AUTOMATED CALIBRATION OF A PARAMETRIC SPRING REVERB MODEL

Hannes Gamper, *

Department of Media Technology
Aalto University
Espoo, Finland
first[dot]last[at]tml.hut.fi

Julian Parker and Vesa Välimäki, †

Department of Signal Processing and Acoustics,
Aalto University
Espoo, Finland
first[dot]last[at]tkk.fi

ABSTRACT

The calibration of a digital spring reverberator model is crucial for the authenticity and quality of the sound produced by the model. In this paper, an automated calibration of the model parameters is proposed, by analysing the spectrogram, the energy decay curve, the spectrum, and the autocorrelation of the time signal and spectrogram. A visual inspection of the spectrograms as well as a comparison of sound samples proves the approach to be successful for estimating the parameters of reverberators with one, two and three springs. This indicates that the proposed method is a viable alternative to manual calibration of spring reverberator models.

1. INTRODUCTION

Spring reverberation is an early method of artificial reverberation, introduced by Laurens Hammond in the 1940s [1]. Its small size and low cost compared to the contemporary methods of artificial reverberation at the time, such as plates or chambers, led to its wide use both in studio applications and within electrical musical instruments and amplifiers. The special sound of the spring reverberator, caused by the highly dispersive nature of wave propagation on the spring, became valued as a musical effect distinct from standard reverberation.

A spring reverberator consists of one or more helical metal springs, connected in parallel, series or in a hybrid configuration. The springs are excited via the use of an electromagnetic coil, which applies a force to a small magnet connected to springs. Varying the signal in the coil produces corresponding vibrations in the spring. The output from the system is taken with a similar configuration. A small magnetic bead oscillates with the spring at another location (usually the opposite end of the spring), and induces current in a nearby electromagnetic coil.

There has been much recent work on modelling the behavior of the spring reverberator digitally for use in a music production environment. The first attempts used an optimisation method to fit the dispersion curve of the spring with a number of allpass filters, and then used these filters within a wave-guide structure [2]. More recent work has modelled the vibration of the spring using finite difference methods [3, 4]. Attempts have also been made to produce a parametric digital spring reverberator which more closely resembles the structure of a traditional digital reverberator [5]. It is this final approach to modelling which we consider in this work.

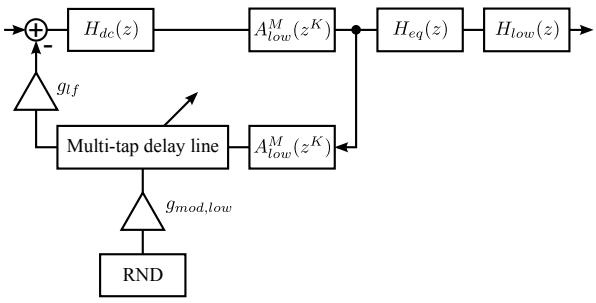

Figure 1: Low-frequency structure of parametric spring reverberation effect, based on [5].

In this work we propose a method for automatically deriving the characteristics of a spring reverb unit from a recorded spring impulse response. These characteristics are used to tune the parameters of the digital spring reverb model, much like previous work has proposed an analogous method that allows the fitting of a digital reverberator to a specific room response [6].

The paper is organised as follows. The parametric spring reverb model used for the automated calibration is briefly introduced in Section 2. Section 3 describes the proposed automated calibration methods for the parametric spring reverb model. The results of the automated calibration are discussed in Section 4. Section 5 concludes the paper.

2. THE PARAMETRIC SPRING REVERBERATION EFFECT

Basis for this work is the parametric spring reverberation effect introduced by Välimäki et al. [5]. Figure 1 shows a block diagram of the feedback structure used to produce the low-frequency chirp sequence of the spring reverberation effect (cf. Figure 2). The filter $H_{\mathrm{dc}}(z)$ is a dc blocking filter with a cutoff frequency at 40 Hz [5]. The filter $A_{\mathrm{low}}^{M}(z^K)$ is a spectral delay filter [7] consisting of M cascaded allpass filters. Each allpass filter section is an interpolated stretched allpass filter, which is composed of a Schroeder allpass filter with an embedded delay line of $K-1$ samples and a first-order fractional-delay allpass filter to implement a delay equal to 1 plus the decimal part of K (in samples). The impulse response of the spectral delay filter imitates the first low-frequency chirp appearing in the response of a spring reverb unit.

* This work has been supported by the [MIDE program] of Aalto university, the Helsinki Graduate School in Computer Science and Engineering (HeCSE), and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [203636].

† This work has been supported by the Academy of Finland, project no. [122815].

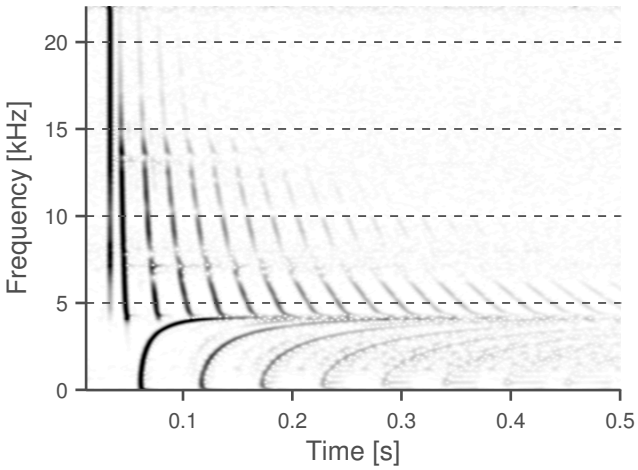

Figure 2: Spectrogram of one spring of the Leem Pro KA-1210 spring reverb. The spectrogram is normalised for each frequency bin, to improve the visibility of the chirp structure.

