OBJECTIVE MEASUREMENT FOR EVALUATION OF HAPTIC AUDIO SIGNALS AND TRANSDUCERS

SG Oxnard Meridian Audio, Huntingdon, UK
LJ Hobden Meridian Audio, Huntingdon, UK
E Stanhope Meridian Audio, Huntingdon, UK
CSE Cotton Meridian Audio, Huntingdon, UK
FL Todd Meridian Audio, Huntingdon, UK, University of York, York, UK

1 INTRODUCTION

The human perception of sound is not restricted to auditory phenomena alone. In many cases, the aural reception of sound is accompanied by significant congruous tactile sensations arising from vibrations transferred to and through the human body. Consider, for instance, the reception of live sound reproduction or any other high sound pressure level (SPL) event, where it is commonly reported that sound can be 'felt' as well as heard. As such, the addition of vibrotactile sensations to the reception of sound culminates in a multi-sensory experience providing an increase in perceived level and sense of immersion in a reproduced sound field. The technology that underpins this outcome, known as "haptics", finds application in a range of industries such as cinema, gaming, automotive and virtual reality (VR). Haptic technologies have also been used to enable individuals with hearing impairments to perceive musical renditions through the delivery of vibrotactile sensations[1].

Broadly, existing haptic products can be categorised as wearables or seat-based solutions with some products currently commercially available[2,3,4]. Recent trends in the development and deployment of such technologies are indicative of an increased interest in the use and experience of haptic audio. However, while haptic transducers (or "shakers") and full haptic reproduction solutions have been researched[1,5] and put to market, there exists a lack of information in related literature on the characterisation, measurement and evaluation of haptic devices and systems. Moreover, the means by which appropriate haptic driving signals are derived from existing audio content for optimal tactile audio reproduction requires further investigation. Recent work[6,7] has demonstrated that a range of different haptic signal generation / derivation processes can be applied to create audio content for tactile signal production. Work to date[6] has revealed that the subjectively evaluated performance of different tactile audio production methods varies depending on personal preference and the genre of input audio programme material. Additionally, objective quantification of tactile transducer characteristics has been shown to be measurable using existing acoustic parameters[7].

The following provides extension to existing work on the objective evaluation of haptic transducer performance by means of physical measurement and analysis of commercially available devices. Influence is drawn from existing acoustic and signal analysis metrics to define objective measures pertaining to the characteristics of driver responses to input analytical signals and specific programme audio. Measurements and derived metrics demonstrate the means by which haptic audio signals and haptic transducers may be evaluated for the purposes of comparison and tuning of tactile audio reproduction systems. Modification of transducer responses was achieved by means of resonance reduction filtering and control of microdynamics of driving signals. Both modification strategies are shown to provide the intended changes to original drive unit responses through comparative objective analysis.

2 HUMAN PERCEPTION OF HAPTICS

In a recent study[8], it was reported that humans can sense sound-induced vibrations from their surrounding environment within the frequency range of 0.3 Hz to 1 kHz. This tactile perception range overlaps with the frequency range of human hearing, which is widely acknowledged to span the range

of 20 Hz to 20 kHz. The perception of haptic information occurs throughout the human body in specialised nerve endings called mechanoreceptors, which are found in cutaneous skin. The detection of stimuli such as temperature, pressure and vibration are all enabled by these nerve endings which are components of the human somatosensory system. The Pacinian corpuscle is the main mechanoreceptor of interest in terms of vibrotactile research as it is responsible for the perception of vibration with an active frequency response range of up to 500 Hz[9].

As there exists an overlapping range of frequencies between the active bandwidths of the human auditory and somatosensory systems (20 Hz – 1 kHz), it is evident that signals existing in the lower region of the audible frequency range should be used to drive haptic transducers. The frequency range of interest can further be informed by the active bandwidth of the Pacinian corpuscle in combination with measured response characteristics of haptic devices under test. As documented in the following section, the selected usable bandwidth of all haptic transducers ranges up to 250 Hz for the purposes of unbiased comparison of measured objective metrics. Given that the upper frequency range of the Pacinian corpuscle is 500 Hz, this upper bound on analysis frequency lies well within the response range of the primary mechanism of the somatosensory system. Moreover, the aforementioned lower frequency limit of the sensation of vibrations (0.3 Hz) nullifies the requirement for a lower frequency bound in terms of analysis. However, a means of DC signal blocking has been applied in the following in the interests of avoiding damage to the transducers under test.

