Electro-Acoustics Group taking it up a level at Reproduced Sound 2020

The 36th annual Reproduced Sound conference, organised by the IOA's Electro-Acoustics Group (EAG), took place online from 17-19 November, 2020. The conference represented the cutting edge of modern audio and acoustics in an informal environment that allowed consultants, manufacturers, contractors, end users, academics and students to mingle and share insights and information.

By Adam Hill

Organisation of the conference was led by EAG Chair, Keith Holland (ISVR, University of Southampton), supported by the 11 committee members and IOA's Linda Canty. Complete online audio-visual support was provided by EAG committee members, John Taylor (d&b audiotechnik) and Ludovico Ausiello (Solent University), with support from student members, Sebastian Duran and Panos Tsagkarakis. d&b audiotechnik have generously provided technical support for Reproduced Sound for many years, to the great benefit of the conference.

Due to the ongoing pandemic, the conference was held virtually over three consecutive afternoons. The delegates numbered nearly 100, representing a good balance between industry and academia, with participants joining the conference from across the globe.

Conference – day one

Prior to the official launch of Reproduced Sound 2020, EAG committee member, John Taylor, introduced the technical team for the conference, explaining to all delegates the logistics for the event, including how to ask questions for each presentation (via Mentimeter) and how to access the virtual break room and lobby (via Jitsi). He then turned the virtual floor over to Keith Holland.

Keith welcomed everyone to the first ever virtual Reproduced Sound, hoping that despite everyone being remote, the conference would still achieve the characteristic friendliness that Reproduced Sound is known for, allowing easy access to a wide and diverse community.

Session 1 – Loudspeakers (Chair, Glenn Leembruggen)

Right: EAG committee member, Glenn Leembruggen

Effect of geometry on cone-driven midrange horns and phase plugs
The first paper of the conference was presented by Lewis MacDonald from the University of Salford, detailing research carried out as part of his undergraduate dissertation under the supervision of Jon Hargreaves. Lewis presented a methodical parameterised study on cone-driven midrange horns and phase plugs, where he specifically inspected the effect of phase plug radius, phase plug length, phase plug curvature, and horn geometry. This was done using FEM alongside a lumped parameter electromechanical model. Performance metrics included average intensity over 60 degrees as well as bandwidth and beamwidth achieved. Results pointed towards clear design considerations for constant directivity with minimal sidelobes. Questions from the audience focused primarily on the model's configuration, offering suggestions for further study.

Infinite waveguide termination by hybrid finite element/ series solution
The session continued with a paper from regular Reproduced Sound contributor, Patrick Macey from PACSYS. Patrick provided a thorough treatment of the underlying mathematics and modelling procedure adopted for this study, with several animated demonstrations to help the audience visualise the issues surrounding problematic terminations in models. The most revealing issues were related to cross mode problems due to an improper termination. A number of questions came from the audience, asking about the relationship between the simulation and reality. Patrick was able to clearly explain the relationship, emphasising the modelling method's accuracy.

Advances in acoustic instrument measurements and system design
The final paper of the session was delivered by Ludovico Ausiello from Solent University. This was a follow-on paper from Ludo's P24

paper at Reproduced Sound 2019, concentrating on furthering the understanding of acoustic instrument design through practical measurement and analysis methods. This work focused on the optimisation of an acoustic guitar's magnitude response, through a parameterised study (performed by carefully and systematically deconstructing an actual guitar). Areas under inspection included string tension, presence of varnish, closing of the body, shaping and position of internal braces, and use of additional braces on the soundboard and back. Through Ludo's analysis, the guitar optimised to exhibit a measurably flatter magnitude response. Many questions came from the audience, indicating considerable interest in the topic. Ludo's enthusiasm on the subject was clear, which made for a very engaging presentation. (See Ludo's technical article on page 30.)

Session 2 – Measurement 1 (Chair, Paul Malpas)
Temporal structure of spectral levels within pre-recorded material
The first conference session on measurement began with a paper from Glenn Leembruggen from Acoustic Directions, Australia, covering a detailed analysis of spectral content for a variety of pre-recorded signals. The data showed material dependent crest factors and spectral content. Unsurprisingly, speech was shown to have the highest crest factor, with metal music exhibiting the lowest. While classical music was relatively consistent across octave bands, jazz music was much less consistent. Ultimately, what this analysis pointed towards was the possibility to save resources on power amplifier requirements if the intended programme material is known. Glenn mentioned, however, that there isn't much published on the audibility of clipping, which should be investigated before committing to a system design process that essentially ignores the peak signal requirements (due to a high crest factor). The presentation prompted numerous questions from the audience. Session Chair, Paul Malpas, was quick to point out that such a system design approach isn't appropriate for hi-fi systems as they require accurate reproduction of peak content. Glenn added that even outside of hi-fi systems, care must be taken before deciding to limit any low frequency content, since this would be audible. This could likely be addressed through use of a multiband compressor.

