Audio Engineering Society
Convention Paper 5378
Presented at the 110th Convention
2001 May 12–15 Amsterdam, The Netherlands

Kautz filters and generalized frequency resolution: Theory and audio applications

Tuomas Paatero and Matti Karjalainen

Helsinki University of Technology
Laboratory of Acoustics and Audio Signal Processing

Espoo, P. O. Box 3000, FIN-02015 HUT, Finland

ABSTRACT

Frequency-warped filters have recently been applied successfully to a number of audio applications. The idea of allpass delay elements replacing unit delays in digital filters allows for focusing of enhanced frequency resolution on lowest (or highest) frequencies and enable good match to the psychoacoustic Bark scale. Kautz filters can be seen as a further generalization where each transversal element may be different, including complex conjugate poles. This enables arbitrary allocation of frequency resolution for filter design, such as modeling and equalization (inverse modeling) of linear systems. In this paper we formulate strategies for using Kautz filters in audio applications. Case studies of loudspeaker equalization, room response modeling, and guitar body modeling for sound synthesis are presented.

INTRODUCTION

In the signal processing context, *rational orthonormal filter structures* were first introduced in the 1950's by Kautz, Huggins and Young [18, 11, 43]. Kautz showed that an orthogonalization process applied to a set of continuous-time exponential components produces orthonormal basis functions having particular frequency-domain expressions. Much earlier, Wiener and Lee [21] proposed synthesis network structures based on some classical orthonormal polynomial expansions [36]. The idea of representing functions in orthonormal components is elementary in Fourier-analysis, but the essential observation in the aforementioned was that some time-domain basis functions have rational Laplace transforms with a recurrent structure, defining an efficient transversal synthesis filter.

Discrete-time rational orthonormal filter structures can be attributed to Broome [2] as well as the baptizing of the *discrete Kautz functions*, consequently defining the discrete-time *Kautz filter*. The point of reference in the mathematical literature is somewhat arbitrary, but a reasonable choice are the deductions made in the 1920's to prove interconnections between rational approximations and interpolations, and the least-square problem, which were assembled and further developed by Walsh [41].

Over the last ten years there has been a renewed interest towards rational orthonormal model structures, mainly from the system identification point of view [9, 38, 42, 26]. The perspective has usually been to form generalizations to the well-established *Laguerre models* in system identification [19, 25, 37] and control [44]. In this context, the *Kautz filter or model* has often the meaning of a two-pole generalization of the Laguerre structure [38], whereupon further generalizations restrict, as well, to structures with identical blocks [10]. Another, and almost unrecognized, connection to recently active topics in signal processing are the orthonormal state-space models for adaptive IIR filtering [29] with some existing implications to Kautz filters [4, 32].

Kautz filters have found very little use in audio signal processing, to cite a rarity [24]. One of the reasons is certainly that the field is dominated by the system identification perspective. The more inherent reason is that there is an independent but related tradition of *frequency warped* structures, which is well-grounded and sufficient for many tasks in audio signal processing. Frequency warping provides a (rough) approximation of the constant Q resolution of modeling [27] as well as a good match with the Bark scale that is used to describe the psychoacoustical frequency scale of human hearing [33].

In our opinion there is a certain void in generality, both in the proposed utilizations of Kautz filters and in the warping-based view to the resolution of modeling. It is a safe guess that Kautz filters are going to find audio applications in acoustic echo cancellation and in adaptive modeling of other audio systems. Here we demonstrate the use of Kautz filters in pure filter synthesis, i.e., modeling of a given target response, which was actually their original usage. Quite surprisingly, to our knowledge there hasn't really been any proposals in this direction on the level of modern computational means in design and implementation.

The objective of this paper is to introduce the concept of Kautz filters in a practical manner and from the point of view of audio signal processing. After the obligatory theoretical part we present various methods for the most essential ingredient in Kautz filter design, the choosing of a particular structure. Audio applications of Kautz filtering, including loudspeaker equalization, room response modeling, and modeling of an acoustic guitar body, are demonstrated and compared with more traditional approaches.

THE KAUTZ FILTER STRUCTURE

For a given set of desired poles $\{z_i\}$ in the unit disk, the corresponding set of rational orthonormal functions is uniquely defined in the sense that the lowest order rational functions, square-integrable and orthonormal on the unit circle, analytic for $|z| > 1$, are of the form [41]

$$G_i(z) = \frac{\sqrt{1 - z_i z_i^*}}{z^{-1} - z_i^*} \prod_{j=0}^{i} \frac{z^{-1} - z_j^*}{1 - z_j z^{-1}}, \quad i = 0, 1, \ldots. \quad (1)$$

The meaning of orthonormality is most economically established in the time-domain: for the impulse responses of (1), $(g_i, g_k) = \sum_{n=0}^{\infty} g_i^*(n) g_k(n) = 0$ for $i \neq k$, and $(g_i, g_i) = 1$. For the remaining conditions it is sufficient to presume stability and causality of the rational transfer function.

Clearly functions (1) form a recurrent structure: up to a given order, i.e., the number of poles, functions associated to the subsets $\{z_j\}_{j=0}^{i}$ of an ordered pole set $\{z_j\}_{j=0}^{N}$ are produced as intermediate substructures defining a tapped transversal system. In agreement with the continuous-time counterpart, a weighted sum of these functions is called a Kautz filter, depicted in Fig. 1.

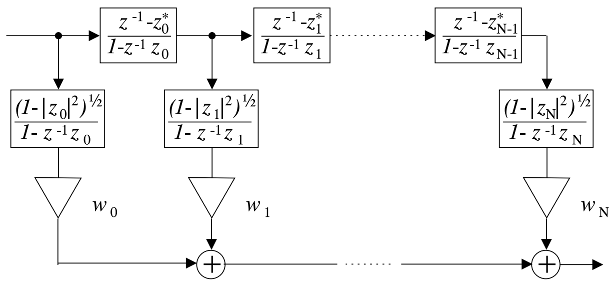

Fig. 1: The Kautz filter. For $z_i = 0$ in (1) it degenerates to an FIR filter and for $z_i = a$, $-1 < a < 1$, it is a Laguerre filter where the tap filters are replaced by a common pre-filter.

Defined in this manner, Kautz filters are merely a class of fixed-pole IIR filters, forced to produce orthonormal tap-output impulse responses. A particular Kautz filter is thus determined by a pole set, consequently defining the filter order, the all-pass filter backbone as well as the corresponding tap-output filters, and a set of assigned filter coefficients $\{w_i\}_{i=0}^{N}$. Kautz filters are clearly related to other orthogonal filter formulations, but here we just exploit this connection by referring to favorable numerical properties implied by the orthonormality [8, 23, 6, 40].

The Kautz model for signals and systems

Defined by any set of points $\{z_i\}_{i=0}^{\infty}$ in the unit disk, functions (1) form an orthonormal set which is complete, or a base, with a moderate restriction on the poles $\{z_i\}$ [41]. The corresponding time domain basis functions $\{g_i(n)\}_{i=0}^{\infty}$ are impulse responses or inverse z-transforms of functions (1). This implies that a basis representation of any causal and finite-energy discrete-time signal is obtained as its Fourier series expansion with respect to the time or frequency domain basis functions. As impulse responses of causal and stable (CS) linear time-invariant (LTI) systems belong to the aforementioned signal space, this is also a valid model structure for input-output-data identification.