3 MEASUREMENT OF HAPTIC TRANSDUCERS

3.1 Transducers under Test

Three commercially available drive units were tested in this work. In the interests of impartiality, the description of drive unit characteristics is restricted to the general properties of each unit. In the following, the transducers are referred to as Driver 1, Driver 2, and Driver 3. All three drive units operate via the principle of electromagnetic transduction with a moving mass suspended within the region of influence of a wound voice coil. The weight and contact surface area of these drive units is shown in table 1.

Table 1: Weight (g) and effective surface area (cm^2) of the three transducers under test.

Parameter	Driver 1	Driver 2	Driver 3
Weight (g)	78.8	340.2	1,232.0
Surface Area (cm^2)	18.85	84.95	132.73

As shown above, the selected drivers increase in both size and weight in line with increasing number (1 – 3). The noted weight is that of the complete unit (including housing and any connected mounting fixtures) and should not be taken as the effective moving mass. The effective surface area is that which is intended to be in contact with the surface to which the driver is affixed in practice. This selected range of drive units is such that the measurement and analysis procedures applied in the following may be said to be general (driver agnostic) for transducers that operate on the principle of electromagnetic transduction.

3.2 Measurement Process

Each driver was measured under a practical approximation of simply supported conditions in an anechoic chamber. In order to minimise the transfer of mechanical energy into the surrounding environment during test, the transducers were placed on a foam bed with the intended bonding surface (that which would be attached to a given surface in practice) facing upwards into free air. In this way, the only significant external force applied to each transducer during test was gravity acting on the moving mass perpendicular to the bonding surface. For the purposes of this study, it was assumed that the suspension of each unit is stiff enough to hold the moving mass in the natural resting position against any external force present in the measurement environment.

A Brüel & Kjær type 4375 accelerometer was attached to the centre position on the bonding surface of each drive unit under test. Varying the position of the accelerometer over the bonding surface was shown to have negligible impact on measured responses in the analysis bandwidth. The accelerometer in turn was connected to a Brüel & Kjær type 2635 charge conditioning amplifier to provide displacement measurements. Drive signals were provided, and responses recorded simultaneously by an Audio Precision (APx) 525[10].

For the purposes of analysis, two forms of driving signal were applied during testing. The first form was an Exponential Sine Sweep (ESS)[11] rendered at a sampling rate of 192 kHz covering a frequency range of 5 Hz – 2 kHz over a time period of 10 s calibrated to dissipate a power of 1 Watt RMS into each transducer, calculated using the nominal resistance in each case. Further driving signals were defined as filtered ESS signals for each drive unit as prescribed in section 3.3. The second form was an audio excerpt consisting of electronic drum samples (available online for reference[12]). Again, further driving signals were derived from this excerpt by means of filtering (see section 3.3) and microdynamics (see section 3.4) strategies. Power levels are given where appropriate.

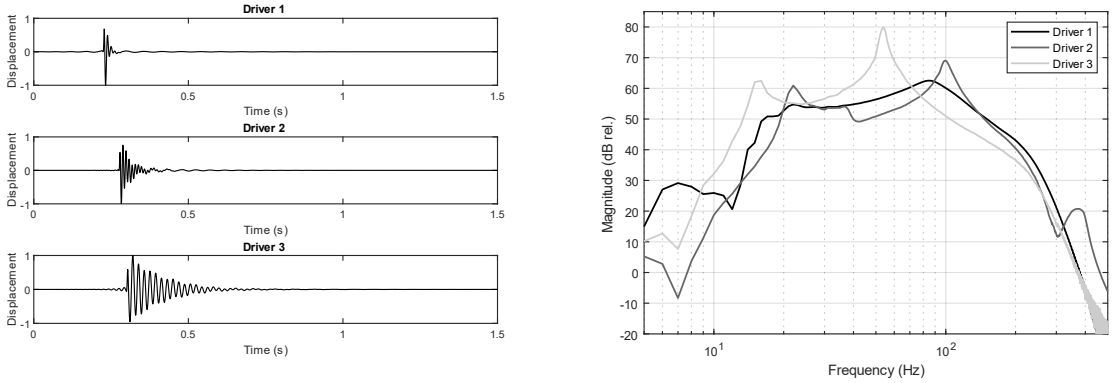

Figure 1: (left) Time domain representation of derived IRs for Drivers 1-3 (note: displacement, y-axis, is normalised). (right) Frequency response of Drivers 1-3 as calculated from IRs (not normalised).

3.3 Initial Response Analysis and Discussion

Impulse responses (IRs) were collected from each drive unit by means of the ESS measurement procedure (via deconvolution). Figure 1 displays the time and frequency domain representations of each derived IR. As shown, it is apparent from rudimentary analysis that the length of temporal response envelopes of the units under test increases with size/weight (see table 1 for reference). In the frequency domain, the cause of this increase in temporal decay envelope length is revealed to be drive unit resonances which become more pronounced as the transducer size increases.