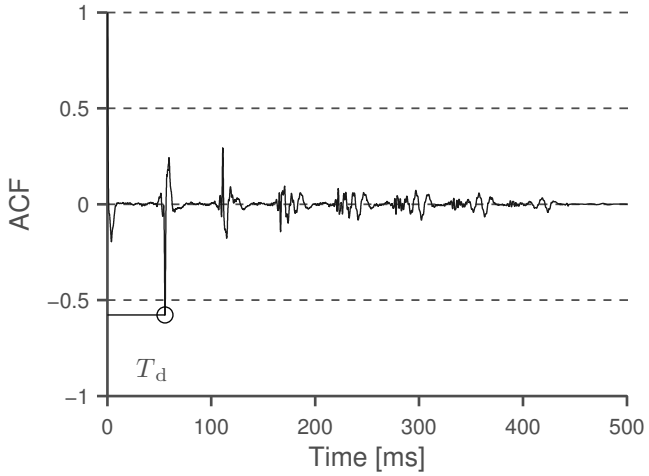

Figure 3: Autocorrelation function. The pulse delay is estimated at the maximum absolute value, $T_\mathrm{d} = 55.6$ ms.

A multi-tap delay line models the propagation and scattering of waves in a helical spring. It contains a long delay line with several extra output taps, which introduce pre-echos in the temporal response and small variations in the loop gain of the system [5]. A random modulation is applied to the delay line to introduce blurring of the response over time [5]. Finally, the output of the model is processed with two filters, $H_\mathrm{eq}(z)$ and $H_\mathrm{low}(z)$, which are a second-order resonator and a low-pass filter, respectively. An allpass cascade A_low is inserted into the feedback path to model the dispersion of each reflection when traversing the spring backwards, as proposed by Parker et al. [8].

To produce the high-frequency chirp sequence in the impulse response of the spring reverberator, a feedback structure similar to the one depicted in Figure 1 is used in the model, as proposed in [5]. It is considered less perceptually important [5]. The allpass cascade in the feedback path is omitted for the high-frequency feedback structure, to reduce the computational complexity of the model. The reader is referred to [5] for a more detailed description of the parametric spring reverb model.

3. AUTOMATED CALIBRATION

The starting point for the automated calibration of the parametric spring reverberation effect is the spectrogram of the impulse response to be modelled. Figure 2 shows the spectrogram of one spring of the Leem Pro KA-1210 spring reverb. It is obtained via an 8192-point short-time Fourier Transform (STFT) using a Blackman window with a hopsize of 8 samples. The spectrogram is normalised at each frequency bin, to enhance the visibility of the high-frequency chirp structure.

The calibration of the model is performed in three steps: First, the pulse delay of the low-frequency chirp structure and the transition frequency are determined. These are the perceptually most important parameters of the model [4]. In the following steps, the parameters for the low- and high-frequency chirp structures are determined separately.

3.1. Determine pulse delay and transition frequency

3.1.1. Pulse delay

The time domain representation of a spring reverb impulse response is dominated by the low-frequency pulse sequence, which contains most of the energy [5]. The pulses recur at regular intervals, given by the pulse delay T_d. It can be derived from the maximum absolute value of the autocorrelation of the impulse response (see Figure 3). For the given impulse response, T_d is estimated as 55.6 ms.

3.1.2. Transition frequency

The transition frequency F_c is defined as the cutoff frequency of the low-frequency pulse series [4]. As can be seen from the spectrogram in Figure 2, the low-frequency pulses overlap in time around the cutoff frequency, and thus cannot be distinguished from one another.

To determine F_c, the normalised autocorrelation of the spectrogram is calculated for each frequency bin along the time axis. The spectrogram can be represented as a matrix $\mathbf{S}$, with rows corresponding to frequency bins and columns corresponding to time instants. The autocorrelation of each row in $\mathbf{S}$ is calculated and normalised to one at lag zero. The result of this calculation is shown in Figure 4 (left). As can be seen, Figures 2 and 4 exhibit a similar pulse structure: For each frequency, the delay caused by dispersion and propagation in the spring is constant between pulses. The autocorrelation around the transition frequency contains no distinct peaks, as around this frequency the pulses overlap in time. Since the autocorrelation is normalised to one at lag zero, the mean value of the autocorrelation calculated over time at each frequency bin is a measure for the periodicity of the impulse response with respect to frequency. The mean value exhibits a peak at the transition frequency, where no distinct pulses are visible in the spectrogram, i.e., the periodicity is lowest (see Figure 4, right). Here, F_c is estimated as 4216 Hz. The method works well also for impulse responses with more than one spring, if the springs share a transition frequency.

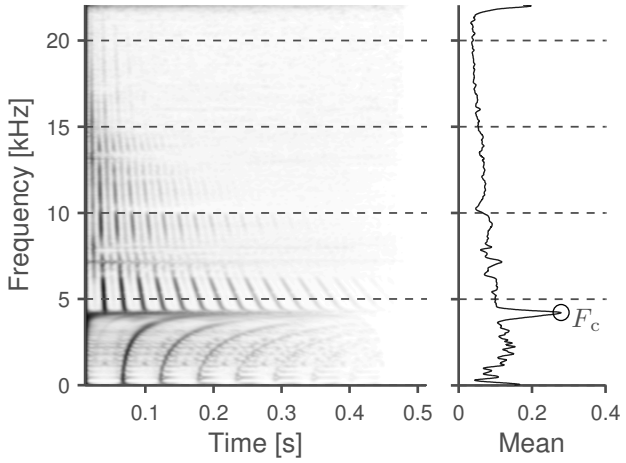

Figure 4: Normalised autocorrelation of spectrogram and mean of autocorrelation. The transition frequency is estimated at the maximum (○), $F_c = 4216$ Hz.

