A LOUDSPEAKER MANAGEMENT SYSTEM WITH FIR/IIR FILTERING

RAINER THADEN[1], SWEN MÜLLER[2], GOTTFRIED BEHLER[3], ANSELM GOERTZ[1], MICHAEL MAKARSKI[1], JOCHEN KLEBER[1]

[1] *Four Audio GmbH & Co. KG, Aachen, Germany*
info@four-audio.com
[2] *Divisão de Acústica, INMETRO, Xerém, Brazil*
smuller@inmetro.gov.br
[3] *Institute of Technical Acoustics, RWTH-Aachen University, Germany*
gkb@akustik.rwth-aachen.de

Most digital loudspeaker processor systems use cascaded IIR filters to perform the crossover function and system equalization. This work presents a multirate platform which allows FIR and IIR filters to be freely combined in each output channel. The pros and cons of both filter types are discussed and some techniques to obtain suitable FIR filter coefficients are described. Dual-range AD conversion satisfies the need for high dynamic input range, while look-ahead peak-limiters in the outputs avoid clipping and mechanical damage without introducing noticeable distortion.

INTRODUCTION

Besides improved audio quality over analog solutions, digital loudspeaker management systems (DLM) offer more flexible audio processing possibilities. Most systems currently on the market use IIR filters to realize the crossover and equalization (EQ) functions. This paper presents a DLM which takes advantage of FIR filters, like controllable group delay, while also offering traditional IIR solutions. Furthermore, the device is equipped with separate peak and RMS limiters and circumvents the input dynamic range bottleneck by using dual-range AD conversion. Techniques for suitable FIR coefficient generation based on loudspeaker transfer-function measurements are shown as well as an approach for integrating equalization for rooms and larger venues into the FIR filters. The limiters and the AD conversion are illustrated in detail.

1 ARCHITECTURE

The motherboard of the DLM was designed to be configurable regarding its input and output sections. Due to its flexibility and small size of 20 x 13 cm (or 11 x 13 cm for a 2-in-4 version), it can be mounted inside the cabinet of an active loudspeaker or in a 19" housing to form a complete DLM.
The input section consists of up to three dual-range or six single-range AD converters, one AES/EBU digital input, and one I²S input as an interface to ADAT, Ethersound, Cobranet, e.g. A sample rate converter can be inserted for input rates other than 96 kHz or to circumvent jitter problems. The output section consists of eight analog and one AES/EBU output. An Ethernet and RS232 port allow remote control.

The signal processing is performed by two Motorola 56321 DSPs with less than 0.4 W power consumption each. They are fed from a flexible routing matrix. Figure 1b shows the signal processing. For the input section, a parametric EQ bank and a main delay block is provided. Four output channels are processed in each DSP. Each output signal chain comprises a delay block for inter-driver arrival time adjustments and optional resampling filters with a sample rate reduction factor up to 32. These can be activated in conjunction with a crossover/EQ-FIR filter, which is handled by the DSP56321 coprocessor (EFCOP). Each signal path also contains an IIR filter bank to allow for IIR crossover and EQ filters and, of course, an individual gain.
At the end of the chain, two separate limiters are provided. A quickly reacting peak limiter is provided to prevent amplifier clipping and mechanical damage of the drivers An RMS limiter with slow reaction time models the loudspeaker's voice coil temperature to protect it from overheating.
Requantization stages from 48 to 24 bit with dither and selectable noise shaper [2] are active behind the input PEQ filter bank and at each channel output.

1.1 48 or 96 kHz ?

The DLM can be operated at either 48 or 96 kHz sampling rates. Regarding audibility, it can be stated that the human auditory system is not capable of perceiving ultrasound. Ultrasound components in the audio signal seem to be detectable only if their intermodulation products fall into the audible band [1]. High-efficiency sound reinforcement systems with their compression drivers and wave-guides are not suited for reproducing ultrasonic frequencies and most digital

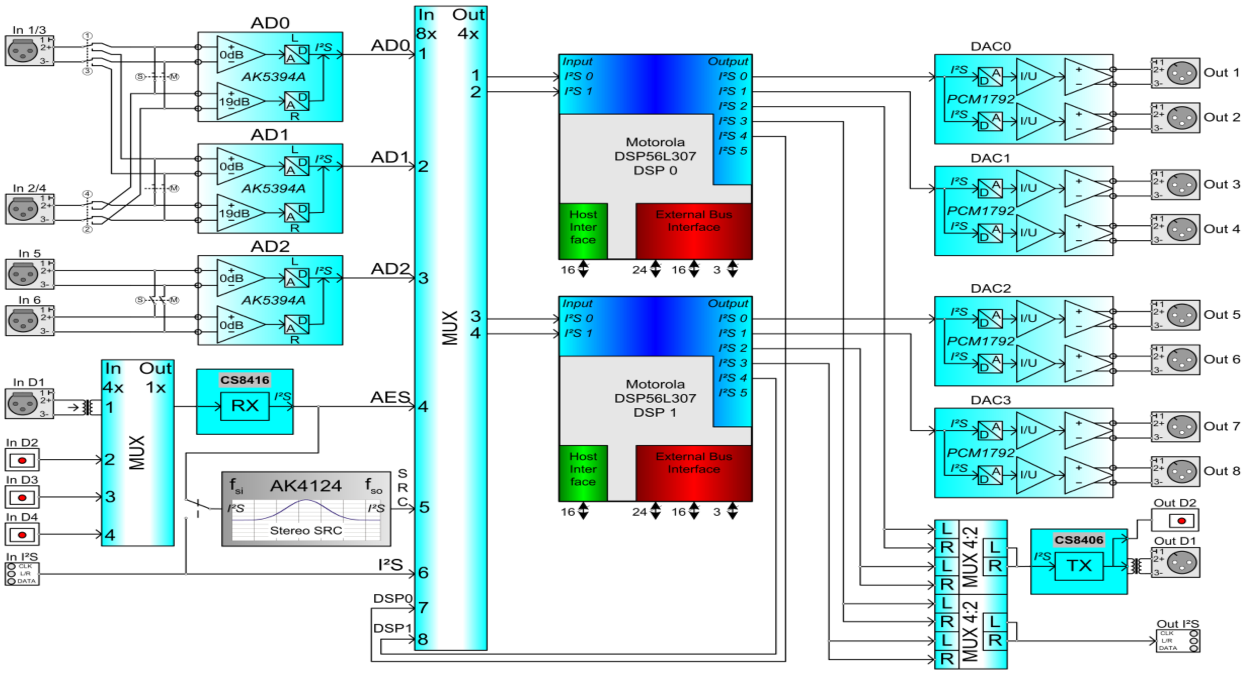

Figure 1: Signal processing components of DLM

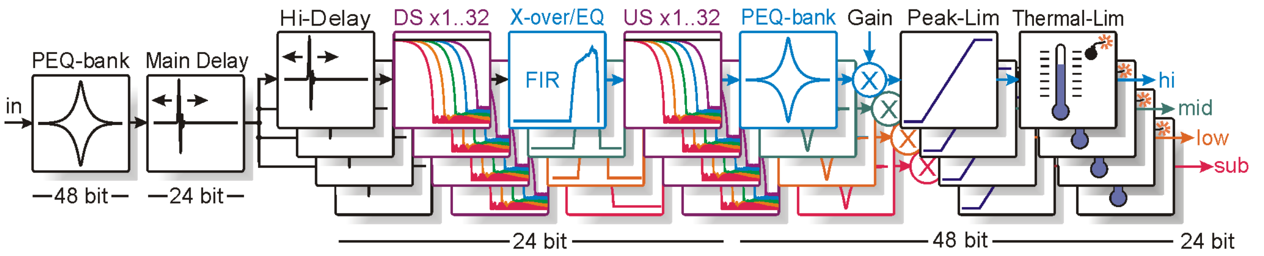

Figure 1b: Signal processing in one DSP for one input and a group of four output channels

mixing desks for live sound run at 48 kHz. So at first glance, for live sound reinforcement, which is the main focus of the presented DLM, 96 kHz seem to be dispensable. In contrast, it has become the standard sampling rate in most studios applications.