Energy-Time Curve (ETC) in electro-acoustic measurement and analysis
Next on the agenda was a paper from James Love and David Gilfillan of Gilfillan Soundwork in Australia. James provided a clear overview of energy-time curves (ETC) including how they are mathematically derived from an impulse response measurement. What quickly became clear in this study was that spectral windowing was essential to get right when using ETCs, otherwise a considerable amount of information could be lost (such as with the use of a Blackmann-Harris window). With proper windowing, however, detailed acoustic information can be extracted from measurements such as low frequency reflection arrivals, identification of closely spaced peaks, and a reduction in time domain aliasing. The signal processing was nicely demonstrated using impulse response measurements from real spaces.

Use of artificial intelligence in room acoustics prediction using a photograph
The final paper of the session was presented by Dan Milne from Solent University and was based on his undergraduate research under the supervision of Lee Davison and Ludovico Ausiello. The research looked into whether there was a reliable method of extracting acoustic information of a space from a 2D photograph. This was explored by utilising a convolutional neural network (CNN) with a focus on estimation of RT60. The CNN was trained with data from 38 classrooms, comprising 24 photographs per room. RT60 measurements were taken in each room using a balloon burst with RT60 at 500 Hz used to characterise the room in this instance. After the CNN training was complete, tests were carried out where the photographs were sent to acoustics experts asking them to predict the RT60 of the room, while the same photographs were sent through the CNN to do the same. The CNN outperformed the experts, where the experts tended to overestimate RT60. There was a potentially unfair advantage for the CNN, however, in that it was limited to guesses between 0.4-1.0 seconds, while the experts could guess any number. Dan usefully identified shortcomings such as this within the test and offered suggestions for further work. The audience had many questions for Dan, with some sceptical about how useful this could be for more general spaces (the study was limited to classrooms), but the interest in the topic was nevertheless high.

Session 3 – System design and modelling

Right: EAG Chair, Keith Holland

(Chair, Keith Holland)
Modern advances in the meeting room ecosystem
The first paper in the system design and modelling session was from John Ellis and Andrew Francis of Shure UK. They presented a thorough overview of the history of meeting room audio-visual technology, leading to the current state of the art, with examples provided in the form of current products offered by Shure. They revealed that past surveys have indicated that most meeting room technology frustrations are related to audio, but that current technology such as networked audio and auto mixing can help to prevent many issues. They were able to give an example very close to home – the IOA's new headquarters in Milton Keynes. The IOA's former meeting room AV system was the source of much frustration. In the new facilities, however, care was taken to implement good acoustic and audio-visual design, which EAG Chair, Keith Holland, was quick to confirm has resulted in much more effective meetings. During Q&A it was emphasised that such systems are still reliant on good room acoustics.

Understanding modern amplification systems
The next paper was presented by Alberto Fueyo Gallego from AMS Acoustics. As with

Glenn Leembruggen's paper from the previous session, Alberto's research focused on power amplifier requirements primarily for the purpose of emergency sound systems. As with Glenn, he suggested that systems may not need so much power due to the high crest factor of speech. He emphasised that clipping and compression don't seem to significantly affect STI, hence designing systems to accurately reproduce sound based on peak levels may be overkill. Alberto's presentation prompted supportive comments, suggesting the need for further work in the industry to consider such ideas, including possible work within standards.

An archaeoacoustics investigation of the Beulieu Abbey

The first day of the conference was concluded with a presentation by Sebastian Duran, Martyn Chambers, and Ioannis Kanellopoulos of Solent University, detailing their research which was under the supervision of Chris Barlow. The paper focused on work to achieve an accurate auralization of Beulieu Abbey, which was destroyed nearly 500 years ago. The modelling was carried out in CATT Acoustic, where the sound sources were modelled to be a priest giving a sermon and monks doing chants. In addition, auralizations and data on STI and clarity were presented, where it was found that the acoustics of the Abbey didn't support speech but would be appropriate for the chants. This aligns with the fact that intelligibility was less important before Vatican II, when mass was conducted in Latin. The students were able to field questions from the audience, which focused on aspects of the acoustic model's design.

