COMMON PITFALLS IN COMPUTER MODELLING OF ROOM ACOUSTICS

I Rees Adrian James Acoustics Ltd, Norwich, UK

1 INTRODUCTION

1.1 With advanced tools, advanced mistakes can be made

The above heading is one of the author’s favourite observations by Bengt-Inge Dalenbäck, developer of CATT-Acoustic. Adrian James Acoustics Ltd has acted as distributor and provided technical support for CATT-Acoustic in the UK and Ireland for many years and particularly over the past five years has delivered regular training courses to consultants and academic staff hailing from throughout the UK, Ireland and sometimes beyond. We have always started off the courses with a smattering of such quotations, partly to give the morning caffeine time to kick in before we start on the heavy technical stuff, but mostly to focus the mind and encourage the right approach to using the software. Indeed, the key aim of our training is not only to explore all the technical nuts and bolts of the software but to encourage users to consider carefully how and why they are using the software, what "Geometrical Acoustics” (GA) means, and to take the time to analyse and challenge the results obtained from calculations.

This paper considers some of the common pitfalls and misconceptions surrounding the use of GA-based room acoustics modelling software that we have encountered and offers some general advice on avoiding these. As our day-to-day experience and expertise is with CATT-Acoustic, technical discussion of programme features relates specifically to this software, but general principles will apply to other GA-based software offering similar technical features.

1.2 Understanding limitations

It is worth highlighting that CATT-Acoustic, Odeon and other acoustic modelling programmes employing processes such as image source modelling and variants of ray or cone tracing fall into a family known as “Geometrical Acoustics” (GA) software. This is a term which seems to have fallen out of common usage in recent years but it is usefully descriptive. GA modelling programmes are essentially energy-based and do not solve the wave equation. These therefore offer only approximations of how sound propagates and interacts within a space, based on the geometric features and surface properties of the space. At the core of the GA limitation is that the effect of object and detail size in relation to wavelength is not handled 'exactly', as with wave-based methods, but using approximations. It is important to have a good understanding of these approximations, both in terms of how the calculations work and the limitations of when they apply, which requires a good practical as well as theoretical understanding of room acoustics.

One of the common misconceptions surrounding CATT-Acoustic and other GA software is that such programmes are design tools which require very little specialist acoustic knowledge or time to learn. It is understandable how such a view can arise, particularly if one has experience of loudspeaker modelling tools such as EASE Focus, Duran Audio DDA, or other loudspeaker manufacturers' direct sound array aiming tools, which will perform a lot of clever automatic alignment routines for complex loudspeaker arrays with relatively little input required from the user. It is of course important to appreciate that these types of software are also based on limited simulations; where reverberation is handled, classical diffuse-field theory is often used where it may not really be applicable.

1.3 GA software as a virtual measurement tool

It is often helpful when teaching new users to recall the origin of GA software as a “virtual” measurement tool, as an alternative to acoustic measurement using physical scale models. This can serve to establish the place of computer modelling within the design process. The first and most crucial step in design is to establish the design aims and then to identify numerical criteria. Concept design then follows, in which the general treatment approach is determined, which is then refined in detailed design. It is at this point that a measurement tool should be employed to test the design and if necessary inform further refinements. There is no value in testing without first establishing clear criteria and a firm design approach. The notion of GA software as a measurement tool also carries a couple of further implications:

1. As acousticians we are familiar with the sound level meter / real-time analyser as a measurement tool. While these are relatively simple devices to operate, it is accepted that specialist skill and knowledge is required to take good, accurate and repeatable measurements (particularly with sufficient source and receiver sampling positions), to analyse the measurement data meaningfully and use these measurements to inform the acoustic design.

2. Acoustic measurement in real spaces carries a degree of uncertainty due to a variety of factors including measurement methodology and room conditions. The same is true of measurements taken via GA modelling. While the selection of source and receiver positions can affect the uncertainty and repeatability of results, the greatest sources of uncertainty are the input data used to create the model, be that the selection / estimation of acoustic surface properties or assumptions made in the creation of the model geometry itself, and errors in the model geometry if this is not checked and debugged Such uncertainties should be carefully considered when creating models and assessing the validity of the results obtained. It needs to be emphasised that even if a model is simplified, due to GA requirements, the model must still be accurate in itself.

