

# A Psychoacoustic Quality Criterion for Path-Traced Sound Propagation

Chunxiao Cao\*

ccx4graphics@gmail.com

Zili An\*

22021151@zju.edu.cn

Zhong Ren\*†

renzhong@zju.edu.cn

Dinesh Manocha‡

dmanocha@umd.edu

Kun Zhou\*

kunzhou@zju.edu.cn

## Abstract

In developing virtual acoustic environments, it is important to understand the relationship between the computation cost and the perceptual significance of the resultant numerical error. In this paper, we propose a quality criterion that evaluates the error significance of path-tracing-based sound propagation simulators. We present an analytical formula that estimates the error signal power spectrum. With this spectrum estimation, we can use a modified Zwicker’s loudness model to calculate the relative loudness of the error signal masked by the ideal output. Our experimental results show that the proposed criterion can explain the human perception of simulation error in a variety of cases.

## 1. Introduction

In virtual reality, computational room acoustics, and various other fields, we need to calculate the sound received by human ears in a virtual acoustic scene. High quality sound rendering can significantly improve the overall user experience in virtual reality [12]. To calculate the received sound, we need to simulate various physical phenomena in the sound generation and propagation process [7]. Simulation of the sound generation provides the user information about sound sources [19, 5, 14], and simulation of sound propagation provides critical information about the sound environment [22, 1, 31].

Sound propagation from the source to the listener can be modeled as a linear time-invariant system, which can be fully described with its impulse response (IR) function  $h(t)$ , a function calculated by the simulator. Given the input audio  $f_i(t)$ , the propagated audio  $f_o$  received at the listener is given by the equation  $f_o(t) = f_i(t) * h(t)$ , where the asterisk stands for convolution. All simulation algorithms inevitably generate an error  $h_e$  over  $h$ , adding an extra sig-

nal  $f_e = f_i * h_e$  to the output audio. The simulation quality can be measured with the energy signal-to-noise rate (SNR) of  $h$ .

While some methods simulate sound propagation by solving the wave equation directly [15, 15]. Speed-critical applications usually rely on methods based on geometrical acoustic (GA) methods [13, 4] or hybrid methods [30]. Among GA methods, Monte Carlo path tracing [3, 24, 23] is one of the most popular methods for real-time simulation. In path tracing,  $h(t)$  is given as an integral calculated by the Monte Carlo method, which uses a large number of random samples called “paths” to simulate the propagation process. An advantage of path tracing is that its computation budget can be easily adjusted to accommodate for different quality requirements. For example, there exists a simple relationship between the number of samples and IR SNR, allowing us to balance between speed and quality easily by controlling the sample count. Other quality control methods have also been proposed in previous works [23, 3].

Since the output audio is usually received by humans, we would like to know the effect of simulation error on receivers’ hearing experiences. Simply put, we need to predict the adequate IR SNR value, under which we can produce plausible outputs for users. However, it turns out that building the relationship between the user experience and IR SNR is not simple. For real-time simulators, where  $h$  is calculated on-the-fly and  $h_e$  changes in every frame, the problem becomes even more complicated. Unfortunately, the change of  $h_e$  between frames is critical for the user to pick up the error of the simulator. There is no analysis on the noise of real-time simulators, let alone its influence on users’ hearing experiences.

In this paper, we will show that it is possible to bridge the gap between IR quality and user hearing experience using established psychoacoustic models. The main results in this paper include:

- A spectral analysis for the error signal  $f_e$  of both static and real-time path-tracing-based simulators. An analytical expression for the power spectrum of  $f_e$  is given, along with practical approximations. With this

\*State Key Laboratory of CAD&CG, Zhejiang University, Hangzhou 310058, China.

†Corresponding author.

‡University of Maryland, College Park, MD 20742, USA.

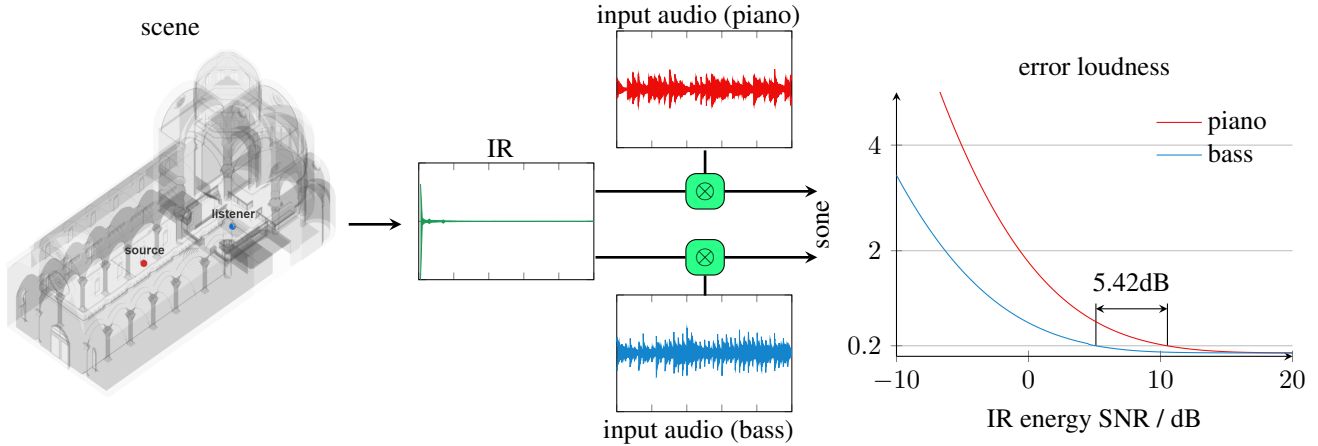


Figure 1: Even in the same acoustic environment, human sensitivity to simulation errors in sound propagation may vary for different input audios. Our criterion exposes this difference and allows the developer to adjust the simulation quality for different cases. The figure above shows that, to achieve the same error loudness level (0.2 sone), the required IR SNR is 5.42dB higher for the “piano” sound when compared to the “bass” sound.

expression, we can estimate  $f_e$  with IR SNR and the spectrum of the input audio.

- A criterion that predicts the audibility of the error signal  $f_e$  based on the widely-used Zwicker’s loudness model [36, 16, 9]. From Zwicker’s original algorithm, which calculates the “absolute loudness” of a signal in a quiet environment, we develop an algorithm to calculate the “relative loudness” of  $f_e$  in the presence of  $f_o$ . This algorithm allows us to predict the “audibility thresholds” below which the error  $f_e$  becomes inaudible, which is valuable for render quality control. Our experiments show the validity of our criterion in various cases.

While the spectral analysis result is dedicated to path-tracing-based simulators, the criterion can be applied to methods other than path tracing if an estimation for their power spectra exists.

## 2. Background and Related Work

To facilitate discussion, we introduce the necessary background knowledge in this section for readers unfamiliar with acoustics and sound simulation.

### 2.1. Numeric Characteristics of Audio Signals

#### 2.1.1 Power and Power Spectrum

In signal processing, the power of a signal  $f$  of duration  $T$  is given by the equation  $p = \frac{1}{T} \int_{\mathbb{R}} f(t)^2 dt$ . For a signal of infinite duration, the power is given as  $\lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-T/2}^{T/2} f(t)^2 dt$ . While power is the simplest

symbol	explanation
$E[X]$	mathematical expectation of variable $X$
$\sigma[X]$	standard variance of $X$ . $\sigma^2[X] = (\sigma[X])^2$
$\text{Cov}[X, Y]$	covariance of $X$ and $Y$
$\mathbb{R}$	set of real numbers
$\mathbb{Z}$	set of integers
$\bar{x}$	complex conjugate of $x$
$\delta$	Dirac delta function
	$\int_{\mathbb{R}} f(t)\delta(t)dt = f(0)$
$\mathcal{F}(f)$	Fourier transform
	$\mathcal{F}(f)(\omega) = \int_{\mathbb{R}} f(t)e^{-2\pi i \omega t} dt$

Table 1: Symbols in this paper. We use square brackets for parameters of random variables, round brackets for functions, and calligraphic letters for operators.

metric for signal intensity, we prefer to know the power distribution of a signal on its frequency domain, as the human hearing system responds to frequential characteristics of incoming sounds. Such a distribution is given by the power spectrum. For a signal  $f$ , the power spectrum is given by the equation  $\frac{1}{T} |\mathcal{F}(f)|^2$ , where  $T$  is the duration of the signal.