Kautz filters provide linear-in-parameter models for many types of system identification and approximation schemes, including adaptive filtering, both for fixed and non-fixed pole structures. Also in the Kautz model case, there are various interpretations, criteria, and methods for the parametrization of the model. Here we address only the "prototype" least-square (LS) approach implied by the signal space description: tap-output signals of a Kautz filter $x_i(n) = G_i[x(n)]$ to input $x(n)$ span an finite dimensional approximation space, providing a LS optimal approximation to any CSLTI system with respect to the basis. The parametrization, i.e., the filter weights, are solutions of the *normal equations* assembled from correlation terms of the tap-outputs and the desired output $y(n)$. For example in matrix form: defining the *correlation matrix* $\mathbf{R}$, $[r]_{ij} = (x_i, x_j)$, the weight vector $\mathbf{w}$ is the solution of the matrix equation $\mathbf{Rw} = \mathbf{p}$, where $\mathbf{p}$, $p_i = (y, x_i)$, is the *correlation vector*. This simply utilizes the previously defined time-domain inner products of the signals, so there are also definitions for the frequency domain correlation terms, and separately for deterministic and stochastic signal descriptions, but this notion is just included to indicate similarities to the FIR model case.

Actually we further limit our attention to the approximation of a given target response since the input-output system identification framework is a bit pompous and impractical for most audio signal processing tasks. Although typical operations such as inverse modeling and equalization are basically identification schemes, they are usually convertible to the approximation problem. Furthermore we are not going to address here the interesting question of the invertibility of a Kautz filter, which would require elimination of delay-free loops, or an implementation method proposed in[13].

For a given system $h(n)$ or $H(z)$, Fourier coefficients provide least-square (LS) optimal parametrizations for the corresponding Kautz model, or synthesis filter, with respect to the pole set. Evaluation of the Fourier coefficients, $c_i = (h, g_i) = (H, G_i)$, can be implemented by feeding the signal $h(-n)$ to the Kautz filter and reading the tap outputs $x_i(n) = G_i[h(-n)]$ at $n = 0$: $c_i = x_i(0)$. This implements convolutions by filtering and it can be seen as a generalization of the rectangular window FIR design. It should be noted that here the LS criterion is applied on the infinite time horizon and not for example in the time window defined by $h(n)$. We use these true orthonormal expansion coefficients because they are easy to obtain, providing implicitly simultaneous time and frequency domain design and powerful means to the Kautz filter structure (i.e., pole position) optimization. Moreover, the coefficients are independent of ordering and approximation order, which makes choosing of poles, approximation error evaluation, and model reduction efficient.

Real Kautz functions

A Kautz filter produces real tap output signals only in the case of real poles. In principle this doesn't in any way limit its potential capabilities of approximating a real signal or system. However, we may want the processing, i.e. internal signals, coefficients and arithmetic operations, to be real valued. A restriction to real linear-in-parameter models can also be seen as a categoric step in the optimization of the structure.

From a sequence of real or complex conjugate poles it is always possible to form real orthonormal structures. Symbolically this is done by applying a block diagonal unitary transformation to the outputs, consisting of ones corresponding to real poles and 2×2 rotation elements corresponding to pole pairs. From the infinite variety of unitarily equivalent possible solutions it is sufficient to use the intuitively simple structure of Fig. 2, proposed by Broome: the second-order section outputs of Fig. 2 are *orthogonal* from which an orthogonal tap output pair if produced in the form $\{x(n) - x(n-1), x(n) + x(n-1)\}$. Normalization terms are completely determined by the corresponding pole pair $\{z_i, z_i^*\}$ and are given by $p_i = \sqrt{(1-\rho_i)(1+\rho_i-\gamma_i)/2}$ and $q_i = \sqrt{(1-\rho_i)(1+\rho_i+\gamma_i)/2}$, where $\gamma_i = -2RE\{z_i\}$ and $\rho_i = |z_i|^2$ can be recognized as corresponding second-order polynomial coefficients. The construction works also for real poles, producing a tap-output pair corresponding to a real double pole, but we use an obvious mixture of first- and second-order sections, if needed.

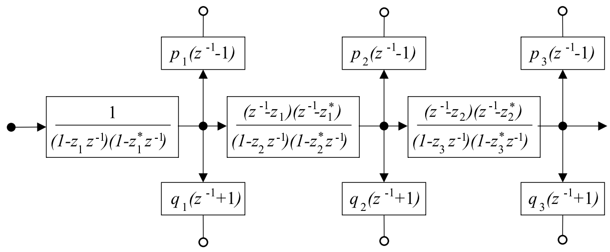

Fig. 2: One realization for producing real Kautz filter tap-output responses defined by a sequence of complex conjugate pole pairs. The transversal all-pass backbone of Fig. 1 is restored by moving the denominator terms one step to the right and compensating for the change in the tap-output filters.

Some of the Kautz filter deductions are made directly on the real function assumption [2]. Moreover, the state-space approach to orthonormal structures with identical blocks, the *Generalized orthonormal basis functions* of Heuberger, is based on balanced realizations of real rational all-pass functions [9].

METHODS FOR THE CHOOSING OF POLES

Kautz filter design can be seen as a two-step procedure involving the choosing of a particular Kautz filter (i.e., the pole set) and the evaluation of the corresponding filter weights. The fact that the latter task is much easier, better defined and inherently (LS) optimal makes it tempting to use sophisticated guesses and random or iterative search in the pole optimization. For the more analytic approach, the whole idea in the Kautz concept is how to incorporate desired *a priori* information to the Kautz filter. This may mean knowledge on system poles or resonant frequencies and corresponding time-constants, or indirect means, such as all-pole or pole-zero modeling.

A practical way to limit the degrees of "freedom" is to restrict to structures with identical blocks, i.e., to use the same smaller set of poles repeatedly. The pole optimization and the model order selection problems are then essentially separated and we may apply the various optimization methods to the substructure. Additionally, for the structure of identical blocks, a relation between optimal model parameters and error energy surface stationary points with respect to the poles may be utilized [5] as well as a classification of systems to associate systems and basis functions [39]. However, in the following we mainly focus on practical methods for choosing a distributed pole set.

Generalized frequency resolution descriptions

As mentioned earlier, frequency-warped configurations in audio signal processing [14] constitute a self-contained tradition originating from *warping effects* observed in analog-to-digital mappings and digital filter transformations [3]. The concept of a warped signal was introduced to compute non-uniform-resolution Fourier transforms using FFT [27] and in a slightly different form to compute warped autocorrelation terms for warped linear prediction [35]. The original idea of replacing a unit delay element with a first-order all-pass operator in a transfer function was restated and applied to general linear filter structures [34, 15, 12].

The warping effect or resolution description introduced by first-order all-pass warping is defined by the all-pass phase function, and the warping parameter (pole location) can be chosen to approximate desired frequency-scale mappings with respect to different error criteria and sampling rates [33]. This is complete in the sense that the first-order all-pass element is the only (rational, stable and causal) filter having a one-to-one phase function mapping. Parallel structures can be constructed to perform any kind of warping [20], but the aim of this chapter is to broaden the concept of frequency resolution to account for the resolution allocation introduced by a Kautz filter.

Warped filters are the optimal choice, if we are happy with the bilinear mapping, and the desired resolution is the only content-dependent thing we are interested in. This has a nice analogy, FIR filters are the optimal choice, if (exponential) stability is the only thing we know about the system. The opposite of the latter is actually used in motivation of the Kautz model [39]. Thus the sometimes stated argument that warped structures are comprehensive in producing a desired modeling resolution allocation is valid, in the sense described above, if we think of filter transformations in general, but for a specific filter design problem this is certainly not the case. To continue with the analogy, our hypothesis is that if the target response has clearly distinguishable features, it is always possible to design a Kautz filter with at least the same (e.g perceptually evaluated) quality, and with better efficiency.