2 BUILDING, TESTING AND DEBUGGING MODELS

2.1 Considerations in building models

It must be stressed that the validity of results obtained from GA models, and indeed any calculation process, is heavily dependent on the quality of data input. It is crucial that the model geometry, absorption and scattering coefficients, source and receiver positions are selected with great care and with good understanding of the likely acoustic effects of these choices. As a good example, a round-robin on room acoustic modelling of a large classroom in Odeon undertaken by the Danish Acoustical Society in 2008[1] demonstrates the wide range of results which can arise when different acousticians are given the same room to model, and estimate surface data for, using the same GA software.

When creating model geometries, simplicity is key. The accuracy of GA modelling relies on the size of surfaces being significantly larger than the wavelengths of sound under consideration. It is important to stress therefore that models featuring very small surfaces will only be valid at high frequencies. While some programmes including CATT-Acoustic have separate calculation processes such as auto edge scattering and diffraction to handle instances where surface sizes are relatively small compared to wavelength, these instances should be seen as exceptional rather than the norm. It should also be noted that auto edge scattering serves to lessen the effects of GA limitations at low frequencies, but these limitations still remain.

As a matter of personal opinion, there is great satisfaction to be gained from finding the simplest possible way of expressing a complex structure sufficiently well for GA prediction and the assessment task at hand. This simplification process needs to be driven by an assessment of features within a space in terms of their acoustic effects. Features such as exposed steelwork, intricate plaster or stone mouldings or blocks of seating or other furniture may be quite complex, but it is important to consider the net effect (e.g. absorption, scattering or focusing) of such features on sound propagation. An understanding of the expected effect of a given feature on incident sound should guide how the feature should be expressed. Similarly, the extent of detail required in a model should be driven by the needs of what the acoustician aims to assess. For instance if using a GA model to assess intelligibility of a sound system, there is a need for room surfaces close to the sources and receivers to be reasonably accurate in their position and orientation as these will have a significant effect on early reflections. However, other parts of the space further from the listening area may not be significant in terms of early reflections and these areas could be modelled much more simply.

To illustrate this, Figures 1 and 2 show two models created in CATT-Acoustic to assess the intelligibility of digitally-steered column loudspeakers in the Nave seating area of a cathedral. The seating area is coloured brown in both models. The first model shown in Figure 1, while visually very pretty, is not an effective model for GA prediction. The majority of fine detail is at high level in the triforium and roof zones, which is well outside of the coverage of the loudspeakers and therefore the exact form of this fine detail is acoustically unimportant. It should also be noted that this model is incomplete, with obviously missing planes, and is therefore acoustically invalid.

In the context of the aims of the model, namely to assess sound system intelligibility in the Nave seating area, the roof, transept and choir merely add diffuse reverberant sound. The second model shown in Figure 2 was therefore created as a much simpler volume. With the exception of the columns flanking the Nave, the majority of fine detail was omitted. The roof area was represented with a single plane with high scattering values. Absorption coefficients for the roof were calibrated to match reverberation times measured in the cathedral. The model therefore fulfils the needs of the assessment task, but uses only 200 planes.