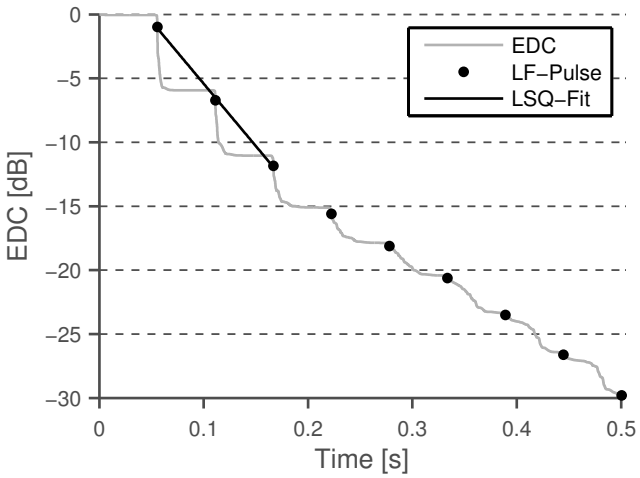

Figure 5: Energy decay curve (grey line) and low-frequency pulse positions (black dots). A least-squares fit (black line) indicates a decay rate of about -5.4 dB per pulse, during the early decay from 0 dB to -10 dB.

3.2. Low-frequency chirps

3.2.1. Gain factor

Figure 5 depicts the energy decay curve of the original impulse response. It is obtained via a backward integration of energy, the so-called Schroeder integration [9]. The black dots are T_d-spaced, indicating the low-frequency pulse positions, which dominate the impulse response [5]. In the parametric model used in this paper [5], the energy decay rate of the impulse response is modelled by applying a constant gain factor g_lf to the pulse series (cf. Figure 1). It determines the attenuation of each reflected pulse with respect to the previous. The value of g_lf is estimated by fitting a line through the energy decay curve.

For the perceived reverberance, the early decay time (EDT) from 0 dB to -10 dB is considered particularly important [10]. To model the EDT, a line is fitted through the energy decay curve at the pulse positions, from the first pulse to the pulse where the impulse response energy decays below -10 dB. The fitting is implemented via the polyfit function in Matlab (cf. black line in Figure 5). The decay per pulse d in dB is obtained as

$$d = mT_\mathrm{d}, \tag{1}$$

where m is the slope of the fitted line and T_d is the pulse delay. From the pulse decay d, the gain factor g_lf is obtained via

$$g_\mathrm{lf} = K10^{(\frac{d}{20})}, \tag{2}$$

with

$$K = \mathrm{sgn}\left\{\max[ACF(x)]\right\}, \tag{3}$$

i.e., the sign of the maximum of the autocorrelation function ACF of the spring impulse response x (cf. Figure 3). To compensate for additional attenuation of the pulses introduced by the modulation and linear interpolation in the delay line, the gain factor g_lf is multiplied by a constant of 1.2. For the given impulse response, this yields $g_\mathrm{lf} = -0.64$.

3.2.2. Spectral delay filter calibration

A crucial step in the calibration of the model is determining the parameters for the spectral delay filter. The goal is to tune the output of the filter to match the dispersion characteristics of the real spring. The transfer function of the stretched interpolated allpass cascade used in the model [5] to generate the chirps is given by

$$A_M(z) = M\frac{a_1 + A_\mathrm{fd}(z)z^{-K_1}}{1 + a_1 A_\mathrm{fd}(z)z^{-K_1}}, \tag{4}$$

where A_fd is a fractional delay allpass filter with

$$A_\mathrm{fd}(z) = \frac{a_2 + z^{-1}}{1 + a_2 z^{-1}}, \tag{5}$$

and

$$K_1 = \mathrm{round}(K) - 1. \tag{6}$$

M is the number of allpass sections in the cascade, a_1 is the allpass coefficient, and K the stretching factor determining the cutoff frequency of the chirps. The fractional delay filter is necessary to implement a nonintegral stretching factor K, to accurately obtain the desired transition frequency of the chirp structure. The parameter a_2 can be derived from the nonintegral part of the stretching factor (for details, see [5]). The parameters M, a_1 and K of the allpass cascade are obtained by iteratively fitting the allpass cascade to the first chirp in the spectrogram. The procedure is applicable to spring reverbs with one or more springs, therefore the general case of N springs is considered in the following.

