

DIRECT DESIGN OF PARALLEL SECOND-ORDER FILTERS FOR INSTRUMENT BODY MODELING

Balázs Bank

Helsinki University of Technology
Laboratory of Acoustics and Audio Signal Processing

ABSTRACT

This paper¹ presents a direct design technique for parallel second-order sections based on a perceptually motivated logarithmic scale, with application to instrument body modeling. Traditional FIR and IIR design techniques work on a linear frequency scale, which is usually not optimal for audio applications. Warped filters and Kautz filters are good candidates for perceptually motivated filter design. However, the direct implementation of warped or Kautz filters is computationally less efficient compared to an IIR filter of the same order. The perceptually motivated frequency resolution can also be achieved by the proposed design method, without the disadvantage of increased computational complexity. The filter design has two phases. First, the pole frequencies of the second-order sections are set to a predefined logarithmic scale, or can be determined from a warped filter design. Then the zeros are found by a simple least squares algorithm, as the optimization problem becomes linear in parameters for a given set of poles. As high-order (up to 1000) filters can be robustly designed, this technique is particularly well suited for instrument body modeling. Moreover, the parallel structure allows flexible modifications to the body transfer function.

1. INTRODUCTION

The basic idea of physics-based sound synthesis is that it models the sound production mechanism of the musical instrument (see [1] for overview). Therefore, the structure of the model follows that of a real instrument. Thus, for string instruments, the model parts are the excitation, the string and the instrument body. The heart of the instrument is the string, as that is the one which produces the periodic vibration as a response to the excitation. The instrument body has two roles from the modeling point of view. First, it produces a filtering to the sound when it converts the string force to sound pressure in a given point in space (“radiation” properties). Second, it acts as a termination to the strings, modifying the partial frequencies and decay times, and implementing a coupling between the different strings (“termination” properties).

As the real-time fully-physical modeling of the instrument body is still not feasible even with today’s fast DSPs, the effect of the instrument body is often treated as a filtering operation acting upon the output of the string model.

Note that this approach separates the “radiation” and “termination” properties of the soundboard. In practice, usually only the radiation effects are implemented by the instrument body model and the termination effects are included in the string model.

The most efficient approach for modeling the radiation properties of the soundboard is the commuted synthesis technique [2, 3]. This is based on the fact that if all the blocks of the model are linear and time invariant, then the order of the blocks can be commuted and the body impulse response can be read from a wavetable. However, as nowadays higher computational power is available, we may seek for other methods that do not require the linearity of the string and excitation models.

The instrument body response can also be implemented as a fast convolution utilizing the FFT algorithm, such as the one presented in [4]. While it is very efficient for long impulse responses, the algorithm becomes quite complicated if low latency is required. Another way of decreasing the complexity of the convolution is the use of the multi-rate approach. In [5] a simple multi-rate algorithm have been presented that splits the input signal into two frequency bands and applies FIR filtering. Naturally, the idea of multi-rate processing can be used not only for FIR filters, but for any other filter structure. Thus, it may be applied as a possible extension to the filter-based body modeling techniques discussed in this paper.

The filtering effect of the instrument body can also be modeled as a reverberation algorithm. However, that approach cannot model the measured responses accurately, and only the overall properties of the instrument body (modal density, decay, spectral shape) can be controlled.

If the transfer function measurement of a real instrument body is available, then body modeling can be considered as a filter design problem. Unfortunately, the transfer function of real instrument bodies exhibits high modal density, making difficulties for standard filter design algorithms. This is mostly because traditional FIR and IIR design algorithms are optimizing on a linear frequency scale. A logarithmic or auditory frequency resolution can be accomplished by warped or Kautz filters, but they require special filter structures to be implemented.

This paper, after reviewing the existing filter design alternatives in Sec. 2, proposes a new method for designing the body filter in the form of parallel second-order sections on the logarithmic frequency scale in Sec. 3. With

¹ Presented at the International Computer Music Conference, Proceedings Vol. I. pp. 458-465, Copenhagen, Denmark, Aug. 2007.

the new method most of the advantages of Kautz filters are retained, while it requires less computation. Examples for piano soundboard modeling are given in Sec. 4. The method can be used not only for instrument body modeling, but for any application that requires nonuniform frequency resolution (i.e., in most audio applications). Some application examples are outlined in Sec. 5 besides describing the advantages of the method.