The Laguerre filter differs structurally from the warped FIR filter only in the all-pole pre-filter. The effect of this distinction can be seen as a mere technicality, which can be compensated if needed, or even as an favorable emphasis in the modeling [35]. Actually we have already referred to three variations of the warping, the two above and the Oppenheim *et al.* construction [27], which is a *bi-orthogonal* analysis-synthesis structure, and this distinction has also an effect on the true frequency resolution mapping. In practice this may not be of importance, but signal and system transformations defined by the Laguerre analysis-synthesis-structure are the only solution where there is a *qualitative* view to, e.g., truncation effects.

The orthonormal *Laguerre transformation* of signals and systems has a generalization to orthonormal (Kautz) structures with identical blocks defined by any all-pass transformation [30, 42], so actually it is just a question of how to interpret the well-defined phase-characteristics as a frequency resolution mapping. Moreover, the same applies to the general Kautz model: the phase function of the transversal all-pass backbone is completely determined by the pole set (as a whole) and we should "just" find a way to decode this information as a frequency resolution allocation.

A trivial way to attain at some level a desired frequency resolution allocation to the Kautz modeling is to use suitable pole distributions. We may place (complex conjugate) poles with pole angle spacings corresponding to any frequency resolution mapping. There is also the choosing of pole radii to be taken to account, in which we can use more sophisticated choices than a constant radius. Motivated from our experiments on producing warping on a logarithmical scale with parallel all-pass (and generalized parallel orthonormal) structures we have used for example pole radii inversely proportional to the pole angles.

Manual fitting to a given response

In this case of approximating a given target response, it is always possible to simply adjust the Kautz filter poles through trial and error to produce a matching Kautz filter response. A Kautz filter impulse response is a weighted superposition of damped sinusoidals, which provide for direct tuning of a set of resonant frequencies and corresponding time constants. In principle this allows for very flexible time-domain design, especially if we use some other weighting than the chosen LS parametrization.

By direct inspection of the time and frequency responses, it is relatively easy to find usable pole sets by selecting a set of prominent resonances and proper pole radius tuning. The choosing of the (complex conjugate) pole angles is more critical than the pole radius selection in the sense that the filter coefficients perform automatic weighting of the sinusoidal components. In practice however, when we design low-order models for highly resonant systems we fine-tune poles very close to the unit circle.

The obvious way to improve the overall modeling with a structure based on a fixed set of resonances is to use the corresponding generating substructure repetitively, producing a Kautz filter with identical blocks. There is no obligation to use the same multiplicity for all poles, but it makes model reduction easier. If the poles are assigned for the substructure, some kind of damping of the pole radii would be advisable. On the other hand the resonance tuning scheme can be made to account for the multiplicity of the poles.

Manual tuning of a large (and possibly in number unspecified) set of poles is obviously a complicated and connected task. One should not however be discouraged by comparing it to resonator or pole-zero model designs since Kautz filter design provides clearly a more robust view to the overall modeling problem. For example in modeling a highly resonant system we may very easily construct Kautz filters with conservative pole choices that are competitive with the FIR model in any measure of complexity and performance.

Methods implied by the orthogonality

We have chosen our model to a given target response $h(n)$ to be the truncated Fourier series expansion,

$$\hat{h}(n) = \sum_{i=0}^{N} c_i g_i(n), \quad c_i = (h, g_i), \tag{2}$$

which makes approximation error evaluation and control easy: for any (orthonormal) $\{g_i\}$ and approximation order N, the approximation error energy satisfies

$$E = \sum_{i=N+1}^{\infty} |c_i|^2 = H - \sum_{i=0}^{N} |c_i|^2, \tag{3}$$

where $H = (h, h)$ is the energy of $h(n)$. Hence we get the energy of an infinite duration error signal as a by-product from finite filtering operations, previously described for the evaluation of the filter weights.

As a more profound consequence of the orthogonality, the *all-pass operator* defined by the (transversal part of the) chosen

Kautz filter induces an *complementary* division of the energy of the signal $h(-n)$, $n = 0, \ldots, M$, [43]. An all-pass filter $A(z)$ is *lossless* by definition and (from the above) it can be deduced that the portion of the response $a(n) = A[h(-n)]$ in the time interval $[-M, 0]$ corresponds to the approximation error E, i.e., the Kautz filter optimization problem reduces to minimization of a finite duration signal.

We have encountered three attempts to utilize the concept of complementary signals in the pole position optimization, at least implicitly [22, 7, 1]. McDonough and Huggins replace the all-pass numerator with a polynomial approximating the denominator mirror polynomial to produce linear equations for the polynomial coefficients in an iteration scheme [22]. Friedman has constructed a network structure for parallel calculations of all partial derivatives of the approximation error with respect to the real second-order polynomial coefficient associated to the complex pole pairs, to be used in a gradient algorithm [7]. Brandenstein and Unbehauen have proposed an iterative method resembling the Steiglitz-McBride method of pole-zero modeling to pure FIR-to-IIR filter conversion [1], which we have adopted to the context of Kautz filter optimization. We have implemented and experimented with all the aforementioned algorithms, but in general just the method of Brandenstein and Unbehauen (BU-method) is found reliable: it genuinely optimizes (in the LS sense) the pole positions of a real Kautz filter, producing unconditionally stable and (theoretically globally) optimal pole sets for a desired filter order. Furthermore, the BU-method works on really high filter orders, e.g., providing sets of 300 distributed and accurate poles, which would be unachievable with standard pole-zero modeling methods.

Hybrids

By trying to categorize various methods for the choosing of a particular Kautz filter we do not in any way intend to be complete or exclusive. There are certainly many other possibilities as well as modifications and mixtures of the presented ones. For example, we have not addressed pole position optimization methods based on input-output-descriptions of the system, which could in some cases be useful even if we have the target response available.

An obvious way to modify the optimization is to manipulate the target response. We may apply time-domain windowing or frequency-domain weighting by suitable filtering, or alternatively divide the optimization in pieces with selective filters. For example the BU-method is based on solving at every iteration a matrix equation with dimensions implied by the duration of the target response and the chosen model order, and therefore pathologicalities can be decreased by truncation. By dividing the frequency range in two (or more) parts we may apply separate methods and allocation of the modeling resolution on the subbands.

In the case where a substructure is assigned and used in a Kautz filter with identical blocks there is always the question of compensating for the repetitive appearance of the poles. We have used constant pole radius damping, individual tuning of the pole radii, and an *ad hoc* method where we simply raise the pole radius to the power of the number of blocks used in the Kautz filter. For poles with dissimilar radii this approach is certainly better-justified than using a constant damping.

The most efficient strategy, in our experience, is to combine various methods with the BU-method. At the least this may mean examining the produced pole set and possibly omitting some of the poles. For a resonant system we get poles really close to the unit circle, which is a cause for checking as such, but there may also occur weak poles that have very little contribution to the model. Furthermore, we don't usually care for real poles because they are by definition either weak or products of the oddities in the target response at frequency band edges. Additionally, sometimes a cluster of poles may be represented with a single conjugate pole pair.

We may also perform a more systematic pruning of the BU-pole set, e.g., by simply omitting the poles outside a specific frequency (pole angle) region, or by spacing out the pole set according to some rule. The fact that a high-order distributed BU-pole set is a good representation of the whole resonance structure can be used to associate poles directly with different choices of prominent resonances. Typically a sharp resonance is represented by a single pole pair, with pole angles corresponding (practically) exactly to the peak frequency, and appropriately tuned pole radii. Because of the latter, this has also proven to be a useful approach in selecting a substructure for the Kautz filter with identical blocks. As an inverse to the pruning we may also add and manually tune poles to the BU-pole set. This is demonstrated in the loudspeaker case study.