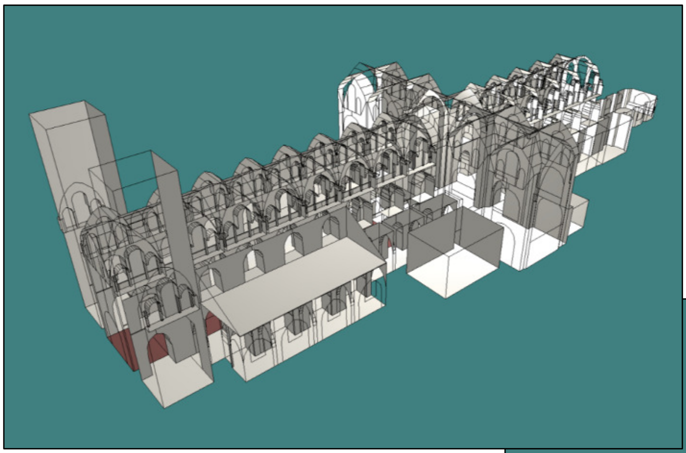

◀ **Figure 1 – Cathedral model Version 1**

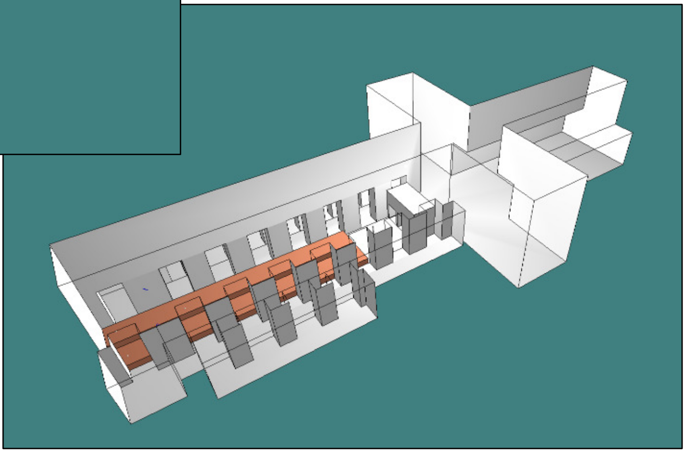

Figure 2 – Cathedral model Version 2 ▶

2.2 Verifying and debugging models

It is often surprisingly necessary to point out to new users of CATT-Acoustic that even if models are to be built in SketchUp and converted using one of the available third-party converters they will not automatically be accurate and correct; knowledge of the structure and syntax of geometry and location files is imperative. This is particularly important to properly check and debug model geometry. Modelling errors such as warped, overlapping, reversed or missing planes are common sources of modelling inaccuracies. Checking and debugging models is an essential, yet often neglected, step in building and verifying models. It is worth noting that debugging functions are not always enabled by default in CATT-Acoustic (although it is possible to reconfigure the default settings) so inexperienced users may easily overlook these. The geometry checking procedures for debugging can be time-consuming in complex models so it may be advisable to turn this option off again once a model has been checked and debugged.

Other commonly overlooked features in CATT-Acoustic are the Interactive RT Estimation module (assuming that the room behaves in a sufficiently diffuse way) and the source and reflector info plots (enabled using the ‘Source info’ and ‘Refl. Info’ settings under the ‘Geometry view/check’ dialogue box), which offer opportunities to make initial coarse adjustments to the model without requiring a full calculation run in TUCT. While it can be tempting to dive into full calculation runs as soon as a model is created, this can quickly lead into an inefficient cycle of repeated iterative adjustments (often informed by little more than trial and error) and lengthy calculation runs. A great deal of adjustment and verification of absorbent finishes, reflector coverage and direct sound system coverage can be undertaken without requiring a full calculation run. Considerable time can be saved by checking that the model is behaving correctly and verifying the validity of initial results before committing to a full calculation run. It is similarly unproductive to always use the more complex Algorithms 2 or 3 in TUCT for cases where these are simply not needed. The differences between TUCT prediction algorithms are discussed in Section 3.2.

3 UNDERSTANDING CALCULATION PROCESSES

3.1 Choosing scattering properties

In GA modelling, surface scattering is very important and serves two purposes. Firstly it is required to describe scattering due to rough (bumpy) surfaces; as discussed in Section 2.1, it is necessary for successful GA modelling to use simple models with large surfaces relative to wavelength and scattering should be used to convey the acoustic effect of fine detail on these surfaces. Secondly, GA modelling uses a finite number of rays or cones to approximate the propagation of curved pressure wave fronts in a fluid medium. At least a small degree of scatter is therefore always necessary even in rooms with very large flat reflecting surfaces (sports halls being a common example).