There are also a couple of technical points in favor of the higher sampling rate. Firstly, noise shapers in requantization stages can shift most of the quantization noise into the inaudible band above 20 kHz.

Moreover, the transfer functions (TF) of IIR filters generated with bilinear transforms without corrections come closer to the shape of their analog counterparts.

The most important argument for 96 kHz is lower latency. Most anti-alias filters in converters have symmetrical impulse responses (IR), causing a delay of typically 1.4 to 2ms at 48 kHz for the AD/DA chain. When the sampling rate is doubled, this value drops by one half. Additionally, some converters, as the DAC PCM1792A used in this project, allow using a relaxed anti-alias filter which reduces the latency for the latter to about 187 µs at 96 kHz.

2 IIR FILTERS

IIR filters are the most commonly used type for equalization and crossover filters. They emulate analog filters and define the frequency and phase response by setting poles and zeroes. Their phase response is related to the frequency response magnitude.

Commonly, filters of second order with two delays and five coefficients (biquads, see Figure 3) are used. They are able to approximate almost all common filters such as high-pass and low-pass with classical characteristics (Butterworth, Bessel, Linkwitz-Riley, etc.). Peak, shelving, and notch filters can also be formed. Filters of higher order are usually realized by cascading second order filters.

The overall transfer-function for one DLM output is composed by serially concatenating a variable number of biquads, realizing the desired crossover bandpass function and performing the passband equalization.

To manually program the individual filters, it might appear that the most straight-forward approach is to first set the desired crossover-frequencies and then equalize the passbands of each individual way. However, reversing the order is often more promising:

1) Flatten the response of each way in the intended passband frequency range and a little beyond.
2) Adjust the output channel gains to bring all passbands to the same level.
3) Now looking at the impulse responses in the time domain, adjust the polarities (0°/180°).
4) Adjust the output channel delays to make all IR peaks coincide.
5) Activate the crossover bandpass filters.
6) Match the phases in the crossover regions by means of all-pass filters.

This procedure avoids corrupting the carefully matched phases of adjacent bands by messing with in-band-PEQs *after* activating the crossover filters.

The whole procedure assumes that a fast-reacting FFT-based measurement system is available to monitor the adjustment results in realtime.

2.1 Implementation

In the DLM presented here, IIR filters can be used for the crossover function as well as for equalization. Instead of deriving the filter coefficients from pre-calculated tables, the DSP is fed with the real filter parameters (type, frequency in Hz, gain in dB, and quality factor) and performs a complete bilinear transform to obtain the digital filter coefficients. Since the DSP56000 family operates with fixed-point numbers, a floating-point library was implemented in order to be able to do the necessary calculations for the transform. A pair of 24-bit numbers forms a floating-point value, with the first one representing the mantissa and the second one the exponent. Cosine, tangent and logarithm calculations are carried out with the help of tables. This allows the user to set nearly every possible combination of filter parameters. The achieved precision of the filter TF, especially for tricky filters such as very narrow low-frequency notch filters, is much better than could be achieved by interpolating through pre-calculated table values.

Since the bilinear transform maps the frequency range of an analog filter (DC - 8) to that of the digital filter (DC - $f_S/2$), its frequency response features become increasingly compressed on the frequency scale when approaching the Nyquist frequency $f_S/2$. Particularly, the common Bell (Peaking)-filters appear with a smaller bandwidth than their analog templates. To remedy this, the quality-factor Q entered in the coefficient calculation can be decreased according to

$$Q_{COR} = Q_{ANA} \cdot \cos\left(p \, \frac{f_0}{f_S} \right) \qquad (1)$$

with f_0 being the center and f_S the sampling frequency. This produces a left wing of a Bell filter nearly identical to the analog model (Figure 2). The right wing has steeper decay because the gain of a common digital Bell filter drops to 0 dB at the Nyquist frequency. The *Orfanidis* method [3] to set a prescribed Nyquist frequency gain, extremely useful for digital filters operating with single audio sampling rates (48 kHz or below), was considered to be dispensable in a DLM operating with 96 kHz.

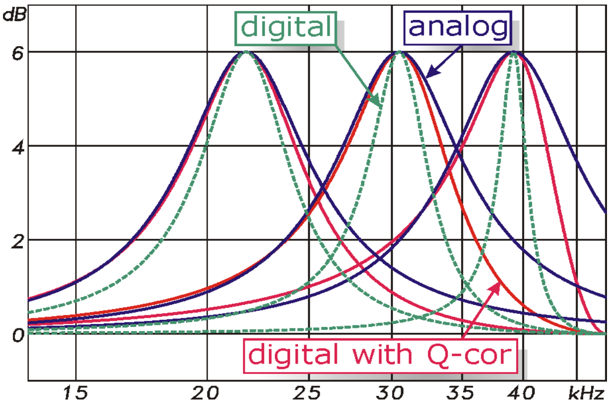

Figure 2: Transfer functions of analog and digital bell filters (6 dB gain, Q=4), with and without Q correction

Since filters of orders > 2 can be realized by cascaded biquads, the parameters frequency, gain, and quality are sufficient to also allow the calculation of these filters. Low-pass and high-pass as well as low-shelf and high-shelf can be modified in their characteristic by altering the quality factor Q.

With the PC remote software, it is also possible to calculate more sophisticated filters on the host and directly transfer the coefficients (instead of the parameters) to the DLM.

The biquad filters can be freely distributed from a common pool across the input and output paths until reaching 100% DSP load. It is thus possible to use them as crossover or EQ in any path.

The DSP shares its processing power between the realtime audio signal processing and the calculation of filter coefficients. The calculation is sliced into 25 pieces. The time needed to calculate the coefficients of one biquad is about 2.5 ms, so about 400 calculations per second are possible. This is fast enough to allow for smooth transitions when adjusting filters at the DLM.

2.2 Audio performance

Due to the feedback loops, the performance in terms of noise and distortion of an IIR filter is more critical than that of a FIR filter, especially on fixed-point platforms. Limit cycles inside the filter and requantization generate noise. A longer word length is therefore required to perform the calculations. Figure 3 shows the noise of an IIR filter on a fixed point DSP with word lengths of 24 bits, 24 bits with error feedback and 48 bits (double precision arithmetic). It can be clearly seen that only double precision *arithmetic* and *data* paths provide satisfying results in terms of the self-generated noise of the filter. In contrast, 24 bits for the *coefficients* provide enough precision to realize very narrow notch filters, even at low frequencies where they are most sensitive to coefficient quantization.

