The Intelligent Loudspeaker

Paul Williams, Mark Bailey

1 Abstract

The technology behind the creation of a new class of loudspeaker, *The Intelligent Loudspeaker* is described, explaining what signal processing is carried out within the loudspeaker and why, for the three distinct phases in the life of the loudspeaker; manufacture, commissioning, and maintaining operation. The processing methodology employed is described in some detail in the areas of automatic equalisation, automatic feedback control, and thermal modeling.

2 Background

Whilst the materials technology used in used in modern loudspeaker systems has continued advancing, there have not been any real revolutions in loudspeaker design for several decades.

Electronic 'Loudspeaker Management' devices incorporating crossover filters and other processing pertaining to the loudspeakers have become very commonplace, liberating the loudspeaker designer somewhat from the constrains of designing within the limits of what simple analogue processing could offer, such as crossover filters with limited slopes, and no programmable equalisation. These have traditionally been de-coupled from the loudspeaker, making it difficult or impossible to optimise the system - since without intimate knowledge of one-another, the devices rarely work in optimal harmony.

The new loudspeaker described here offers a breakthrough in ease of installation, and improved ongoing performance by using Digital Signal Processing (DSP) technology, tightly coupled with the amplifiers and loudspeaker drivers. An optimal system is created at the time of manufacture, further optimisation being possible at installation according to the environment, and further dynamic optimisation is possible depending on the changing environment and changing driver characteristics.

The Intelligent Loudspeaker utilizes Digital Signal processing (DSP) for the following functions:

Static processing (As manufactured)
- Crossover filtering
- Fixed driver equalisation
- All-pass filters and delays for inter-driver phase correction

Setup processing (Installation)
- Delay measurement and compensating delay placement
- Automatic room measurement and equalisation
- Fixed feedback reduction

Dynamic processing (Maintaining operation)
- Protection Limiting
- Dynamic feedback reduction
- Thermal modeling for thermal limiting and power compression reduction

Fig. 1 shows the general arrangement of The Intelligent Loudspeaker. It comprises three drivers, driven by three amplifiers, and a Digital Signal Processing (DSP) card performing all the processing, this being fed from a single analogue input.

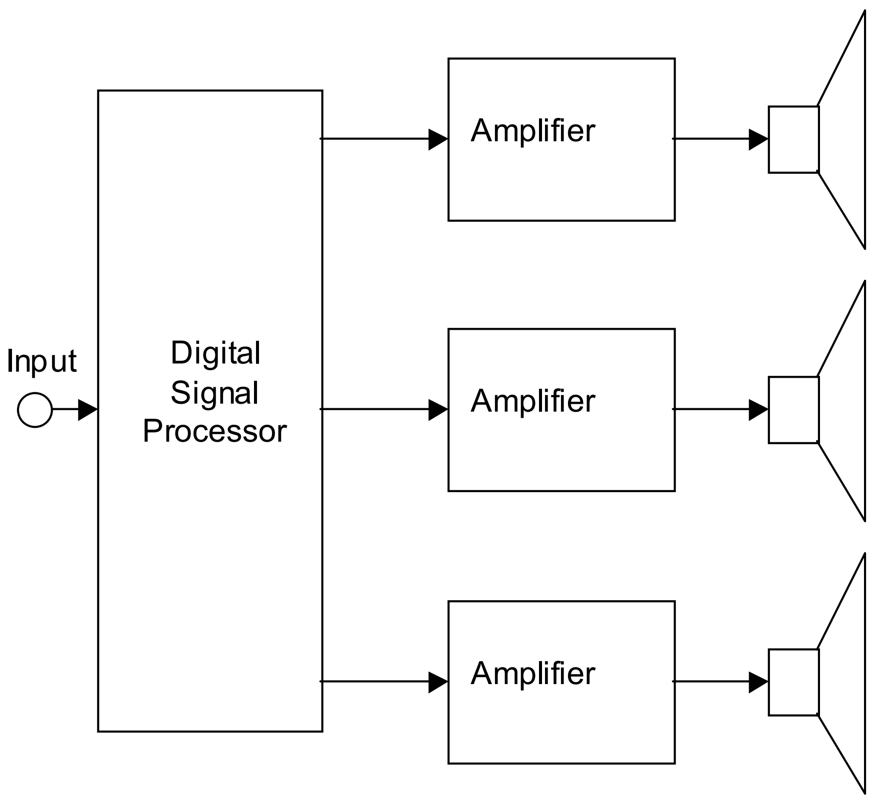

Fig. 1 General arrangement of the Intelligent Loudspeaker

Some of these processing elements are now discussed further.