Conference – day two

Session 4 – Signal processing (Chair, Ludovico Ausiello)

Analysis of neural network architectures for audio signal processing

Day two of Reproduced Sound commenced with a paper from Vlad Paul of ISVR, University of Southampton, with a focus on his research on neural network architectures under the supervision of Philip Nelson. Vlad pointed out that while the research community has seen widespread adoption of neural networks for various purposes over the past few years, most projects utilise these systems as black boxes, where details of specific parameters during the training process are unknown. To illustrate the usefulness of having sight of these parameters, Vlad utilised multilayer perceptron (MLP) to determine the hidden layer variables. Spectral difference detection was used an as example application, whereby the spectra under analysis contained magnitude peaks spaced too closely to accurately resolve with conventional analysis techniques. Through this process, Vlad was able to demonstrate how the neural network can be optimised when the often-hidden variables are revealed. Questions from the audience focused on potential applications for this neural network optimisation approach, while other commented that it was nice to see such a topic presented in an accessible manner to non-experts in the audience.

Automated delay estimation and time alignment in a reflective environment

The next paper in the signal processing session was from Reproduced Sound regular, Ambrose Thompson from Martin Audio and Alexander Holt from the University of Surrey. They presented their work on time-alignment of sound reinforcement systems in reflective environments. They described an optimisation routine that they developed to determine the optimal delay parameters to achieve an ideal magnitude response (characterised by summing all available magnitude responses without phase information). As the system contained a single variable, an exhaustive search was possible. This was shown to be resilient in the face of room modes and strong reflections. Questions focused on the practicality of the approach, given that the optimisation was based on a single measurement location, despite having a wide audience area.

Portable synthesizer on embedded system

The session on signal processing was concluded with a presentation by João Davi de Campos, which described his undergraduate research under the supervision on William D'Andrea Fonseca at the Federal University of Santa Maria in Brazil. After giving the audience an overview on the history of analogue and digital synthesizers, João went on to describe the synthesizer implementation he developed using an MDI keyboard and a Raspberry Pi 3B+, which was running Zynthian software (an open source synthesizer platform). Measurement examples from the device's output were shown along with an impromptu live demo during the Q&A.

Session 5 – Measurement 2 (Chair, Adam Hill)

Development and use of a low cost acoustic flow resistance meter

The second measurement session of the conference began with a presentation from Camille Hanrahan-Tan of Acoustic Directions, Australia. Camille detailed her work on developing and testing a low cost acoustic flow resistance meter, which cost approximately one-tenth of a similar commercially available device. She gave ample evidence of the device's performance, with example measurements from a variety of materials. The measurements were in good agreement with predictions and any anomalies were clearly explained. Several questions were received from the audience, focusing on material properties and methods to convert between various metrics. Overall, the impression was a positive one, where it was encouraging to see developments to make specialist equipment more accessible within the industry.

Suitability of hi-fi loudspeakers measuring reverberation time in domestic rooms

The session continued with a paper from Chris Adair of Adair Acoustic Design, which focused on using typical hi-fi loudspeakers to measure reverberation times in domestic environments.

P26

Below: Reproduced Sound technical crew (clockwise from top left: John Taylor, Ludo Ausiello, Sebastian Duran, Panos Tsagkarakis)

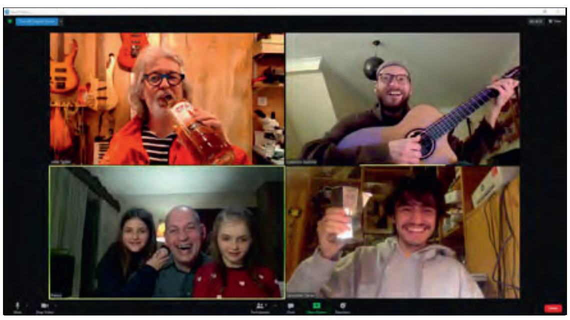

Chris emphasised the balance between acoustics and reality, primarily questioning whether we need to measure all spaces with dodecahedral loudspeakers. Considering home theatre systems, why not measure with the in-situ sound system? He noted that there is no requirement within ISO 338-2 for an omnidirectional speaker for surveying and engineering applications. Comparisons of results between conventional and dodecahedral speakers were given, with good agreement at high frequencies and a small amount of variation at low frequencies. Chris's experiment shows that the hi-fi speakers could give sufficient accuracy for RT60 measurements in small rooms, with measurement errors all falling below the just noticeable difference of 10%. The audience responded with many questions about future applications for this, indicating broad support for this concept.