We can also describe the signal intensity with pressure or power levels. For a signal of power  $p$ , the power level  $L_{p/p_r}$  is given by the equation

$$L_{p/p_r} = 10 \log_{10} \frac{p}{p_r} \text{ dB.} \quad (1)$$

Here  $p_r$  is an appointed reference value. For example, decibels relative to full scale (dB FS) is the case where  $p_r =$

1. Sound pressure level (dB SPL) is the case where the average sound pressure is  $2 \times 10^{-5} \text{Pa}$  at  $p_r$ .

For the error signal  $f_e$  overlapped on the signal  $f$ , we assume that  $E[f_e] = 0$  and that its Fourier spectrum  $\mathcal{F}(E[f_e])$  is also zero. Thus we have

$$\begin{aligned} E\left[\frac{1}{T}|\mathcal{F}(f_e)|^2\right] &= \frac{1}{T}E[|\mathcal{F}(f_e) - \mathcal{F}(E[f_e])|^2] \\ &= \frac{1}{T}E[|\mathcal{F}(f_e) - E[\mathcal{F}(f_e)]|^2] \\ &= \frac{1}{T}\sigma^2[\mathcal{F}(f_e)]. \end{aligned} \quad (2)$$

Thus, the power spectrum of the error is also the spectrum of its variance. This  $\sigma^2[\mathcal{F}(f_e)]$  is what we are trying to calculate in [Sect. 3](#).

### 2.1.2 Energy SNR

A useful numerical metric for signal error is the energy signal-to-noise rate (SNR). For a signal  $f$  of limited duration overlapped with an error signal  $f_e$ , the energy SNR is

$$\text{SNR}(f) = \frac{\int_{\mathbb{R}} f(t)^2 dt}{\int_{\mathbb{R}} f_e(t)^2 dt}. \quad (3)$$

When the  $f_e$  is a random function,

$$\text{SNR}(f) = \frac{\int_{\mathbb{R}} f(t)^2 dt}{E[\int_{\mathbb{R}} f_e(t)^2 dt]}. \quad (4)$$

Energy SNR is also frequently given in the form of SNR levels:

$$L_{\text{SNR}}(f) = 10 \log_{10} \text{SNR}(f) \text{ dB}. \quad (5)$$

### 2.2. Mechanisms Influencing Auditory Perceptibility

A good way to understand the various factors that influence human sound perception is to classify them by their effects on resultant neuron signals. In this way, we can divide hearing mechanisms into excitation and suppression mechanisms.

The most remarkable feature of the excitation mechanism is frequency selectivity: sound components of different frequencies are separated by the auditory system and perceived by different receptors. Frequency selectivity has been well-studied from both anatomical and psychological perspectives. It has been observed that components of different frequencies resonate with different positions on the basilar membrane (BM) of the cochlea. In psychoacoustics, a rough relationship between frequency and BM position is given by frequency scales like the Bark scale [33]. On further investigation, one finds that a single point on the BM responds to a range of frequencies rather than a single frequency. We use band-pass filters called auditory

filters to describe these frequency response patterns. There are several models describing the characteristics of these filters, with parameters like critical bandwidth (CB) [33, 35] or equivalent rectangular bandwidth (ERB) [8]. With an appropriate scheme, one can divide the audible frequency range into several bands for which the excitations minimally interfere.

The suppression mechanism of human hearing is not as well understood as excitation. An important related discovery is otoacoustic emission [10], which shows that the cochlea adjusts its sensibility not only neurally, but also mechanically. The suppression mechanism affects frequency sensitivity in a highly nonlinear and nonlocal way and drives many effects such as the ‘‘upward spread’’ of masking [20], where a low-frequency sound can suppress the sensibility of another sound with much higher frequency, and auditory adaptation [11], where the continued presence of a sound will be regarded as ‘‘background’’ and ignored by the hearing system. Other studies show that the frequency selectivity of hearing is also a combined effect of activation and suppression mechanisms [17].

### 2.3. Masking Effect

The masking effect describes the change of perceptibility of a sound (signal) in the presence of another sound (masker). Due to its relationship with various hearing mechanisms, masking has long been a useful tool in the study of human hearing. The perceptibility of the signal is non-linearly related to its intensity level: it decreases much faster below a certain threshold than above it [27]. This threshold is called the mask threshold. Naturally, this threshold should be above the absolute hearing threshold, under which a sound is imperceptible. The absolute threshold can also be considered a special case of mask threshold, with the masker being the internal noise generated by physiological activities [34]. Here, we note the absolute threshold as  $p_0$  and the mask threshold as  $\max\{p_0, p_t\}$ , as  $p_t$  can usually be described with simple models.

Narrow-band masking occurs when a narrow-band masker masks a signal (usually a sinusoidal one) within the band. This phenomenon is closely related to frequency selectivity. The relationship between  $p_t$  of narrow-band masking and the power spectrum of the masker is generally assumed to be linear [6]:

$$p_t = \int_{\mathbb{R}} p_m \cdot f d\omega \quad (6)$$

where  $p_m$  is the power spectrum of the masker and  $f$  is a function that shows the shape of the auditory filter.

Broad-band masking occurs when the signal is masked by a broad-band masker. When the signal is sinusoidal and the masker is white noise, [34] approximates the mask threshold with the relationship  $p_t(\omega_t) = k(\omega_t) \cdot p_m(\omega_t)$ ,

where  $w_t$  is the signal frequency. With this relationship, a uniform masking noise is developed in [34] that makes  $p_t$  equal for all frequencies.

### 2.3.1 Loudness

The perceived intensity of the signal under a certain masker is called its loudness. The “loudness” in the general sense is the loudness when the masker is the internal noise. Loudness models are closely related to auditory filter models. Different auditory filter descriptions lead to different loudness algorithms. The criterion in this paper is related to the loudness algorithm proposed by Zwicker, which uses CB to describe auditory filters.

The general unit of loudness is sone. Like power and SNR, there is a corresponding unit for loudness level named phon. Given an audio signal, different loudness algorithms may result in different sone values. Since our criterion is a metric for relative loudness, we also use sone as the unit of its result values.

### 2.4. Sound Propagation and IR

The goal of sound propagation simulation is to calculate the IR of the sound environment. As its name indicates, IR represents the sound received at the listener’s position when the sound source emits an impulse signal.

#### 2.4.1 The Structure of IR

The impulse response can be broken down into three parts: direct contribution, early reflections and late reflections. The direct contribution represents the sound propagated from the source to the listener directly. It may not exist when there is an occluding object between the listener and the source. This part usually has an analytical expression and is computed directly in most path-tracing-based simulators. The early reflections represent the sound reflected by the scene in the first few times. It is related to the geometry and acoustic characteristics of objects near the listener. The late reflections represent the sound reflected multiple times in the environment. It is related to the general volume and sound absorption rate of the scene. Late reflections are important to user experience because they give the listener a “sense of volume” of the scene [1]. In path-tracing-based simulators, both early and late reflections are generated by random sampling and are thus not immune to numerical error.

#### 2.4.2 Psychoacoustics in IR Computation

While path tracing is already one of the most efficient methods to calculate IR, the calculation process is still costly, especially for real-time applications and low-end devices. It is possible to use psychoacoustic knowledge to improve

the simulation efficiency. For example, [25] uses a psychoacoustic method to suppress IR noises. [21] reduces the computation cost using the relationship between the scene changes and the audibility of IR changes. [32] uses the A-weighting scheme to estimate the error significance and the appropriate computational budget for simulation. However, the error analyses in these works are usually based on oversimplified models that ignore some critical factors, like the masking effect and the frame-to-frame IR variance. In contrast, these factors are all taken into consideration in our new criterion.