2. EXISTING FILTER DESIGN ALTERNATIVES

2.1. Traditional filter design techniques

The most straightforward approach to design the body model as an FIR filter is to apply windowing to the measured impulse response. In the case of the acoustic guitar, filter orders lower than 1000 do not produce satisfactory sound [6]. For modeling the piano soundboard at a sampling rate of $f_s = 44.1$ kHz, the present author has found that 2000 tap filters give a good overall sound quality. However, for synthesizing the characteristic knocking sound of the soundboard for high tones, even higher orders (≈ 10000) are required.

Another straightforward choice for body modeling is applying one of the standard IIR design techniques. In [6], two IIR filter design methods were compared for modeling the guitar body. It has turned out that IIR filters perform almost the same as FIR filters for the same computational cost. As noted in [6], minimum-phase equalization is not preferable, since it destroys the reverberant character of the response.

Traditional FIR and IIR techniques are not optimal for instrument body modeling (and for most of the other audio applications), because their frequency resolution is linear, as opposed to the quasi-logarithmic resolution of the human hearing. In the particular case of body modeling, this is emphasized by the fact that usually the lower body modes have smaller bandwidth compared to the higher ones. The linear frequency resolution is inherent in FIR filters, and while IIR filters could have a higher pole density in the low frequencies in theory, the traditional filter design techniques tend to place the filter poles uniformly around the unit circle. Unfortunately, the logarithmic frequency scale is so highly distorted compared to the linear scale, that even weighted filter design cannot give satisfactory results [7]. In the following sections such filter structures and design techniques will be reviewed which use a perceptually motivated frequency resolution.

2.2. Warped filters

Warped filters are particularly well suited for audio applications because their frequency resolution can closely match that of the human hearing [7, 8]. The basic idea of warped filters is that the unit delay z^{-1} in the traditional FIR or IIR filters is replaced by an allpass filter

$$z^{-1} \leftarrow D(z) = \frac{z^{-1} - \lambda}{1 - \lambda z^{-1}}. \quad (1)$$

By a particular choice of the λ parameter, it is possible to match the Bark or ERB scale closely [9].

The design of warped filters starts with warping the target impulse response $h_t(n)$, e.g., by the use of an allpass chain. Then, warped FIR (WFIR) filters can be obtained by truncating or windowing the warped target response $\tilde{h}_t(n)$. The WFIR filters have a similar structure as FIR filters, but the unit delays are replaced by the allpass filter $D(z)$. That is, the WFIR filter is an allpass chain, where the signals between the first-order allpass blocks are tapped and weighted by b_k . For warped IIR (WIIR) filters, standard IIR filter design techniques are applied to the warped target response $\tilde{h}_t(n)$. However, for WIIR filters, the replacement of unit delays by $D(z)$ leads to delay-free loops, and the filter structure has to be modified for practical implementation [8].

The application of frequency warping produces large savings in terms of filter order for instrument body modeling. For modeling the body of the classical guitar, while the required order is around 1000–2000 for FIR and around 500–1000 for IIR filters, WIIR filters are capable of producing the same quality by an order of 100–200.

The shortcoming of warped filters is that special filter structures are required for their implementation, which lead to higher computational complexity compared to traditional FIR or IIR filters of same order. For low orders (<20 – 30), WFIR and WIIR filters can be converted to traditional (direct form) IIR filters, but for higher orders this is not possible because of numerical inaccuracies during conversion [8].

2.3. Kautz filters

In the case of warped filters the frequency resolution is controlled by one parameter (i.e., λ). A more flexible allocation of frequency resolution can be achieved by Kautz filters. Kautz filters can be seen as the generalization of WFIR filters, in the sense that now the allpass filters in the chain do not have to be identical (for overview and audio applications, see [10]). The Kautz transfer function can be written as

$$\begin{aligned} H(z) &= \sum_{k=0}^N w_k G_k(z) \\ &= \sum_{k=0}^N w_k \left(\frac{\sqrt{1 - p_k p_k^*}^{k-1} z^{-1} - p_j^*}{1 - p_k z^{-1} \prod_{j=0}^{k-1} \frac{z^{-1} - p_j^*}{1 - p_j z^{-1}}} \right), \quad (2) \end{aligned}$$