This trend was also revealed in recent work[7] and the application of filtering strategies for resonance reduction was shown to decrease the decay time in resulting IRs. The same procedures are applied in this work to reinforce the claim that filtering to reduce resonance is a key component in altering and controlling haptic transducer performance. To this end, 2nd-order biquadratic peaking filters were applied with parameters as per table 2 to alleviate transducer resonances in each instance of unit under test. Note that one filter for Driver 1 has positive gain – this is to flatten the overall magnitude response after application of the resonance reduction filter applied at Fc = 82 Hz. Measurements relating to the resonance reduction filters are given in section 4.1.

Finally, it is noted from figure 1 (right) that the response of each transducer exhibits significant roll-off beyond f = 200 Hz and that Driver 2 presents a response node in the region of ~300 Hz. Hence, it is apparent that the response of Driver 2 begins to deviate from that of the other two above 250 Hz. Based on these findings, and with reference to the aforementioned tactile response of the somatosensory system, the bandwidth for comparative measurements was given an upper limit of 250 Hz in this work.

Table 2: Parameters defined for resonance reduction 2nd-order biquadratic peaking filters for Drivers 1-3.

Driver	Fc (Hz)	Q	Gain (dB)
1	40	1.5	2
1	82	2.4	-3
2	21	4.0	-22
2	102	2.6	-16
3	17	4.5	-13
3	54	6.0	-23

3.4 Microdynamics Modification

The transient or steady-state portions of the selected drum loop driving excerpt can be exaggerated by means of modulating the original signal with a modified envelope follower derived from the audio excerpt itself. The envelope follower applied in this work was implemented using the asynchronous full-wave envelope detector method[13].

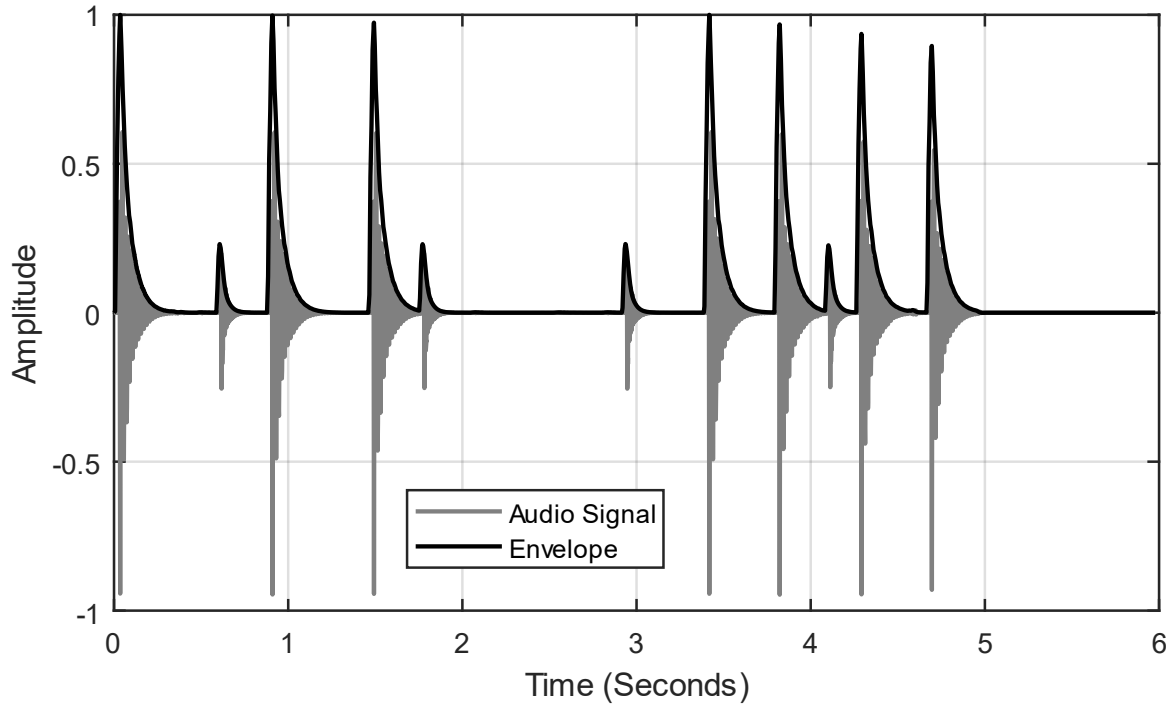

Figure 2: A plot of the drum loop excerpt used in this work overlayed with the detected envelope using the asynchronous full-wave envelope detector method[13].