First, peaks in the spectrogram are extracted at each frequency bin, up to the transition frequency. After low-passing the rows of the spectrogram matrix $\mathbf{S}$, which correspond to frequency bins (cf. Section 3.1.2), the peak locations at each frequency bin are detected as zero crossings of the derivative of each row. Next, the peaks in $\mathbf{S}$ are grouped to connected segments, by identifying sequences of peaks that form a connected line in the spectrogram. Finally, these peak segments are grouped to chirps in an iterative process, using the output of the allpass cascade as a model for the chirps:

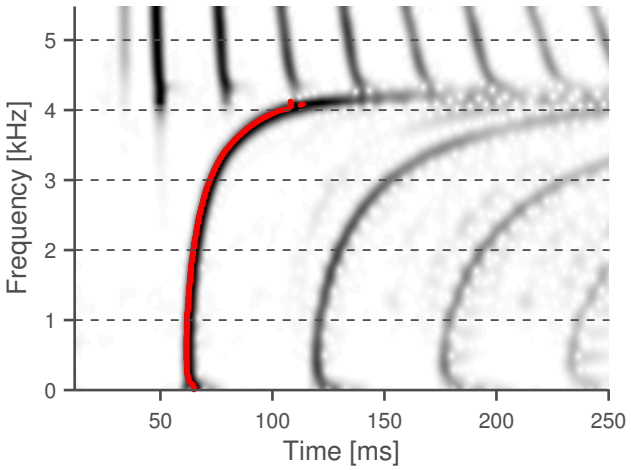

Figure 6: The first chirp (red) is identified through iterative fitting of an allpass cascade to peak segments of the spectrogram.

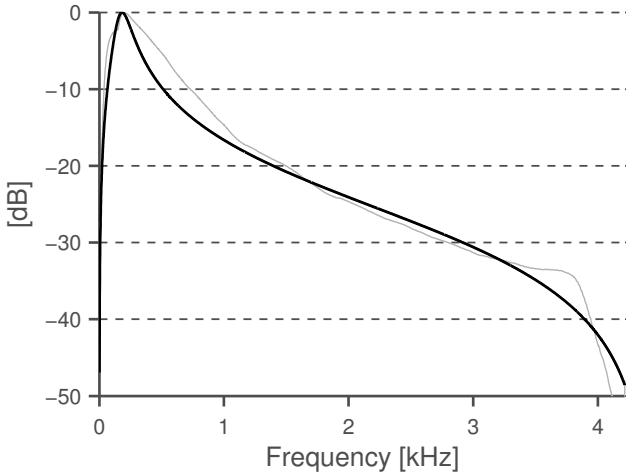

Figure 7: Spectrum of first pulse (light grey) and fitted equalisation filter: $F_{\mathrm{peak}} = 183Hz$, $B = 146Hz$, $K = 5$.

1. Fit the output of an allpass cascade to the longest peak segment in the spectrogram, using nonlinear least squares fitting via the lsqnonlin function in Matlab. This yields a rough estimate $\hat{c}_i$ of the i-th chirp in the spectrogram.
2. Find peak segments that lie close to $\hat{c}_i$, update $\hat{c}_i$ to include those segments and re-fit the allpass cascade to the new $\hat{c}_i$.
3. Repeat step 2 while peak segments close to $\hat{c}_i$ are found.
4. Remove all peak segments allocated to any chirp estimate $\hat{c}$, and repeat steps 1–3, until all chirps $\hat{c}_{1\ldots\mathrm{N}}$ are extracted.

This process identifies the chirps in the spectrogram produced by the N springs and directly yields the parameters of the allpass cascades to model their estimates $\hat{c}_{1\ldots\mathrm{N}}$. The pulse delay $T_{\mathrm{d},i}$ of each chirp is given as

$$T_{\mathrm{d},i} = T_{\mathrm{d}} + 2\Delta T_{0,i}, \tag{7}$$

where T_{d} is the pulse delay (cf. Section 3.1.1), and $T_{0,i}$ is the offset of the i-th chirp from the first chirp visible in the spectrogram. The result of the chirp extraction for the given impulse response is shown in Figure 6.

3.2.3. *Chirp equalisation*

The spectrum of the first chirp is obtained by calculating the Fourier Transform of the impulse response from 0 to $2T_{\mathrm{d}}$. To approximate the spectral shape of the chirp, a second-order IIR filter is used. Its transfer function is stretched by replacing the unit delays of the filter structure with a delay line of length K [5]. The filter parameters consist of the stretching factor K, the frequency F_{peak} of the peak in the transfer function, and the -3 dB bandwidth B of the peak. Nonlinear least-squares fitting is used to fit the frequency response of the equalisation filter to the spectrum of the first chirp (see Figure 7).

3.3. High-frequency chirps

The high-frequency chirp sequence is considered less perceptually important [5]. Therefore, a simpler approach towards calibration can be taken than for the low-frequency structure. Based on the assumption of a regular pulse structure, the chirps can be extracted from the autocorrelation of the spectrogram matrix $\mathbf{S}$ (cf. Section 3.1.2). This procedure emphasises periodic energy components in the spectrogram, whilst suppressing nonperiodic components, such as the first pulse appearing in the spectrogram above 5 kHz (see Figure 2).

To detect the first chirp of the high-frequency pulse series in the autocorrelation of the spectrogram, the locations of the two largest peaks are determined for each frequency bin, using the method described in Section 3.2.2. To eliminate all peaks not belonging to the periodic pulse series, a simple check is performed: All peaks in the autocorrelation belonging to a periodic chirp in the pulse series must have a corresponding peak at twice the delay (cf. Figure 8). The peaks in the autocorrelation passing this check indicate the dispersion and delay induced to a periodic pulse after traversing the spring twice: Dividing the locations of the peaks by two yields the form of the first chirp in the spectrogram, at half the delay between adjacent chirps. It is modelled by fitting an allpass cascade via nonlinear least-squares fitting.

For simplicity, the decay rate of the high-frequency chirps is obtained by multiplying the decay rate of the low-frequency chirp sequence with a constant factor. Based on inspection of the spectrograms of spring reverb impulse responses, we chose 1.3 for the factor, assuming that the decay rate per pulse is about 30% lower for the high-frequency than for the low-frequency chirp sequence. This yields a high-frequency gain factor $g_{\mathrm{hf}} = -0.83$.