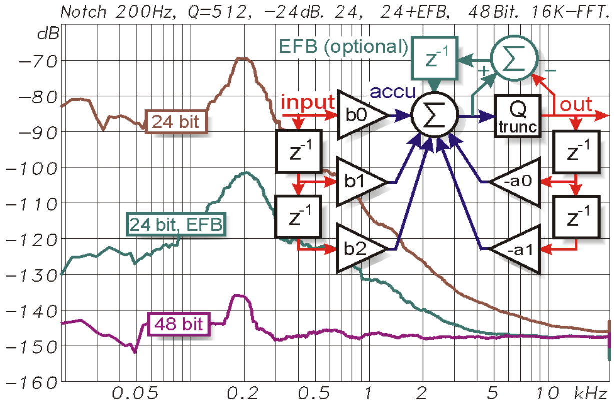

Figure 3: Topology of a biquad in direct form I with optional error feedback, and self-generated noise of a notch filter with single and double precision processing.

While 48-bit double-precision biquad calculation on a 24-bit fixed point DSP is cumbersome, involving shift-and-add operations of partial results and taking about three times more processor cycles to execute than a single-precision biquad, it compares favourably to 32-bit floating-point processing in terms of noise and distortion. For very faint signals, the noise floor of 32-bit float processing is typically a bit lower than that of 48-bit integer processing. However, the noise floor of floating-point processing scales with the signal amplitude, while the noise floor of integer processing remains on a very low fixed level for all signal amplitudes. This means that 48-bit integer processing produces considerably less noise and distortion than 32-bit floating-point processing for medium and high signal amplitudes.

3 FIR FILTERS

FIR filters allow a radically different approach for the crossover and system equalization task. While the desired TF of an IIR based DLM is usually realized with generic filter blocks concatenated from a building set, one single FIR filter can replace the whole IIR filter chain, performing the roles of the crossover and equalizer at once.

Since the IR of an FIR filter is identical to its coefficients, it is possible to realize filters with phase and group delay characteristics which can range from minimum-phase through linear-phase up to maximum-phase (which would be the time-inverted minimum-phase-filter). The necessary length and hence number of coefficients of a FIR filter to reach satisfactory spectral selectivity scales with wavelength. If not handled by overlapped FFT convolution techniques, an FIR filter dedicated for the mid, low or sub band has to be run at a reduced sample rate (dealt with in 3.4) to spare DSP resources.

3.1 Minimum phase

The behaviour of a minimum phase FIR filter closely resembles that of a chain of minimum-phase IIR filters as described in section 2, that is, phase distortion is introduced but with the benefit of the lowest physically possible overall group delay. A minimum phase FIR filter can be easily derived from the following two cases (3.2 and 3.3) by means of the Hilbert transform.

3.2 Linear phase

Linear phase filters used to realize the crossover function in the mid and high frequency ranges are an interesting alternative to the conventional IIR designs as they avoid any phase distortion introduced by the filter itself, regardless of the steepness of the transition band. This potentially eliminates problems with dips near the crossover frequency caused by phase mismatch between the two involved drivers. Since the impulse response of the filters is symmetric, they add an overall group delay of half the filter length to the system.

A linear phase FIR filter can be easily derived from the previous (3.1) and the next case (3.3) by setting the group delay of its TF to a constant value and performing an IFFT. Some DLM permit to use linear phase FIR filters to perform the crossover functionality, while the passband equalization is done conventionally with a chain of IIR-PEQs. This possibility also exits in the DLM presented in this work. However, the more typical application is the “one-FIR-filter-does-it-all” approach.

3.3 Complex-equalizing

Instead of opting for linear phase crossover/EQ filtering, it is even more enticing to let the FIR filter not only equalize the loudspeaker’s amplitude response, but also linearize its phase response, yielding a constant, frequency-independent group delay. In other words, the loudspeaker’s dispersion is removed, resulting in a faithful reproduction of the audio signal’s temporal waveform. However, the resulting perfect transmission behavior is usually restricted to the proximity of the reference point where the TF measurement was made. Averaged over a larger area, the quality does not necessarily improve compared to a well-done conventional processing with IIR filters.

Indeed, the very detailed equalization of even the finest TF dips and peaks possible with FIR filters of sufficient length is a powerful, but also potentially dangerous tool. Many of these irregularities stem from interference of the direct sound with reflected or refracted components. As the complex-equalizing filter is time-selective, a delayed component in the loudspeaker’s non-equalized IR will be faithfully cancelled out through the equalizing filter by a component of opposite polarity (and a subsequent train of level-decaying peaks of alternating polarity). When observed from a different listener angle, the path length between the direct sound and an interfering component varies and instead of cancellation, a second peak and ringing appears in the IR. A similar effect can be observed for very steep crossover transitions which only recombine perfectly at the measurement microphone position. Away from this point, the extended pre- and post-ringing inherent in these highly frequency-selective filters can become visible (but not necessarily audible) in the IR. However, this problem has diminished considerably with the advent of line-array configurations that produce an almost cylindrical wave originating from the vertical acoustical center line in the middle of the box.

Other major concerns are reflections in horns and cone-break-up in compression drivers. It is obviously counterproductive to boost narrow frequency ranges where forward and backward moving parts of the diaphragm cancel, producing deep dips in sound pressure. Moreover, the exact location and shape of the spectral features of a driver's TF are slightly different for every individual sample from the assembly line.

If a driver/horn combination is plagued by multiple reflections due to acoustical impedance discontinuities, then complex-equalizing can become a "mission impossible". In this case, a simple high-shelf correction and some accompanying IIR-PEQs to flatten the response often yield better results.

For these reasons, the FIR generation software provides a couple of tools to optionally pre-process the measured TFs, filling dips and smoothing amplitude and group delay courses. These in the end will produce filter-TFs which ironically are more similar to those obtained with chains of conventional IIR-PEQs.

Complex-equalizing filters suffer from the same problem as linear phase filters, namely the high group delay of the overall system due to the principle of causality. To illustrate this, let's consider a 3-way system as depicted in Figure 4 and Figure 5. The magnitude responses of the overall system after equalization are shown in the top graph. An equalization using minimum phase filters, and complex- equalization was carried out. In Figure 5, the phase response (after subtracting the overall group delay) is shown.

Complex-equalizing produces almost no phase distortions except for very low frequencies, which could be corrected by using a longer filter. Not addressing the loudspeaker's phase response in the minimum phase approach results in considerably more phase shift. The resulting overall basic group delay is 15 ms for the minimum phase approach and 80 ms for the complex-equalization. For pure playback situations this is not a problem, but this latency is prohibitive for live usage. A remedy for this is the use of mixed minimum-phase and complex-equalizing. Since the woofer path normally determines the maximum delay, it is equalized using a minimum-phase filter, thus its phase is not linearized for the sake of lower group delay.

It is also possible to "morph" a FIR filter's group delay between minimum and complex-equalizing phase.

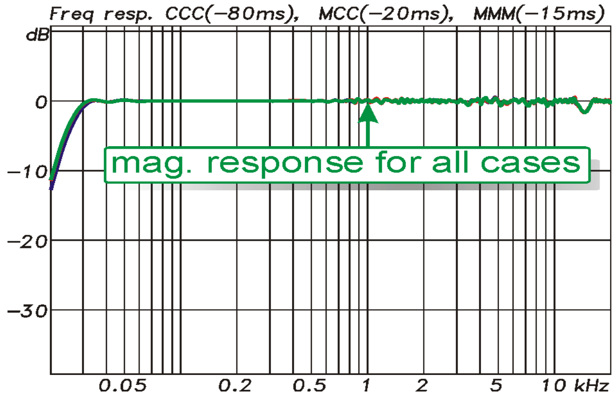

Figure 4: Magnitude responses of equalized LS.

