

Phonak Insight

April 2016

SoundRecover2 – the adaptive frequency compression algorithm

More audibility of high frequency sounds

Phonak led the way in modern frequency lowering technology with the introduction of SoundRecover in 2008. Since then, extensive worldwide field studies with adults and children have found increased detection, distinction and recognition of high frequency sounds, better speech understanding and significant improvement in intonation and overall voice quality for users. For those with more extreme severe-to-profound losses, however, including left corner audiograms and ski-slope losses, the benefits have been limited due to the restricted audible bandwidth in which frequency compression could be applied. The new SoundRecover2 algorithm aims to restore the audibility of relevant high frequency sounds while leaving intact the low frequency structures important for good sound quality. The new frequency lowering scheme retains the essence of the original SoundRecover, utilizes an adaptive algorithm and an additional cut-off frequency in order to successfully extend these benefits to those with more extreme severe-to-profound hearing loss.

Introduction

The remapping of frequencies for extending the perceptual auditory bandwidth of hearing aid users has been commercially available for approximately 10 years now. Phonak introduced SoundRecover, non-linear frequency compression, with the first Naída in 2008, offering a solution for restoring audibility of conventionally unaidable high frequency sounds.

In order to achieve distortion-free amplification of the input signal, SoundRecover exploits the fact that vowels are dominated by greater energy in the low frequencies whereas voiceless fricatives are dominated by greater energy in the high frequencies. For this reason, SoundRecover was designed with a cut-off frequency which is the starting point of compression. For inputs below the cut-off frequency, the input signal is not subject to frequency compression. All inputs above the cut-off frequency are subject to frequency compression. Hence, the output area below the cut-off frequency is unchanged, while the output area between the cut-off frequency and the upper edge frequency is compressed with a constant compression ratio¹. The upper-edge frequency corresponds to the maximum output frequency and is set according to the audible bandwidth of the individual audiogram. The cut-off frequency is

limited to a minimum value of 1.5 kHz. This ensures that important vowel structures are left unchanged, and at the same time allows presentation of the noise-like high frequency components of speech in the individual audible range. A more detailed discussion of the SoundRecover frequency lowering system is provided in the paper by McDermott 2010.¹

Experience has shown that the SoundRecover frequency lowering scheme works very well for speech and for high frequency sounds such as bird song or environmental sounds. It can be successfully applied to hearing losses where audible hearing is available above 1.5 kHz allowing the high frequencies to be compressed into an audible region. However, the fitting of more extreme severe-to-profound losses, like left corner audiograms and ski slope audiograms, where users have only low frequency hearing thresholds, has presented unique challenges. These losses require more aggressive parameter settings (lower cut-off, stronger compression ratio) than those admissible with SoundRecover due to concerns about impact on sound quality.

In order to widen the reach of SoundRecover, the new SoundRecover2 algorithm is designed to allow operation with lower cut-off frequencies and weaker compression ratios, thereby extending the benefits of frequency compression to a broader audience of children and adults.

¹ The compression ratio is constant if viewed on logarithmically scaled frequency axes.

Functional Description of SoundRecover2

Basic functional principle

SoundRecover2 is an adaptive frequency lowering scheme based on the original SoundRecover framework. The new algorithm protects vowels which are dominated by lower frequency energy and compresses voiceless fricatives which are mainly composed of energies in the higher frequency range. Figure 1 depicts the spectral distribution of different English phonemes.




Figure 1: The speech banana showing the frequency content of speech sounds in English.

The significant difference with SoundRecover2 is that the extent of frequency lowering, i.e. the area of protection and the starting point of compression is not fixed, but is instead set adaptively as a function of the input signal. This adaptive nature is realized by the use of two cut-off frequencies, of which only one is active at any moment in time. Based on the momentary energy distribution in the input signal, the system determines instantaneously which one of the two cut-off frequencies is applied. Thus, the functional principle of SoundRecover2 is similar to that in SoundRecover; with the increased sophistication that it now switches automatically between two possible starting points of compression, respectively between a "lower" and an "upper" cut-off frequency. As in SoundRecover, frequency lowering is always carried out with a predefined constant compression ratio regardless of which cut-off frequency is momentarily active.

Hence, SoundRecover2 instantaneously maps input components to the output depending on their energy content. This adaptive frequency lowering processing is accomplished by simply recognizing the different energy distributions of tonal and noise-like structures of the input signal. In case of more low frequency content, frequency compression takes place with the upper (higher) cut-off frequency in order to "protect" the low frequency sounds from being compressed. In case of more high frequency content, frequency compression takes place with the lower cut-off frequency to restore audibility of the high frequency sounds. When applied to speech signals, this strategy leaves vowels intact while allowing compression of important high frequency information in fricatives down to sufficiently low output frequencies. The output curve shown in Figure 2 illustrates this adaptive behavior schematically.




Figure 2: Sample output curve for SoundRecover2. Depending on the energy distribution of the input signal, the frequency compression starts either at the lower or at the upper cut-off frequency.