### 3. Spectrum Analysis of the Error Signal

The spectrum of the ideal output  $f_o = f_i * h$  is given by the equation  $\mathcal{F}(f_o) = \mathcal{F}(f_i) \cdot \mathcal{F}(h)$ . In this section, we will estimate the power spectrum of the error signal  $f_e$ . The calculation process is similar to the proof of the Wiener-Khinchin theorem [29]. When dealing with discrete signals that are sampled from continuous signals, we use the equation

$$\mathcal{F}(f)(\omega) = \sum_{t \in \mathbb{Z}} \frac{1}{s} f(st) e^{-2\pi i \omega st} \quad (7)$$

for Fourier transform of discrete functions. Here  $s$  is the sample rate. This equation still satisfies  $\mathcal{F}(f \cdot g) = \mathcal{F}(f) * \mathcal{F}(g)$  and  $\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$ . Under this definition, we have the discrete version of Parseval’s theorem [2]:

$$\sum_{t \in \mathbb{Z}} \frac{1}{s} f(st)^2 dt = \int_{-s/2}^{s/2} |\mathcal{F}(f)(\omega)|^2 d\omega. \quad (8)$$

#### 3.1. Static Case

In path tracing, the IR function is estimated by accumulating a series of independent random samples in the form  $a\delta(t - t_0)$ . The accumulated result  $h'$  can be regarded as a stochastic process with expectation  $E[h'] = h$ . In practice,  $h'$  is a discretely sampled function, and  $E[h'(t)] = \int_{t-1/2s}^{t+1/2s} h(t) dt$ . When the number of samples is sufficiently large,  $h'(t_0)$  and  $h'(t_1)$  become approximately independent when  $t_0 \neq t_1$ . This independence is critical to our analysis below.

The sample error  $h_e = h' - E[h']$  is also a stochastic process. From the independence assumption, we have  $E[h_e] = 0$  and  $\forall t_0 \neq t_1, \text{Cov}[h_e(t_0), h_e(t_1)] = E[(h_e(t_0) - h(t_0))(h_e(t_1) - h(t_1))] = 0$ . Therefore,  $h_e$  is a modulated white noise process:

$$h_e(t) = \sigma[h_e(t)] \cdot \nu(t). \quad (9)$$

Here  $\nu(t)$  is a white noise with  $E[\nu(t)] = 0, \sigma[\nu(t)] = 1$  and  $\forall t_0 \neq t_1, \text{Cov}[\nu(t_0), \nu(t_1)] = 0$ .

Now consider the Fourier spectrum  $\mathcal{F}(h_e)$ . We immediately know from  $E[\nu(t)] = 0$  that  $E[\mathcal{F}(h_e)(\omega)] = 0$ , and




Figure 2: Energy spectrum of IR and its noise on the audible range, produced by our simulator in the Sibenik scene (see [subsubsection 5.1.1](#)). The deviation of IR variance from the estimation given by [Eq. 11](#) is about 5dB at low frequencies and less than 2dB at middle frequencies. However, this deviation is minimal compared to the IR energy spectrum itself, which is  $|\mathcal{F}(h)|^2$ .

$$\begin{aligned}
& \text{Cov}[\mathcal{F}(h_e)(\omega_0), \mathcal{F}(h_e)(\omega_1)] \\
&= E[\mathcal{F}(h_e)(\omega_0) \overline{\mathcal{F}(h_e)(\omega_1)}] \\
&= \frac{1}{s^2} \sum_{t,u \in \mathbb{Z}} E[h_e(st)h_e(su)] e^{-2\pi i s(\omega_0 t - \omega_1 u)} \\
&= \frac{1}{s^2} \sum_{t \in \mathbb{Z}} \sigma^2[h_e(st)] e^{-2\pi i(\omega_0 - \omega_1)st} \\
&= \frac{1}{s} \mathcal{F}(\sigma^2[h_e])(\omega_0 - \omega_1).
\end{aligned} \tag{10}$$

Specifically,

$$\sigma^2[\mathcal{F}(h_e)(\omega)] = \frac{1}{s} \mathcal{F}(\sigma^2[h_e])(0) = \frac{1}{s^2} \sum_{t \in \mathbb{Z}} \sigma^2[h_e(st)]. \tag{11}$$

This is an interesting result: no matter what shape the IR noise envelope  $\sigma^2[h_e]$  is, it is always equivalent to a white noise: the energy spectral density  $\sigma^2[\mathcal{F}(h_e)(\omega)]$  is equal for all  $\omega$ . In actual path tracers, the dependence between different parts of  $h_e$  may exist, and the density spectrum may not be entirely flat. However, the fluctuation is largely negligible, especially when compared with the IR spectrum (see [Fig. 2](#)).

Given the input audio  $f_i$ , the energy spectral density of the output error  $f_e$  is

$$\begin{aligned}
\sigma^2[\mathcal{F}(f_e)] &= \sigma^2[\mathcal{F}(f_i * h_e)] \\
&= \sigma^2[\mathcal{F}(f_i) \mathcal{F}(h_e)] \\
&= |\mathcal{F}(f_i)|^2 \cdot \sigma^2[\mathcal{F}(h_e)].
\end{aligned} \tag{12}$$

Thus, we have the power spectrum estimation of the error signal in the static case.

Now we'll look at the relationship between the error power spectrum and IR energy SNR. We discretize the integral of SNR in [Eq. 4](#), then use [Eq. 8](#) and [Eq. 11](#):

$$\begin{aligned}
\text{SNR}(h) &= \frac{\sum_{t \in \mathbb{Z}} h(st)^2}{E[\sum_{t \in \mathbb{Z}} h_e(st)^2]} \\
&= \frac{\sum_{t \in \mathbb{Z}} h(st)^2}{\sum_{t \in \mathbb{Z}} \sigma^2[h_e(st)]} \\
&= \frac{\int_{-s/2}^{s/2} |\mathcal{F}(h)(\omega)|^2 d\omega}{s \cdot \sigma^2[\mathcal{F}(h_e)(\omega)]}.
\end{aligned} \tag{13}$$

Now with [Eq. 12](#), we have

$$\sigma^2[\mathcal{F}(f_e)] = |\mathcal{F}(f_i)|^2 \cdot \frac{\int_{-s/2}^{s/2} |\mathcal{F}(h)(\omega)|^2 d\omega}{s \cdot \text{SNR}(h)}. \tag{14}$$

This is the relationship between the power spectrum of the error signal and the energy SNR of IR.

### 3.2. Dynamic Case

In this subsection, we discuss the case of real-time renderers, where the IR is recalculated in every frame (hence the name “dynamic”). We assume that the scene remains unchanged during rendering. The IR difference between frames is thus only due to calculation error. We can regard the different IRs as different sample paths of the same variable  $h_e$ , the spectrum of which is given in the previous section. The resulting audio is produced by stitching multiple audio frames together. If we note the IR error of the  $k$ -th frame as  $h_k$ , then the output error signal is

$$f_e(t) = \sum_{k \in \mathbb{Z}} ((f_i * h_k) \cdot w(t - k\Delta t_w)), w(t) \in [0, 1]. \tag{15}$$

Here  $w(t)$  is the window function for frame interpolation and  $\Delta t_w$  is the frame length. Let  $g_k = f_i * h_k$ , then  $g_k$  is a sample path of  $g = f_i * h_e$ , and

$$\mathcal{F}(f_e) = \sum_{k \in \mathbb{Z}} (\mathcal{F}(g_k) * \mathcal{F}(w) e^{2\pi i \omega k \Delta t_w}). \tag{16}$$

Let  $g_k^* = (f_i * h_k) \cdot w(t - k\Delta t_w)$ , then

$$\mathcal{F}(g_k^*)(\omega) = \int_{\mathbb{R}} \mathcal{F}(g)(x) \mathcal{F}(w)(\omega - x) e^{2\pi i(\omega - x)k\Delta t_w} dx, \tag{17}$$

and

$$\begin{aligned}
\sigma^2[\mathcal{F}(g_k^*)(\omega)] &= \int_{\mathbb{R}^2} (\text{Cov}[\mathcal{F}(g)(x), \mathcal{F}(g)(y)] \cdot \\
&\quad \mathcal{F}(w)(\omega - x) \mathcal{F}(w)(\omega - y) \cdot \\
&\quad e^{2\pi i(y-x)k\Delta t_w}) dx dy.
\end{aligned} \tag{18}$$



Substituting  $y - x$  with  $y$ , we have

$$\sigma^2[\mathcal{F}(g_k^*)(\omega)] = \int_{\mathbb{R}^2} \left( \text{Cov}[\mathcal{F}(g)(x), \mathcal{F}(g)(x+y)] \cdot \mathcal{F}(w)(\omega-x) \mathcal{F}(w)(x+y-\omega) \cdot e^{2\pi i y k \Delta t_w} \right) dx dy, \quad (19)$$

and

$$\begin{aligned} \sigma^2[\mathcal{F}(f_e)(\omega)] &= \sum_{k \in \mathbb{Z}} \sigma^2[\mathcal{F}(g_k^*)(\omega)] \\ &= \int_{\mathbb{R}^2} \left( \text{Cov}[\mathcal{F}(g)(x), \mathcal{F}(g)(x+y)] \cdot \mathcal{F}(w)(\omega-x) \mathcal{F}(w)(x+y-\omega) \cdot \sum_{k \in \mathbb{Z}} e^{2\pi i y k \Delta t_w} \right) dx dy. \end{aligned} \quad (20)$$