4. DISCUSSION

Figure 9 allows a visual comparison of the spectrogram of a real spring reverb unit and the spectrogram of the digital parametric model after automated calibration.

The main perceptual parameters, i.e., the transition frequency and pulse delay of the low-frequency chirp sequence are modelled quite accurately. The form of the low-frequency chirps is captured well, although the group delay close to the transition frequency appears to be larger in the real impulse response than in the model. As a result of the linear approximation of the early decay time (EDT, cf. Figure 5), the energy of the low-frequency chirps decays

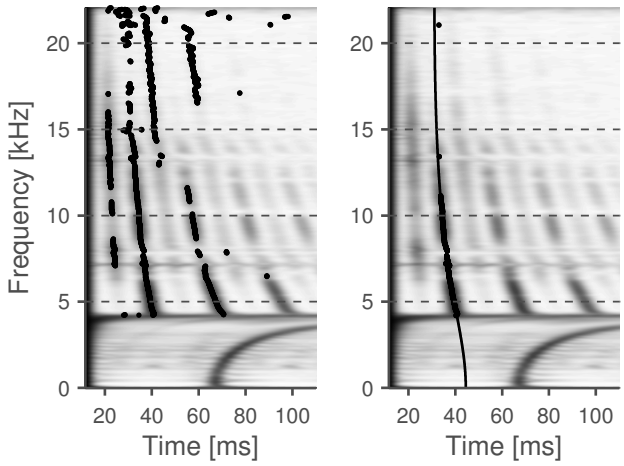

Figure 8: Left: First and second peaks in normalised autocorrelation of spectrogram. Right: Peaks of first pulse after outlier removal (○) and fitted allpass cascade (–): $a_1 = -0.34$, $M = 189$.

slower in the real response than in the model. An auditory comparison reveals a marked difference in timbre between the real and modelled low-frequency chirp sequence. This may partly be due to the coarse approximation of the chirp spectrum (cf. Figure 7), which results in a discrepancy between the desired and modelled spectrum, particularly below 1.5 kHz and above 3.5 kHz. Using a higher-order equalisation filter in the digital model might help tackle this problem. Furthermore, the acoustic quality of the increasing diffuseness of successive chirps could be modelled more accurately, for example by replacing the delay line modulation with an automatically calibrated digital reverberator [5].

The main characteristics of the high-frequency chirp structure are successfully reproduced by the model. The form of the first high-frequency chirp is modelled accurately. The following chirps are more strongly dispersed in the real impulse response than in the model. To model the dispersion characteristics more accurately, an allpass cascade could be inserted in the feedback path of the high-frequency structure of the model. It is omitted here since it is not considered perceptually important and reduces the computational load of the model considerably [5]. The decay rate of the chirps seems to be slightly lower in the real impulse response than in the model. This is a result of the linear approximation of the decay rate in the parametric model.

The automated calibration was performed without constraints in terms of the computational complexity of the digital parametric model. The computational load is dominated by the allpass chains. To lower the computational load of the model, an upper limit can be set to the length of the allpass chains fitted by the optimisation algorithms described in Sections 3.2.2 and 3.3. As an example, we set the maximum length of the low-frequency allpass cascade to $\max\{M_{\mathrm{low}}\} = 100$. Table 1 presents an overview of the calibrated model parameters for the Leem Pro KA-1210 spring reverb. The first column contains parameter values obtained through manual and semi-automated calibration [5]. The middle column contains the values obtained using the automated calibration proposed in this paper, without constraints in terms of computational complexity. The last column presents the values obtained with an upper limit on the length of the low-frequency allpass chain. There