3 Automatic room measurement and equalisation

Automatic room equalisation is performed when the product is first installed, to correct for room response anomalies.

3.1 The acoustic environment

The principles uncovered by the seminal paper by Toole [1] showed us that:

- The best sounding loudspeakers are those with a flat on-axis response
- Low frequency response (below about 500Hz) is dictated by the reverberant field, determined entirely by the acoustics of the room
- The high frequency response is dictated almost entirely by the direct radiated on-axis response of the loudspeaker, and is not impacted significantly by the room.

Since we are striving to preserve a good on-axis response, it follows that any equalisation applied to correct any deficiencies in the off-axis response, which may be evident at higher frequencies, will be to the detriment of the perceived sound quality. In a well-designed loudspeaker, care will have been taken to preserve a good off-axis response, the designer being aware that deficiencies here cannot be corrected without detrimentally affecting the all-important on-axis response. So, even though in spectral analysis we are able to measure any deficiencies in frequency response flatness at higher frequencies, we should deliberately ignore these. Our automatic equalisation will therefore concentrate only on those artifacts considered to be due to room effects, which are those below 500Hz or so.

The nature of the room can have very significant effects on the flatness of the reproduced spectrum. Reflections from hard surfaces for example can alter the reproduced levels by as much as 12dB. Standing waves can make further contributions to this.

The automatic equalization process breaks down into two sub-processes: measurement and correction.

3.2 Measurement

Measurement is performed using a form of the Maximum Length Sequence (MLS) technique, as described by Schroeder [2] and further discussed and developed by Borish & Angell [3], Rife & Vanderkooy [4] and others. Here, a Pseudo-Random Binary Sequence (PRBS) of maximum length is used to stimulate the environment, and the acoustic results captured by a suitably placed microphone, the recovered sequence being processed to discover the characteristics of the acoustic environment. The MLS technique benefits from good tolerance of poor Signal-to-Noise Ratio (SNR), so a high level stimulus is not required. Pink filtering is applied to the stimulus to adequately stimulate the LF end of the spectrum without over-driving HF drivers.

The stimulus is 'played' over the system, and is synchronously recorded via the measurement microphone into the DSP. Several sequences are repeated, and the results averaged to further improve the SNR, as discussed by Nielsen [5].

The properties of the PRBS sequence exploited by the MLS technique are that the circular autocorrelation of the PRBS sequence is an impulse, and the circular cross-correlation of the captured sequence with the original PRBS sequence is the impulse response of the measured system:

$$\Omega_n = \frac{1}{L+1} \cdot \sum_{k=0}^{L-1} s_k \cdot y_{n+k} \tag{1}$$

Where L is the length of the sequence, s is the PRBS sequence, and y is the recovered sequence.

With knowledge of the parameters of the pink filter etc, an impulse response for the room can be calculated using equation 1, from which the delay and magnitude response parameters can be obtained by taking the Discrete Fourier Transform (DFT) of the impulse response using FFT techniques. The recovered impulse also allows us to determine the time taken for the signal to traverse the measured environment, enabling us to determine any required correction delays.

3.3 Correction

We can now determine what equalisation is required to make the measured response look more like the desired target response, limited to the range of frequencies that we will allow to be corrected for room equalisation (up to about 500Hz). We also have to be mindful of the dangers of trying to correct significant nulls, and significant resonances, either of which would require unduly large amounts of equalisation. It is clearly not practical to attempt to correct for complete nulls caused by standing waves, and it may be undesirable to fully correct for very large peaks in response.

The parameters of a bank of parametric equalisers are then manipulated to give a best-fit to the required response. The parameters for the equaliser bands are determined by an iterative search-and-fit algorithm, which applies equalisation filters at those frequencies where the measured response departs most significantly from the target response.