where $G_k(z)$ are the orthonormal Kautz functions determined by the pole set p_k . From the design point of view, Kautz filters can be considered as fixed-pole IIR filters. The advantage of the orthonormality of $G_k(z)$ functions is that the weights w_k can be determined from the target response $h_t(n)$ by a scalar product

$$w_k = \sum_{n=1}^{\infty} g_k(n) h_t(n), \quad (3)$$

where $g_k(n)$ are the inverse z transform of $G_k(z)$.

For determining the poles p_k of the Kautz filter, several methods are discussed in [10]. The simplest is to use a logarithmic pole distribution

$$\vartheta_k = \frac{2\pi f_k}{f_s} \quad (4)$$

$$p_k = R^{\vartheta_k/\pi} e^{\pm j\vartheta_k}, \quad (5)$$

where ϑ_k are the pole frequencies in radians determined by the logarithmic frequency series f_k and the sampling frequency f_s . The pole magnitudes form an exponentially damping sequence approximating a constant Q resolution. The pole magnitude at the Nyquist rate is set by the damping parameter R . Another way of finding the poles is an iterative algorithm, like the Brandenstein-Undehaunen method discussed in [10]. For better results in the case of body modeling, warped pole positioning can be applied. Adding some poles manually can also help in focusing on the critical regions of the transfer function. For modeling the guitar body, a Kautz filter with warped Brandenstein-Undehaunen pole positioning produced satisfactory results already at an order of 120 [10]. This is comparable to the orders required for warped IIR filters.

It is impractical to implement Kautz filters by a series of complex first-order allpass filters as in Eq. (2), and combining the complex pole pairs to second-order sections yields lower computational complexity [10]. However, the combined cascade-parallel nature of the filter still requires more computation compared to filters implemented in direct or cascade form on DSP. This can be overcome by converting the Kautz filter to a direct form IIR filter, but this is possible only for low filter orders due to numerical problems during conversion.

3. PARALLEL SECOND-ORDER SECTIONS

We have seen that Kautz filters provide more flexibility in the distribution of frequency resolution compared to warped filters. This is because the resolution is controlled by the entire pole set and not only by one parameter. It would be beneficial to find a filter structure that retains this property of Kautz filters, while enables more efficient implementation. This is achievable by parallel second-order sections.

Implementing IIR filters in the form of parallel second-order sections have been used traditionally because its better quantization noise performance [11] and the possibility of code parallelization. The parameters of the second-order sections are usually determined from the direct form IIR filters, by, e.g., the partial fraction expansion or a similar algorithm [12]. However, this approach is not applicable for our case, as traditional IIR filter design techniques are not well suited for audio applications because of the reasons discussed in Sec. 2.1. For converting the better performing warped or Kautz filters to parallel form, they should be converted to direct form IIR filters first, which is not possible due to numerical inaccuracies. Therefore, a different approach has to be taken for the parameter estimation of the second-order sections.

The algorithm presented here designs the parallel second-order filter-bank directly, without the intermediate direct-form IIR filter. After determining the poles, it uses the outputs of the second-order sections (i.e., exponentially decaying sinusoidal functions) as basis functions of a linear-in-parameter model. In this sense the concept is similar to that of Kautz filters, but by giving up the orthonormality of the basis functions, the scalar product of Eq. (3) cannot be used for parameter estimation. On the other hand, the new method results in a simpler structure and significant computational savings.

3.1. Problem formulation

Every transfer function of the form $H(z^{-1}) = B(z^{-1})/A(z^{-1})$ can be rewritten in the form of partial fractions:

$$H(z^{-1}) = \sum_{k=1}^K c_k \frac{1}{1 - p_k z^{-1}} + \sum_{m=0}^M b_m z^{-m}, \quad (6)$$

where p_k are the poles, forming either conjugate pairs or real valued, if the system has a real impulse response. The second sum in Eq. (6) is the FIR part of order M . If the order of $A(z^{-1})$ and $B(z^{-1})$ is the same, then it reduces to a constant coefficient b_0 . Note that Eq. (6) is valid only if no multiple poles are present. In the case of pole multiplicity, terms of higher order also appear.