Selection of appropriate surface scattering coefficients is critical to the success of GA prediction. In many geometrically non-mixing spaces, variations in scattering can have a larger effect on the accuracy of results than similar variations in absorption[2]. While selection of scattering properties should of course be informed by the size, shape and surface roughness of the surface, it is useful to consider the space as a whole and have a good understanding in advance of whether the space is expected to be ‘mixing’ or not. Although academics may balk at the suggestion, as consultants we generally prefer to design to a reasonable worst case to give ourselves at least a little margin for error. If a non-mixing space is anticipated, for instance in sports halls or unusually-shaped atria, it is perhaps preferable to err on the side of slightly underestimating rather than overestimating the scattering of anticipated ‘problem surfaces’. Of course, this should always be informed by an understanding of what is ‘realistic’ and a healthy sense of moderation; implausible extremes of absorption or scattering should normally not be used even if the resulting numbers meet expectations (or hopes).

It can be useful to compare the results of initial quick calculations run with 100% scattering and 0% scattering on all surfaces to 'get to know the room' (TUCT2 has a special export function to make the process very simple). If results with 0% scattering are already close to the Sabine RT then the model will not be very sensitive to scattering (the geometry is mixing and the absorption distribution is fairly even). Results with 100% scattering should normally be very close to the Sabine RT; if these are not then there is likely to be a model issue such as overlapping planes or a very unusual shape where GA prediction may not work well.

In practice, in typical classroom-sized rectangular rooms, where walls are regularly interrupted by windows, doors and furniture, we would typically select scattering coefficients of around 20%. In auditoria with more curved and broken surfaces, a typical scattering coefficient of 30% tends to provide more representative results. In either case, the frequency-dependence of scattering coefficients must be carefully considered; this will be influenced by the size of surfaces and details in comparison to the wavelength. The auto-edge scattering function of CATT-Acoustic provides useful assistance in estimating this size effect and is discussed in Section 3.3.

In large spaces such as sports halls, which typically have large areas of parallel hard surfaces, we should know from our experience as acousticians that flutter echoes and other non-diffuse reflections between walls will have a very substantial effect on the measured reverberation time; Sabine calculations rarely hold true when acoustic absorption is concentrated in one plane, typically the ceiling. Armed with this knowledge, the acoustician should assign low scattering coefficients for the large reflecting surfaces within the room, typically around 10%, tending to 5% for exceptionally large and flat surfaces such as sports hall floors.

For surfaces with significant surface relief, typically in the order of one or two inches or more, for example planes used to represent stairs or seating blocks, the frequency-dependence of scattering becomes acoustically very significant. In these instances, CATT-Acoustic can provide an estimation of frequency-dependent scattering for a given depth of surface roughness using the ESTIMATE function (inserted in place of scattering coefficients in the ABS surface properties statement). This can provide a useful starting point when estimating scattering but coefficients may require some manual adjustment to suit the situation, which again should be guided by the acoustician’s experience. It is important to emphasise again that fine surface details should not be modelled ‘as is’. As noted earlier, a model using very small planes will only be valid at high frequencies where the wavelength is substantially larger than the plane dimensions. It is always important to remember that one is creating an acoustic model for GA prediction and not a visual model.

3.2 CATT-Acoustic TUCT prediction algorithms

TUCT (The Universal Cone Tracer) was introduced as an add-on module to CATT-Acoustic during the later releases of Version 8 so that users could familiarize themselves with what was to come. In CATT-Acoustic v9 and above it is the only prediction engine. TUCT offers a choice of three different algorithms for detailed source-receiver prediction, which vary by their handling of high order specular and diffuse reflections. Randomised tail correction (RTC) of previous CATT-Acoustic versions is no longer used in TUCT; the full-length impulse response is calculated using a mix of image source modelling for first-order specular reflections, and cone and ray tracing first order diffuse and later reflections. It is important to understand the basic workings of each algorithm so that the appropriate calculation method can be selected for a given task.