Impact of the adaptive behavior

The impact of the adaptive behavior of SoundRecover2 is substantial. The adaptively determined starting point of compression ensures that input signal components are lowered, to a high extent, only when they have significant high frequency energy. As a result, the lower cut-off can be set well below the current limit of 1500 Hz increasing the area of compression and allowing in turn, weaker compression ratios than possible with the original SoundRecover. The value of the upper cut-off remains limited however, and can be set quite high, because the upper cut-off is applied only when significant low frequency energy is present. This extends the output area below the upper cut-off frequency where the signal is protected and left intact. Thus it is guaranteed that tonal structures and other low frequency components are protected and not compressed, and in particular, it is ensured that important vowel formants will not be adversely affected.

Outcomes

The adaptive nature of SoundRecover2 allows frequency lowering with a lower overall cut-off frequency and weaker compression ratio than was available with the original SoundRecover. The high value of the upper cut-off yields a better sound quality with more naturalness and familiarity, and less distortions of tonal components. At the same time the lower possible value of the lower cut-off provides for extended access to high frequency sounds for all kinds of hearing losses. Thus, the fitting range of SoundRecover2 is expanded to include subjects with a very restricted audible bandwidth as commonly seen with those who have greater severe-to-profound hearing loss, left corner audiograms and ski slope losses. Put simply, more users will be able to benefit from frequency lowering than ever before with the new SoundRecover2 algorithm. The use of a weaker compression ratio results in less alteration and thus an improved preservation of the spectral shape of mid to high frequency inputs, this in turn allows for better recognition of speech and environmental sounds and leads to even better spontaneous acceptance than seen with SoundRecover.

Figure 3 presents the spectrograms of the sample sentence "my name is asa" (a) without frequency lowering, (b) with SoundRecover and (c) with SoundRecover2 processing.

Figure 3(a) shows pronounced formant structures up to 5.5 kHz at 0.2 seconds to 0.5 seconds and two high frequency /s/ phonemes at 1.2 seconds and 1.9 seconds.




Figure 3(a): Spectrogram of the sample sentence "my name is asa" without frequency lowering.

In Figure 3(b), showing SoundRecover with a cut-off frequency of 1500 Hz and compression ratio of 2.1, the audible bandwidth extends up to approximately 4000 Hz. The /s/ phonemes at 1.2 and 1.9 seconds are compressed down into a frequency area between 2.5 and 4 kHz. Note that the fine-spectral structures above the cut-off frequency of 1500 Hz at the beginning of the sentence are not fully preserved at this maximum setting.




Figure 3(b): Spectrogram of the sample sentence "my name is asa" with SoundRecover (cut-off: 1500 Hz, compression ratio: 2.1).

In Figure 3(c), showing SoundRecover2 with a lower cut-off frequency of 1479 Hz, upper cut-off frequency of 3600 Hz, and compression ratio of 1.4, the audible bandwidth extends up to approximately 4000 Hz as well. Note the preservation of the spectral fine structures up to the upper cut-off of 3600 Hz at the beginning of the sentence at 0.2 to 0.5 seconds, and the remapping of the two significant high frequency /s/ phonemes at 1.2 and 1.9 seconds down into a frequency area as low as between 2000 and 3000 Hz.




Figure 3(c): Spectrogram of the sample sentence "my name is asa" with SoundRecover2 (lower cut-off: 1479 Hz, upper cut-off: 3600 Hz, compression ratio: 1.4).

Evidence of Benefit

Table 1 summarizes the theoretical benefits of SoundRecover2 in comparison to the original SoundRecover, namely the extension of the fitting range and audiological benefits in terms of audibility, distinction and sound quality. Moreover, a smooth transition from existing frequency lowering technology is expected for most users. In particular, users of the original SoundRecover feature should be able to switch over and acclimatize to SoundRecover2 without concerns.

Extended fitting range	<ul style="list-style-type: none"> fits even more severe-to-profound losses with acceptable sound quality better inclusion of left-corner and ski slope audiograms
Improved audiological performance	<ul style="list-style-type: none"> better audibility of high frequency sounds. See internal evidence below improved distinction, detection and recognition of compressed high frequency components maintains sound quality, more specifically preservation of familiarity and naturalness in general and in particular for low and mid frequency components better awareness of environmental sound and as a result potentially better spontaneous acceptance and reduced acclimatization time

Table 1: Theoretical benefits of SoundRecover2 relative to the original SoundRecover.

Study results

Research has confirmed the theoretical benefits outlined in Table 1. An early study tested the performance of 14 children with severe-to-profound high frequency sensorineural hearing loss on a very mature prototype of SoundRecover2 relative to the original SoundRecover feature (Wolfe et al. 2016). The study showed improved word recognition in quiet and improved recognition of plurals. No detriment in the detection or recognition of consonants was seen, and long-term users of the original SoundRecover feature were able to switch to SoundRecover2 without a long period of acclimatization.