The amplitude envelope, as per figure 2, was defined as a lowpass filtered version of the absolute amplitude. The applied lowpass filter was defined as a 2nd-order Butterworth filter with a cutoff frequency of 20 Hz. A gain can be applied to the resultant envelope, though its values must be bounded between 0 and 1. This amplitude envelope can then be modified to achieve either transient suppression or exaggeration. In the case of transient suppression (exaggeration of steady state signal components), the envelope was subtracted from a value of one prior to modulation with the audio signal. For transient exaggeration, the envelope can be directly modulated with the incoming data.

4 RESULTS

4.1 Resonance Reduction

The purposes of applying the resonance reduction filters (see table 2) to the driving signals for each transducer are to achieve less timbral response colouration and to gain further control of transient response characteristics. As shown via initial analysis, it is apparent that the lighter and smaller the transducer, the less pronounced response resonances are and, therefore, the quicker the transient response. This is again demonstrated in figure 3 which shows the normalised squared displacement yielded from IRs with (post) and without (pre) resonance reduction filters. As shown, the impact of the resonance reduction filters becomes more apparent as drive unit size and weight increase. In the case of the smallest unit, Driver 1, it is difficult to ascertain whether or not the resonance reduction filter provides a faster transient response. In contrast, for the cases of Drivers 2 and 3, it is clear from visual inspection of figure 3 that the post ring present in each IR is reduced.

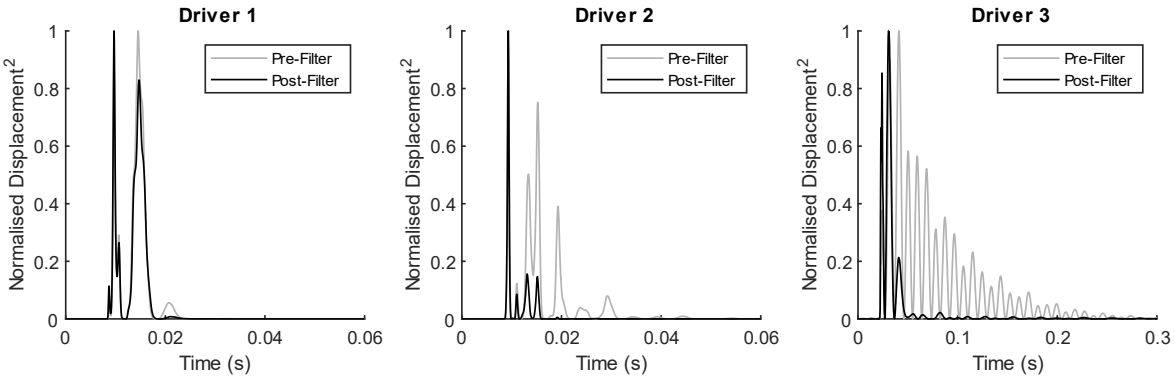

Figure 3: Normalised squared displacement calculated from the IR of each driver before (grey) and after (black) resonance reduction filtering. Note there are different timescales for Drivers 1 & 2 compared to Driver 3.

The Early Decay Time (EDT) and Early-to-Late Index (C_{50}) parameters[14] can be applied to yield objective data to support or disprove the claim that the application of resonance reduction filters is producing faster IR decay characteristics as intended. Values for EDT and C_{50} as calculated from the IRs recorded pre- and post-filtering for Drivers 1-3 are given in table 3.

Table 3: EDT and C_{50} values calculated from the measured IRs for Drivers 1-3 pre- and post-resonance reduction filtering.

	Pre-Filtering		**Post-Filtering**	
Driver	**EDT (s)**	**C_{50} (dB)**	**EDT (s)**	**C_{50} (dB)**
1	0.0762	21.8159	0.0658	19.4078
2	0.2195	12.5518	0.0620	19.9526
3	0.7343	1.9613	0.1335	11.7398

Given the nature of the natural transient response of Driver 1 and the comparatively subtle resonance reduction filter parameters, it is reasonable to expect little difference in measured characteristics between pre- and post-filtered responses. However, even in this case, the EDT of the IR has been shown to reduce as a result of the filtering. Yet, at the same time, the C_{50} decreases, which is contrary to expectations. This highlights a measurement case in which the application of such acoustic parameters for objective evaluation of haptic driver IRs leads to seemingly contradictory results. It is for this reason that two additional metrics have been utilised in this study to refine the objective evaluation of haptic transducer performance.

In the cases of Drivers 2 and 3, it is readily apparent that the resonance reduction filters are achieving the intended outcome in terms of reducing EDT and increasing C_{50}. An encouraging outcome is attained for Driver 2 specifically, whereby the resonance reduction filters have provided an IR decay profile that is objectively consistent with that of Driver 1. This suggests that haptic transducers of quite dissimilar physical properties may be tuned to produce similar response characteristics.