TUCT uses two different methods of generating diffuse reflections in cone and ray tracing; deterministic split ray scattering and random scattering, illustrated schematically in Figure 3. Each of the three prediction algorithm uses different combinations of deterministic split ray scattering and random scattering.

Algorithm 1 primarily uses random scattering with optional ray splitting for up to 2nd order reflections. Algorithm 1 is only suitable for closed rooms with good mixing (TUCT will not allow Algorithm 1 to run if an open model is detected) and is not ideal for identifying late reflection effects such as flutter echoes unless a very large number of rays are used. 'Random' scattering is the method traditionally used in older versions of CATT-Acoustic and other similar software, where each incident ray generates a single reflected ray. This will either be specular or scattered, the probability of which is determined by the scattering coefficient. Random scattering requires a very large number of rays as no additional rays are created down the line. If too few rays are used then impulse responses generated using Algorithm 1 may not have sufficient reflection density in the late part of the decay to be suitable for auralisation. Also, because diffuse reflections are randomly determined, the late impulse response may vary significantly between successive calculation runs, particularly if a relatively small number of initial rays are selected. This can result in 'phantom' late echoes which in extreme cases can lead to large variations in calculated T30 between successive runs.

Figure 3 – Schematic comparison of deterministic split-ray scattering and random scattering

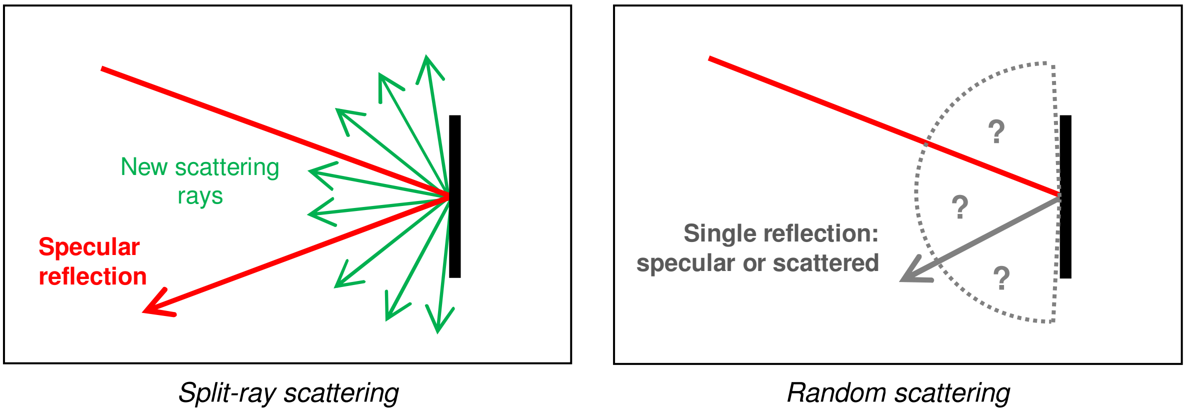

Split-ray scattering *Random scattering*

Algorithms 2 and 3 use ray-splitting for all reflections along the specular path of each ray or cone. Algorithm 3 adds further levels of ray-splitting for scattered rays. Split ray scattering generates a specular reflection ray <u>and</u> many new rays to represent diffuse reflections. Ray splitting requires a relatively low number of initial rays, but generates a very large number of diffuse rays (each with relatively low energy) after a few orders of reflection, resulting in an increase in reflection density as the sound decays. It is important to note that calculation time in a closed model can increase enormously if Algorithms 2 or 3 are used. In the interests of efficiency, selection of the appropriate algorithm therefore needs to be carefully considered; a substantially longer calculation time does not necessarily result in significantly more accurate results. Algorithm 2 is generally a good choice for auralisation in difficult, non-mixing spaces, particularly to identify and auralise flutter echoes well. To an extent, the choice between Algorithms 1 and 2 is self-regulating; if an Algorithm 2 run takes much longer than a good Algorithm 1 run in a closed model (with a suitably large number of rays selected) then Algorithm 2 is probably not needed.