We have also developed and applied successfully a combination of the warping and the BU-method. By using the warped target response in the BU-method we may optimize the pole positions corresponding to Kautz modeling on a warped frequency scale. We get the poles in the original domain by simply mapping them according to the inverse all-pass transformation. This is at the present maybe the best way to incorporate perceptually justified allocation of the modeling resolution to the Kautz filter, but warping can also be utilized to focus the resolution in a more technical fashion. It should also be noted that our interpretation of warping has a definite effect on the poles we get. However, the possible "tilt" in the Kautz model magnitude response is due to the choice of poles and it is not a property of the model, in contrast to some warped designs. A peg leg version of the proposed method would be to down-sample the response prior to the BU-method and to map the produced poles back to the original frequency domain.

APPLICATION CASES

We have tested the applicability of Kautz filter design in three audio-oriented applications. The first one is loudspeaker equalization task where frequency resolution is distributed both globally and locally. The second case, room impulse response, is modeled at low frequencies to capture the modal behavior as a robust filter structure. In the third case we use Kautz filters to model the body response of the acoustic guitar, where also, the lowest frequencies are of primary interest.

Case 1: Loudspeaker equalization

An ideal loudspeaker has a flat magnitude response and a constant group delay. Simultaneous magnitude and phase equalization would be achieved by modeling the response and inverting the model, or by identifying the overall system of the response and the Kautz equalizer, but here we demonstrate the use of Kautz filters in pure magnitude equalization, based on an inverted target response.[1] The measured loudspeaker magnitude response and a derived equalizer target response are included in Fig. 3. The sample rate is 48 kHz.

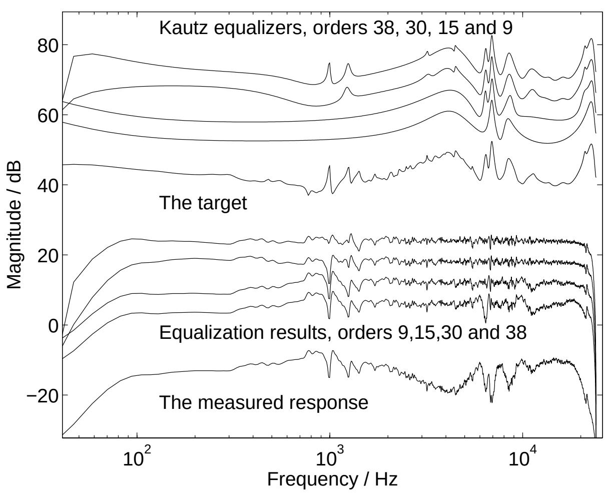

Fig. 3: Kautz equalizers and equalization results for orders 9, 15, 30 and 38, compared to the measured loudspeaker response and the equalizer target response.

Magnitude response equalization consists typically of compensating for three different types of phenomena: slow trends in the response, sharp and local deviations, and correction of roll-offs at the band edges. This makes "blind equalization" methods ineffective. We propose that Kautz filters provide a useful alternative between "blind" and hand-tuned "parametric" equalization.

As is well known, FIR modeling has an inherent emphasis on high frequencies on the auditorily motivated logarithmic frequency scale. Warped FIR (or Laguerre) [14] filters release some of the resolution to the lower frequencies, providing a competitive performance with 5 to 10 times lower filter orders than with FIR filters [16]. However, the filter order required to flatten the peaks at 1 kHz in our example is still high, of the order 200, and in practice Laguerre models up to order 50 are able to model only slow trends in the response. In a recent publication [31], general (real) all-pass cascades were proposed for low-frequency equalization, but they were found too difficult to design. Here we demonstrate efficient design methods for the orthonormal and (possibly) complex counterpart.

The simplest choice of Kautz filter poles in this case is to focus on the 1 kHz region with a single (multiple) pole pair. By tuning the pole radius we trade-off between the 1 kHz region and overall modeling. Quite interestingly, as a good compromise, we end up with a radius close to a typical warping parameter at this sampling rate (e.g., $\lambda = 0.76$) [33] and we get surprisingly similar results for the Laguerre and two-pole Kautz equalizers for filter orders 50–200. Actually this simply means that perceptually motivated warping is also technically a good choice for the flattening of the 1 kHz region.

The obvious way to proceed would be to add another pole pair corresponding to the 7 kHz region. In search for considerably lower Kautz filter orders, compared to the Laguerre equalizer, we however utilize directly the BU-method. It provides stable and reasonable pole sets for orders at least up to 40. In Fig. 3 we have presented Kautz equalizers and equalization results for orders 9, 15, 30 and 38. These straightforward Kautz filter constructions are very compatible with the FIR and Laguerre counterparts, but we may further lower the filter orders by omitting some of the poles. For example for orders above 15, the BU-method produces poles really close to z=1 (because of the low frequency boost in the target) and omitting some of these poles actually tranquilize the low frequency region. We obtain, for example, quite similar equalization results for orders 28 and 34, from the sets with 30 and 38 poles, respectively (Fig. 4).

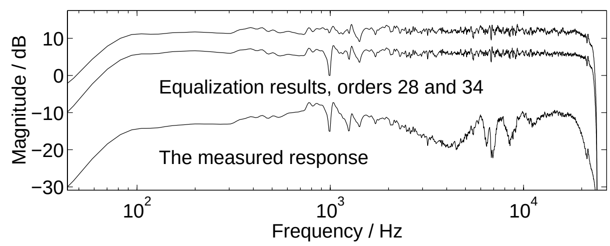

Fig. 4: Kautz equalization results for orders 28 and 34, with pruned BU-poles.

To improve the modeling at 1 kHz, we added three to four manually tuned pole pairs to the BU-pole sets, corresponding to the resonances in the problematic area. This is actually not too hard since the 1 kHz region is isolated from the dominant pole region, which allows for undisturbed tuning. As starting points we used the 15th and 30th order Kautz equalizers of Fig. 3, omitting three pole pairs in the latter case. Three pole pairs were tuned directly to the three prominent resonances

[1] This kind of detailed equalization of a loudspeaker is not necessarily practical, except maybe for avoiding acoustic feedback. Here our aim was primely to use this specific loudspeaker as a test case of a demanding equalization task.

and one pole pair was assigned to improve the modeling below the 1 kHz region. Results for final filter orders 23, 32 and 34 are displayed in Fig. 5, where the last two differ only in the optional compensating pole pair.

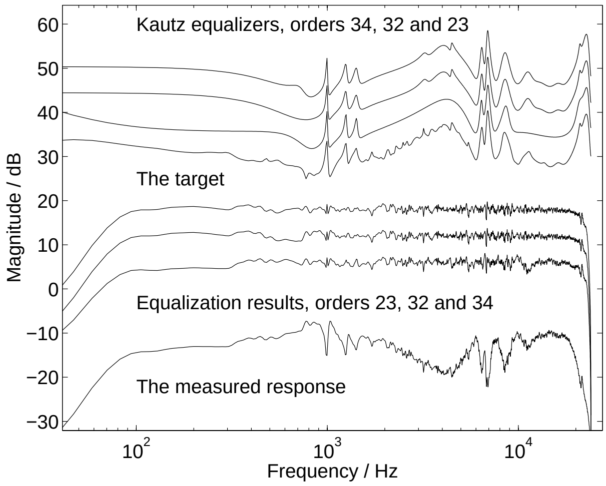

Fig. 5: Kautz equalizers and equalization results for orders 23, 32 and 34, with combinations of manually tuned an BU-generated poles.