Further to analysis using acoustical metrics, two additional derived metrics are applied to demonstrate the impact of the resonance reduction filters. This marks a deviation from previous work[7] in which the changes in transient properties were evaluated through calculation of EDT and C_{50} only. Given that the systems under test in this work are largely independent of room interaction effects, the two additional derived metrics are defined to better reflect the nature of the systems. To this end, a combination of signal processing and acoustic analysis techniques culminate in the use of "Fall Time" (ms) and definition of "Transient-to-Late Energy Ratio" (TLER).

In the following, the Fall Time is derived from the normalised energy decay curve, *E(n)*, for each discrete drive unit IR, *h(n)*:

$$E(n) = \sum_{k=n}^{N} h(k)^2 \quad (1)$$

for discrete time sample index, '*n*', and IR length, '*N*' (samples). This represents the discrete form of the continuous room impulse response decay curve as specified in the related standard documentation[14]. The Fall Time is then derived from the signal *E(n)* as the time interval between the instant the normalised energy is at a value of 0.9 (90%) and the instant the normalised energy is at a value of 0.1 (10%). This type of metric is commonly encountered in analogue signal processing and control system engineering and is able to shed light on how quickly a system is able to slew between two states. In the case of haptics transducers, the two states are the onset of motion to cessation of motion. Hence, the lower the Fall Time, the faster the transient response.

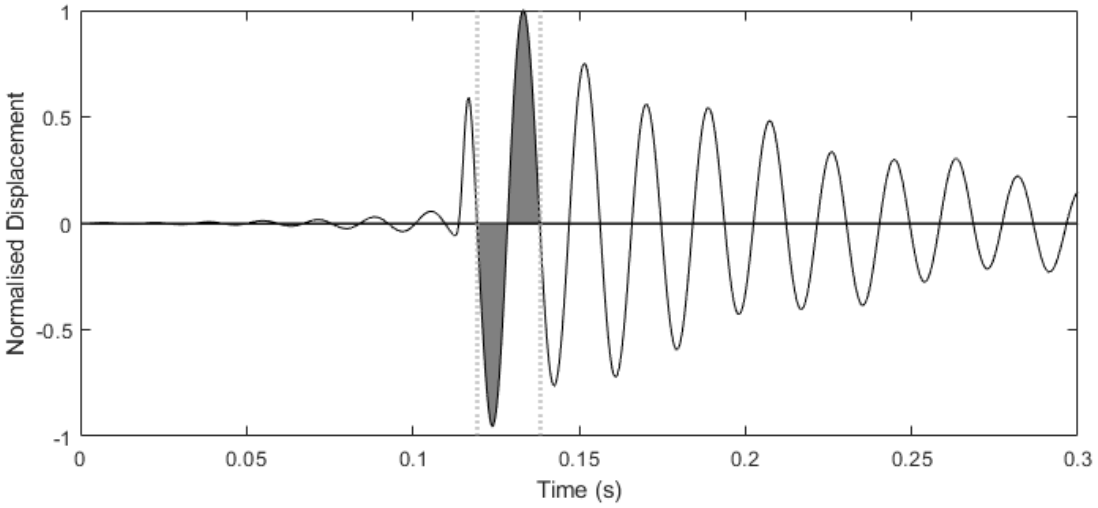

Figure 4: Graphical depiction of the definition for the transient temporal region as applied to the calculation of TLER (shown in the case of Pre-Filtered Driver 3 IR).

TLER is defined as similar to C_{50} with the distinction that the time period over which the early energy contribution is calculated is restricted to the initial transient event present in a given IR. In turn, this initial transient event is defined as spanning a temporal region around the maximum absolute displacement value present in the IR. The temporal region of interest is that which includes two zero crossing points before and one zero crossing point after the time instant of maximum absolute displacement. An example is shown in figure 4 where the initial transient event of the unfiltered Driver 3 response is defined. The late energy contribution is then calculated from the IR from the end of the

transient event to the time instant at which *E(n)* (as per Eqn. 1) reaches a value of -60 dB below the logarithmic value of $E(n)_{n=1}$ (i.e., the initial energy level (dB)). Note that the transient and late energies are defined as the squared displacement summed over the respective temporal regions. Taking the ratio of the transient and late energies calculated in this way gives the final TLER value for a given IR.

Table 4: Fall Time and TLER values calculated from the measured IRs for Drivers 1-3 pre- and post-resonance reduction filtering.