Audio signal statistics revisited: A homomorphic separation approach
Building on an emerging theme of music and speech signal statistical analysis at this year's conference, Jamie Angus-Whiteoak from the University of Salford presented her paper looking into a homomorphic separation approach to audio signal analysis. The issue at hand was that signals that have been combined in a multiplicative manner can't be separated, since original spectral content may not be present as distinct content. To overcome this, Jamie explained how the combined signals can be transformed through an invertible function which converts the signal to an additively combined signal, allowing for straightforward separation. As an example, Jamie applied the process to a speech signal, which in turn confirmed the findings of some previous papers in the conference, showing that speech and music signals aren't gaussian, where these signals spend most time at lower levels – adding weight to the argument for lower power amplification requirements for efficient (non hi-fi) systems. Questions from the audience focused on the signal analysis technique's applications within dynamic processing such as compressors, limiters and even wireless audio systems.

Determining the source of coherence reduction using playback level of M-Noise
The final paper of this session on measurement was from Roger Schwenke of Meyer Sound, USA, where he presented a hands-on look into how sources of audio signal coherence degradation can be determined with M-Noise. Signal coherence is often affected by background noise, distortion, and the acoustics of a venue. With the underlying theory explained, Roger turned to a live demo of M-Noise being used to measure coherence of a simulated system. The results were systematically analysed while adjusting key parameters of the simulated system to show the tell-tale signs of each form of coherence reduction. Overall, this presented a useful toolbox for system engineers to use while on-site tuning sound systems to avoid incorrect analysis of data. Session chair, Adam Hill, noted that the paper session literally stretched around the globe, starting in Australia, stopping off in the UK, and concluding in California.

Electro-Acoustics Group AGM
The AGM of the Electro-Acoustics Group was held prior to the close of day two of the conference. The meeting was chaired by Keith Holland and was attended by 25 delegates, including eight EAG committee members. Keith delivered the chairman's report, describing all activities of the group over the past year, the central focus being the organisation of this conference. It was noted that this year required a very steep learning curve in order to adjust Reproduced Sound to operate effectively in an online space. Notably, Keith highlighted the ongoing efforts of John Taylor and Ludo Ausiello in overcoming the technological challenges this has presented and permitting the delivery a very enjoyable and successful conference. It was agreed that many lessons have been learned over the course of this year and the committee will consider what to roll forward to benefit future Reproduced Sound conferences.

Conference – day three

Session 6 – Virtual audio (Chair, Jamie Angus) WHAM: Webcam head-tracked ambisonics
The final day of this year's Reproduced Sound conference was kicked off by Mark Dring and Bruce Wiggins from the University of Derby who presented their work on WHAM – Webcam head-tracked ambisonics. This research was necessitated by the national lockdown in March, where there was no access to their research group's specialist facilities. To overcome this, a workaround was developed to allow for head-tracked ambisonics to be available to anyone with a webcam and a pair of headphones, where the webcam was used to track head movements alongside presentation of binaural room impulse responses (BRIR), where the software included the necessary support for asymmetrical filtering. In this first instance, only horizontal head rotations were considered (where the head was rotated within the ambisonics system, not the virtual room). Live demos were given via a purpose-built web tool hosted at **brucewiggins.co.uk/WHAM**. Due to the current capabilities of web-based audio, the system was limited to 7th order ambisonics (although Bruce noted that this restriction may be eased soon). The demo gave clear evidence that the webcam-based tracking was effective, giving high-quality ambisonics-based sound reproduction over headphones and (critically) allowing subjective evaluations to continue despite the current restrictions.

Rendering binaural signals for moving sources
The second and final paper of this session on virtual audio was presented by a trio consisting of Lucas Gomes, William D'Andrea Fonseca, Davi Carvalho, all from the Federal University of Santa Maria, Brazil. The work focuses on accurately rendering binaural signals for moving sources. Critically, the Doppler effect had to be considered in these situations, which was implemented within an image source acoustical model. Interpolation was required in this case to account for the variable time of arrival of successive samples from moving sources. P28

Below: Paper from Mark Dring and Bruce Wiggins

With the moving source model in place, an investigation was carried out looking into the perception of the rendered sound sources in terms of the perception of speed, distance, position, realism, and number of simultaneous sources. It was found that participants were able to correctly identify relative distance and speed, but further work was required to achieve better externalisation over headphones. A key question from the audience revealed that the ground reflection was omitted from the model, where it was agreed that this should be included as part of further work.