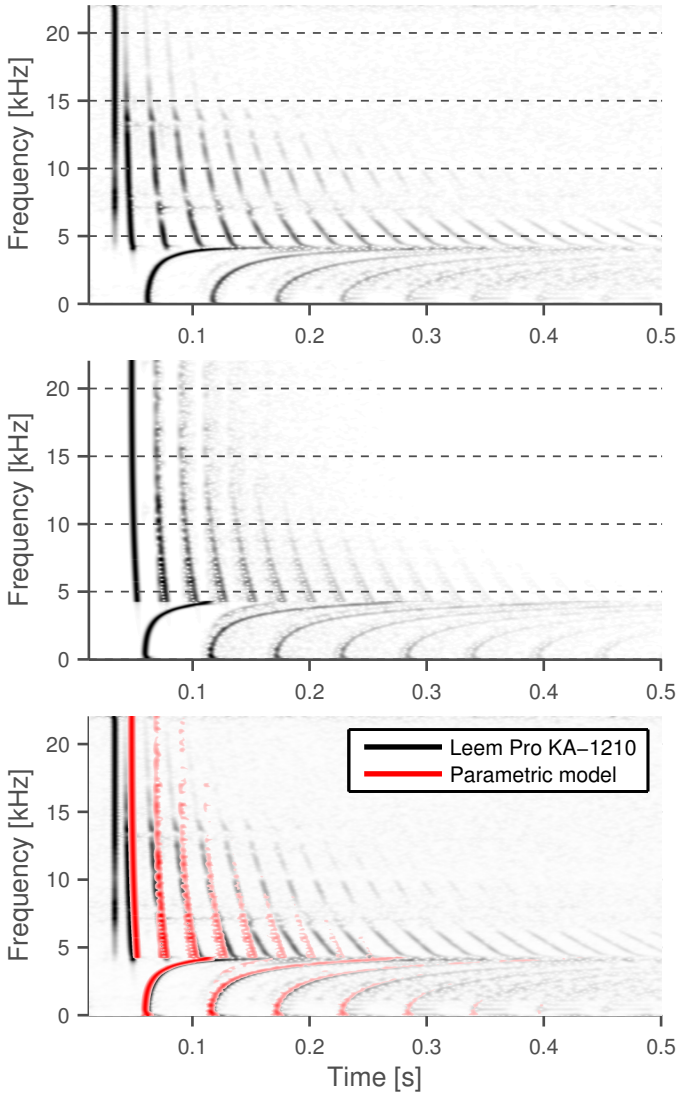

Figure 9: Spectrogram of the impulse response of the Leem Pro KA-1210 (top) and of the parametric model with automated calibration (middle). The bottom graph shows the model spectrogram as an overlay (red) onto the spectrogram of the Leem Pro KA-1210 (greyscale).

is good correspondence between the manual and automated calibration of the values T_{d} and F_{c}. The form of the modelled chirps is determined by the parameters a and M, as well as $F_{\mathrm{c,lf}}$ in case of the low-frequency chirp structure. If an upper limit is set to M, the values of a and $F_{\mathrm{c,lf}}$ change such as to obtain an optimal fit of the modelled chirp. The spectrogram of the model with limited allpass cascade length is shown in Figure 10. A visual comparison reveals no substantial differences to the spectrogram of the model with unconstraint cascade length, indicating that the restriction is a viable option to reduce computational load without major perceptual impact. However, the form of the first chirp is modelled less accurately with the reduced allpass chain length. Figures 11 and 12 demonstrate the usage of the automated calibration for spring

	Leem Pro KA-1210, Spring 1		
	Manual [5]	Automated	Automated*
T_{d}	0.056	0.056	0.056
F_{c}	4300	4216	4216
$F_{\mathrm{c,lf}}$	4300	4980	4275
M_{low}	100	318	100
a_1	0.62	0.69	0.63
g_{lf}	-0.8	-0.64	-0.64
M_{high}	200	189	189
a_{high}	-0.6	-0.34	-0.34
g_{hf}	-0.77	-0.83	-0.83

* With $\max\{M_{\mathrm{low}}\} = 100$.

Table 1: Comparison between manual and automated calibration of the parametric spring reverberation effect. The transition frequency F_{c} refers to the transition frequency of the impulse response, whereas $F_{\mathrm{c,lf}}$ refers to the transition frequency of the fitted allpass cascade. The optimised calibration results were obtained by setting an upper limit to M_{low}.

reverbs with two and three springs. In both cases the main features of the spectrogram of the real impulse response are captured successfully by the proposed automated calibration, and reproduced by the parametric model. However, the modelling accuracy is inferior to the single-spring example, since the parameters of each spring are extracted from a single response of the whole unit.

5. CONCLUSION

In this paper, an approach was proposed to automatically calibrate parameters of a spring reverberation model. A slight modification of a previously presented spring reverberation model is used [5], which produces impulse responses containing the same basic features appearing in responses measured from spring reverb units.

Signal processing methods to estimate the values of several parameters of the spring reverb model were suggested. All estimations are based on a measured impulse response of a spring reverberation unit or its spectrogram. First, the time delay between repetitive pulses appearing at low frequencies was estimated by detecting the first peak in the absolute value of the autocorrelation function of the measured response. Next, the transition frequency, which corresponds to the cutoff point of the low-frequency chirp sequence, was estimated from the autocorrelation function of the spectrogram. The maximum of the mean of the autocorrelation function appeared to indicate the transition frequency. The value of the feedback gain factor was computed from the early decay time, which we estimated as the difference of time instants, where the energy decay curve passes the -10 dB level.

A spectral delay filter consisting of a chain of first-order allpass filters imitates the shape of the first chirp in the impulse response. The number of cascaded allpass filters and their filter coefficient value, which is the same for all filters in the cascade, were chosen by fitting the output of the allpass cascade via an iterative procedure employing the nonlinear least squares method. The spectral delay filter was equalised by fitting the parameters of a resonant second-order filter to the magnitude spectrum of the chirp, using the nonlinear least-squares method.

Model parameters for the high-frequency chirp sequence are also extracted from the autocorrelation of the spectrogram, although the data is noisier than that containing the low-frequency chirps. No equalisation is performed for the high-frequency chirps.

The proposed calibration methods can be applied to a response of a single-spring unit or to one with several parallel springs. The calibrated parameter values were compared against the manually and semi-automatically calibrated values presented in [5]. It was observed that slightly different but similar values are obtained. The optimal number of allpass sections in the spectral delay filter for the low-frequency chirp can become very large, such as about 300, leading to a high computational load in the implementation. For this reason a constrained optimisation was tested in which the maximum number of filter sections is limited to 100. This can lead to a sufficiently good fit and to a reasonable number of filtering operations per sample.

Recently, Parker has proposed multirate and subband techniques to reduce the computational cost of the parametric spring reverberation model [11]. These ways to improve the computational efficiency are suggested to be applied after parameter values have been calibrated using methods proposed in this paper.

6. REFERENCES

[1] L. Hammond, "Electrical musical instrument," US Patent No. 2230836, 1941.

[2] J. S. Abel, D. P. Berners, S. Costello, and J. O. Smith, "Spring reverb emulation using dispersive allpass filters in a waveguide structure," in *Audio Engineering Society Convention 121, San Francisco*, 2006.

[3] S. Bilbao and J. Parker, "A virtual model of spring reverberation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 4, pp. 799 –808, 2010.