	Pre-Filtering		Post-Filtering	
Driver	**Fall Time (ms)**	**TLER**	**Fall Time (ms)**	**TLER**
1	10.94	4.27	9.86	4.69
2	31.65	0.68	14.54	7.37
3	112.98	0.48	20.21	3.43

Table 4 gives the calculated Fall Times and TLERs for each measured drive unit with and without resonance reduction filtering applied. The same trend of increase in reduction of the temporal measure (Fall Time) and increase in the energy ratio metric (TLER) is consistent with expectations and largely agrees with trends taken from EDT and C_{50} (specifically in the cases of Drivers 2 and 3). Again, the impact of the resonance reduction filters as applied to Drivers 2 and 3 is immediately apparent with significant reductions in measured Fall Time and increases in TLER. However, the aforementioned alignment of transient behaviour in post-filtered response metrics between Driver 1 and Driver 2 is no longer suggested using these metrics. The TLER of Driver 2 is greater than that of Driver 1 for post-filtered IRs and, therefore, the TLERs of both could be better aligned by relaxing the resonance reduction filter applied to Driver 2. However, the Fall Time of Driver 2 is longer than that of Driver 1 and so equalising the two drive units in terms of these objective measures may not be possible. It is interesting to note at this point that the perception of Fall Time / TLER, in terms of Just Noticeable Differences (JNDs) for example, remains to be investigated. Finally, it is shown that the refined metrics reveal the expected trends for the smallest unit, Driver 1, suggesting that these additional objective measurements may be better suited to the detailed objective measurement of haptic transducer transient response than EDT / C_{50}. In all measurement cases, the trends in calculated Fall Time and TLER match expectations.

4.2 Evaluating Dynamics Modification

Modifications to the dynamics of a signal were evaluated to verify the validity of the Loudness Dynamic Range (LDR)[15] and Crest Factor (CF) metrics for quantifying the dynamic properties of haptic audio reproduction. As described in section 3.4, the microdynamics of the selected electronic drums excerpt were exaggerated and suppressed by means of amplitude modulation using an envelope follower algorithm. This resulted in two versions of the input audio to be used for measurement. One version exhibited exaggerated transient events, while the other contained enhanced steady state content. Figure 5 (overleaf) shows a comparison of the steady state and transient extraction of a single drum hit compared to the original signal before this process. Each version of the input audio was filtered using the resonance reduction filters for each drive unit under test with dynamics modification applied before the filtering stage. The series of test signals was as follows: original drum sample excerpt; drum sample excerpt with respective resonance reduction filters applied; drum sample excerpt with respective resonance reduction filters and transient enhancement; drum sample excerpt with respective resonance reduction filters and steady state enhancement. This series of test signals enabled comparison of objective metrics measured from input audio that is more representative of common audio programme material (compared to ESS signals) with and without the additional processing step of microdynamics manipulation.

Each of the generated signals were applied to the respective drive unit under test as per the measurement setup and conditions given in section 3. Each signal was amplified to drive Drivers 1-3 within their operable power range to circumvent any potential issues arising from noise when taking measurements (1 W, 50 W and 75 W RMS respectively).

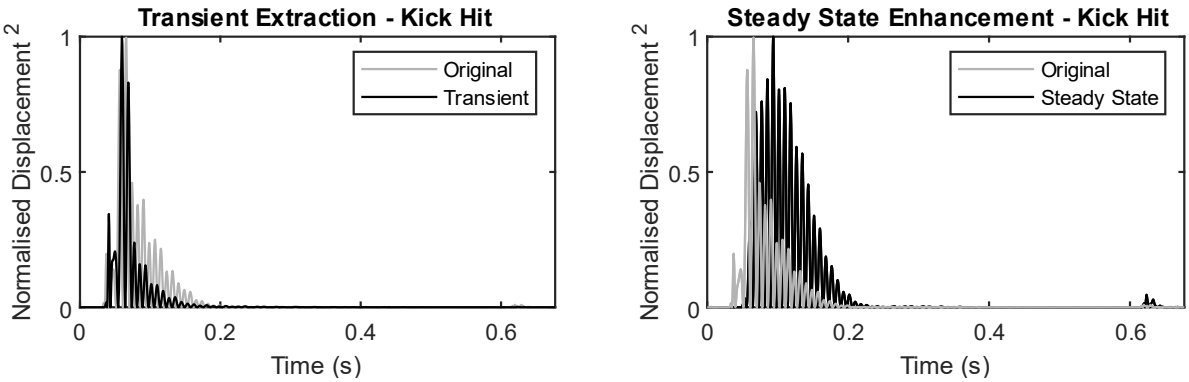

Figure 5: Normalised squared displacement of the transient extraction and steady state enhancement versions of a single kick drum hit from the drum sample audio excerpt compared to the unmodified version (original).