The sum  $\sum_{k \in \mathbb{Z}} e^{2\pi i y k \Delta t_w}$  is equal to  $\text{III}(\Delta t_w y)$ , where  $\text{III}(x)$  is the shah function [2]:

$$\text{III}(x) = \sum_{n \in \mathbb{Z}} \delta(x - n). \quad (21)$$

Therefore,

$$\begin{aligned} \sigma^2[\mathcal{F}(f_e)(\omega)] &= \\ \frac{1}{\Delta t_w} \sum_{k \in \mathbb{Z}} \int_{\mathbb{R}} &\left( \text{Cov}[\mathcal{F}(g)(x), \mathcal{F}(g)(x + \frac{k}{\Delta t_w})] \cdot \mathcal{F}(w)(\omega-x) \mathcal{F}(w)(x + \frac{k}{\Delta t_w} - \omega) \right) dx. \end{aligned} \quad (22)$$

For a narrow-band  $w$ ,

$$\sigma^2[\mathcal{F}(f_e)(\omega)] \approx \frac{\sigma^2[\mathcal{F}(g)] * |\mathcal{F}(w)|^2}{\Delta t_w}. \quad (23)$$

Here  $\sigma^2[\mathcal{F}(g)]$  is equal to  $\sigma^2[\mathcal{F}(f_e)]$  in Eq. 14. So we also have the relationship between  $\sigma^2[\mathcal{F}(f_e)]$  and  $\text{SNR}(h)$  for the dynamic case. In practice, the audio quantization error will also affect the spectrum of  $f_e$ . The effect of quantization adds an additional term  $\frac{q^2}{3}$ , where  $q$  is the quantization interval.

One may expect an even simpler estimation for  $\sigma^2[\mathcal{F}(f_e)]$  if the bandwidth of  $\omega$  is completely ignored. If we approximate  $\mathcal{F}(w)$  with a Dirac delta function, we can replace Eq. 23 with

$$\sigma^2[\mathcal{F}(f_e)(\omega)] \approx \sigma^2[\mathcal{F}(g)] \cdot \frac{\int_{\mathbb{R}} w(t)^2 dt}{\Delta t_w}. \quad (24)$$

However, this turns out to be an oversimplification. Fig. 3 shows the results of various estimation formulae applied to a simulated dynamic rendering process. We can see

that both Eq. 22 and Eq. 23 fit the real data very well, while Eq. 24 fails to predict the noise power at low frequencies. This is where the dynamic case differs from the static case — the spectral distortion brought by the window function  $w$  cannot be ignored.




Figure 3: Comparison of the noise signal spectrum generated by a dynamic sound rendering process and its estimations.

## 4. Quality Criterion

Since our quality criterion is developed based on Zwicker’s loudness, we need to investigate Zwicker’s loudness algorithm first. Zwicker proposed two different loudness algorithms in [9], one for “stationary” sounds and another for “time-varying” sounds. The procedures of the two algorithms are largely similar. In this section, we assume the input audios to be stationary. That is, the spectral characteristics of the input audios do not change significantly in their durations.

A simple diagram for the loudness algorithm for stationary sounds is given in Fig. 4. We first filter the input audio  $f$  by 28 band filters and calculate the power of filtered audios. The results are expressed by a vector  $p$ . After this, we convert  $p$  into the specific loudness  $l$  through a series of processes.  $l$  is a function over the Bark scale, which maps the audible frequency range into a number between 0 and 24. The range  $[n-1, n]$  on the scale is called the  $n$ -th critical band. The mapping scheme of the Bark scale is designed to make different critical bands approximately independent from each other in auditory perception. In the end, we integrate  $l$  over the Bark scale to achieve the total loudness, which is Zwicker’s loudness. Refer to [9] for the detailed implementation of Zwicker’s algorithm.

One can see from Fig. 4 that our criterion requires two input audios instead of one as in Zwicker’s algorithm. Zwicker’s algorithm takes an audio  $f$  as input and outputs its loudness  $S$  relative to the internal noise. Our criterion requires a masker  $f_m$  and a full audio  $f$ , which is the sum

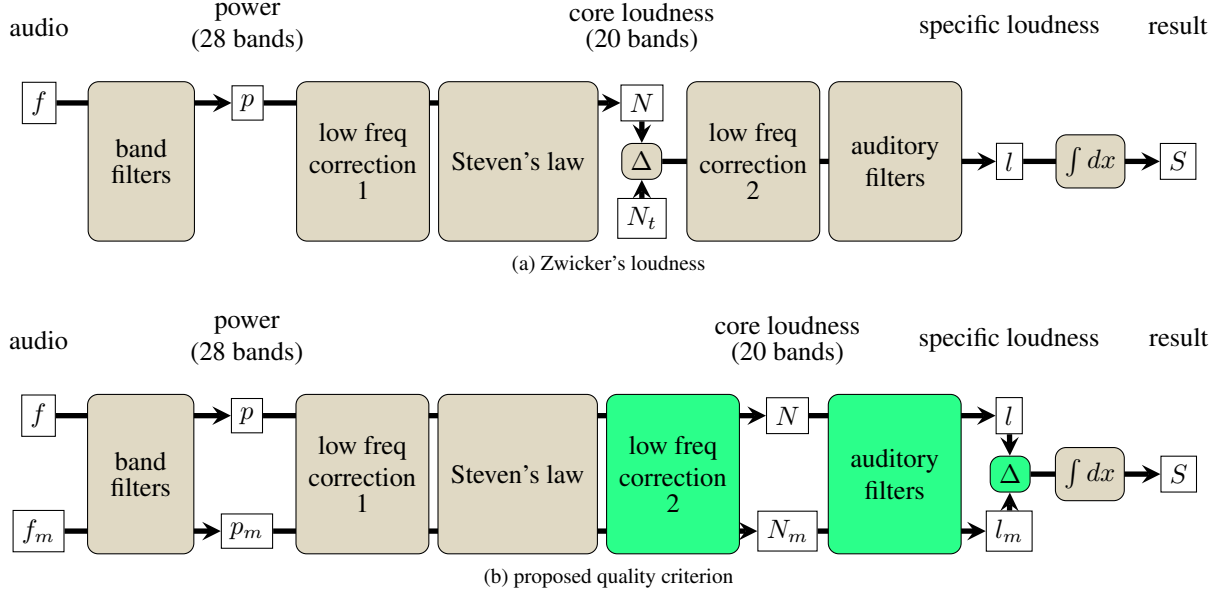


Figure 4: Diagrams for algorithms of Zwicker's loudness and our quality criterion. We use  $\Delta$  and  $\int dx$  to represent difference and integral operations. Green blocks signify operations different from Zwicker's original algorithm.

of the masker and the masked signal. In the context of simulation quality assessment,  $f_m = f_o$ ,  $f = f_o + f_e$ , where  $f_o$  and  $f_e$  is the ideal output and the error signal. We have already discussed the power spectrum of  $f_o$  and  $f_e$  in the previous section. One can use this information to calculate the value of  $p$  and  $p_m$  on different bands.

In Zwicker's algorithm, the loudness  $N$  of each critical band (core loudness) is an important intermediate variable. This variable is calculated using Steven's power law [26], which gives the relationship between audio power  $p$  and sensation magnitude [36]:

$$N(p) = Cp^a \quad (25)$$

where  $C$  and  $a$  are parameters. Zwicker uses  $C = 0.03175\sqrt{2}$  and  $a = 0.25$  in his algorithm. The loudness of a certain band is the difference between the actual sensation magnitude and the sensation threshold:

$$\Delta N = \max\{N(p+p_m) - N_t, 0\}, N_t = N(p_t+p_m) \quad (26)$$

where  $p$ ,  $p_m$  and  $p_t$  are power of input audio, internal noise and hearing threshold, respectively. Zwicker used  $p_t = \frac{1}{3}p_m$  in his algorithm. At low frequencies, the above equation no longer reflects the actual human perception correctly, so Zwicker introduced a few nonlinear corrections into the algorithm to mitigate this problem.