Now let us assume that not only the target response, but the poles of our IIR filter are known (the pole positioning problem will be outlined in the next section). In this case Eq. (6) becomes linear in parameters c_k and b_m , thus, they can be estimated by a simple least squares algorithm to match the required response. This can also be done in the frequency domain by the substitution $z^{-1} = e^{-j\vartheta}$, where ϑ is the digital angular frequency. Then the right-hand side of Eq. (6) is evaluated at those frequencies ϑ_i , where the target response $H(e^{-j\vartheta_i})$ is known, and the parameters c_k and b_m are set by least squares optimization so that the mean squared error between the target and the implemented frequency response is minimal.

However, the simplest way for finding the coefficients c_k and b_m is to compute the impulse responses of the complex first-order filters of Eq. (6) in the time domain, up to the length of the target impulse response. The impulse response of the k th term $1/(1 - p_k z^{-1})$ becomes $u_k(n) = (p_k)^n$, where n is the discrete time step. Then the columns of the modeling signal matrix \mathbf{M} are composed of the vectors $u_k(n)$. Moreover, for the FIR part, a set of delayed impulses $\delta(n - m)$ have to be added. The parameter vector \mathbf{p} is composed of the c_k and the b_m parameters $\mathbf{p} = [c_1 \dots c_K, b_0 \dots b_M]^T$. Thus, the resulting filter impulse response \mathbf{h} can be calculated as

$$\mathbf{h} = \mathbf{M}\mathbf{p}. \quad (7)$$

Because of the nature of the linear-in-parameter problems, the error function is quadratic in \mathbf{p} and has a unique minimum, as opposed to nonlinear optimization problems. The

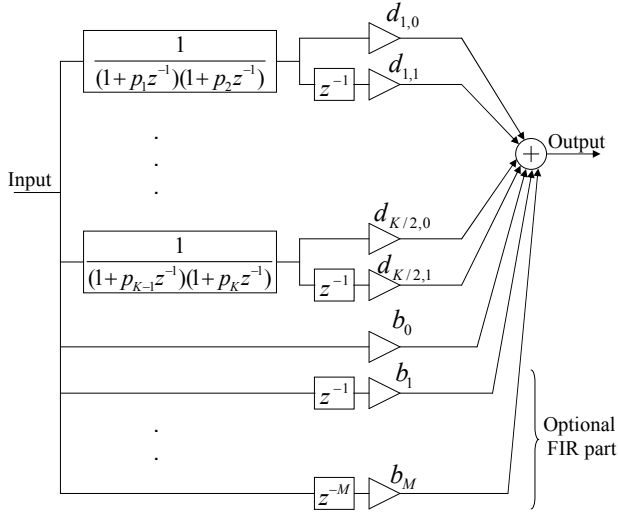


Figure 1. The structure of the parallel second-order filter.

mean squared error between the target response \mathbf{h}_t and the resulting filter response \mathbf{h} is minimized by the expression

$$\mathbf{p}_{\text{opt}} = (\mathbf{M}^H \mathbf{M})^{-1} \mathbf{M}^H \mathbf{h}_t, \quad (8)$$

where the vector \mathbf{p}_{opt} contains the optimal parameters, and \mathbf{M}^H is the conjugate transpose of \mathbf{M} .

The resulting filter can be implemented directly as Eq. (6) forming parallel first-order complex filters. However, this is not optimal from the computational point of view. On the other hand, if the complex pole pairs are brought to a common denominator, we obtain

$$\begin{aligned} \frac{c_k}{1 - p_k z^{-1}} + \frac{c_{k+1}}{1 - p_{k+1} z^{-1}} &= \\ \frac{c_k(1 - p_{k+1} z^{-1}) + c_{k+1}(1 - p_k z^{-1})}{(1 - p_k z^{-1})(1 - p_{k+1} z^{-1})} &= \\ \frac{d_{k,0} + d_{k,1} z^{-1}}{1 - (p_k + p_{k+1})z^{-1} + p_k p_{k+1} z^{-2}}, \end{aligned} \quad (9)$$

where p_k and p_{k+1} , and c_k and c_{k+1} are complex conjugate pairs. This is a second-order section with real valued coefficients, which can be implemented much more efficiently. Those fractions of Eq. (6) that have real poles are either implemented as first-order IIR filters, or combined with other real poles to form second-order IIR filters, yielding a canonical structure. This is depicted in Fig. 1. In the case of multiple poles, higher order sections also appear. However, in our applications, multiple poles do not occur in practice (see Sec. 3.2).