The differences between Algorithms 2 and 3 in closed models is generally insufficient to justify the additional calculation time required for Algorithm 3. Algorithm 3 is best suited to very open models such as amphitheatres, where the majority of rays are lost very quickly and hence where it is important to squeeze as much detail as possible, as deterministically as possible, from the first few reflections. Due to the low reflection density in open models, there are often insufficient reflections in a given echogram time segment for meaningful 'average' results to be obtained from random scattering.

3.3 Auto edge scattering

Auto edge scattering is an extremely useful function for estimating low-frequency scattering of small planes, so as to avoid unrealistic low-frequency specular reflections from such planes. In practice where the wavelength is large compared to the surface size, sound will be diffracted by the surface, rather than reflected specularly. Edge scattering uses frequency-dependent scattering to approximate the diffracting effects of smaller planes at low frequencies on the ‘illuminated’ side (i.e. the side receiving the incident sound). In contrast to the early sound diffraction method implemented in recent versions of CATT-Acoustic (which is discussed further in Section 3.4), auto edge scattering is applied to all reflections (early and late). Edge scattering is therefore a more appropriate method than diffraction for handling low-frequency scattering of elements such as suspended ceiling rafts or baffles, which may be relatively far away from the source and therefore beyond the significant influence of the direct or first-order reflected sound.

Edge scattering is however a function which is often overlooked or misused. Edge scattering needs to be enabled for specific planes, usually on specific protruding or free-standing features such as suspended reflectors, columns, balcony fronts and office screens; it is not advisable to indiscriminately implement edge scattering on all planes. It is important to remember that to use edge scattering in CATT-Acoustic, the ‘Surface + Edge’ diffuse reflection option needs to be enabled in the ‘General settings’ dialogue box. This is useful since it can be switched on and off as a whole to see the effect.

Figure 4 demonstrates the effect of auto edge scattering in a comparison of SPL maps at 125Hz with a source placed close to a free-standing plane 1 metre wide (considerably narrower than the 2.6 metre wavelength at 125Hz), with a totally absorbent ground surface. Without auto edge scattering applied, a clear specular reflection can be seen from the plane. With edge scattering applied, scattering in the 125Hz octave band is 89%, and hence the specular reflection component becomes insignificant in this band.

Figure 4 – Comparison of edge scattering in 125Hz SPL maps close to a free-standing plane