With the core loudness known, we can calculate specific loudness by applying acoustic filters. The specific loudness curve is calculated as follows:

$$l(x) = \max_{1 \leq i \leq 20} \{AF_i(\Delta N_i)(x)\} \quad (27)$$

where  $\Delta N_i$  is the core loudness of the  $i$ -th band and  $AF_i(\Delta N_i)$  is the corresponding specific loudness curve. The value of  $AF_i$  is equal to  $\Delta N_i$  inside the  $i$ -th band, but also extends to higher frequencies.

The main difference between our criterion and Zwicker's loudness algorithm is that the mask is no longer predetermined. In our criterion,  $p_m$  is the power of the input mask audio  $f_m$  plus the internal noise power, and  $p_t$  is no longer a static value. This leads to the possibility of upward spreading: When the mask audio is intensive on a certain band, its masking effect will also extend to higher frequencies like the function  $AF_i$ . Because of this, we calculate the difference of specific loudness instead of core loudness in our criterion.

From Fig. 4, we can see that the position change of the difference operation has some effect on two other operations. "low frequency correction 2" multiplies the core loudness of the first critical band by a factor  $c$ :

$$c = \min\{0.4 + 0.32\Delta N^{0.2}, 1\}. \quad (28)$$

To achieve a similar effect, we replace  $c\Delta N$  with  $c(N - N_t) + N_t$  and  $c(N_m - N_t) + N_t$  in our criterion, where  $N_t$  is the threshold in Zwicker's algorithm. For auditory filters, we use a linear approximation to replace  $AF_i(\Delta N_i)$  in Eq. 27:

$$AF_i(N) \approx \frac{N}{30} \cdot AF_i(30). \quad (29)$$

Using this approximation, we have  $AF_i(N - N') \approx AF_i(N) - AF_i(N')$ , making the difference operation and the application of auditory filters exchangeable. Fortunately, the original  $AF_i$  is already very close to a linear function, and this approximation does not introduce too much error.

#### 4.1. Soft Threshold

While the relationship  $p_t = \frac{1}{3}p_m$  is used in Zwicker's algorithm, experiments have shown that the ratio  $k_p = p_t/p_m$  is not a fixed value but is highly related to frequency and relatively independent from  $p_m$  [18]. Therefore, we can model  $k_p$  as a function of frequency. However, the function  $k$  may also vary for different individuals, and we should take this variance into account if possible.

Our solution to the abovementioned problem is to use a "soft threshold" in our criterion. Suppose that we calculate  $\Delta a$  with the equation below:

$$\Delta a = \max\{a - a_t, 0\}. \quad (30)$$

The key point of the soft threshold is to regard  $a_t$  as a distribution instead of a fixed value. we consider  $\Delta a$  as a function of  $a$ , then we have

$$E[\Delta a](a) = \int_{-\infty}^a \text{CDF}(a_t) dx. \quad (31)$$

Here  $\text{CDF}(a_t)$  is the cumulative distribution function of  $a_t$ . In our case, we assume that  $a_t$  conforms to a logistic distribution. When  $E[a_t] = E_a$  and  $\sigma[a_t] = \sigma_a$ , we have

$$E[\Delta a] = s \ln(e^{\frac{a-E_a}{s}} + 1), s = \frac{\sqrt{3}\sigma_a}{\pi}. \quad (32)$$

We will take  $E[\Delta a]$  as the difference under the soft threshold.

In our algorithm, we are calculating the difference of specific loudness  $\Delta l$  instead of  $\Delta N$  or  $\Delta p$ . Like Zwicker, we assume that  $k_l = l_t/l_m$  is independent from  $l_m$ . The values of  $E[k_l]$  and  $\sigma[k_l]$  are measured by experiments.

To compare  $k_l$  with  $k_p$  values measured in previous papers, we ignore the low-frequency correction and derive the relationship of  $p$  and  $l$  from Eq. 25 and Eq. 29:

$$l \propto C((p + p_m)^a - p_m^a). \quad (33)$$

Thus, we have the approximate conversion formula

$$k_l = C((1 + k_p)^a - 1). \quad (34)$$

## 5. Experiments and Results

In this section, we will show the results of our experiments. We first measured the values of  $E[k_l]$  and  $\sigma[k_l]$  and tested the independence between  $k_l$  and  $l_m$ . Then we test

our criterion on real outputs. We achieve these goals by measuring the audible power threshold of different signals under masking. The details are presented below.

### 5.1. Experiment Procedure

#### 5.1.1 Stimuli

In our experiment, we ask the subjects to listen to different audio clips with Beyerdynamic DT770 headphones in a quiet room. The background noise of the experiment environment is between 30-45 dB SPL. The audio clips are played by a web application on a regular PC computer. All the audio clips are generated and played in the 16-bit, 48kHz stereo format. The content of the left and right channels are all identical.

When measuring  $k_l$ , we use sinusoidal waves of different frequencies as signals and pink noise as maskers. The power spectrum of a pink noise signal is inversely proportional to frequency, which is similar to many natural sounds. In our experiment, we use pink noise at 55 dB SPL generated by the stochastic Voss-McCartney algorithm [28]. The frequencies of sinusoidal waves are the midpoints of the 24 critical bands defined in [34], which is a bit different from the specification in [9]. We change the definition of the first band from 0-100Hz to 20-100Hz in our experiment. Each band corresponds to a test case, and both are indexed by the upper bound Bark value of the bands. This gives 24 test cases in total.

To validate the independence between  $k_l$  and  $l_m$ , we measured the audible threshold of sinusoidal signals on band 1, 7, 9, and 19 with maskers of varied intensities. We used 6 different masker power levels equally spaced between 50-70 dB SPL. This validation process also requires 24 test cases.

When testing our criterion on real outputs, we generate the audio clips similar to real sound renderers:

1. We simulate sound propagation using a path tracing-based simulator in different scenes. Here we tried to cover all the "typical varieties" of sound environments. The reverberation length and the presence of direct contribution (DC) are two of the most prominent factors that affect user perception. Thus, we used 4 scene configurations with 2 different geometry models (see Fig. 5). The configurations are listed below:
  - Sibenik-DC: Interior model of a cathedral with a long reverberation time. The audio source and the listener are mutually visible so that the direct contribution (DC) component is present in the resultant IRs.
  - Sibenik-noDC: Same as above, but now the source is not visible from the listener.



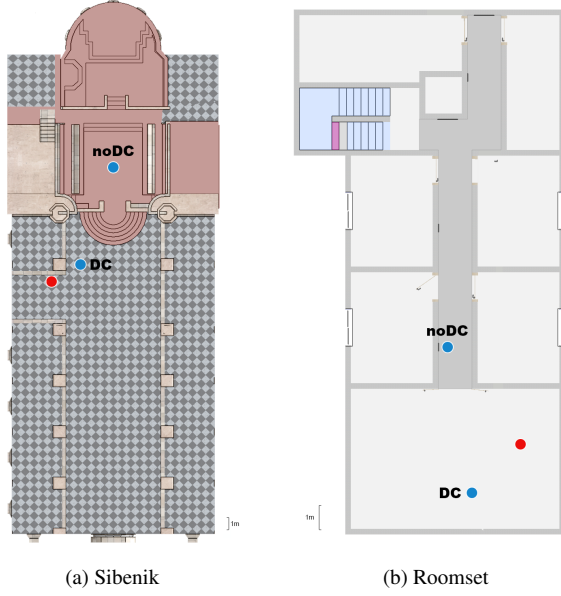


Figure 5: Floor plans of the geometry models used in our experiment. The position of the sound source and listener are marked with red and blue dots.

- Roomset-DC: Interior model of a two-story guesthouse with multiple small rooms and narrow corridors. The reverberation time is relatively short. The source is visible from the listener.
- Roomset-noDC: Same as above. The source is now invisible.

We computed 2000 IRs in each configuration. Each IR is computed from 48000 sampled paths, and the maximum number of reflections in a path is 40. We take the average of the 2000 IRs as the “reference IR” and the difference between each IR and the reference IR as its numerical error.

2. We generate the masker audio clip by convoluting the input audio with the reference IR. We scaled the masker audio so that its power level is around 65 dB SPL when being played. The signal clip is generated by convoluting the input with IR errors according to Eq. 15. Here we set  $w$  to a 512-sample Hanning window and  $\Delta t_w = 256$ . The IR error of each frame is randomly chosen from the 2000 candidates.

We used 6 different input audios in our experiment:

- bass: A music piece played by synthetic bass. The power spectrum concentrates at around 70Hz.