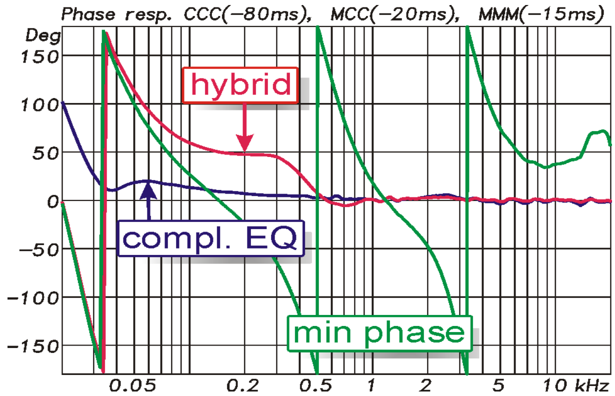

Figure 5: Phase response for complex equalization, overall minimum phase and combined minimum phase and complex equalization

Regarding audibility, group delay distortion caused by loudspeakers themselves seems to be inaudible even under critical listening conditions. An exception is bandpass-type subwoofers. Their high group delay ranging up to some tenths of ms often causes a retarded "slugging" bass response not in line with the first transient of a bass drum or a slapped bass guitar. Complex-equalizing is able to dramatically change the reproduction of such a device, yielding a very tight, compact bass response strictly synchronized with the mid and high range sound. However, the flip side of the coin is an overall latency of 50 - 100 ms, depending on the particular construction and filter parameters.

So complex-equalizing suffers a dilemma: In the mid and high range, where it is feasible with only a moderate latency penalty, constant group delay and detailed equalization of even the finest spectral features do not necessarily improve the perceived sound quality. In contrast, a considerable improvement is possible for the subwoofers, but only at the expense of an overall latency not compatible with live sound requirements.

3.4 Multirate processing

If opting for FIR bandpassing and equalization by direct convolution in the time domain, the only viable way to process the mid and low bands with reasonable computational burden and memory usage is by previous sample rate reduction. This implies a series of problems:

1) The necessary low-pass filters in the down/up sampler stages introduce additional latency.
2) The low-pass filters have to be carefully designed to avoid any aliasing (which would fall into the audio band and thus be very audible).
3) The roll-off of the low-pass filters has to be equalized by the inserted FIR filter.
4) Given a fixed DSP word length and signal bandwidth, the SNR drops by 3dB with every halving of the sample rate.

For the design of a suitable low-pass filter, a compromise between passband flatness, filter slope and stop-band attenuation vs. increased latency has to be chosen. The following criteria have influenced the design of the anti-alias filter in the down/upsampler stages:

1) Lowest possible delay was considered mandatory in order to not jeopardize live sound applications and to reduce the computational burden.
2) Stop band attenuation when using the highest sample rate reduction factor should be at least 120 dB.
3) The desired stop band attenuation should be fully reached at the new Nyquist frequency to eliminate any danger of aliasing.
4) The time response of the filter should be free of overshoot (which could cause internal clipping at the filter output even if the passband gain does not exceed 0 dB), e.g. its IR should only contain positive values.
5) Passband flatness was considered less important as it can be easily compensated by the equalizing FIR filter.
6) Possible sample rate reduction factors should include 2, 4, 8, 16 and 32.

These demands were satisfied with a FIR filter of 321 coefficients. It simply consists of a sampled Albrecht 5-cosine-term window function [4] suitably stretched in time to let the first null of its spectrum fall exactly on the new Nyquist frequency (1.5 kHz @ 96 kHz system sample rate) for the highest possible sample rate reduction factor (32).

The lower reduction factors 16, 8, 4 and 2 are easily accomplished by only considering every 2nd, 4th, 8th or 16th coefficient, respectively, scaling up the filter's frequency response accordingly. Because the number of coefficients has been chosen to be $n = k * Hrf + 1$, with k being an integer and Hrf the highest possible reduction factor, the two outer coefficients and the central peak of the 321-taps window function are always included in the convolution. The computational burden of a little less than one MIPS for each stage is practically equal for all supported reduction factors because the product of the reduced sample rate and the number of utilized coefficients is nearly the same.

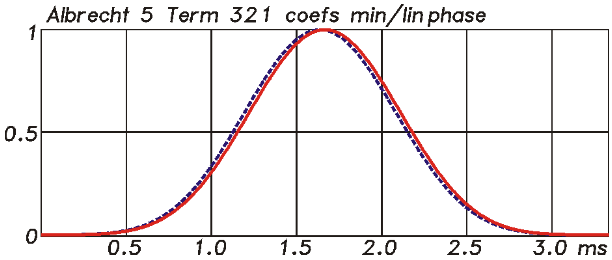

Figure 6: IR of anti-alias low-pass.
Solid line: linear phase, dotted line: minimum-phase.

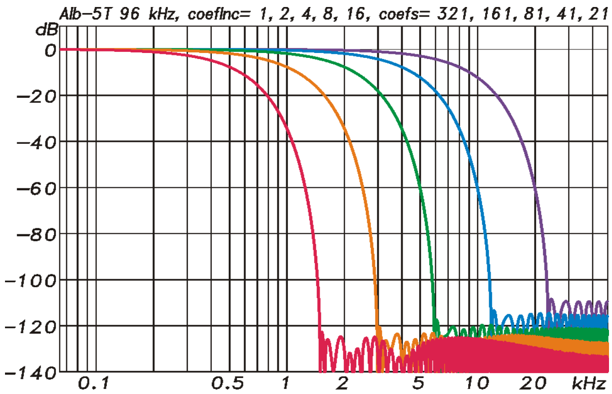

Figure 7: Frequency response of anti-alias low-pass for sample rate reduction factors of 2, 4, 8, 16 and 32

Other non-2^N integer reduction factors between 2 and 32 are possible, but not covered by the scheduler implemented in the DSP software to homogeneously distribute processing power. The highest reduction factor normally used for the sub and low frequency channel is only 16, yielding a usable bandwidth up to approximately 1kHz and 1.66 ms overall delay through the resamplers.

Instead of the linear phase TF inherent in the window's symmetrical IR, a minimum phase version could be derived by means of the Hilbert transform in an attempt to further reduce latency. However, Figure 6 shows that the difference between both versions is negligible.

The very short IR of this low-pass filter corresponds to a broad roll-off starting well below the Nyquist frequency (Figure 7). As has been stated, the roll-off can be easily compensated with the equalizing FIR filter. However, to avoid excessive boost, the sample rate reduction should be moderate enough to ensure that the new Nyquist frequency is at least three times higher than the upper corner frequency of the desired passband.

Some loudspeaker combinations cannot be handled with one single FIR filter per way because its relative bandwidth would be too large. Typical and very frequent cases are full-range speakers with passive crossover to be combined with subwoofers in a two-way active configuration. To assure sufficient spectral selectivity over the full bandwidth, two or more channel outputs can be added up to form a "multipath" way. By letting the involved FIR filters have the same internal crossover frequencies as the passive network in the cabinet, individual limiting (treated in detail in chapter 6) for each loudspeaker chassis becomes possible even though the cabinet is fed from one single amplifier channel. The summed multipath signal is passed through a separate fifth peak limiter programmed to the amplifier limits to avoid clipping of the recomposed signal.