Fig 3.3.1 Auto EQ Response and Correction

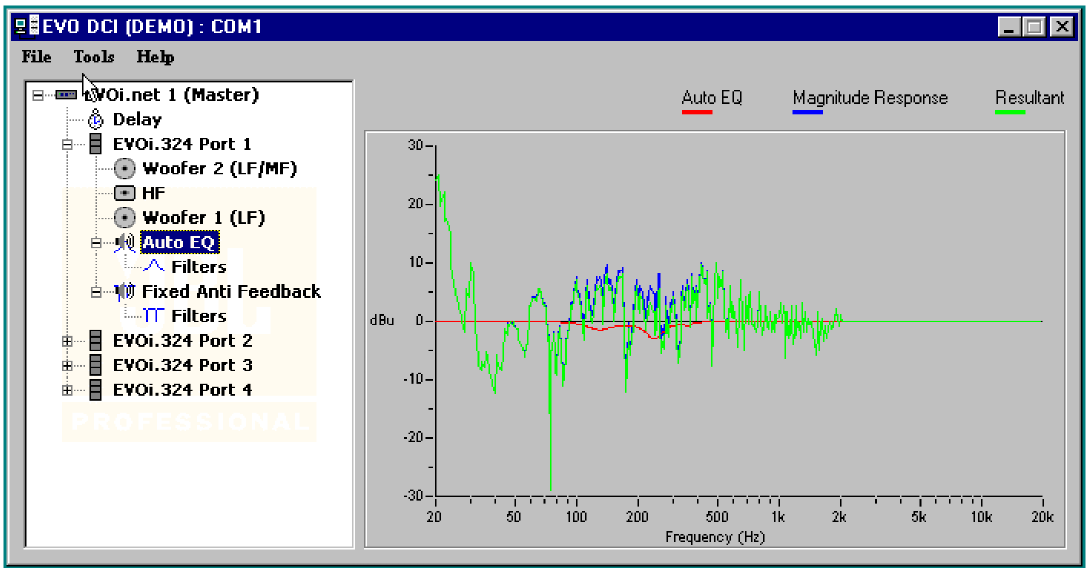

Fig 3.3.2 Auto EQ data

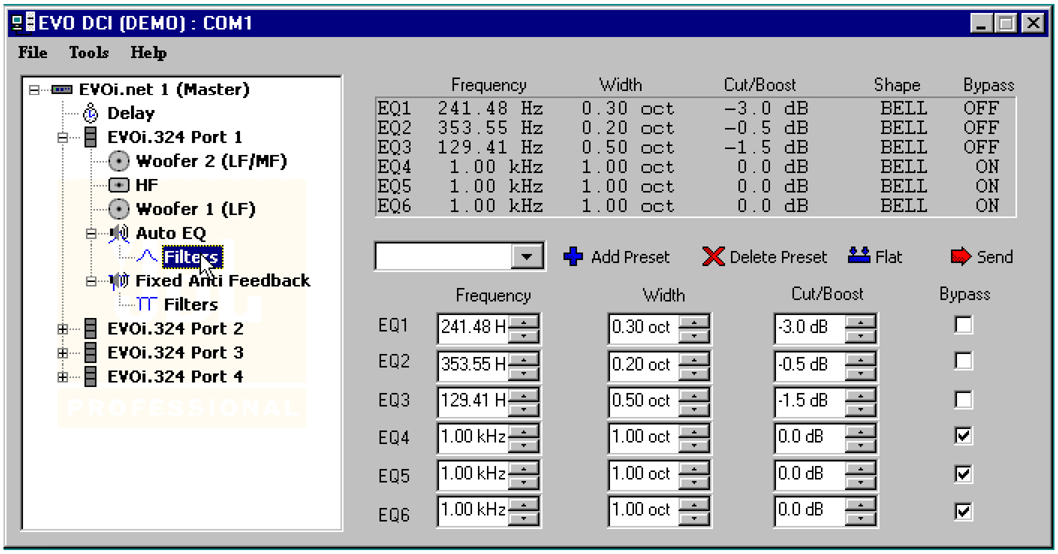

4 Automatic Feedback Control

Regenerative feedback via the loudspeaker-microphone is clearly a common problem, and is a problem which can be tackled at two distinct stages:

1. When the system is installed, taking care of resonances caused by the room, which can be catered for at setup time.
2. When the system is in use in changing circumstances, such as differing sizes of audience, and roving microphone position, both of which need a more dynamic approach.