- piano: A piano piece whose energy distributes broadly in the 350-1000Hz range.
- voice: Speech of a male narrator. The energy concentrates on the 60-600Hz range.
- drums: Intense outdoor percussion music. The energy distributes broadly on the 100-2500Hz range.
- drips: Sound of dripping water. The energy distributes mainly between 600-3000Hz.
- nature: Environmental noise recorded on the Patagonian plain with rain sounds and occasional animal callings. The power spectrum is very similar to pink noise.

The above list covers most types of input audios that may appear in virtual acoustic simulations. We consider the first three audios as “melodic” because they contain pitch-related semantic information, which is usually sensitive to spectrum distortion. The last three audios are considered “noisy” for the absence of such information. The combination of 6 input audios and 4 scene configurations gives 24 test cases in total.

After generation, we crop all the signal and masker clips to the first 10 seconds and apply an exponential fading to the beginning and the end of each clip. The fading duration is 0.5s.

### 5.1.2 Method

We use the following trial to determine if a masked signal is audible to a subject. We present 4 candidate clips to the subject and the masker clip for reference. 1-3 of these candidates are the masked signal clip, and the rest are masker only, leading to 14 different possible combinations. The probability for each combination is the same. We require the subject to listen to the candidates and select all the clips containing the signal or press “give up” if they cannot tell a difference. All the candidate clips are replayable until the subject has submitted their answer. If the answer is correct, the signal is considered audible under masking.

Given a masker, the audible power threshold of a signal is determined by a bisection search. We set the initial width of the search range to 60 dB on the logarithmic scale and choose the upper bound empirically. For sinusoidal signals masked by pink noise, the upper bound is 30-40 dB above the power of the masker. We set the signal power to the midpoint of the search range and test its audibility. The midpoint is used as the new upper bound if the masked signal is audible or the new lower bound otherwise. We repeat this process until the width of the search range is below 1 dB and take the midpoint as the measured threshold.

### 5.1.3 Subjects

Our main experiment results are produced by 17 subjects (8 female and 9 male, aged between 19-27). The age group is similar to the one for measuring the absolute threshold of hearing [34] used in Zwicker’s original algorithm. To examine the subjects’ hearing condition, We performed pure tone audiometry tests in the range of 250-8000 Hz. For every subject, the measured absolute hearing threshold is less than 25 dB HL on the whole frequency range, which indicates normal hearing. We divide the test cases into 3 groups according to their purposes (measurement of  $k_l$ , validation of independence, validation on real outputs). 10 subjects are allocated to each group and their audible thresholds are measured under all 24 cases. Each group of tests takes 1-2 hours to finish.

Additionally, we validate our criterion on real outputs with 5 subjects (3 female and 2 male, aging between 28-40). According to audiometry results, 2 out of 5 among these subjects have their absolute hearing threshold between 25-40 dB HL on certain frequencies, which indicates mild hearing impairment.

None of the subjects has previous experience in psychoacoustic experiments.

## 5.2. Result

### 5.2.1 Measurement of $k_l$

The influence of sinusoidal signals is always concentrated on the influence range of one core loudness band. Thus we can determine from measured audibility thresholds the value of  $k_l$  on every single band. Fig. 6a shows the value of measured  $k_l$  from sinusoidal band 1-24 mask tests.

To use the measured data in our criterion, we need to fit  $E[k_l]$  and  $\sigma[k_l]$  with functions on the Bark scale. In the actual criterion, we use the following piecewise polynomials:

$$E[k_l] = \begin{cases} 0.0013(4.4 - x)^4 + 0.0052, & x < 4.4 \\ 0.0052, & 4.4 \leq x \leq 18.1 \\ 0.0011(x - 18.1)^2 + 0.0052, & x > 18.1 \end{cases} \quad (35)$$

$$\sigma[k_l] = \begin{cases} 0.0013(4.4 - x)^4 + 0.004, & x < 4.4 \\ 0.004, & 4.4 \leq x \leq 18.1 \\ 0.002(x - 18.1)^2 + 0.004, & x > 18.1 \end{cases} \quad (36)$$

To compare with previous data, we convert  $k_p$  values measured by Glasberg and Moore in [16] and  $k_p = 1/3$  in Zwicker’s algorithm into  $k_l$  curves on the bark scale. We use the conversion formula Eq. 34 and the frequency-to-Bark mapping function from [37]. The result is displayed in Fig. 6b, with curves defined by Eq. 35 and Eq. 36. We can see that Eq. 35 matches well with previously measured

band index	1	7	9	19
$r$	-0.5084	-0.0384	-0.0298	-0.2593

Table 2: Pearson correlation coefficient  $r$  of  $k_l$  and  $l_m$  on different bands.

data at central frequencies and shows a similar trend at low and high frequencies.

For the independence assumption of  $k_l$  from  $l_m$ , we calculated the Pearson correlation coefficient between  $k_l$  and  $l_m$  on four bands with measured  $k_l$  values under varied  $l_m$ . The result is given in Table 2.  $r < 0$  implies that the value of  $k_l$  tends to decrease when the masker loudness increases. This trend is negligible for bands 7 and 9, but not for bands 1 and 19. Hence the independent assumption holds for middle frequencies but is not entirely valid for low and high frequencies. This will not affect the accuracy of our criterion in most cases, but it may have some effect when the input sound has a very intense low/high-frequency component. We will not fix this problem in the current criterion for the sake of simplicity.

### 5.2.2 Application on Real Outputs

When applied to real outputs, our criterion gives interesting results and becomes a powerful tool for analyzing noise audibility. The experiment results are shown in Fig. 7 and Fig. 8.

Fig. 7 shows that our criterion can help in the choice of IR SNR for the actual rendering process. The figure shows the loudness-SNR relationship curves under a certain scene configuration. We can see that the error audibility thresholds cluster around the “knee” part of the relationship curves, where the loudness goes up quickly when the SNR decreases and stay invariant when the SNR increases. Our criterion can explain the difference of error significance between different input audios very well. The average IR SNR required for the error to be imperceptible is 2.09 dB when the input is “bass”, and 9.32 dB when the input is “piano”, which is much higher. While it is known that human hearing is less sensitive to the difference of low-frequency sounds, our criterion can show this difference in a quantitative manner.

One may have concerns about the robustness of our criterion to different inputs. Ideally, we expect the measured error audibility thresholds to be “nearly equally loud” in all cases. Fig. 8 shows that the thresholds are indeed largely stable. The average threshold loudness is 0.7813, and 94% of the results fall between 0-2 sone. Even for older and mildly hearing-impaired subjects, the result of our criterion remains valid (See Fig. 9).

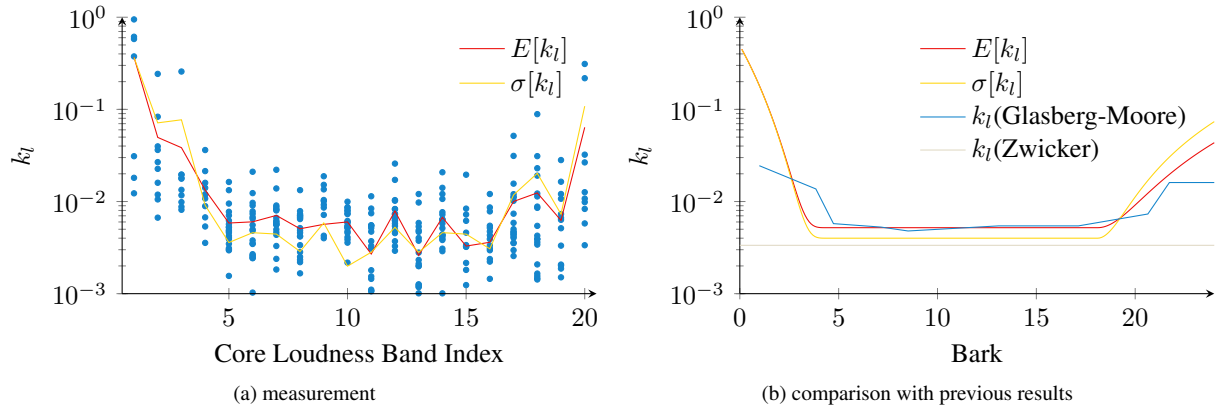


Figure 6: Visualization of various information about  $k_l$ . Figure (a) shows the values of  $k_l$  on different core loudness bands measured from our experiment. Figure (b) shows the comparison of the fitting curve for  $k_l$  given in Eq. 35 and Eq. 36 with previous results.