Although $d_{k,0}$ and $d_{k,1}$ can be determined uniquely from c_k and p_k by the help of Eq. (9), the optimization problem is still linear in parameters if the second-order sections are used instead of the complex first-order filters. With that approach the feedforward parameters $d_{k,0}$ and $d_{k,1}$ can be estimated directly with the least squares algorithm, and there is no need to calculate them from the c_k and p_k values of Eq. (6). In that case, the matrix \mathbf{M} contains not only the outputs of the two-pole resonators but

also their delayed versions by one temporal sample. Then, the optimal parameter set \mathbf{p}_{opt} is obtained by Eq. (8), and now $\mathbf{p}_{\text{opt}} = [d_{1,0}, d_{1,1}, \dots, d_{K/2,0}, d_{K/2,1}, b_0 \dots b_M]^T$.

3.2. Pole positioning

The poles of the second-order sections can be determined by all the methods suggested for the case of Kautz filters [10], some of them were briefly outlined in Sec. 2.3. Positioning the poles logarithmically is the simplest way. In this case, pole multiplicity is naturally avoided. In general, the pole set does not have to be strictly logarithmic, but can also focus to specific frequencies by increasing the pole density in that region.

Another simple way of finding the poles of the second-order filters is first designing a warped IIR filter to the warped target response $\tilde{h}_t(n)$. Then, the poles of this warped filter have to be found. This is in practice possible even for filter orders of 1000, as the poles are spread approximately uniformly along the unit circle in the warped domain. The warped poles \tilde{p}_k are converted back to linear frequency scale by an expression similar to Eq. (1):

$$p_k = \frac{\tilde{p}_k + \lambda}{1 + \tilde{p}_k}, \quad (10)$$

where λ is the warping parameter with which the WIIR filter was designed. The poles p_k can be directly used for filter design. It has to be noted that the author has never experienced the appearance of multiple poles with any practical design method, meaning that the canonical structure of Fig. 1 can usually be retained.

4. DESIGN EXAMPLES

The performance benefit of the warped or logarithmic frequency scale methods strongly depends on the phase properties of the target response. Warping tends to stretch the high-frequency components of the target response in time, while compresses the signals at low frequencies [8]. Therefore, if the target response is non-minimumphase, thus, having maximum power at, e.g., 5 ms, warping will move this to 10–20 ms in the high frequencies, making them out of reach for the filter design algorithm for low filter orders. On the other hand, minimum-phase target responses have their power maxima in the beginning of the response for all frequencies, thus, warped or logarithmic frequency resolution filters can perform optimally.

Body responses are non-minimumphase in the strict sense. However, closely miked responses behave almost like minimum-phase responses from the design point of view. Generally, the literature [8, 10] uses closely miked responses and sometimes even makes them completely minimum-phase artificially. In this section, a closely miked (20 cm) piano soundboard response is used in general, but as an example of a largely non-minimumphase target, a far-miked soundboard response is presented in Sec. 4.4.

Examples of warped or Kautz designs are not shown in the next sections. Kautz filters perform in principle in




Figure 2. Comparison of frequency responses achieved by different filter structures at an order of 200. From the top to bottom: the proposed parallel filter, the Kautz filter, the warped filter, and finally the target specification.

the same way as the parallel second-order sections for the same pole set, and warped filters have a very similar performance as the parallel filter, when the pole set is determined from a warped filter. This is illustrated in Fig. 2, where first a 200 order warped IIR filter was fitted to the target specification. Then, the poles of the warped filter were used for designing the Kautz and parallel filters.