Finally, we abandon the pole sets proposed by the BU-method and try to tune 10 pole pairs manually to the target response resonances. The design is based on 10 selected resonances, represented with 10 distinct pole pairs, chosen and tuned to fit the magnitude response. This is of course somewhat arbitrary, but it seems to work! In Fig. 6 along with the equalizer and target responses are vertical lines indicating pole pair positions. This is clearly one form of "parametric equalization" with second order blocks since each resonance is represented with a single pole pair. However, with Kautz filters we have, at least to some extent, separated the choice of the resonance structure and the fine-tuning produced by the (linear-in-parameter) model parametrization.

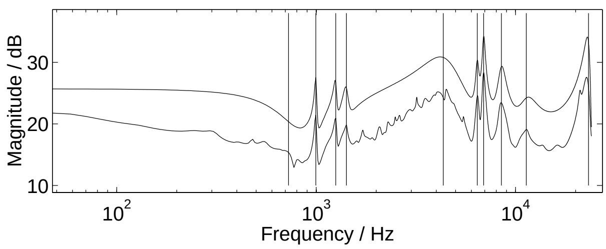

Fig. 6: Manually tuned 20th order Kautz equalizer and target responses, with lines indicating pole pair positions.

Figure 7 compares some of the Kautz equalization results to those achieved with FIR and Laguerre equalizers of orders 200 and 100, respectively. If we are only concerned with the amount of arithmetic operations at runtime, we should actually compare Laguerre and Kautz equalizers at about the same filter orders. The actual complexity depends on many things, but we have in any case achieved very low order equalizers. Furthermore, the low filter orders enable in principle filter transformations to other structures, possibly more efficient or otherwise preferable.

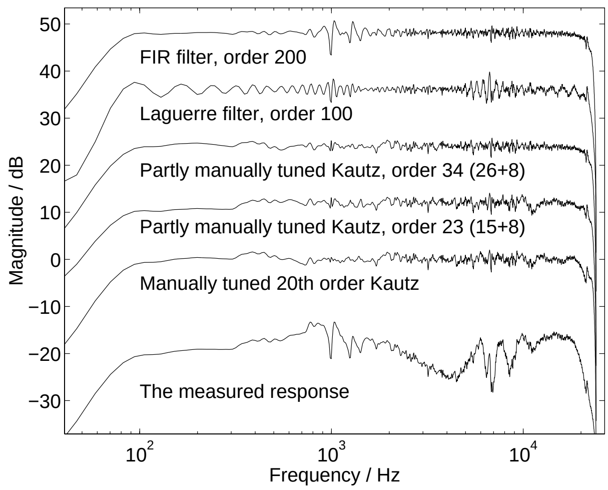

Fig. 7: Comparison of FIR, Laguerre, and Kautz equalization results.

Case 2: Room response modeling

Models for the transfer function from a sound source to an observation location are used for many purposes in audio signal processing, typically as a part of a larger system. These room response models, understood in a broad sense, may however constitute a major computational burden, both because of the complexity of a particular model and the number of distinct models needed to describe the spatial information.

The obvious difficulties in modeling a room transfer function are that the duration of the target response is usually long and that the frequency content is distributed, both in frequency and in time. Physically speaking, there are low-frequency modes determined essentially by the room dimensions and on the other hand a reverberation structure produced by the multitude of reflections. There are many methods proposed for taking into account for the various time and frequency domain modeling aspects, for example for reverb design, usually with complicated parallel and feedback structures, so it is interesting to see how a straightforward transversal linear filter can perform the same task.

We have chosen as an example a measured room (impulse) response (re)sampled at 11025 Hz. The target response for modeling in Fig. 8 is composed from the original signal by omitting the early delay and by truncation to 8192 samples, which yields a duration of 743 ms. Clearly this is not an ideal Kautz modeling task since the (early) response is not a superposition of *coincident* damped exponentials. However, if the Kautz filter is long enough, it should be able to model all the temporal details, though with an potential inefficiency, e.g., in producing pure delays. After all, functions (1) define an exact representation of any (finite-energy) signal, a generalized z-transform, so this is brute, FIR-type design, but with a more delicate choice of basis functions.

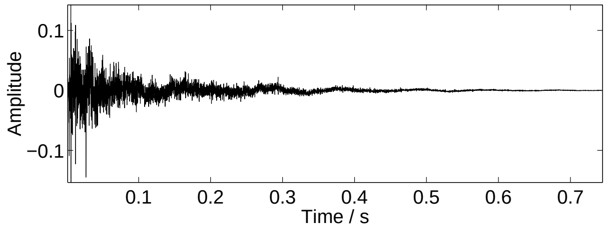

Fig. 8: A measured room impulse response.

In Fig. 9 we have simply used two different pole distributions, logarithmically spaced poles with radus 0.99 and a combination of logarithmically and linearly spaced angles with radius 0.97. The filter orders are 236 and 320, respectively, and the magnitude responses are plotted against the target magnitude response. These are bare and purely methodologically justifiable examples. For the former the fit at low frequencies is complete, but the poor performance in modeling the dense resonance structure at higher frequencies is evident. The latter model works better in the high frequency part, with the expense of the low-frequency modeling.

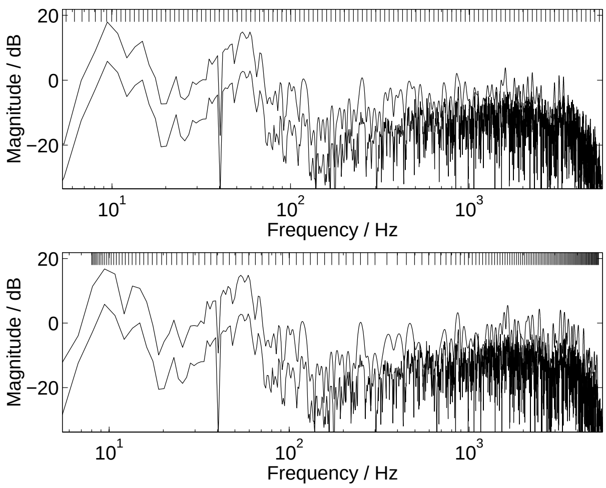

Fig. 9: Two Kautz models plotted with upward offset to the target magnitude response, Kautz model orders 236 (top) and 320 (bottom), with pole pair positions indicated by vertical lines.

There is a certain dilemma in Kautz modeling the high frequency region of the target response. We would like the pole distribution to reflect the relevant resolution in modeling a diffuse sound field, but the Kautz model tries to model individual resonances and a mismatch between the chosen poles and the resonances produces effectively a tilt in the response. One way to overcome this problem would be to use the Kautz filter tap-outputs assigned to the high frequency modeling as a tunable resonator, i.e., replace the LS coefficients with appropriate gains. Another possibility would be to carefully choose representatives with desired spacings from the actual resonance structure. Maybe the most attractive solution would be to directly replace the high frequency part in the target response with a desired resonator structure. However, here we just proceed in modeling, as a bare example, a given target response in the stated blind LS filter synthesis sense.

The proposed iterative pole optimization method, the BU-algorithm, is in great trouble because of the nature and duration of the target response. Even at very high (time-consuming and unreliable) filter orders 200–300, the method is not capable of producing poles for the modeling of lowest frequencies. However, the frequency region above 500 Hz is modelled much better than in the previous examples, i.e., the method genuinely finds many of the prominent resonances. A 300th order Kautz model is presented in top pane of Fig. 10.