[4] J. Parker and S. Bilbao, "Spring reverberation: A physical perspective," in *Proceedings of the International Digital Audio Effects Conference, Como, Italy*, 2009.

[5] V. Välimäki, J. Parker, and J. S. Abel, "Parametric spring reverberation effect," *Journal of the Audio Engineering Society*, vol. 58, no. 7/8, pp. 547–562, 2010.

[6] J.-M. Jot and A. Chaigne, "Digital delay networks for designing artificial reverberators," in *Audio Engineering Society Convention 90, Paris*, 1991.

[7] V. Välimäki, J. S. Abel, and J. O. Smith, "Spectral delay filters," *Journal of the Audio Engineering Society*, vol. 57, no. 7/8, pp. 521–531, 2009.

[8] J. Parker, H. Penttinen, S. Bilbao, and J. S. Abel, "Modeling methods for the highly dispersive slinky spring: A novel musical toy," in *Proceedings of the International Digital Audio Effects Conference, Graz, Austria*, 2010.

[9] M. R. Schroeder, "New method of measuring reverberation time," *The Journal of the Acoustical Society of America*, vol. 37, no. 3, pp. 409–412, 1965.

[10] J.S. Bradley, "Review of objective room acoustics measures and future needs," *Applied Acoustics*, vol. 72, no. 10, pp. 713 – 720, 2011.

[11] J. Parker, "Efficient dispersion generation structures for spring reverb emulation," *EURASIP Journal on Advances in Signal Processing*, vol. 2011, 2011, Article ID 646134, 8 pages.

7. APPENDIX

This appendix presents the spectrograms of the real impulse response and of the parametric model with automated calibration for spring 1 of the Leem Pro KA-1210 spring reverb (see Figure 10), a no-name spring reverb containing two springs (see Figure 11), and a Mesa Boogie spring reverb containing three springs (see Figure 12). Sound samples are available for download at `http://www.tml.tkk.fi/~hannes/DAFx2011/`.

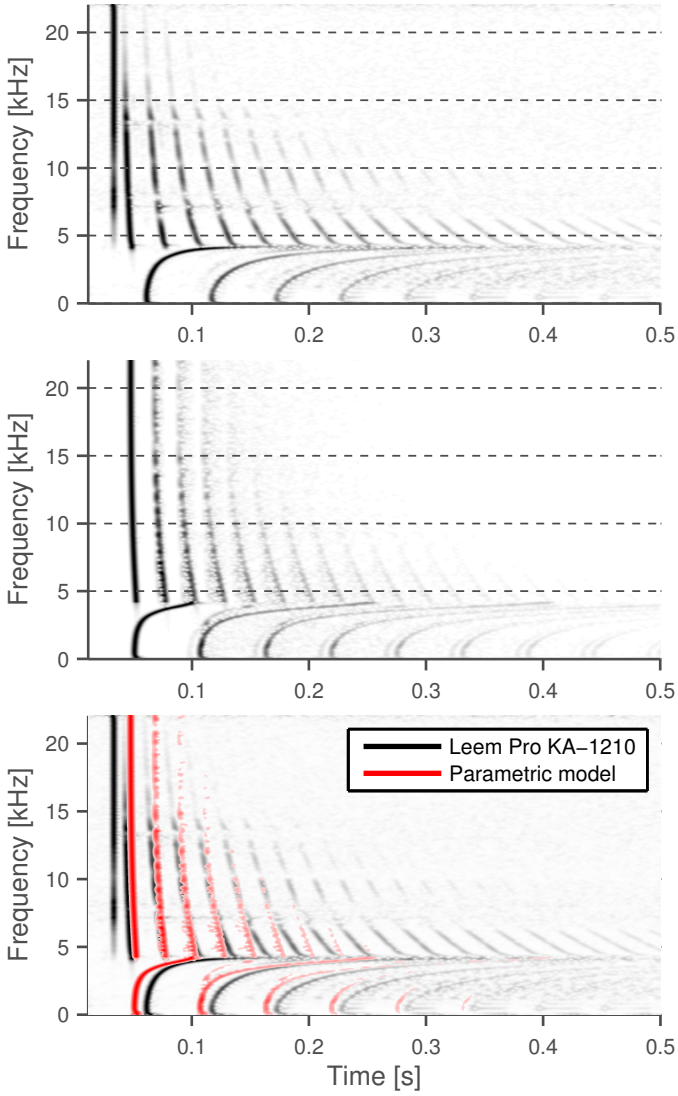

Figure 10: Spectrogram of impulse response of spring 1 of Leem Pro KA-1210 (top) and of the parametric model with automated calibration (middle), with length of the low-frequency allpass cascade limited to $\max\{M_{\text{low}}\} = 100$. The bottom graph shows the model spectrogram as an overlay (red) onto the spectrogram of the real unit (greyscale). Although the main perceptual aspects of the low-frequency chirp sequence are captured, the limited length of the allpass chain deteriorates the modelling accuracy of the form of the low-frequency chirps. This is a trade-off for reducing the computational complexity by limiting the length of the allpass cascade.