4.1. Reference IIR filter design

Figure 3 (a) shows the frequency response of two IIR filters designed by the Prony's method. The target response is displayed by the bottom curves of Figure 3, and has a length of 10000 taps at $f_s = 44.1$ kHz. Filter orders of 200 are needed for acceptable sound (top curves), while the filter response starts to be perceptually indistinguishable from the target only at an order of 800. It can be seen from Fig. 3 (a) that the design wastes too much effort for the perceptually less significant high frequencies. The impulse responses of the filters are depicted in Fig. 3 (b). They produce an excellent match for the first few hundred samples (until the length of their FIR part, they are perfect by principle), but they are unable to follow the long-ringing low modes of the target response.

4.2. Parallel second-order sections with logarithmic pole positioning

Figure 4 (a) shows the frequency response of two parallel second-order designs, where the poles are set on a logarithmic scale. For the top curve, 25 logarithmically spaced resonators are used from 60 Hz to 20 kHz yielding a filter order of 50. The poles are calculated by Eqs. (4) and (5) with $R = 0.9$. The middle curve uses 100 logarithmically spaced poles with $R = 0.98$, giving a filter order of 200. The order of the FIR part M was set to zero in




Figure 3. Frequency (a) and time-domain (b) responses of traditional IIR filters of order 200 (top) and 800 (middle), and the target impulse response (bottom).

both examples, resulting in a canonical structure containing second-order IIR sections only. The time-domain responses of Fig. 4 show that the parallel filter can follow the long-ringing modes better than traditional IIR filters even at lower filter orders. Naturally, this can be already anticipated from the more precise fit in the magnitude response at low frequencies. The author has found in general that around four times lower order filters are required for the same sound quality compared to traditional IIR design.

4.3. Parallel second-order sections with warped-filter based pole positioning

Next, the pole set of the parallel second-order filter is estimated by warped filter design. For the top curve of Fig. 5 a 50th order WIIR filter was designed on the warped target response with $\lambda = 0.75$. Then, the poles of the warped filter were determined and transferred to the normal (unwarped) domain by Eq. (10). The middle curve of Fig. 5 shows a similar example but with an order of 200. Similarly to the previous examples, no parallel FIR part was implemented. The warping-based pole positioning gives




Figure 4. Frequency (a) and time-domain (b) responses of parallel second-order sections designed with logarithmic pole positioning for orders of 50 (top) and 200 (middle), and the target impulse response (bottom).

comparable perceptual results to the logarithmic positioning for the same filter order.

4.4. Fitting on a largely non-minimum phase response

In this example, a piano soundboard impulse response is used which was recorded at 2 m distance from the piano in a fairly damped rehearsal room. This means that the target response $h(n)$ (the lowest curve in Fig. 6 (b)) is far from being minimum-phase: it reaches its maximum around the 200th sample. Indeed, a 200th order parallel filter (middle curve in Fig. 6 (b)) is unable to follow the onset accurately. It can be seen that the resulting response contains high-frequency components only in the first part, which is a general shortcoming of warped or logarithmic filter designs. In this case, the pole set is determined from a warped design, but similar results are obtained for logarithmic pole positioning. Designing a filter of an order of 400 cures the problem, but at the expense of doubled computational complexity.

It is a better solution to include a parallel FIR chain




Figure 5. Frequency (a) and time-domain (b) responses of parallel second-order sections with a pole set obtained by warped filter design for orders of 50 (top) and 200 (middle), and the target impulse response (bottom).

in the filter structure, as depicted in the bottom of Fig. 1. For the top curves of Fig. 6, a 100th order IIR part and a 200th order FIR part was used, which have the same computational complexity in terms of additions and multiplications as the 200th order parallel IIR filter displayed in the middle. The pole set was again determined from a warped design, but now with a $\lambda = 0.9$ giving more emphasis on low frequencies, as the high frequency part is taken care by the FIR part anyway. The top curve of Fig. 6 (b) shows that the time-domain response is much more accurate (actually, perfect up to the length of the FIR part), and the magnitude response has also improved in the high frequencies.

In this example the linear frequency resolution of the FIR filter is combined with the auditory scale of the parallel filter. This approach often results in the lowest computational complexity even in the case of closely miked responses, but the canonical structure of having only second-order filters in parallel is lost, meaning that some of the additional benefits of the method (discussed in Sec. 5.1) cannot be utilized.