4 FIR FILTER COEFFICIENT GENERATION

The fundament for the FIR coefficient generation consists of separate measurements of the complex TF of each loudspeaker way and its respective DLM channel. For the latter, the desired sample rate reduction factors are chosen and a neutral response (single Dirac pulse) is loaded into the inserted FIR filters.

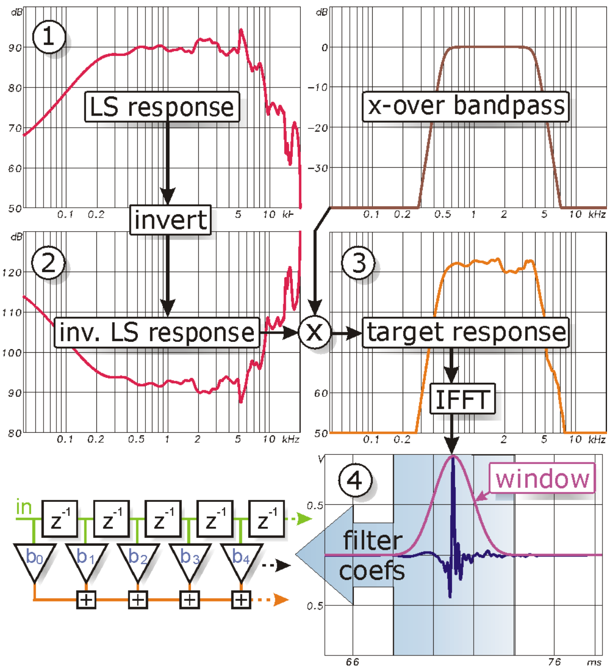

Figure 8: Simplified principle of complex EQ FIR filter coefficient generation for mid-range loudspeaker

The filter generation process for *complex-equalizing* (both amplitude and phase are equalized as explained in 3.3) is fairly straight-forward:

1) Each measured loudspeaker TF is multiplied with the respective DLM channel self-response.
2) Each resulting TF is inverted (the magnitude is turned upside down and the phases are negated).
3) Each inverted TF is multiplied with its respective prototype bandpass (see below).
4) The TFs are transformed into IRs by IFFT.
5) Each IR is treated with a window whose length equals to the product of final coefficient number and sample rate reduction factor.
6) The IR is sampled according to the sample rate reduction factor (only every nth coefficient is considered, with n being the reduction factor) and normalized to the DSPs 24 bit format.

If *minimum*-phase filters according to section 3.1 or linear-phase filters as described in 3.2 are to be generated, the phase information is deleted before step 4), yielding a symmetrical IR which is then windowed to the desired length. In case of the *minimum*-phase filter, a Hilbert transform is then performed.

Accompanying the whole process, group delay and level information are raised to set up suitable channel pre-delays and establish an internal gain structure which makes maximum use of the FIR filter's 24 bit dynamic range while avoiding any possible internal clipping.

4.1 Bandpass prototypes

The prototype bandpass slopes can be derived from traditional analog filters or be set as brickwall for highest selectivity. However, the final windowing to confine the filter's IR and extract the coefficients will obviously have a smoothing effect on the filter slopes which is constant on a linear frequency scale. As the filter responses have to become longer for the lower bands to permit sufficient spectral resolution, the slopes between two adjacent bands become asymmetrical (L1 and H in Figure 9).

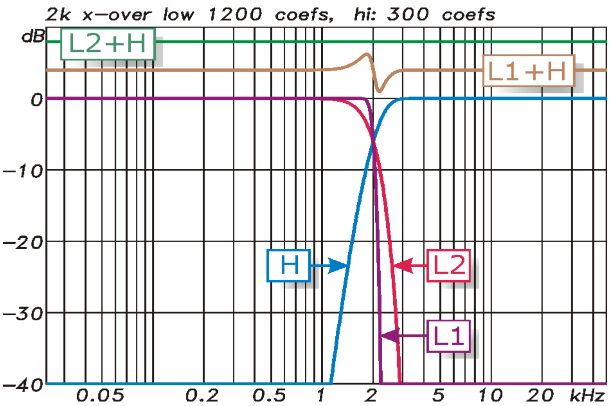

Figure 9: Effect of adding asymmetrical slopes (L1+H) and correction by pre-windowing the IR (L2+H)

So the upper slope of the lower band has to be artificially made sloppier. This can be easily achieved by the following steps:

1) The lower band IR is built by IFFT.
2) The *quotient function* of the higher band window through the lower band window is built.
3) The IR is multiplied with this quotient function.
4) The IR is back-transformed to the spectral domain to yield the corrected lower band TF.

Later, prior to extracting the coefficients in the coefficient generation process, the lower band IR is windowed one more time with its associated window. As this window has been the denominator in step 2), its contribution on the filter slope cancels out and only the influence of the narrow higher band window prevails. Thus the resulting slope (L2 in Figure 9) is symmetrical to the higher band and both add up perfectly.

Instead of using the same crossover-frequency for two adjacent bands, it is also possible to make them overlap. This tool, interesting to maximize SPL and/or to control directivity, works particularly well with line-arrays in conjunction with complex-equalizing filters.

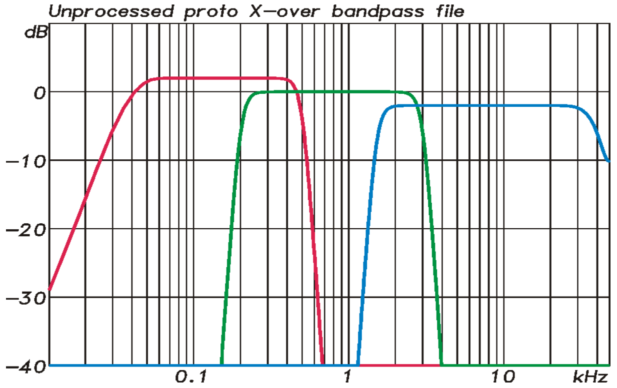

Figure 10: Example of overlapping prototype bandpass

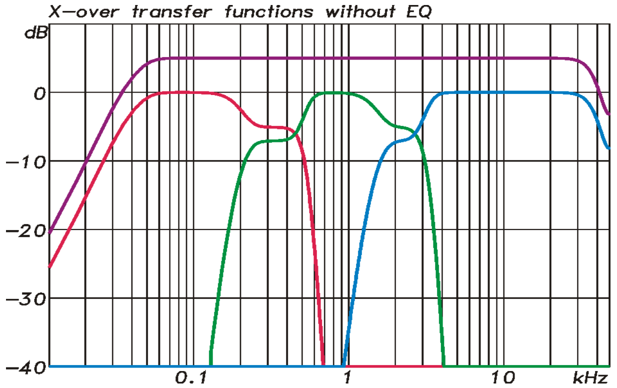

Figure 11: Resulting FIR filter and overall responses

The user can freely select level differences, slopes and individual corner frequencies (as in the example in Figure 10). A possible sequence to derive suitably corrected bandpass responses is:

1) All user-defined bandpass TFs are summed up.
2) The TFs are divided through this sum.
3) The TFs are multiplied with the desired overall frequency response, if different from unity-gain.
4) Now a consecutive subtraction scheme is used to construct the corrected upper slope of each bandpass, in which the sum of all higher bands is subtracted from the desired overall response.