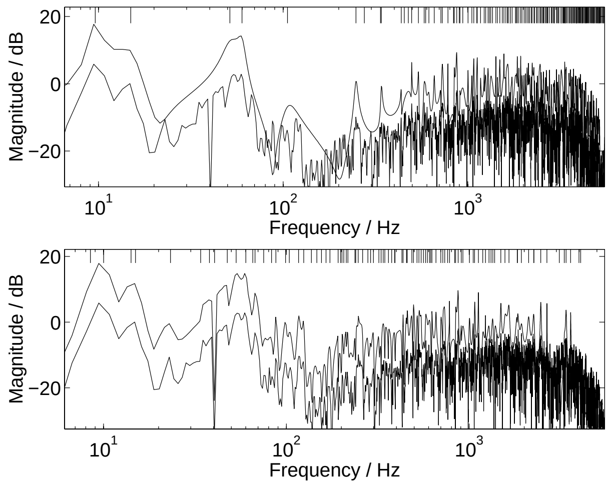

Fig. 10: Kautz models with BU-poles, with filter order 300 (top) and by warping and back-mapping at the filter order 200 (bottom).

Finally we demonstrate the use of warping and downsampling combined with the BU-method. In addition to the perceptual motivation, low-frequency emphasizing warping has an effect of shortening the low-frequency components in the signal, which in this case shortens the effective length of the target response, and makes the BU-method work better. In (bottom pane of) Fig. 10 we have used a warping parameter $\lambda = 0.8$, in producing the warped response and in the back-mapping of the poles, which is much higher than the corresponding "Bark-warping" value $\lambda = 0.48$ [33]. This is however the amount of warping that is needed to capture reasonably well the details at low frequencies at Kautz filter order 200.

In Fig. 11 we have simply composed a 50th order Kautz model for the low frequencies by decimating the target response (by 40) prior to the BU-method and then mapped the poles back to the original design domain according to $z \mapsto z^{0.025}$. We could also apply warping to the decimated response, but e.g. the optimal Bark scale mapping is almost linear at these sampling rates [33]. This and other efficient low order Kautz models for the low-frequency part introduce a steep low-pass cut-off (see Fig. 11). The other possibility is to produce an accurate low-frequency model and an envelope-like high frequency model. These two approaches provide Kautz models for the low-frequency region, to be used for example with an artificial reverberation structure, both for parallel and cascaded constructions.

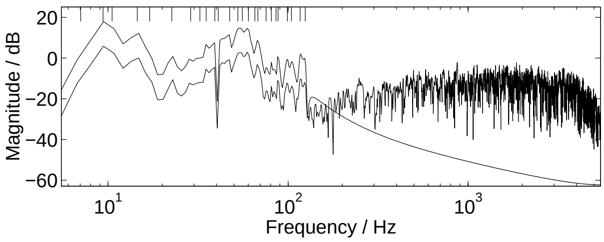

Fig. 11: A 50th order Kautz model for accurate modeling of the lowest frequencies.

As a mare curiosity, we conclude this chapter by attaching a pole distribution (with radii 0.995), corresponding to the *ERB rate scale* [33] spacing, to the low-frequency model of Fig. 11, displayed in Fig. 12.

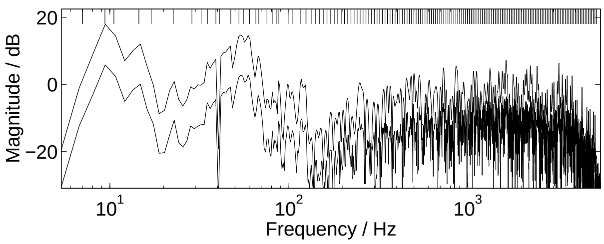

Fig. 12: A 326th order Kautz model produced by adding poles (indicated by vertical lines) to the pole set of the model in Fig. 11.

It should be noted that this was not a detailed case study on room response modeling and that especially the temporal evolution of the model should be evaluated more carefully. However, as hinted previously, it would be interesting to design a Kautz filter using the response of a complicated, highly integrated and perceptually accurate (artificial) room response model as the target response. After all, the building blocks of such systems are usually delay lines, all-pass filters and resonators, which are inherent in a Kautz filter.

Case 3: Guitar body modeling

As another example of high-order distributed-pole Kautz modeling we approximate a measured acoustic guitar body response sampled at 24 kHz (Fig. 13). The impulse response was obtained by hitting (with an impulse hammer) the bridge of an acoustic guitar, with strings removed. This type of measurement produces a response with emphasis on low frequencies and in fact a more accurate impulse response would be achieved by extracting it by deconvolution from an identification setup using spectrally rich real playing of the acoustic guitar as excitation [28]. This more truthful guitar body response is typically like the response of a "small room", but since the density of modes is proportional to the volume, it is actually a much better justified Kautz modeling task than the above room response modeling case. For the sake of diversity we, however, use the response of Fig. 13 as the target response.

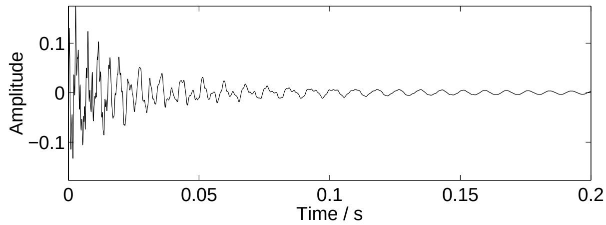

Fig. 13: Measured impulse response of an acoustic guitar body.

The obvious disadvantage of a straightforward FIR filter implementation is that modeling of the slowly decaying lowest resonances requires a very high filter order. All-pole or pole-zero modeling are the traditional choices in improving the flexibility of the spectral representation. However, model orders remain problematically high and the basic design methods seem to work poorly. A significantly better approach is to use separate IIR modeling for the slowly decaying lowest resonances combined with the FIR modeling [28]. Perceptually motivated warped counterparts of all-pole and pole-zero modeling pay off, even in technical terms [17], but here we want to focus the modeling resolution more freely.

Direct all-pole or pole-zero modeling were found to produce unsatisfactory pole sets $\{z_i\}$ for the Kautz filter, even in searching for a low-order substructure. On the other hand, it is relatively easy to find good pole sets by direct selection of prominent resonances and proper pole radius tuning. In Fig. 14 we have simply used five pole pair, corresponding to the resonances indicated by vertical lines, with multiplicity eight, i.e.,the Kautz filter order is 80.

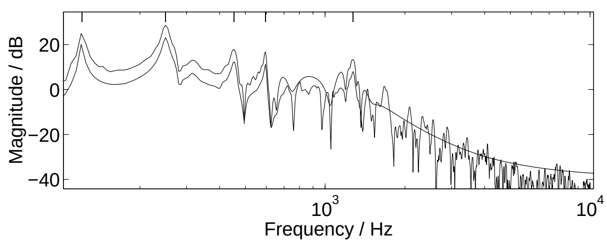

Fig. 14: Displayed with upward offset, magnitude responses of the target and a 80th order Kautz model.

Figure 15 demonstrates that the proposed pole position optimization scheme, the BU-method, is able to capture essentially the whole resonance structure. The Kautz filter order is 102 and the poles are obtained from the 120th order BU-pole set, omitting some poles close to $z = -1$. In general, the BU-method works quite well at least up to order 300 and the lower limit for finding the chosen prominent resonances is about 100. Based on the fact that some of the BU-poles are good representatives for individual resonances can then again be used to form Kautz filters with identical blocks, as in Fig. 14, corresponding to different choices of prominent resonances.

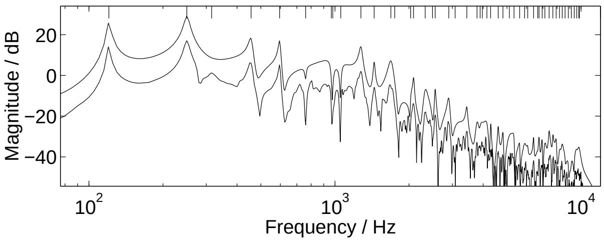

Fig. 15: A 102th order Kautz model and vertical lines indicating poles, obtained by the BU-method.