The Automatic Feedback Control (AFC) employed in The Intelligent Loudspeaker thus has the concept of fixed (Fixed AFC) and dynamic (Live AFC) filters that may be deployed to reduce the regenerative gain at problem frequencies to reduce the gain around the system below unity. Live filter settings are considered temporary, and so are not saved during power cycles.

The feedback reduction system employed in this loudspeaker therefore has two such distinct phases, but both using a similar method.

FFT techniques are used to dynamically capture the changing spectrum of the signal being applied to the input of the loudspeaker system. The spectral data is subjected to sophisticated statistical analysis to identify any resonant peaks in the signal that are not due to any wanted music or speech signal. From this, we are able to deduce the frequency at which any feedback is occurring, and automatically place a narrow notch filter at that frequency to attenuate the feedback.
We then have measures that ensure that the spectral feature thus attenuated is definitely not due to a wanted signal as a double check. If it is, we are able to then remove the filter.

The loudspeaker allows a 'ring out' setup phase to be carried out, whereby the user lifts the faders on the mixing console, raising the gain in the system until feedback occurs, at which point the feedback reduction system will apply a notch filter at the appropriate frequency. The user keeps raising the system gain, allowing the feedback reduction system to apply more notch filters until either there is sufficient gain without feedback, or all the setup notch filters are used up.

Dynamic feedback reduction is exactly the same except a different bank of notch filters are used, and also, depending on how many dynamic notch filters are free, the filters are slowly removed over time to free one or more up, the assumption being that the offending feedback was a due to an non-ideal short-term microphone position.

This method is the subject of a patent application.

5 Thermal Modeling

It is well known that the impedance of the voice-coil of a loudspeaker changes significantly with temperature, and such variation causes not only changes in the reproduced acoustic signal level, but also the spectral content, because of the changes that occur in complex impedance. The results are the reduction of system gain at high power levels (commonly known as *power compression)*, and dynamic timbral shifting as the power level varies.

In this loudspeaker, we compensate for both of these effects by dynamically adjusting a compensating equaliser bank depending on the voice coil temperature. It is of course difficult to monitor the temperature of the voice coil directly, so we use thermal modeling to predict the voice coil temperature merely by examining the signal being applied to the driver.

By implementing this model, we are able to predict the voice coil temperature with sufficient accuracy that we can confidently apply compensating equalisation to substantially eliminate power compression and timbral shifting.

An additional benefit is that we can use the temperature information to cause the gain of the system to be reduced if any driver voice-coil is predicted to be exceeding what is considered to be a safe working temperature.

Fig. 2 shows how the sensitivity of a typical driver at different frequencies varies with voice-coil temperature.

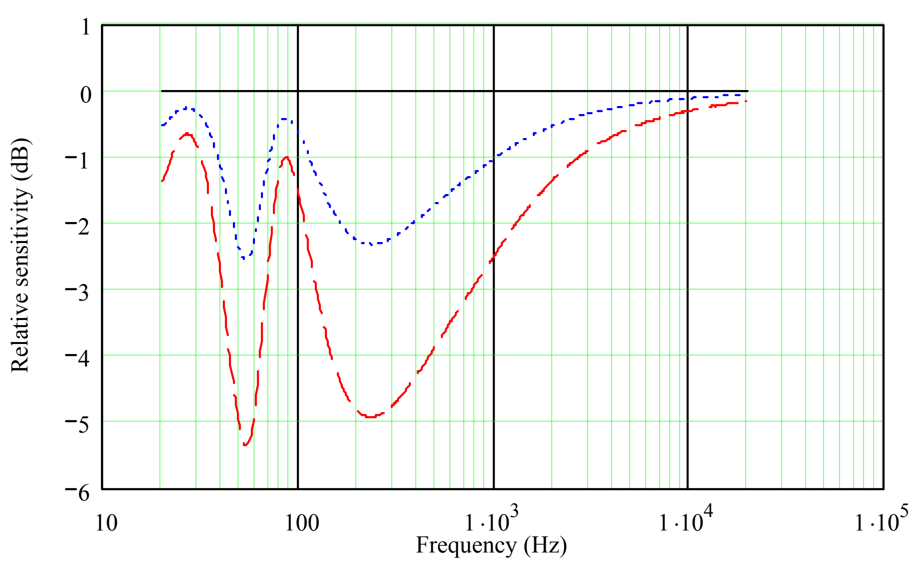

Fig. 2 Sensitivity of a typical driver at voice-coil temperatures of 20 Centigrade (solid line), 120 Centigrade (dotted line) and 270 Centigrade (dashed line).