This last step is performed after multiplying the associated IR of each TF with a quotient window as described earlier. Later, in the coefficient generation process, the filter TFs will add up perfectly after the final windowing has been applied to their IRs (see Figure 11).

4.2 Iterative filter response refinement

After calculating the FIR-coefficients, each filter TF can be multiplied with the self-response of the associated DLM channel and the respective measured loudspeaker response. This yields the simulated bandpassed and equalized loudspeaker response for each way of the cabinet. Summing them all up will eventually result in the simulated overall frequency response. Depending on the chosen filter sizes and window types, it will deviate more or less from the desired overall response previously defined. A considerable improvement of the overall magnitude response is possible through the following iterative scheme:

1) The desired overall response is divided through the simulated one.
2) The result is slightly stretched (the deviation in dB is multiplied with a factor slightly over 1, individual for positive and negative deviation).
3) The phase is deleted and the processed deviation multiplied with a copy of the desired overall response.
4) The whole calculation is repeated with this modified copy of the desired overall response.

The result converges rapidly in three to five iterations , yielding an overall response with substantially reduced deviation from the desired target response.

4.3 DC-elimination in FIR filter response

Treating bandpass IRs with narrow windows often produces an odd phenomenon: The transfer function “leaks“ towards lower frequencies. Instead of continuously decreasing below the lower cut-off frequency, the TF takes a horizontal course. This highly undesired behavior, particularly annoying for the subwoofer filter, but also often seen in loudspeaker TFs obtained from windowed IRs, occurs because the sum of all FIR coefficients generally deviates from 0 after applying a window. However, the balance between positive and negative coefficients can be easily reestablished with the following simple procedure:

1) The sum of all windowed coefficients is calculated.
2) The sum of all the coefficients of a window of same length (and normally same type) as used to window the IR is calculated.
3) This window function is now scaled with the quotient 1) through 2) and *subtracted* from the FIR coefficients.

The result is that the IR is gently bowed to the opposite direction of the coefficient “center of gravity”, now producing a coefficient sum of exactly 0, yielding a first-order high-pass characteristic in its TF.

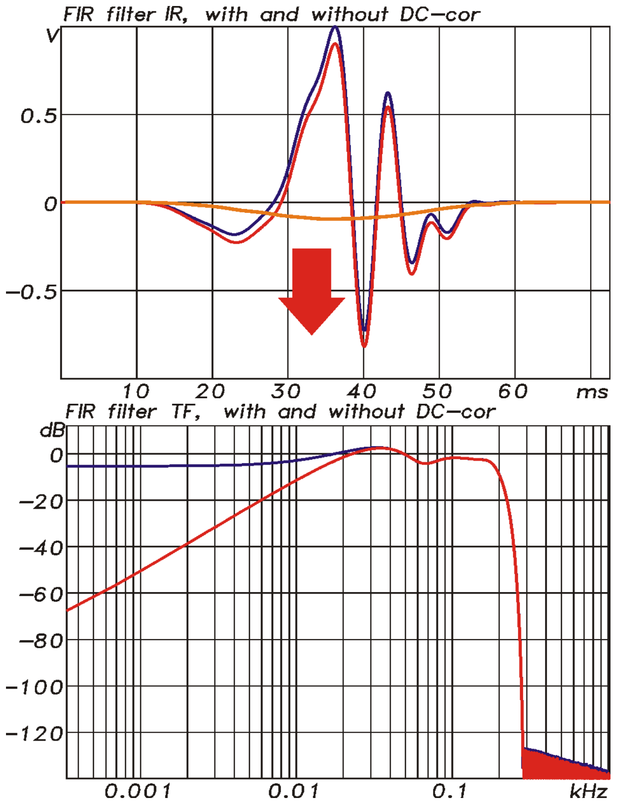

Figure 12: Original and DC-corrected filter IR and TF

4.4 Adjustments at large venues

A cluster of loudspeaker cabinets behaves quite differently than the single box for which a particular equalization was made. That's why checking the response at various points over the covered area is indispensable to fine-adjust the average response of a sound system. As an example, Figure 13 presents measurements at various rows in a football stadium refurbished with a new line array system. The energetic average (discarding the phases) of all those curves is shown in Figure 14. Based on the gentle broadband roll-off to higher frequencies, a "desired response" curve is constructed. With the help of the input PEQ bank, the averaged TF is approximated to this desired response, filling dips and cutting peaks.

The resulting *modulus* of the PEQ bank frequency response curve can be multiplied with the desired overall spectrum previously defined in the FIR filter coefficient generation process and then the whole coefficient generation is repeated. This imprints the PEQ bank's amplitude response to the FIR filter responses without introducing its phase distortion. ..This way, the final EQ adjustment work can be shifted into the FIR filters, freeing the complete PEQ bank for future adjustments.

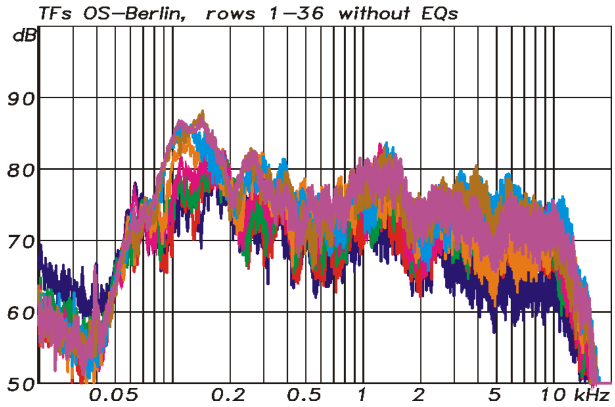

Figure 13: TFs measured in Olympia stadium, Berlin.

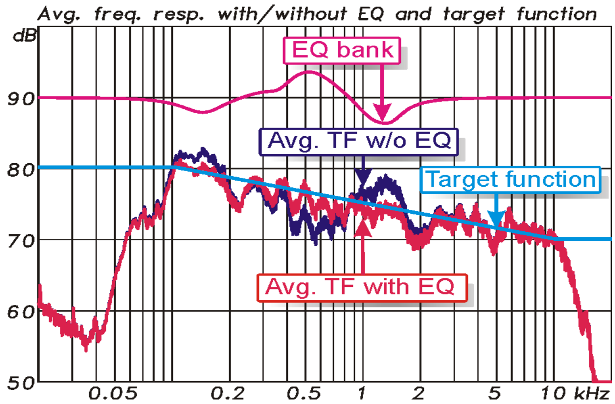

Figure 14: Averaged TFs without and with EQ, target response, and EQ curve

5 DUAL-RANGE AD CONVERSION

A DLM is typically located between the mixing desk and the power amplifiers. If the mixer is digital, the connection to the DLM is of course best done digitally via AES3, avoiding additional latency and preserving the dynamic range. However, many consoles for live use are still purely analog. Typically equipped with opamps fed with up to ±18V supply voltage, their symmetrical outputs are capable of delivering around +28 dBu (27.5 Vpp). In contrast, when sliding the master faders down to mute the signal, only the output buffer noise remains, letting the noise level easily drop to below -100 dBu. The difference of both values yields a dynamic range of approximately 130 dB, not attainable by today's AD converters. For this reason, most DLM have a gain control potentiometer or switches in their input stages, tasking the sound engineer with the decision as to whether high headroom or lowest noise are more important.

Ideally, the DLM inputs should not clip before the mixing desks in order to not diminish headroom needed for the limiters. At the other end of the input level range, the DLM input circuit should not contribute significantly to the overall system noise when only a moderate volume is needed.