Values for LDR and CF (dB) were calculated from recordings of the input audio series as applied to each haptic transducer. The calculations were based on the full excerpt as reproduced under test by Drivers 1-3. Given the definition of LDR and CF (dB), it is expected that the values of each will change depending on the nature of the input audio signal (i.e., with/without resonance filtering and/or dynamics modification). In particular, both LDR and CF are expected to increase for resonance filtered and transient enhanced versions of the original audio and decrease in the case of steady state enhancement (effectively, transient reduction). Note that the measures of LDR were calculated using fast and slow integration times of 50 ms and 2 s respectively. The values of LDR stated in this work are the peak values (100th percentile) calculated in each measurement case.

Table 5: Results of the LDR and Crest Factor (CF) of a drum loop with/without resonance reduction filtering and/or dynamic processing. Values are given for Drivers 1-3 in dB.

	Original		Filtered		Steady State		Transient	
Driver	**LDR**	**CF**	**LDR**	**CF**	**LDR**	**CF**	**LDR**	**CF**
1	14.9204	7.7815	14.8470	7.7328	12.4718	5.8198	16.1880	8.8004
2	15.8818	8.5585	16.4777	10.3885	15.6277	5.2343	16.2552	9.9288
3	13.6738	5.2175	13.2022	6.8601	12.4539	4.6568	15.2442	7.5266

Table 5 shows measured objective metric values for each input signal and drive unit combination. It is evident that the steady state exaggeration is successfully reflected by both LDR and CF metrics for all drive units. This is shown by the reduction in values for both metrics when compared to those calculated from original and filtered measurements.

In the case of Driver 1, the application of resonance reduction filters has produced a very small degradation in dynamic range when compared to the original input measurement. Recall from the analysis of section 4.1 that Driver 1 exhibited the most marginal change in terms of measured transient response characteristics after the application of resonance reduction filtering. Moreover, the related C_{50} value decreased after filtering. Therefore, the slight decrease in LDR and CF for the filtered versus original is not entirely unexpected. Nevertheless, it can be seen that transient enhancement of the input audio signal has resulted in an increase in LDR and CF for Driver 1 compared to both original and filtered measurements. This indicates that even when the driver has a naturally transient response, pre-processing of the audio content can increase transientness.

The results for Driver 2 are particularly interesting. Firstly, it is noted that the LDR and CF values for the original and filtered measurements are higher than that of Driver 1. This is in contrast to the EDT and C_{50} measurements calculated from the impulse response measurements of the drivers. It is surmised that this is due to the interaction between the drivers' natural responses, their corrective filters, and the specific audio content used in this study. For instance, a resonant frequency range may be energised more in the case of Driver 1 where the same frequency range has little or no impact

on the resonance of Driver 2. This notion requires further investigation through measurement of a wider range of input programme material.

Another interesting result to note about Driver 2 is that the transient enhancement has resulted in a decrease in both LDR and CF, which is in contrast both to the other drivers, and contrary to what is expected. It is unclear exactly why this has occurred, but it should be noted that each drive unit measurement will contain a different noise floor levels and proportions of non-linear behaviour which could both affect these metrics. This is in contrast to the measurements of IRs which are assumed to be linear and have largely suppressed noise floor levels.

Finally, it is encouraging to note the results calculated in the case of Driver 3 from which it can be discerned that the processes of resonance reduction, steady state enhancement and transient exaggeration are all acting to yield the expected performance changes. Given that Driver 3 exhibits the most resonant (least ideal) natural response, these results give support to the claim that haptics transducers of comparatively poor response characteristics can be tuned / optimised and that the optimisation is quantifiable through analysis of programme material playback as well as IR measurements.

4.3 Summary of Results

The natural and modified responses of a series of three different haptic transducers have been measured and objectively quantified using existing and derived metrics. The behaviour of the EDT and C_{50} metrics that were examined in previous work[7] were discussed. As can be seen from the results shown in table 3, the metric of EDT behaves as expected for the drive units under test and seems to be useful when paired with C_{50} as an indicator facilitating calibration to match the transient profiles of differing drive units. It is noted that the temporal regions of energy calculations for C_{50} may cause the metric to not reflect expectations when used to examine transducers with especially short decay profiles, as seen from the results of Driver 1.

New derived metrics (Fall-time and TLER) are also evaluated in this work (see table 4) to analyse the effect of resonance reduction filtering on transient performance. These new metrics largely agree with the changes observed when the EDT and C_{50} metrics are deployed for analysis on signals before and after the application of resonance reduction filtering. Furthermore, it is demonstrated that these two metrics consistently align with expectations given the physical and response properties of the units under test.