5.1 Thermal model

Button [6] has described a method of modeling the thermal circuit of a driver by using an electrical circuit involving two resistors and two capacitors to represent the major components involved. The thermal model based on the equivalent electrical circuit shown in Fig 3. When driven by a voltage equating to power delivered to the driver, this circuit will cause a current to flow, which is equivalent to the temperature.

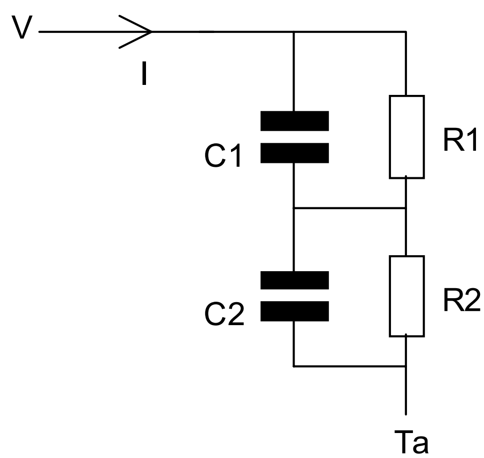

Fig.3 An electrical circuit model of the thermal circuit of a driver

Current, I, is equivalent to the power in the thermal circuit
Voltage, V is equivalent to the temperature above the reference (ambient, Ta).
R1 is equivalent to the thermal resistance from coil to magnet.
R2 is equivalent to the thermal resistance from magnet to ambient.
C1 is equivalent to the thermal capacity of the coil.
C2 is equivalent to the thermal capacity of the magnet.

To model the temperature in the voice-coil for a given power level, we solve for I, given V in this network:

$$Z(s) = \frac{\frac{1}{C1 \cdot s} \cdot R1}{\frac{1}{C1 \cdot s} + R1} + \frac{\frac{1}{C2 \cdot s} \cdot R2}{\frac{1}{C2 \cdot s} + R2} \tag{2}$$

In The Intelligent Loudspeaker, we implement this model in the Z-plane in a DSP, enabling us to predict the voice coil temperature within the DSP.

By applying the Bilinear Transform to equation 2, we arrive at a z-plane equation, which we can run in the DSP:

$$Z(z) = \frac{z \cdot T \cdot (R1 \cdot z \cdot T + R1 \cdot R2 \cdot C2 \cdot z - R1 \cdot R2 \cdot C2 + R2 \cdot z \cdot T + R1 \cdot R2 \cdot C1 \cdot z - R1 \cdot R2 \cdot C1)}{z^2 \cdot T^2 + z^2 \cdot T \cdot R1 \cdot C1 + z^2 \cdot R2 \cdot C2 \cdot T + z^2 \cdot R2 \cdot C2 \cdot R1 \cdot C1 - z \cdot R2 \cdot C2 \cdot T - 2 \cdot z \cdot R2 \cdot C2 \cdot R1 \cdot C1 - z \cdot T \cdot R1 \cdot C1 + R2 \cdot C2 \cdot R1 \cdot C1} \quad (3)$$

This boils down to a second order IIR filter. From equation (3), the coefficients for this filter may be calculated for a given set of R and C values in the equivalent electrical circuit.

There are however complications which we cannot ignore. The movement of the voice-coil and diaphragm in air help to cool the voice-coil. To account for the cooling effect of the moving coil and diaphragm, the thermal resistance values are modified dynamically, but since this cooling effect changes with frequency a cooling EQ filter is used to spectrally shape the cooling signal, whose RMS level is used to modify the thermal resistance values. Also, the estimated driver power cannot be derived from voltage alone, and must take account of driver coil resistance, which also changes with temperature. The calculated temperature value is therefore used to calculate the current coil resistance. Fig. 4 shows the building blocks making up the thermal modeling process.