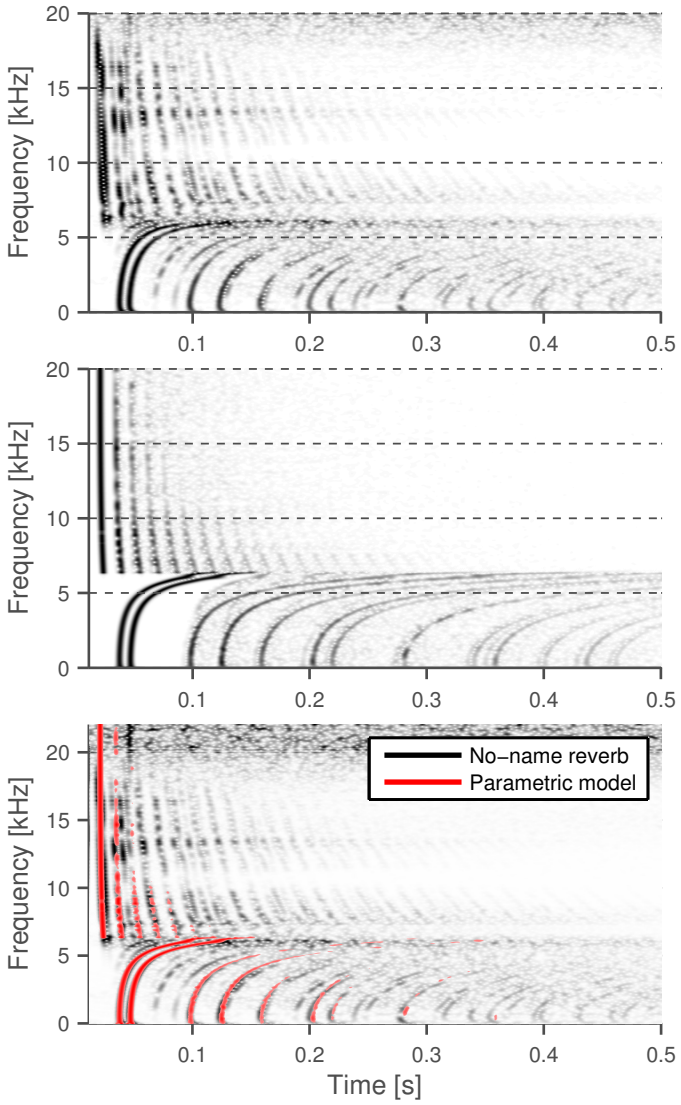

Figure 11: Spectrogram of impulse response of a no-name spring reverb with two springs (top) and of the parametric model with automated calibration (middle). The bottom graph shows the model spectrogram as an overlay (red) onto the spectrogram of the real spring reverb unit (greyscale). The main features of the real impulse response are captured well. However, additional low-frequency chirp reflections and some details of the high-frequency chirp sequence are not reproduced by the model. The modelling accuracy could presumably be improved if individually measured responses of both springs were used for the automated calibration. The same holds for the model response shown in Figure 12. It is more difficult to derive parameters of all springs from a single response measured of the whole unit than using individually measured responses for each spring.

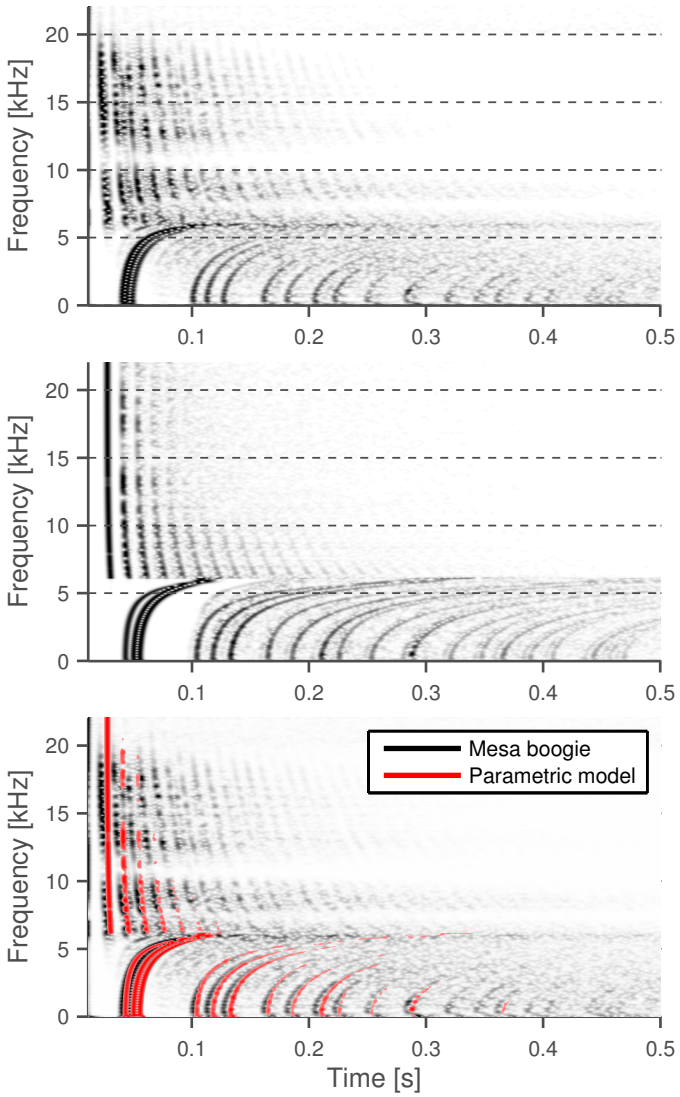

Figure 12: Spectrogram of impulse response of Mesa Boogie spring reverb with three springs (top) and of the parametric model with automated calibration (bottom). The bottom graph shows the model spectrogram as an overlay (red) onto the spectrogram of the real spring reverb (greyscale). All three low-frequency chirp sequences are modelled relatively accurately, although there are discrepancies regarding the form of the first modelled chirps and the details of the high-frequency chirp sequence.