The dual-range conversion principle depicted in Figure 15 solves the problem, generously boosting the dynamic range by using a stereo AD-converter in a mono configuration. The input signal is treated by two preamplifiers, one with unity gain and the other one with a gain equal to the desired dynamic range expansion. The circuit must be carefully designed to avoid crosstalk to the other channel when the preamplifier with the higher gain clips. The entailing leap of distortion can be kept satisfactorily small, as the red curve in Figure 16 documents.

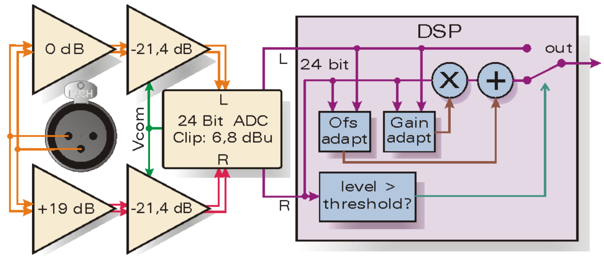

Figure 15: Dual-range AD conversion principle

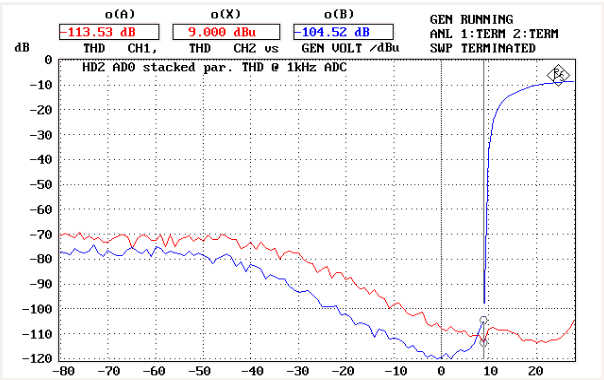

Figure 16: THD vs. level of the two input channels of the dual-range AD converter

At levels below clipping, the digitized higher-gain ADC signal, attenuated by the gain difference and corrected for the DC offset difference between both channels, is elected to be the source signal. When the clipping level is reached, the lower-gain ADC input becomes the source.

Some precautions must be taken to achieve seamless and inaudible switching. First of all, the gain difference and the DC offset difference between both channels are prone to slight drift due to circuit warm-up and other factors. To keep the gain mismatch and offset error sufficiently low, both values have to be tracked adaptively. The adaptation speed is dependent on the signal level: At high levels, the gain detection is fast and the DC-offset detection is slow and vice versa. When clipping occurs, the adaptation is stopped and the last values are frozen. The gain detection error is usually less than ±0.005 dB and the DC offset error less than 20 µV.

When switching from one channel to the other, any remaining tiny discontinuity is handled by adding a small offset to the new channel so that old and new samples are exactly the same. This added offset is smoothly driven to zero with a slow exponential decay during the next milliseconds.

While these precautions already guarantee that the transitions are inaudible under any circumstance, some further improvements can be made to drastically reduce the possible number of transitions per second, making use of psycho-acoustic properties of our hearing sense. After switching to the lower-gain signal, it is kept as the source for at least 20 ms. Even if the level falls well below the transition threshold during this “hold” period, the higher noise level of the lower-gain input channel is completely masked by the previous loud sound. Additionally, the thresholds for switching up and down are not equal (-2 and -5 dBFS). This hysteresis prevents switching when the input signal has almost constant level.

The technique described here obviously does not increase the *instantaneous* SNR of the ADC. At a given moment, only one of the two channels is active, so the value of the single ADC channel rules. When switching from the high to the low gain channel, the SNR drops from 118 dB to a little less than 100 dB. However, this noise level continues to be masked even in case of the most critical signals, such as very low frequency tones.

So the dual-range ADC input can be considered as a providential dynamic range extension. The high gain ADC channel assures a conversion with very low input noise level (-108 dBu) up to about +8 dBu – slightly above the typical levels encountered at the mixer output when reaching maximum volume. From there on, the other ADC channel takes over, allowing for 20 dB of headroom to be used by the limiters, now described in detail.

6 LIMITERS

Limiter design has been the last holdout of analog-addicts. However, digital signal processing allows devising better sounding limiters with much improved precision, no added noise and considerably reduced distortion as compared to analog designs.

6.1 Look-ahead peak limiter

Digital peak limiters can make use of a simple yet effective tool not available to analog limiters. A small delay inserted in the peak limiter’s signal path (Figure 17) allows its control logic to gradually reduce the gain (with a constant dB/s slope) when triggered by a peak. This limits the transient to exactly the programmed threshold when reaching the output (Figure 18). This way, clipping in the subsequent signal processing stages is safely avoided. The main benefit, however, is that the gain reduction signal (corresponding to the VCA control voltage in an analog solution) has significantly lower high frequency content, effectively reducing the unpleasant and aggressive sounding distortion introduced if high frequency components modulate the audio signal.

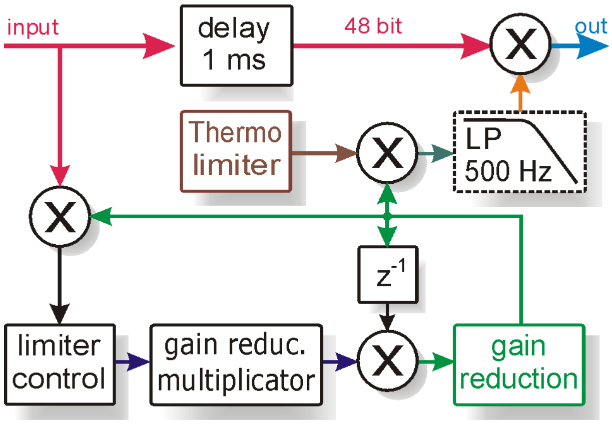

Figure 17: Block diagram of look-ahead peak limiter

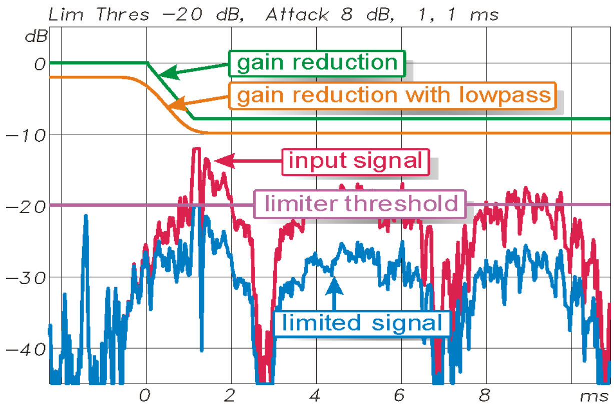

Figure 18: Limiter reaction to transient above threshold

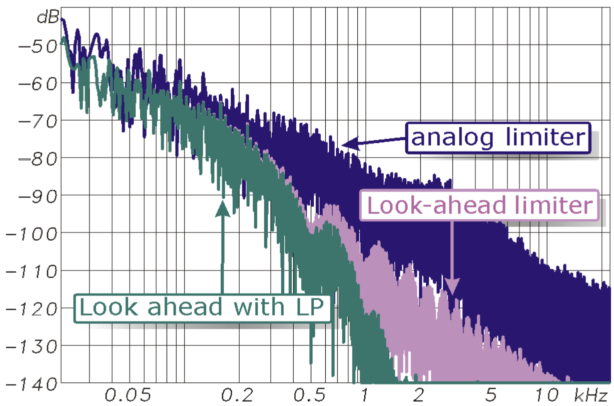

Figure 19: Gain reduction spectrum in conventional analog peak limiter compared to look-ahead limiter with and without 500 Hz low-pass.