The modification of signal dynamics is also evaluated by means of application of LDR and crest factor metrics. From the results in table 5, it is shown that both LDR and crest factor increase in line with the application of transient exaggeration and decrease when steady-state enhancement is deployed. This is when applied in conjunction with resonance reduction filtering and when compared to the response to unmodified input audio content. It is noted that of the two metrics evaluated, crest factor follows the expected trends more closely than LDR. Crest factor largely increases when resonance control filtering is applied. The range of LDR results suggest that further investigation be conducted to determine the effectiveness of this metric. The effectiveness of LDR as an objective measure in this context may be increased by selecting a different set of integration times or by taking the final values from a lower percentile.

5 CONCLUSION

This work has demonstrated the means by which haptic transducer response characteristics may be modified by means of manipulating input driving signals through simple filtering and dynamics modification strategies. Measured results taken from existing commercially available haptic drive units have shown that it is possible to alter and tune such devices, particularly in terms of transient response performance. Findings drawn from time and frequency domain analysis of captured haptic transducer IRs have informed the definition of resonance reduction filters and microdynamics modification algorithms. Application of filtering and dynamic control has been objectively evaluated

using existing (EDT, C_{50}, LDR, CF) and derived (Fall Time, TLER) metrics for the cases of three different haptic reproduction devices. In doing so, the means of objective performance evaluation of haptic transducers applied in this work have been proven to provide good insight to the reproduction characteristics of each unit under test. Such results can then be applied to the optimisation, tuning and performance evaluation of a haptic device when deployed as part of a multisensory sound reproduction system.

6 REFERENCES

1. S. Nanayakkara, E. Taylor, L. Wyse, and S. H. Ong, "An enhanced musical experience for the deaf: Design and evaluation of a music display and a haptic chair," Conf. Hum. Factors Comput. Syst. - Proc., 337–346 (2009).
2. S. Hashizume, S. Sakamoto, K. Suzuki, and Y. Ochiai, "LIVEJACKET: Wearable Music Experience Device with Multiple Speakers," 359-371. Springer, (2018).
3. Subpac, “SUBPAC X1 - SUBPAC,” https://www.subpac.com/x1/.
4. A. Haynes, J. Lawry, C. Kent, and J. Rossiter, "Feelmusic: Enriching our emotive experience of music through audio-tactile mappings," Multimodal Technol. Interact. 5(6), (2021).
5. M. Karam, C. Branje, G. Nespoli, N. Thompson, F.A. Russo, and D.I. Fels, ‘The emoti-chair: An interactive tactile music exhibit,’ Proc. Conf. Hum. Factors Comput. Syst. 3069–3074. (April 2010).
6. E. Stanhope, S.G. Oxnard, and L.J. Hobden, ‘Tactile Signal Derivation Part 1: Evaluation of Bandwidth Extension Methods,’ Presented at the 154th AES Conv. 1-10. Helsinki (2023).
7. S.G. Oxnard, E. Stanhope, and L.J. Hobden, ‘Tactile Signal Derivation Part 2: Extracting Haptic Information from Traditional Audio,’ Presented at the 154th AES Conv. 1-10. Helsinki (2023).
8. B. Remache-Vinueza, A. Trujillo-León, M. Zapata, F. Sarmiento-Ortiz, and F. Vidal-Verdú, ‘Audio-tactile rendering: A review on technology and methods to convey musical information through the sense of touch,’ Sensors 21(19), (2021).
9. D. M. Birnbaum and M. M. Wanderley, "A systematic approach to musical vibrotactile feedback," Int. Comput. Music Conf. ICMC 2007 397–404 (2007).
10. Audio Precision, “APx52x Series Datasheet,” https://www.ap.com/download/apx52x-series-datasheet-2/.
11. A. Farina, ‘Simultaneous Measurement of Impulse Response and Distortion with a Swept-Sine Technique,’ Presented at the 108th AES Conv. 1-24. Paris (2000).
12. E. Stanhope, S.G. Oxnard, L.J. Hobden, F.L. Todd, and C.S.E Cotton, ‘Objective Measurement for Evaluation of Haptic Audio Signals and Transducers,’ https://osf.io/6nvez/.
13. R. Lyons, “Digital Envelope Detection: The Good, the Bad, and the Ugly,” https://www.dsprelated.com/showarticle/938.php/.
14. International Organization for Standardization, “Acoustics - Measurement of room acoustic parameters - Part 1: Performance spaces (ISO 3382-1:2009),” (2009).
15. E. Skovenborg, ‘Measures of Microdynamics,’ Presented at the 137th AES Conv. 1-9. Los Angeles (2014).