Session 7 – Case studies and regulation (Chair, Mark Bailey)

The acoustic design of the Qingdao Oriental Movie Metropolis Grand Theatre

The final session began with a presentation by Shenzhi Su from CSP Acoustics, where she walked the audience through the electroacoustic design of the Qingdao Oriental Movie Metropolis Grand Theatre in China. The 1,970-seat theatre was specifically designed to host a film festival but would have to be flexible to accommodate a range of different events after the festival. Interestingly, it was stated that the theatre is the only one in the world with a Dolby Atmos system as well as an electronic room acoustic enhancement system. During the design stages of the theatre, it was determined that the venue's use would be as follows: movie projection (80% of all events), award ceremonies, and symphonic performances, each requiring a different acoustic. The venue's construction was explained, with a particular focus on the incorporation of the necessary absorption and diffusion. The theatre was initially evaluated by twenty acousticians, sound engineers, and musicians. All groups rated the sound quality as excellent with natural sounding reverberation. Quite a few questions came from the audience, indicating a good deal of interest in the project.

Impact of sound level regulations on sound engineering practice

The second paper of the session was jointly presented by Adam Hill and Jon Burton from the University of Derby. The work focused on a case study looking into the effect of local sound level regulations on live sound engineering practice. A dataset was used encompassing roughly five years of touring and festival dates (130 in total) from a popular UK-based band. For each event, one minute equivalent continuous sound levels were recorded with A and C weighting, allowing for later conversion to any required measurement time window greater than one minute. A number of statistically significant points were drawn from the data. Room acoustics play an important role in terms of overall sound level in the audience. Venues with poor acoustics resulted in louder shows. Additionally, small venues were louder than large venues, since loud stage levels are more significant within the audience area in small venues. Events with an LAeq limit in place were roughly 2 dB quieter overall, but this was only for events with limits at or below 101 dBA (indicating the engineer's natural mixing level to be around 100 dBA). In terms of the time measurement window, it was clear that short times (less than 10 minutes) reduced the dynamic range of the event, due to the engineer having to constantly adjust to maintain compliance with the limit, reducing the possibility for strong dynamics in the reinforced music. In the few cases where an LCeq limit was in place, there were clear issues with the engineer struggling to maintain compliance due to arbitrarily defined limits.

UK and international guidance for the control of noise from outdoor events

Next was an overview of noise regulations in the UK and internationally, presented by Peter Wheeler from Vanguardia. Peter presented a detailed comparison between sound level regulations across the UK, with clear evidence of confusion, either through misinterpretation of the UK's Noise Council Code of Practice on Environmental Noise Control at Concerts or basing limits on best practice and not from guidance documents. While there is clear evidence of the steady ride in low-frequency content within popular music, there are few regulations with low-frequency (LCeq) limits. Of the few areas with such limits, there is little agreement in specific limits and applications, which agrees with the findings from Adam Hill and Jon Burton's paper. Peter explained that this points to a need for updates to the existing guidance documents (work which is currently underway). Many questions came from the audience, primarily focusing on the issue surrounding low-frequency limits. It was clear that the regulations are lagging reality.

The Sound of SoFi – amplifying the experience

The final paper of Reproduced Sound 2020 was presented by Jim Burdette (JBL/Harman), Kevin Day (WJHW), and Demetrius Palavos (Pro Media) on the sound system design for the new SoFi Stadium in Los Angeles, USA. The overarching design criteria for the sound system was to fit most of the sound system elements into the scoreboard, the largest of its type in the world, which was flown centrally above the playing field. This approach went against the current trend of implementing distributed systems in such venues. Kevin Day took the audience through development and testing of the system in EASE, describing key considerations along the way. All presenters made it clear that it was critical to work closely with other departments to make sure nothing interfered with the sound system's operation once it was installed (such as using acoustic mesh to obscure the speakers within the scoreboard to not affect them acoustically). The presentation generated many questions and comments from the audience, unsurprisingly, due to the large and unique nature of the project.

Conference close

EAG chair, Keith Holland, closed the formal proceedings of Reproduced Sound 2020, by thanking all those who were involved with the organisation and running of the event as well as the delegates for attending. Keith specifically thanked John Taylor and Ludo Ausiello, along with assistants Sebastian Duran and Panos Tsagkarakis, for developing and running the technical aspects of the conference, allowing for a smooth and engaging conference experience for all attendees and presenters. He also expressed thanks to the IOA's Linda Canty for her continued support and guidance. With that, Keith formally closed the conference and expressed hope that we'd all be able to see each other in person for Reproduced Sound 2021.