Figure 7: Relationship between IR energy SNR and noise loudness of input audio ‘piano’ and ‘bass’ under the ‘Sibenik-DC’ scene configuration. The measured error audibility thresholds of 10 subjects are marked on the loudness curves.

Fig. 8 also shows some interesting phenomena. The sensibility of different subjects to simulation noise varies a lot, but we can still find similar features in their results. The presence of DC lowers the threshold significantly. In other words, the simulation error is much more distinguishable with the help of DC. When the DC is absent, the noise of input ‘nature’ is the least perceptible one among all cases. This is probably due to that the input audio itself is similar to a ‘noise signal’. However, for most input audios, the threshold loudness is not very different. The influence of reverberation strength is minimal. The results in the Roomset and the Sibenik scene are almost indistinguishable. The difference between ‘melodic’ and ‘noisy’ input audios doesn’t seem to have much influence on the loudness either.

## 6. Conclusion and Future Work

In this paper, we demonstrated that an algorithm that calculates the auditory effects of sound propagation simulation error can be built on a solid theoretical basis. Experimental results show that our criterion is valid for a variety of sounds and people with different hearing abilities, and the result matches results from previous research.

Our error analysis for path tracing relies only on the assumption of sample independence. Thus it can be applied to other path traced methods like bidirectional path tracing (BDPT) with sample strategies like multiple importance sampling (MIS) [3]. Our analysis is also helpful in the optimization of sound propagation simulators. For example, Eq. 22 shows us the connection between the shape of the interpolation window and the error spectrum. Many audio renderers use linear interpolation between frames, which is equivalent to using a triangular window. Our analysis shows that a better window function, like the Hanning window, could improve the final quality of the output audio.

There is still a lot of work to do related to the current criterion. We will list a few points here.

We have mentioned that Zwicker’s loudness algorithm has another version for ‘time-varying’ sounds, the spectra of which vary significantly during playback. One can see from Fig. 8 that the noise loudness for input audio ‘drips’ seems to be very different from other audio. The audio ‘drips’ contains intermittent water dripping sounds and is not as ‘stationary’ as other inputs. This probably indicates that we need to modify our algorithm when dealing with time-varying input audios.

While our ‘soft threshold’ has taken the variance of human hearing abilities into consideration, the variance model is probably oversimplified. We chose the logistic distribu-

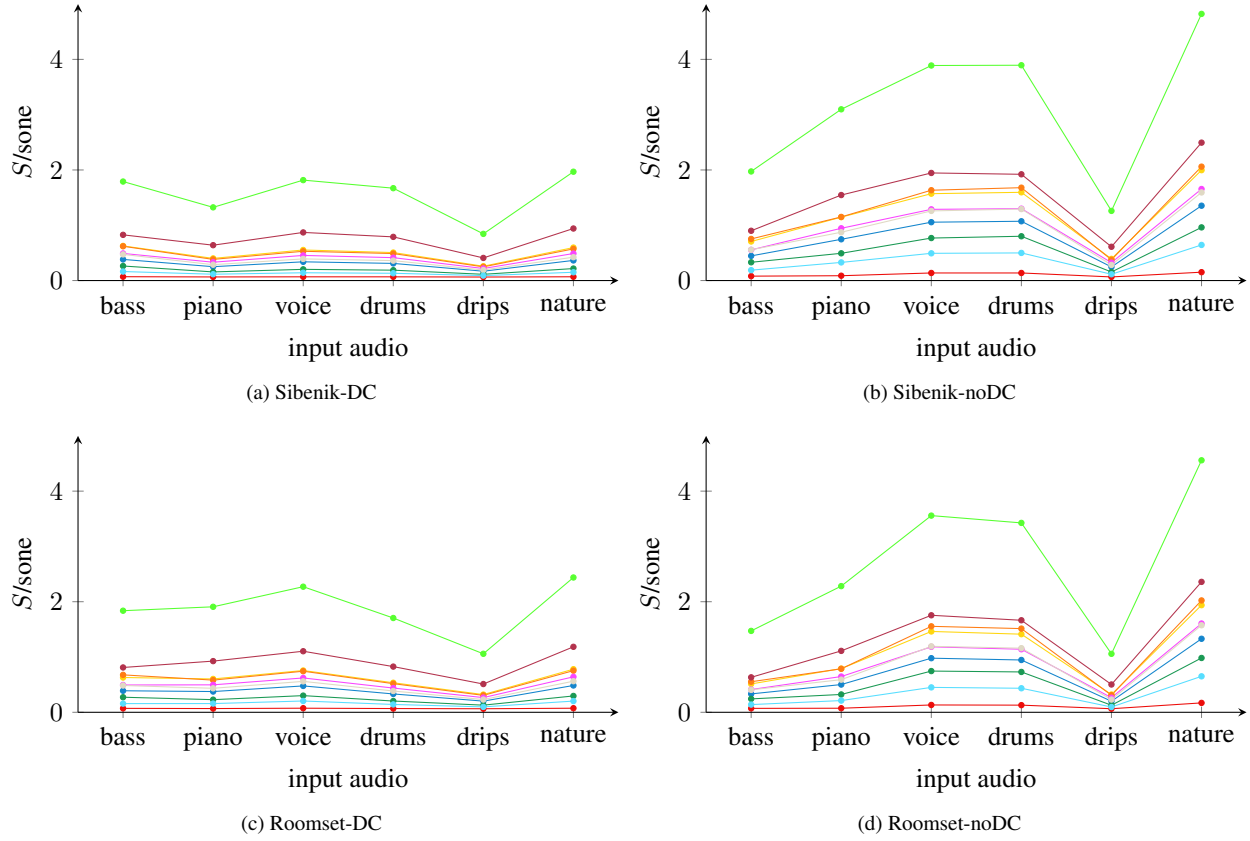


Figure 8: Error loudness at measured audibility thresholds in different test cases according to our criterion (age group 19-27). A low threshold indicates that the error is easily perceptible by the listener. The threshold loudness values of each single test subject are connected with lines. Different test subjects are marked with different colors.

tion in subsection 4.1 for its mathematical simplicity. Ideally, for a better description of  $k_l$ , we need a nonsymmetric distribution defined on  $[0, +\infty)$  that is relatively uniform on the logarithmic scale. Unfortunately, distributions satisfying such conditions usually have very complex CDFs with no analytical integral expression. This problem could be solved with numerical approximation, from which we may derive a soft threshold that better explains human hearing.

Finally, while we expect our criterion to be able to improve the efficiency of existing sound simulators, especially the real-time ones, it is not as simple as it seems. Our criterion uses signal statistics and operations on the spectral domain. The statistics require knowledge of multiple IR frames, and the Fourier transform can be time-consuming. To integrate our criterion into real applications, we need an efficient implementation that works well with real-time simulators, which we will try to develop in the future.

## References

- [1] J. Blauert and W. Lindemann. Auditory spaciousness: some further psychoacoustic analyses. *The Journal of the Acoustical Society of America*, 80 2:533–42, 1986.
- [2] R. N. Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
- [3] C. Cao, Z. Ren, C. Schissler, D. Manocha, and K. Zhou. Interactive sound propagation with bidirectional path tracing. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016.
- [4] A. Chandak, C. Lauterbach, M. Taylor, Z. Ren, and D. Manocha. Ad-frustum: Adaptive frustum tracing for interactive sound propagation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1707–1722, 2008.
- [5] G. Cirio, D. Li, E. Grinspun, M. A. Otaduy, and C. Zheng. Crumpling sound synthesis. *ACM Transactions on Graphics (TOG)*, 35:1 – 11, 2016.
- [6] B. Glasberg and B. Moore. Derivation of auditory filter shapes from notched-noise data. *Hearing Research*, 47:103–138, 1990.
- [7] S. guang Liu and D. Manocha. Sound synthesis, propagation, and rendering: A survey. *ArXiv*, abs/2011.05538, 2020.
- [8] T. Houtgast. Auditory-filter characteristics derived from direct-masking data and pulsation-threshold data with a rippled-noise masker. *The Journal of the Acoustical Society of America*, 62(2):409–415, 1977.