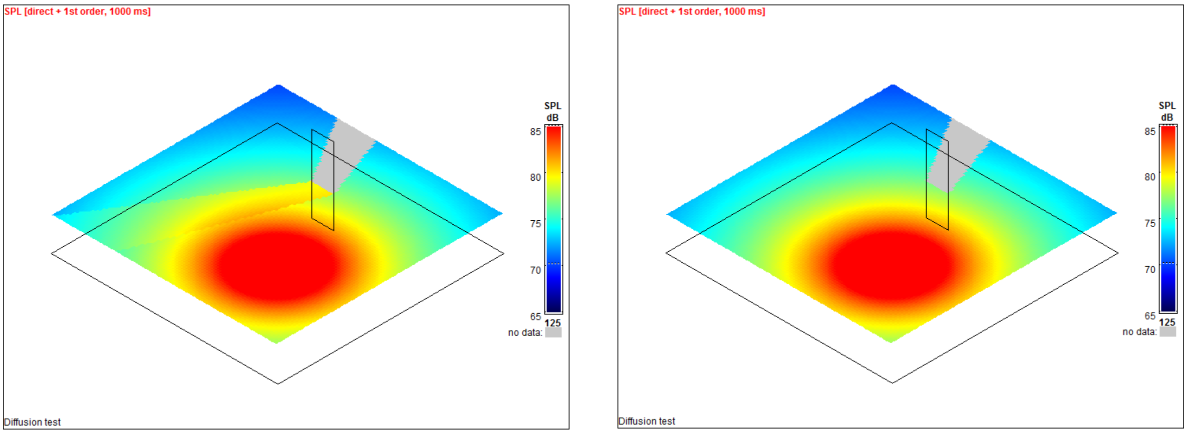

125Hz SPL map without edge scattering

125Hz SPL map with edge scattering

It should be noted in Figure 4 that even with edge scattering applied, there is still an unrealistically hard shadow visible behind the upright plane. In a closed reverberant room, such shadows may be masked by reflections from other surfaces, but in open models or acoustically dry spaces, these may be very apparent. CATT-Acoustic versions 9.0c and later include separate diffraction calculation processes that deal with this shadow zone more realistically and that also include interference with the direct sound on the ‘illuminated’ side.

3.4 Diffraction

CATT-Acoustic handles diffraction by generating new secondary ‘edge sources’ along the diffracting edge. The process is very elaborate and Bengt-Inge Dalenbäck has produced a very comprehensive white paper on diffraction[4] which is accessible via the CATT-Acoustic software manual and help file. Acknowledging that the method cannot be done adequate justice in a meagre few sentences of description, this shall not be reproduced here but it is worth highlighting some of the key limitations as well as the limitation of GA. As with ray-splitting, the process of diffraction adds new secondary sources and thus can add significantly to the calculation time. It is therefore important to understand when it is appropriate and worthwhile to use diffraction; employing this indiscriminately may extend calculation times unnecessarily while bringing no significant benefit to prediction accuracy.

Perhaps the most important limitation to note regarding diffraction is that it is handled in CATT-Acoustic solely as an early reflection process; diffraction is only applied to direct or first-order incident sound (depending on options selected) and diffracted sound rays are not permitted more than a first order reflection, if any (again depending on options selected.) Diffracted sound does not therefore contribute to reverberant energy in the model.

Diffraction is by default (again this default setting can be changed) only applied to edges shared by two planes where the average scattering or absorption coefficients of the two planes at 125Hz are 20% or less (excluding auto edge scattering). This is on the basis that in practice a very rough, scattering or very absorbing edge is unlikely to diffract cleanly according to the hard-edge diffraction theory applied. Because the checking routine considers only the scattering coefficients defined in the surface properties (ABS) definition, it is imperative that auto edge scattering is enabled in conjunction with diffraction, rather than manually defining high surface scattering coefficients at low frequencies in the ABS definition.

The effects of diffraction on the scenario introduced in Figure 4 are demonstrated in Figure 5. The hard shadow area behind the column is now filled in and we can see a recognisable diffraction grating interference pattern radiating from the column edges. It should also be noted that diffraction also creates a spread of reflected energy on the source side, rather than a linear specular reflection as shown in Figure 4. The spread of this reflection is frequency-dependent (resulting in a tighter reflection 'beam' at high frequencies.) Diffraction can therefore also be beneficial in predicting frequency-dependent spreading or focusing of first or second order reflections from curved reflectors, for instance in orchestra shells.