Figure 6. Frequency (a) and time-domain (b) response comparison of different filter designs: 50 parallel second-order sections and a 200 tap FIR filter in parallel (top), and 100 parallel second-order sections (middle). The target impulse response is displayed at the bottom.

5. DISCUSSION

The new method requires around four times less number of additions and multiplications compared to traditional IIR filter designs for piano soundboard modeling. This decrease is even larger if the target response is made minimum-phase before filter design.

If the pole set of the parallel filter is determined from a warped design, the parallel filter performs the same way as the WIIR filter of the same order. However, the parallel filter yields a simpler DSP implementation and provides more flexibility in the distribution of frequency resolution.

As the design of the parallel second-order filter is very similar to that of the Kautz filter, the parallel filter performs the same way as the Kautz filter for a given pole set. On the other hand, it requires smaller number of multiplications and additions and its parallel structure is better suited for DSP implementation compared to the mixed cascade-parallel structure of the Kautz filter. The proposed method has two disadvantages compared to Kautz

filters, but none of them cause any problems in our case. First, in the case of multiple poles, the canonical structure of Fig. 1 has to be extended, while for Kautz filters the structure is not dependent on pole multiplicity (in practice, multiple poles do not appear in body modeling). Second, because the basis functions of the parallel filter are not orthonormal, filter design takes longer time and can be in theory numerically more sensitive, as the filter coefficients have to be determined by the matrix operations of Eq. (8) instead of a scalar product. However, in practice the filters used in this study (or even orders up to 1000) could be robustly designed without any numerical problems.

5.1. Additional advantages

The structure of the parallel filter is well suited for code parallelization, which can lead to even higher computational savings on some DSPs. Moreover, the parallel second-order structure has a lower quantization noise compared to direct form or cascade implementations [11].

Another benefit for physical modeling is that the parallel filter structure can have a physical interpretation: it can be related to modal synthesis, where the vibration of a structure is combined from that of its modes [13]. This analogy to the physical reality leads to interesting parameter modifications. The resonance frequencies of the second-order filters (pole angles) can be changed to simulate a bigger or smaller instrument body. The decay times of the body modes can be influenced by varying the pole radii. Changing the overall magnitude response can be accomplished by scaling the feedforward coefficients $d_{k,0}$ and $d_{k,1}$ of Fig. 1, without an additional filter. Morphing of different body responses is also possible.

The same method can be utilized for scalability in terms of sampling rate. The parameters of the second-order sections have to be determined for the highest possible sampling rate, and when the system has to run at a lower rate f'_s , the modes above $f'_s/2$ are dropped and the remaining ones are stretched appropriately. Multiple outputs (e.g., for stereo effects) can be efficiently achieved by using the same set of poles for the different channels, thus, only the output coefficients of Fig. 1 have to be implemented for the channels separately.

5.2. Future research

The construction of an adequate perceptual model for the quantitative evaluation of body modeling techniques would be beneficial, but it is far from being a straightforward task. This is because the perception of the body response is highly dependent on the input signal, thus, looking at transfer functions or impulse responses gives only a rough measure of quality.

The proposed method may find applications in all those cases, where Kautz filters have been proven to be effective, as it produces comparable performance at lower computational cost. Trivial examples are audio applications such as loudspeaker or room equalization. The application of the technique for the source-filter modeling of mu-

sical instrument sounds and comparison with previous parameter estimation techniques, such as the ones used in [14, 15], can also be a part of future research.

As the method has been found to be robust even in the order of a thousand, it might be applicable for room response modeling. While not as efficient as fast convolution algorithms [4], it gives larger flexibility, since all the parameters of the late reverberation could be changed as outlined in Sec. 5.1. Room modeling requires significantly higher filter orders compared to instrument body modeling. Therefore, combining the approach with multi-rate techniques could be favorable.

6. CONCLUSION

In this paper, a simple and robust method has been proposed for the direct design of a parallel set of second-order filters. While the idea of using such a filter structure is not new, the proposed direct design method makes it possible to use the parallel structure in an auditory (or logarithmic) frequency scale, which is better suited for audio applications. Design examples have been presented for the case of piano soundboard modeling. The method produces comparable results to warped and Kautz filters, while does not require the implementation of special structures. Compared to traditional IIR design, around four times lower filter orders are obtained for the same sound quality. Moreover, the parallel structure is well suited for DSP implementation and produces lower quantization noise compared to cascade or direct structures. A significant benefit for physics-based sound synthesis is that the parameters of the parallel filter can be changed in a meaningful way, leading to the real-time variability of body size and decay, or to the morphing of different responses. The method can not only find applications in instrument body modeling, but also in other applications where nonuniform frequency resolution is preferable.