Though already reduced by more than 20 dB compared to an analog limiter with similar parameters and the same input signal (Figure 19), the high-frequency components can be further attenuated by low-pass filtering the "VCA control voltage" (dashed block in Figure 17). A 6th order Bessel filter with 500 Hz cut-off frequency has been found suitable for this purpose. Its constant passband group delay of nearly 1ms can be easily compensated by adding it to the length of the signal delay line. Though providing a very subtle increase in sound clarity, the additional latency may be objectionable in live sound applications, so a switch allows bypassing the filter.

To summarize, the attack *rate* is adapted to the amplitude of the detected transient, while the attack *time* is constant. Its value of 1ms has been chosen to be independent of the processed band. This means that in the high-band, the limiter has time to reduce the gain over at least one half-wave. In contrast, a slowly developing half-wave triggering the limiter of the subwoofer channel is chamfered over a one-millisecond period before reaching the threshold level and remains there until reaching the cusp. This behavior, in conjunction with the "controlled overshoot" feature described in 6.2, preserves the "kick" while avoiding continued clipping in the subsequent half-waves.

The limiter control is rounded out by a retriggerable hold phase of 50 ms and a release rate adjustable by the user from 10 dB/s to 200 dB/s, just as in a conventional peak limiter. However, the actual release rate is not static. If the calculated output signal level remains near the limiter threshold, the release process is continuously slowed down. This eliminates a potential problem for slowly decaying sounds like a stroke on a crash cymbal: If the decay rate was slower than the static release rate, release and hold phases would alternate, giving the sound a sawtooth-like modulation.

It must be stressed here that although the look-ahead concept reduces distortion by keeping the gain reduction signal cleaner and by safely avoiding clipping in subsequent stages, the action of the peak limiter is of course not inaudible. Feeding an audio signal with a level well above the threshold and choosing a very fast release will result exactly in the same very dense and yelling loud sound unfortunately overused by many mastering studios and radio stations.

6.2 Controlled overshoot

The peak limiter's purpose is to avoid amplifier clipping and mechanical damage of the drivers. If the connected loudspeaker easily handles the maximum amplifier output (often true for subwoofer arrays), a peak limiter programmed to the amplifier's rated long-term power could mean squandering its short-term power capability. In order to remedy this waste of 2 or 3dB (or even more) for transients, the peak limiter concept was modified to simulate the power supply voltage drop occurring when power is drawn. Two additional parameters are needed to program this "controlled overshoot" feature: The "surge" value indicating the peak power available when the power supply electrolytic capacitors are fully charged, and the "duration" value which is the time when the output power has dropped to the continuous level (observing the initial declination of the power drop). Both values can be easily measured by loading both amplifier channels, applying a burst signal that slightly drives the amplifier into clipping, and analyzing the resulting waveform on an oscilloscope (Figure 20). If not available, a "surge" value of 2dB and a "duration" value of 30 ms harmonize well with most common power amps.

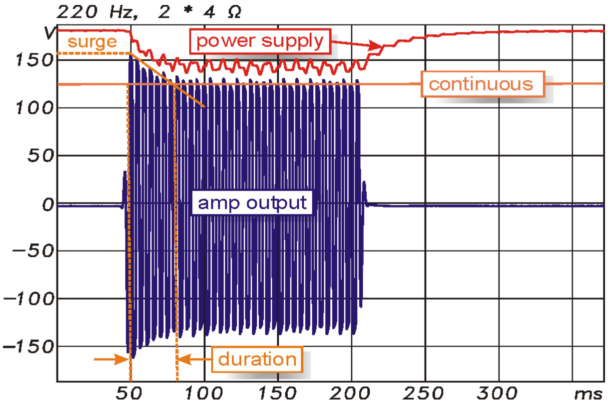

Figure 20: Determination of parameters for "controlled overshoot" limiter threshold adaptation.

The power drawn from the amplifier is modeled by squaring the signal and feeding an RC-like accumulator whose output shifts the limiter threshold downwards.

6.3 RMS limiter

The independent RMS thermal limiter is programmed with the continuous power rating of the loudspeaker and a time-constant which is the intersection of the initial temperature rise with the temperature limit upon applying the rated power. While the first value is available from the loudspeaker manufacturer, the second one has to be estimated according to the thermal capacity of the driver. The limiter algorithm also includes a second thermal circuit, simulating the heat transfer from the voice coil to the magnet, but due to the absence of reliable data, this feature has not yet been used.

The power transformed into heat is modeled by squaring the output signal, assuming a constant load impedance. The fact that both the impedance and the excursion and thus the main cooling mechanism of a woofer's voice coil are strongly frequency-dependent makes this assumption fairly imprecise. Exciting a vented box with rated power at its tuning frequency (where excursion and heat convection are minimal) will make the thermal limiter fatally fail in protecting the woofer from overheating. However, just as the rated power information is usually evaluated with AES noise (a pink noise with controlled crest factor), the limiter model is based on the broad spectral statistics of a musical signal. The RMS limiter acts independently of the peak-limiter and both gain reductions (in dB) are added (Figure 17). This means that peaks are limited to a lower output value when the RMS limiter is active, a desired protection feature.

7 CONCLUSIONS

The idiosyncrasies of FIR and IIR filtering for use in loudspeaker management systems have been illustrated based on a newly developed DLM. FIR filters allow using steep slope crossover filters without introducing phase distortions, reducing interference effects between transducers. Since the filters are based on the measured complex transfer functions of the loudspeakers, a constant group delay can be achieved for the whole reproduction chain. For live applications, the high group delay introduced by this approach can be reduced by using minimum phase filters in the low frequency paths.

The problem of high calculation power for low frequencies is being dealt with by using a multi-rate approach.

Besides this, it is possible to use IIR filters additionally, or exclusively. They can be used as crossover filters as well as for EQs. A versatile IIR filtering section was implemented by performing the necessary filter calculations inside the fixed-point DSP. Double precision arithmetic guarantees audio signal processing with very low self-generated filter noise.

The illustrated dual range AD conversion yielding an enhanced dynamic range relieves the sound engineer of fitting the AD range of the DLM to the output range of the mixing desk.

The separate limiters for peak and average power preserve the "kick" while simultaneously preventing the loudspeakers from mechanical damage or overheating.

Concluding, it can be said that the presented DLM platform opens up enhanced features in loudspeaker equalization while also providing standard techniques as already known.

REFERENCES

[1] Ashihara Kaoru, Kiryu Shogo, "Detection Threshold for tones above 22 kHz", 110th AES Convention, Amsterdam, paper 5401.

[2] John Vanderkooy, Stanley Lipshitz, Robert Wannamaker, "Minimally Audible Noise Shaping", *J.AES*, vol. 39, 1991, pp. 836-852

[3] Sophocles J. Orfanidis, "Digital Parametric Equalizer Design With Prescribed Nyquist-Frequency Gain", J.AES, p. 444, June 1997

[4] Hans-Helge Albrecht, "A Family of Cosine-sum Windows for High-Resolution Measurements", Proceedings. (ICASSP '01) 2001 IEEE