.

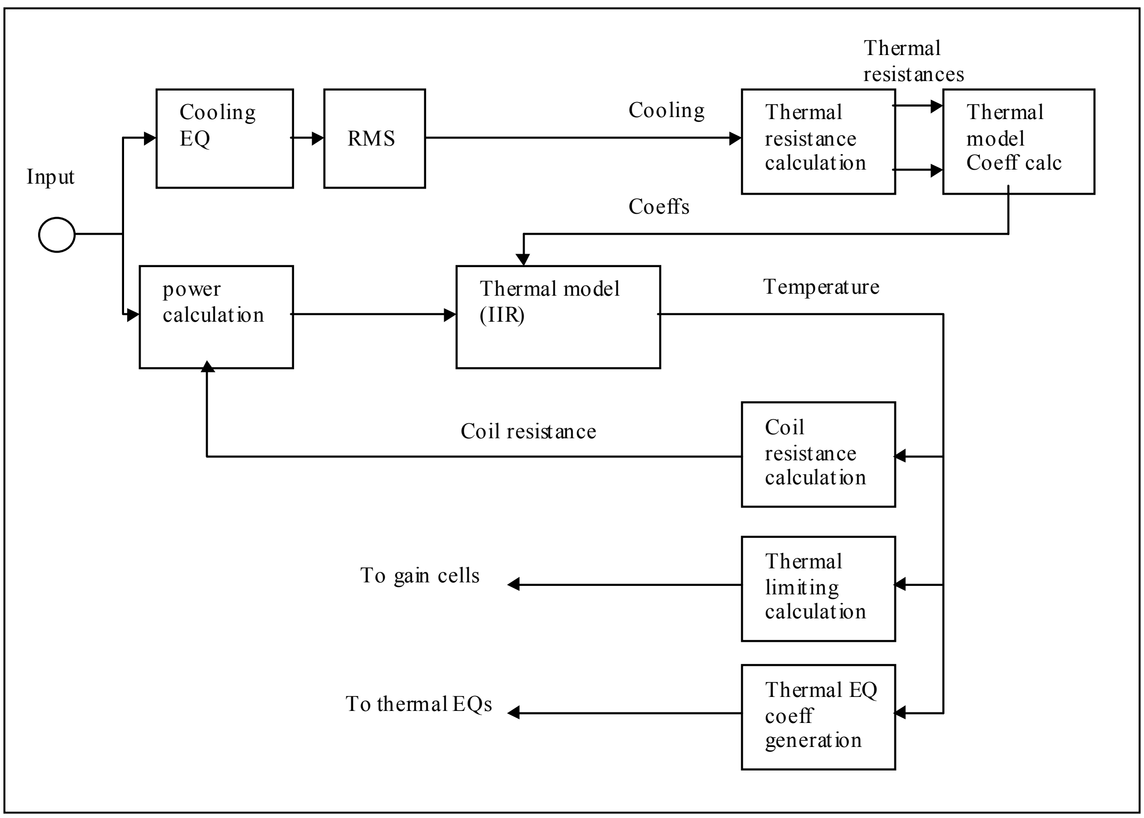

Fig. 4. The thermal modeling processing elements.

The heart of the engine is the thermal model itself, a 2nd order filter that produces a signal representing the temperature value. The resulting temperature from the thermal model is used for re-evaluating the voice coil resistance, recalculating the required driver equalisation, and for controlling gain cells for absolute thermal limiting. The power calculator uses the instantaneous 'voltage' as seen by the power amplifier, and the coil resistance, to produce an instantaneous value proportional to the power delivered to the driver. This value drives the thermal model. The coefficients for the thermal model are calculated using equivalent R and C values. The C values are fixed, but the R values vary with diaphragm activity, as determined by the cooling filter, together with the RMS meter. The resulting cooling effect value from the cooling meter is used to calculate the thermal resistance R values.

The modeled temperature of a typical driver is shown in Fig.5.