Especially in this case of a target response dominated by the low-frequency part, it is possible to compose very low order Kautz models with the proposed combination of warping and BU-method. By altering the warping parameter and the desired filter order we may focus the resolution of the modeling between the prominent low-frequency resonances and the overall spectral model. For example using the corresponding Bark scale warping ($\lambda = 0.66$) [33] and Kautz filter order 100 produces a very detailed model for frequencies up to 3 kHz (Fig. 16).

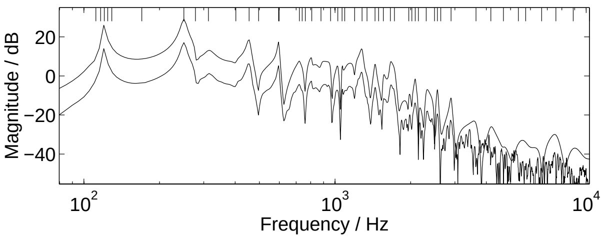

Fig. 16: A 100th order Kautz model.

If we truly concentrate on the lowest resonances we may obtain acceptable models with very low-order Kautz filters, using the warped target response in the pole position optimization. In Fig. 17 are presented the magnitude responses of the attained Kautz models for orders 10, 16, 20 and 40. We used warping parameter $\lambda = 0.7$. It is quite surprising that the BU-method found the five chosen prominent resonances at model order 10, i.e., with exactly five complex conjugate pole pairs, in contrast to the unwarped case where we needed the filter order 100.

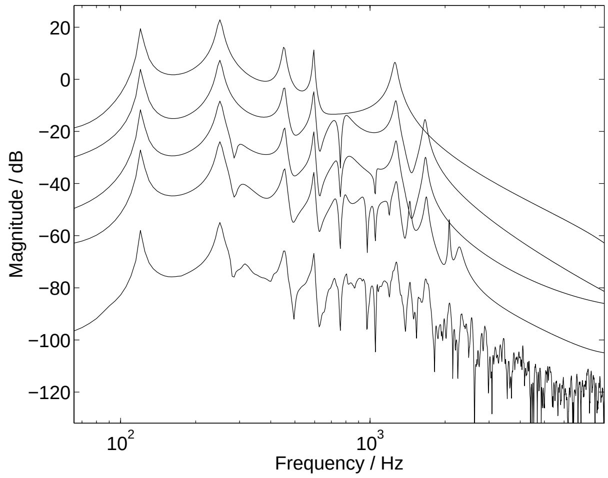

Fig. 17: Displayed with offset from top to bottom, Kautz models of orders 10, 16, 20 and 40, and the target magnitude response.

In Fig. 18 we demonstrate that good fitting to the five prominent resonances of the 10th order Kautz filter of Fig. 17 means also good matching in the time-domain (compared to Fig. 13).

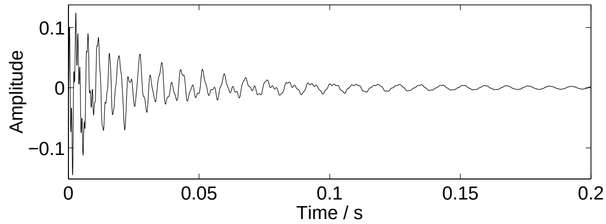

Fig. 18: The impulse response of a 10th order Kautz filter modeling the response in Fig. 13.

For comparison, we conclude with a Kautz model for a more realistic acoustic guitar body impulse response, discussed in the beginning of this chapter. The 160th order Kautz model of Fig. 19 is obtained by applying the BU-method to the warped ($\lambda = 0.7$) target response, and by mapping the poles back to the original frequency domain.

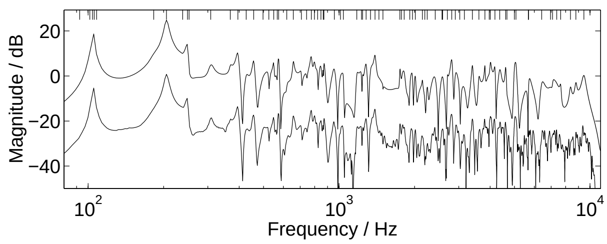

Fig. 19: Magnitude responses for a 160th order Kautz model of another acoustic guitar body impulse response, displayed together with upward offset.

DISCUSSION AND CONCLUSIONS

The above cases of audio equalization and modeling were taken as challenging examples in order to show the applicability of Kautz filters. Many specific questions, such as audio relevance of the modeling details, perceptual aspects of the designs, as well as computational robustness and expense have been addresses briefly or not at all. Thus call for further investigations.

The aim of this study was to show that it is possible to achieve good modeling or equalization results with lower Kautz filter orders than with warped (Laguerre) or traditional FIR and IIR filters. In the loudspeaker equalization case Kautz filter orders of 20–30 can achieve similar results of flatness as warped IIR models of order 100–200, or much higher orders with FIR equalizers. This reduction is due to well controlled focusing of frequency resolution on both global shape and particularly on local resonant behavior.

Especially the room response modeling case should be taken as a mere methodological study, except possibly regarding the low-frequency models. In the case of guitar body response modeling the low-frequency modes are important perceptually, and relatively low-order Kautz filters can focus sharply on them, showing advantage over warped, IIR, and FIR designs, especially when focus is on the separate low-frequency modes of the body response.

The basic flexibility of Kautz filter designs doesn't come without complications. In this paper we have, e.g., used hand-tuned poles, in addition to the more obvious and automatized choices of poles, to yield superior modeling with low orders. There are numerous possible techniques and strategies to search for an optimal model for a given problem, and different tasks may be solved best with different approaches. The cases investigated here just hint general guidelines, and fully automated search for optimal solution even in the present cases requires further work. However, we have demonstrated the potential applicability of Kautz filters. They are found flexible generalizations of FIR and Laguerre filters, providing IIR-like spectral modeling capabilities with well-known favorable properties resulting from the orthonormality. The competitiveness compared to Laguerre modeling is based on the fact that the generalization step imposes little or no extra computation load at runtime, even if the design phase may become more complicated.

MATLAB scripts and demos related to Kautz filter design can be found at `www.acoustics.hut.fi/software/kautz`.

ACKNOWLEDGEMENT

This work has been supported by the Academy of Finland as a part of project "Sound source modeling".

REFERENCES

[1] H. Brandenstein and R. Unbehauen, "Least-Squares Approximation of FIR by IIR Digital Filters", *IEEE Trans. Signal Processing,* vol. 46, no. 1, pp. 21–30, 1998.

[2] P. W. Broome, "Discrete Orthonormal Sequences", *Journal of the Association for Computing Mashinery,* vol. 12, no. 2, pp. 151–168, 1965.

[3] A. G. Constantinides, "Spectral Transformations for Digital Filters", *Proc. IEE*, vol. 117, no. 8, pp. 1585–1590, 1970.

[4] J. E. Cousseau, P. S. R. Diniz, G. B. Sentoni and O. E. Agamennoni, "An Orthogonal Adaptive IIR Realization for Efficient MSOE Optimization", *Proc. IEEE Int. Conf. Electronics, Circuits and Systems*, 1998, vol. 1, pp. 221–224.

[5] A. den Brinker, F. Brenders and T. Oliveira e Silva, "Optimality Conditions for Truncated Kautz Series", *IEEE Trans. Circ. and Syst.,* vol. 43, no. 2, pp. 117–122, 1996.

[6] E. Deprettere and P. Dewilde, "Orthogonal Cascade Realizarion of Real Multiport Digital Filters",*International Journal of Circuit Theory and Applications*, vol. 8, pp. 245–272, 1980.