Figure 5 - Comparison of diffraction in SPL maps close to a free-standing plane

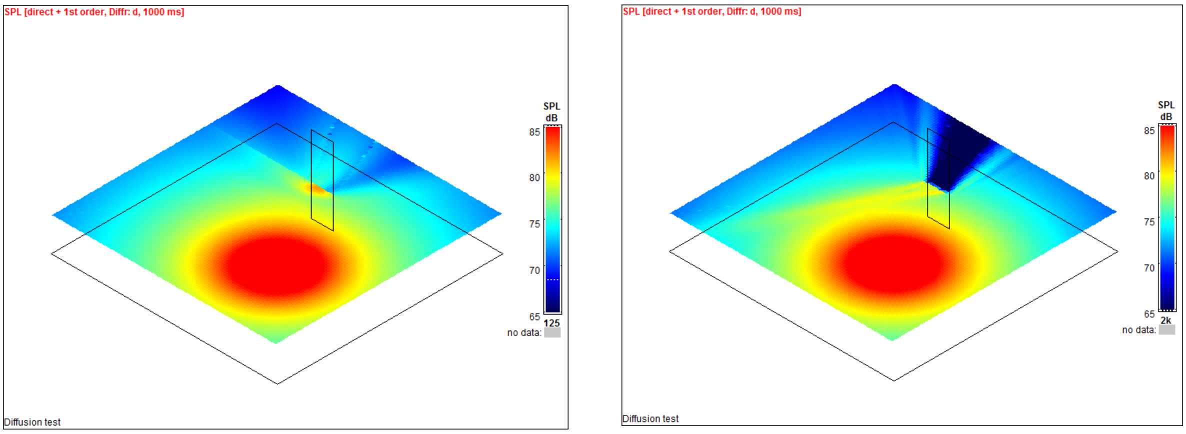

125Hz SPL map with diffraction

2000Hz SPL map with diffraction

4 ANALYSING RESULTS

Once a model has been created, there is often a temptation to rush through calculations and obtain results – any results – as quickly as possible. It is important however to remember that creation of the room geometry is not the end of the process but merely the beginning. It is important to allow due time to digest and analyse results and, if these are not as expected, to investigate why; lessons learned in this process may also prove to be invaluable for future projects. As discussed in Section 2.2, it is advisable to make use of the Interactive RT Estimation module and similar peripheral modules to quickly make initial coarse adjustments to the model before committing to a full calculation run.

It is also important to consider which of the bewildering array of measures available is most informative to the acoustic design and which tell the story most clearly for the benefit of a non-technical client. There is clearly much more to be written on this topic than can be crammed into the closing paragraphs of a short paper but a key point to note is that it is easy in consultancy to become over-reliant on meeting simple numerical criteria; finer analytical tools such as echograms and impulse responses (particularly the exceptionally powerful rotating sector microphone impulse response feature of TUCT) are often overlooked. Auralisation can also be a powerful tool for analysing results, as well as for demonstration to clients. As a practice we have found auralisation especially useful in combination with the aforementioned rotating microphone feature for locating problematic directional reflection patterns, such as flutter echoes and other strong late reflections. It is of course important to ensure that an appropriate calculation algorithm and number of rays / cones is selected for auralisation, as discussed in Section 3.2, and that the validity of late-part impulse responses is considered before attempting to use auralisations for analysis or demonstration.

5 CONCLUSION

Geometrical Acoustics (GA) modelling software can provide powerful tools for virtual measurement and prediction of room acoustics but ‘with advanced tools, advanced mistakes can be made.’ GA modelling programmes do not solve the wave equation and therefore offer only approximations of how sound propagates and interacts within a space, based on the geometric features and surface properties of the space. It is vital to have a good understanding of these approximations, both in terms of how the calculations work and the limitations of when they apply, which requires a good practical as well as theoretical understanding of room acoustics.

It must be stressed that the validity of results obtained from GA models, and indeed any calculation process, is heavily dependent on the quality of data input. It is crucial that the model geometry, absorption and scattering coefficients, source and receiver positions are selected with great care and with good understanding of the likely acoustic effects of these choices. This paper has presented a brief discussion of considerations surrounding simplification of models, selection of surface properties and prediction algorithms in CATT-Acoustic, and analysis of results.

We clearly cannot provide all the answers in a brief paper but it is our intention that users will have been prompted to ask all the right questions when creating and running models for GA prediction, and when analysing the results.

6 REFERENCES

1. C.L. Christensen, G.B. Nielsen and J.H. Rindel, Danish Acoustical Society Round Robin on room acoustic computer modelling, www.odeon.dk/pdf/Classroom%20RR.pdf (2008

2. B.I. Dalenbäck, S.G. Brown, Characterizing rooms regarding reverberation time prediction and the sensitivity to absorption and scattering coefficient accuracy, ASA Baltimore (2010)

3. B.I. Dalenbäck, The Rediscovery of Diffuse Reflection in Room Acoustics Prediction, ASA Cancoun (2002)

4. B.I. Dalenbäck, White paper regarding Diffraction (v5), [available to CATT-Acoustic users] (2012)