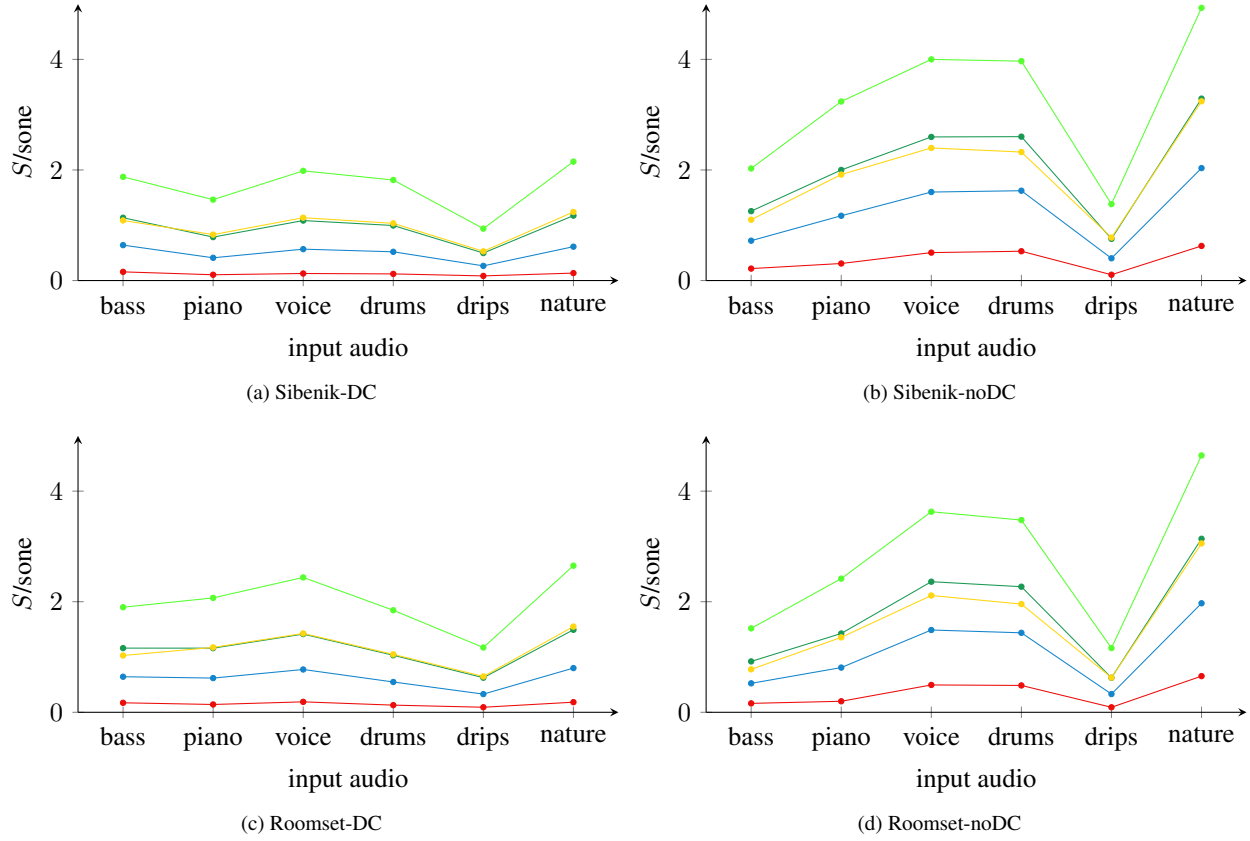


Figure 9: Error loudness at measured audibility thresholds for the 28-40 age group with hearing-impaired subjects. The result is similar to Fig. 8. Different subjects are marked with different colors.

- [9] Acoustics — Methods for calculating loudness — Part 1: Zwicker method. Standard, International Organization for Standardization, Geneva, CH, June 2017.
- [10] D. T. Kemp. Stimulated acoustic emissions from within the human auditory system. *The Journal of the Acoustical Society of America*, 64(5):1386–1391, 1978.
- [11] A. J. King and K. M. M. Walker. Listening in complex acoustic scenes. *Current Opinion in Physiology*, 18:63–72, 2020. *Physiology of Mammalian Hearing*.
- [12] P. Larsson, D. Västfjäll, and M. Kleiner. Better presence and performance in virtual environments by improved binaural sound rendering. 2002.
- [13] C. Lauterbach, A. Chandak, and D. Manocha. Interactive sound rendering in complex and dynamic scenes using frustum tracing. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1672–1679, 2007.
- [14] D. Li, Y. Fei, and C. Zheng. Interactive acoustic transfer approximation for modal sound. *ACM Transactions on Graphics (TOG)*, 35:1 – 16, 2015.
- [15] R. Mehra, A. Rungta, A. Golas, M. Lin, and D. Manocha. Wave: Interactive wave-based sound propagation for virtual environments. *IEEE transactions on visualization and computer graphics*, 21(4):434–442, 2015.
- [16] B. Moore and B. Glasberg. A revision of zwicker’s loudness model. *Acta Acustica united with Acustica*, 82:335–345, 03 1996.
- [17] B. Moore, B. Glasberg, R. Hess, and J. Birchall. Effects of flanking noise bands on the rate of growth of loudness of tones in normal and recruiting ears. *The Journal of the Acoustical Society of America*, 77(4):1505–1513, 1985.
- [18] B. Moore, R. Peters, and B. Glasberg. Auditory filter shapes at low center frequencies. *The Journal of the Acoustical Society of America*, 88(1):132–140, 1990.
- [19] W. Moss, H. Yeh, J.-M. Hong, M. C. Lin, and D. Manocha. Sounding liquids: Automatic sound synthesis from fluid simulation. *ACM Transactions on Graphics (TOG)*, 29(3):1–13, 2010.
- [20] A. J. Oxenham and C. J. Plack. Suppression and the upward spread of masking. *The Journal of the Acoustical Society of America*, 104(6):3500–3510, 1998.
- [21] A. R., N. R., R. K., and D. M. P-verber: Perceptual characterization of early and late reflections for auditory displays. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 455–463, 2019.
- [22] A. Rungta, S. Rust, N. Morales, R. L. Klatzky, M. C. Lin, and D. Manocha. Psychoacoustic characterization of propa-



- gation effects in virtual environments. *ACM Transactions on Applied Perception (TAP)*, 13:1 – 18, 2016.
- [23] C. Schissler and D. Manocha. Interactive sound propagation and rendering for large multi-source scenes. *ACM Transactions on Graphics (TOG)*, 36(4):1, 2016.
  - [24] C. Schissler, R. Mehra, and D. Manocha. High-order diffraction and diffuse reflections for interactive sound propagation in large environments. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014.
  - [25] J. V. D. Schuitman and D. D. Vries. Applying cochlear modeling and psychoacoustics in room acoustics. *journal of the audio engineering society*, 2008.
  - [26] S. Stevens. The direct estimation of sensory magnitudes: Loudness. *The American Journal of Psychology*, 100(3/4):664–689, 1987.
  - [27] S. Stevens and M. Guirao. Loudness functions under inhibition. *Attention Perception & Psychophysics*, 2:459–465, 10 1967.
  - [28] L. Trammell. Improved Pink Noise Generator Algorithm. <https://www.ridgerat-tech.us/pink/newpink.htm>.
  - [29] N. Wiener. *Time series*. MIT press, 1949.
  - [30] H. Yeh, R. Mehra, Z. Ren, L. Antani, D. Manocha, and M. Lin. Wave-ray coupling for interactive sound propagation in large complex scenes. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013.
  - [31] P. Zahorik, D. S. Brungart, and A. W. Bronkhorst. Auditory distance perception in humans : A summary of past and present research. *Acta Acustica United With Acustica*, 91:409–420, 2005.
  - [32] H. Zhou, Z. Ren, and K. Zhou. Adaptive geometric sound propagation based on a-weighting variance measure. *Graphical Models*, 116:101109, 2021.
  - [33] E. Zwicker. Subdivision of the audible frequency range into critical bands (frequenzgruppen). *The Journal of the Acoustical Society of America*, 33(2):248–248, 1961.
  - [34] E. Zwicker and H. Fastl. *Psychoacoustics: Facts and models*, volume 22. Springer Science & Business Media, 2013.
  - [35] E. Zwicker, G. Flottorp, and S. Stevens. Critical band width in loudness summation. *The Journal of the Acoustical Society of America*, 29(5):548–557, 1957.
  - [36] E. Zwicker and B. Scharf. A model of loudness summation. *Psychological review*, 72:3–26, 1965.
  - [37] E. Zwicker and E. Terhardt. Analytical expressions for critical-band rate and critical bandwidth as a function of frequency. *The Journal of the Acoustical Society of America*, 68(5):1523–1525, 1980.