Sound examples and MATLAB functions can be found at the webpage <http://www.acoustics.hut.fi/go/icmc07-parfilt>.

7. ACKNOWLEDGEMENTS

This work has been supported by the EC 6th Framework Programme Marie Curie Intra-European Fellowships contract no. 041924 “Nonlinear Effects in String Instruments: Perception and Modeling”. The author is thankful for the helpful comments of Prof. Matti Karjalainen, Prof. Vesa Välimäki, and Dr. Tuomas Paatero.

8. REFERENCES

- [1] Vesa Välimäki, Jyri Pakarinen, Cumhuri Erkut, and Matti Karjalainen. Discrete-time modelling of musical instruments. *Reports on Progress in Physics*, 69(1):1–78, Oct. 2006.
- [2] Matti Karjalainen and Vesa Välimäki. Model-based analysis/synthesis of the acoustic guitar. In *Proc. Stockholm Music Acoustics Conf.*, pages 443–447, Stockholm, Sweden, July 1993.
- [3] Julius O. Smith. Efficient synthesis of stringed musical instruments. In *Proc. Int. Computer Music Conf.*, pages 64–71, Tokyo, Japan, Sep. 1993.
- [4] William G. Gardner. Efficient convolution without input–output delay. *J. Audio Eng. Soc.*, 43(3):127–136, Nov. 1995.
- [5] Balázs Bank, Giovanni De Poli, and László Sujbert. A multi-rate approach to instrument body modeling for real-time sound synthesis applications. In *Proc. 112th AES Conv., Preprint No. 5526*, Munich, Germany, May 2002.
- [6] Matti Karjalainen, Vesa Välimäki, Heikki Räisänen, and Harri Saastamoinen. DSP equalization of electric film pickup for acoustic guitar. *J. Audio Eng. Soc.*, 48(12):1183–1193, 2000.
- [7] Michael Waters and Mark B. Sandler. Least squares IIR filter design on a logarithmic frequency scale. In *Proc. IEEE Int. Symp. on Circuits and Systems*, pages 635–638, May 1993.
- [8] Matti Karjalainen and Julius O. Smith. Body modeling techniques for string instrument synthesis. In *Proc. Int. Computer Music Conf.*, pages 232–239, Hong Kong, Aug. 1996.
- [9] Julius O. Smith and Jonathan S. Abel. Bark and ERB bilinear transform. *IEEE Trans. Speech Audio Proc.*, 7(6):697–708, Nov. 1999.
- [10] Tuomas Paatero and Matti Karjalainen. Kautz filters and generalized frequency resolution: Theory and audio applications. *J. Audio Eng. Soc.*, 51(1–2):27–43, Jan./Feb. 2003.
- [11] Wei Chen. Performance of cascade and parallel IIR filters. *J. Audio Eng. Soc.*, 44(3):148–158, 1996.
- [12] M. Price, S. Holden, and M. Sandler. Accurate parallel form filter synthesis. *IEE Electronics Letters*, 32(22):2066–2067, Oct. 1996.
- [13] Jean-Marie Adrien. The missing link: Modal synthesis. In Giovanni De Poli, Aldo Piccialli, and Curtis Roads, editors, *Representations of Musical Signals*, pages 269–297. The MIT Press, Cambridge, Massachusetts, USA, 1991.
- [14] Jean Laroche and Jean-Louis Meillier. Multichannel excitation/filter modeling of percussive sounds with application to the piano. *IEEE Trans. Speech Audio Proc.*, 2(2):329–344, Apr. 1994.
- [15] J. P. Mackenzie, I. Kale, and G. D. Cain. Applying balanced model truncation to sound analysis/synthesis models. In *Proc. Int. Computer Music Conf.*, pages 400–403, Banff, Canada, Sep. 1995.