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 26, no. 10, pp. 1897–1912, Oct. 2018.

[23] T. Kailath, *Linear Systems*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1980.

[24] L.-J. Brännmark, "Robust sound field control for audio reproduction: A polynomial approach to discrete-time acoustic modeling and filter design," Ph.D. dissertation, Dept. Eng. Sci., Uppsala Univ., Uppsala, Sweden, 2011.

[25] Y. Haneda, S. Makino, and Y. Kaneda, "Common acoustical pole and zero modeling of room transfer functions," *IEEE Trans. Speech Audio Process.*, vol. 2, no. 2, pp. 320–328, Apr. 1994.

[26] Y. Haneda, Y. Kaneda, and N. Kitawaki, "Common-acoustical-pole and residue model and its application to spatial interpolation and extrapolation of a room transfer function," *IEEE Trans. Speech Audio Process.*, vol. 7, no. 6, pp. 709–717, Nov. 1999.

[27] K. Öhrn, A. Ahlén, and M. Sternad, "A probabilistic approach to multivariable robust filtering and open-loop control," *IEEE Trans. Autom. Control*, vol. 40, no. 3, pp. 405–418, Mar. 1995.

[28] L.-J. Brännmark, "Robust audio precompensation with probabilistic modeling of transfer function variability," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust.*, New Paltz, NY, USA, Oct. 2009, pp. 193–196.

[29] M. Grant and S. P. Boyd, CVX: Matlab software for disciplined convex programming. 2014. [Online]. Available: http://cvxr.com/cvx/

[30] O. Kirkeby, "Reproduction of acoustic fields," Ph.D. dissertation, Inst. Sound Vib. Res., Univ. Southampton, Southampton, U.K., 1995.

[31] H. Fastl and E. Zwicker, *Psychoacoustics*, 3rd ed. Berlin, Germany: Springer, 2007.

[32] D. H. Raab, "Forward and backward masking between acoustic clicks," *J. Acoust. Soc. Amer.*, vol. 33, no. 2, pp. 137–139, 1961.

[33] A. Barkefors, S. Berthilsson, and M. Sternad, "Extending the area silenced by active noise control using multiple loudspeakers," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2012, pp. 325–328.

[34] M. Sternad and A. Ahlén, "A novel derivation methodology for polynomial-LQ controller design," *IEEE Trans. Autom. Control*, vol. 38, no. 1, pp. 116–121, Jan. 1993.

[35] V. Kučera, *Discrete Linear Control: The Polynomial Equation Approach*. Chichester, U.K.: Wiley, 1979.

[36] A. Ahlén and M. Sternad, "Wiener filter design using polynomial equations," *IEEE Trans. Signal Process.*, vol. 39, no. 11, pp. 2387–2399, Nov. 1991.

Simon Widmark, photograph and biography not present at the time of publication.