[7] D. H. Friedman, "On Approximating an FIR Filter Using Discrete Orthonormal Exponentials", *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP–29, no. 4, pp. 923–831, 1981.

[8] A. H. Gray and J.D. Markel, " A Normalized Digital Filter Structure", *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 23, no. 3, pp. 268–277, 1975.

[9] P. S. C. Heuberger,*On Approximate System Identification with System Based Orthonormal Functions,* Ph.D dissertation, Delft University of Technology, The Netherlands, 1991.

[10] P. S. C. Heuberger, P. M. J. Van den Hof and O. H. Bosgra, "A Generalized Orthonormal Basis for Linear Dynamical Systems", *IEEE Transactions on Automatic Control*, vol. 40, no. 3, pp. 451–465, 1995.

[11] W. H. Huggins, "Signal Theory", *IRE Transactions on Circuit Theory*, vol. CT–3, pp. 210–216, 1956.

[12] M. Karjalainen, U. K. Laine and A. Härmä, "Realizable Warped IIR Filter Sructures", *Proc. IEEE Nordic Signal Processing Symposium, NORSIG'96*, Espoo, Finland, 1996, pp. 483–486.

[13] A. Härmä, "Implementation of Recursive Filters Having Delay Free Loops", *Proc. IEEE ICASSP'98,* Seattle 1998, vol. III, pp. 1261–1264.

[14] A. Härmä, M. Karjalainen, L. Savioja, V. Välimäki, U. K. Laine and J. Huopaniemi, "Frequency Warped Signal Processing for Audio Applications", *J. Audio Eng. Soc.*, Nov. 2000.

[15] S. Imai, "Cepstral Analysis Synthesis on the Mel Frequency Scale", *Proc. IEEE ICASSP'83*, Boston, MA, April 1983, pp. 93–96.

[16] M. Karjalainen, E. Piirilä, A. Järvinen and J. Huopaniemi, "Comparison of Loudspeaker Equalization Methods Based on DSP Techniques", *J. Audio Eng. Soc.,* vol. 47, no. 1/2, pp. 15–31, 1999.

[17] M. Karjalainen and J. O. Smith, "Body Modeling Techniques for String Instrument Synthesis", *Proc. Int. Computer Music Conf.*, Hong Kong, Aug. 1996, pp. 232–239.

[18] W. H. Kautz, "Transient Synthesis in the Time Domain", *IRE Trans. Circuit Theory*, vol. CT–1, pp. 29–39, 1954.

[19] R. E. King and P. N. Paraskevopoulos, "Parametric Identification of Discrete-time SISO Systems", *International Journal of Control*, vol. 30, no. 6, pp. 1023–1029, 1979.

[20] U. K. Laine, "FAMlet to be or not to be a wawelet", *Proc. IEEE-SP International Symposium on time-frequecy and time-scale analysis,* Victoria, British Columbia, Canada, October 1992, pp. 335–338.

[21] Y. W. Lee, *Statistical Theory of Communication.* John Wiley and Sons, New York, 1960.

[22] R. N. McDonough and W. H. Huggins, "Best Least-Squares Representation of Signals by Exponentials", *IEEE Trans. Aut. Cont.,* vol. 13, no. 4, pp. 408–412, 1968.

[23] C. T. Mullis and R.A. Roberts, "Synthesis of Minimum Roundoff Noise Fixed Point Digital Filters", *IEEE Transactions on Circuits and Systems*, vol. 23, no. 9, pp. 551–562, 1976.

[24] L. S. H. Ngia and F. Gustafsson,"Using Kautz Filters for Acoustic Echo Cancellation", *Conf. Rec. 33th Asilomar Conf. Signals, Systems and Computers*, 1999, vol. 2, pp. 1110–1114.

[25] Y. Nurges, "Laguerre Models in Problems of Approximation and Identification of Discrete Systems", *Automation and Remote Control*, vol. 48, pp. 346–352, 1987.

[26] T. Oliveira e Silva, " Rational Ortonormal Functions on the Unit Circle and the Imaginary Axis, with Applications in System Identification", 1995. `ftp://inesca.inesca.pt/pub/tos/English/rof.ps.gz`

[27] A. V. Oppenheim, D. H. Johnson and K. Steiglitz, "Computation of Spectra with Unequal Resolution Using Fast Fourier Transform", *Proceedings of the IEEE*, vol. 59, pp. 299–301, 1971.

[28] H. Penttinen, V. Välimäki and M. Karjalainen, "A Digital Filtering Approach to Obtain a More Acoustic Timber for an Electric Guitar", *Proc. EUSIPCO 2000*, Tampere, Finland, September 4–8, 2000.

[29] P. A. Regalia, "Stable and Efficient Lattice Algorithms for Adaptive IIR Filtering", *IEEE Transactions on Signal Processing*, vol. 40, no. 2, 375–388, 1992.

[30] R. A. Roberts and C. T. Mullis, *Digital Signal Processing*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1987.

[31] M. Tyril, J. A. Pedersen and P. Rubak "Digital Filters for Low-Frequency Equalization", *J. Audio Eng. Soc.*, vol.49, no. 1/2, pp. 36–43, 2001.

[32] L. Salama and J. E. Cousseau, "Comparison of Orthonormal Adaptive FIR and IIR Filter Realizations", *Proc. IEEE Int. Symp. Advances in Digital Filtering and Signal Processing*, 1998, pp. 77–81.

[33] J. O. Smith and J.S. Abel, "Bark and ERB Bilinear Transform", *IEEE Trans. Speech and Audio Processing,* vol. 7, no. 6, pp. 697–708, 1999.

[34] K. Steiglitz, "A Note on Variable Resolution Digital Filters", *IEEE Trans. Acoustics, speech and Signal Processing*, vol. ASSP–28, no. 1, pp. 111–112, 1980.

[35] H. W. Strube, "Linear Prediction on a Warped Frequency Scale", *J. Acoust. Soc. Am.*, vol. 68, no. 4, pp. 1071–1076, 1980.

[36] G. Szegö, *Orthonormal Polynomials.* American Mathematical Society, New York, 1939.

[37] B. Wahlberg, "System Identification Using Laguerre Models", *IEEE Transactions on Automatic Control*, vol. 36, no. 5, pp. 551–562, 1994.

[38] B. Wahlberg, "System Identification Using Kautz Models", *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp. 1276–1282, 1994.

[39] B. Wahlberg and P Mäkilä, "On Approximation of Stable Linear Dynamical Systems Using Laguerre and Kautz Functions", *Automatica*, vol. 32, no. 5, pp. 693–708, 1996.

[40] P. Vaidyanathan, "A Unified Approach to Orthogonal Digital Filters and Wave Digital Filters, Based on LBR Two-Pair Extraction", *IEEE Transactions on Circuits and Systems*, vol. CAS-32, no. 7, pp. 673–686. 1985.

[41] J. L. Walsh, *Interpolation and Approximation by Rational Functions in the Complex Domain*, 2nd Edition. American Mathematical Society, Providence, Rhode Island, 1969.

[42] P. M. J. Van den Hof, P. S. C. Heuberger and J. Bokor, "System Identification with Generalized Orthonormal Basis Functions", *Proceedings of the 33rd Conference on Design and Control*, Lake Buena Vista, FL, December 1994. pp. 3382–3387.

[43] T. Y. Young and W.H. Huggins, "'Complementary' Signals and Orthogonalized Exponentials", *IRE Transactions on Circuit Theory*, vol. CT–9, pp. 362–370, 1962.

[44] C. Zervos, P. R. Bélanger and G. A. Dumont, "On PID Controller Tuning Using Orthonormal Series Identification", *Automatica*, vol. 24, no. 2, pp. 165–175, 1988.