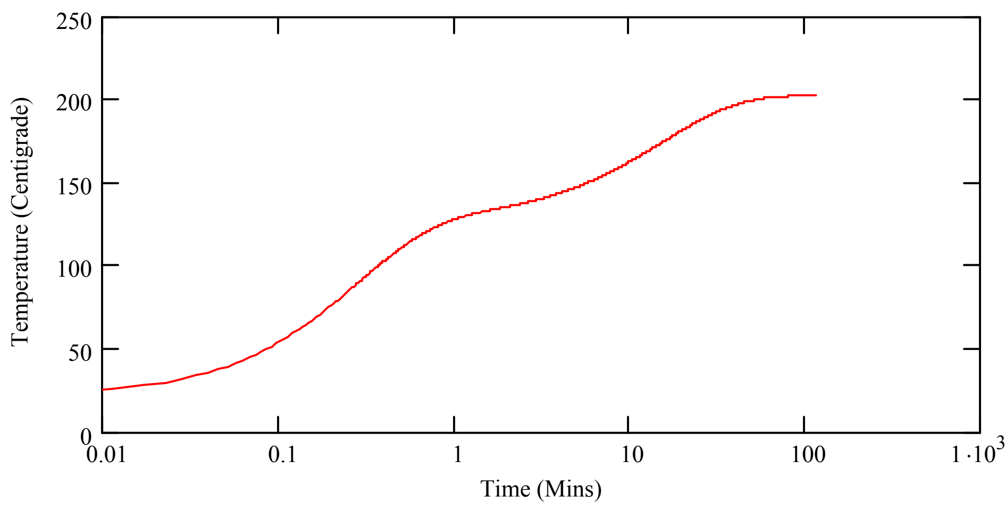

Fig. 5. Typical output from the thermal model.

5.2 Thermal Correction

Once the temperature of the voice-coil is known, it is then possible to calculate parameters for a bank of correction equalisers which invert the known spectral patterns of the driver, as shown in Fig.2. This equalisation can correct for both spectral shifts, and changes in sensitivity.

We can also now prevent thermal damage to the drivers by observing when the predicted temperature exceeds a predetermined limit, whereupon we can cause the gain in all bands to be reduced.

This method is the subject of a patent application.

6 Conclusions

There are many aspects of sound system design that can be readily predicted and calculated for, the ***majority*** of these can be accommodated by operator assisted computer control in the ***majority*** of cases.

It is clear that in many system realizations today, much of the 'basic' system engineering is performed to an insufficient standard to get the best from the equipment installed. This is often due more to a lack of time or resource than a lack of knowledge, though the latter can play a part.

The areas where assistance is most often required are: System Commissioning, System Maintenance and System Operation.

In utilizing the *Intelligent* loudspeaker design, we can create a sound system that will produce:

- A *more accurate result* at the outset by computer assisted commissioning.
- A *consistent result in operation.* Attained by monitoring operational conditions such as power compression, feedback etc.
- A *simple* indication of what the fault is if something goes wrong.

In the imperfect world of sound system installation, *The Intelligent Loudspeaker* offers a real benefit to all who have cause to install, use, or even just listen to, sound systems in a professional environment.

7 References

[1] F.E. Toole, "Loudspeaker Measurements and Their Relationship to Listener Preferences", J. Audio Eng, Soc., vol. 34, pt.1 pp.227-235 (1986 April), pt. 2, pp. 323-348 (1986 May).

[2] M. R. Schroeder, "Integrated-Impulse Method Measuring Sound Decay Without Using Impulses", J. Acoust. Soc. Am., vol. 66, pp. 497-500 (1979 Aug).

[3] J. Borish and J. B. Angell, "An Efficient Algorithm for Measuring the Impulse Response Using Pseudorandom Noise", J. Audio Eng. Soc., vol. 31 pp. 478-488 (1983 Aug.).

[4] D. D. Rife and J. Vanderkooy, "Transfer-Function Measurement with Maximum-Length Sequences", J. Audio Eng. Soc., vol37, pp. 419-443 (1989 June).

[5] J. Nielsen, "Improvement if Signal-to-Noise Ratio in Long-Term MLS Measurements with High Level Non-stationary Disturbances", J. Audio Eng. Soc., vol 45, pp1063-1066 (1997 Dec.).

[6] Douglas J Button, "Heat Dissipation and Power Compression in Loudspeakers", J. Audio Eng. Soc., vol. 40, pp.32-40 (1992 Jan).