Another study carried out at Phonak headquarters compared the audiological performance of the original SoundRecover with SoundRecover2 on 8 male hearing impaired adults (average age: 56.8 years) with profound symmetrical sensorineural or mixed average loss higher than 90 dB over the frequencies from 250 Hz to 8 kHz). For this profound group, the phoneme perception test (Schmitt et al. 2016) revealed significantly better detection thresholds for 3 out of 4 tested stimuli (Figure 4), and a significantly better recognition threshold for 1 out of 4 stimuli (Figure 5).




Figure 4: Phoneme perception test: median detection thresholds for profound hearing losses, Naída Q SR (original SoundRecover)² vs Naída V SR2 (SoundRecover2). Detection of 3 out of 4 stimuli tested (sh5, s6, s9) was significantly better with SoundRecover2.




Figure 5: Phoneme perception test: median recognition thresholds for profound hearing losses, Naída Q SR (original SoundRecover)² vs Naída V SR2 (SoundRecover2). Recognition of 1 out of 4 stimuli tested (Asha5) was significantly better with SoundRecover2.

The SoundRecover2 Fitting Concept

Objective

From extensive research studies with patients, it is known that achieving good audibility and distinction are among the most important goals for providing better hearing. At the same time, good sound quality is important for spontaneous acceptance and for hearing comfort. The objective in fitting SoundRecover2 is to preserve, and when necessary, restore the audibility of high frequency sounds and the discrimination of compressed high frequency sounds, and to maintain the familiarity of the entire frequency range.

² Where detection or recognition thresholds were not measurable with the original SoundRecover (s9, Asa6 and Asa9 stimuli), the thresholds were set to 75 dB.

Perceptual trade-offs in fitting frequency lowering solutions

In fitting any kind of frequency lowering scheme, a perceptual trade-off must be managed between the balance of audiological benefit and overall sound quality. Most contemporary frequency lowering approaches manage this trade-off sufficiently well for varying degrees of high frequency hearing loss. However, for some extreme severe-to-profound hearing losses, such as those with no aidable hearing beyond 2000 Hz, the application of frequency lowering is a particular challenge. When frequency lowering is applied to present high frequency sounds in the strongly limited audible range of these patients, sound quality may be compromised and thereby limits the benefit of frequency lowering algorithms.

This challenge is illustrated in the trade-off diagram of Figure 6 which qualitatively shows the effect of adjusting the strength of the frequency lowering scheme along the perceptual dimensions of audibility and sound quality.




Figure 6: Perceptual trade-off curves for extreme severe-to-profound hearing loss with contemporary frequency lowering solutions showing the unsatisfactory compromise between audiological benefit and sound quality. Improved audibility is only achieved at the expense of sound quality, and vice versa.

SoundRecover2 not only offers the technical means to overcome the above described limitations, but also a way for hearing care professionals to better manage the involved perceptual trade-offs.

Basic principle of SoundRecover2 fitting

The performance of frequency lowering schemes can be characterized in an intuitive and simple way by means of perceptual dimensions. Thus the fitting approach for SoundRecover2 is based on the following three important perceptual dimensions:

- **Audibility** of high frequency sounds like phonemes /s/, /f/ and /th/,
- **Distinction or discrimination** of lowered high frequency sounds like /s/ and /sh/ and
- **Sound quality** of low and mid frequency sounds, like vowels /a/, /e/, /i/.

These three interrelated perceptual dimensions can be depicted graphically with the "trade-off triangle". Figure 7 illustrates three possible configurations of the trade-off triangle, a default setting, a weaker setting and a stronger setting.




Figure 7: The SoundRecover2 trade-off triangle for a default SoundRecover2 setting (green), a setting with stronger audibility (blue) and another setting with stronger distinction (purple). Note the effects of modifications on one perceptual dimension onto the other two dimensions.

SoundRecover2 aims to optimize the balance between these interdependent perceptual dimensions by:

- 1) **making use of the individual's audible bandwidth in the best possible way**
 - the maximum output frequency is set to the upper limit of the individual's audible bandwidth in order to maximize stimulation of the hearing nerve where audibility can be achieved without frequency lowering
 - no risk of deprivation
- 2) **protecting the mid and low frequencies in the best possible way**
 - the upper cut-off frequency is set high enough so that speech which is audible without compression remains unaffected
 - no risk for distortion of mid and low frequencies
- 3) **presenting the compressed sounds in an optimally selected frequency area**
 - the starting point of compression (the lower cut-off) is set at the lowest possible frequency yielding a wide area of compression with weak compression ratio
 - increases the overall strength of frequency compression for those extreme severe-to-profound losses extending the benefits of frequency lowering to more users.

Pre-calculation

Taking into account these three elements and based on subject data collected over several years, an optimized pre-calculation and starting point for fitting has been carefully developed. The resulting default setting for all fittings is the best possible balance across the three perceptual dimensions, in particular, allowing for good audibility of high frequency sounds while enabling sufficient distinction between compressed sounds, and at the same time yielding an acceptable overall sound quality.

SoundRecover2 fitting procedure

Fine-tuning allows for individual adjustment, if required, of the two dimensions audibility and distinction. An adjustment for greater audibility results in a decreased lower cut-off frequency, while an

adjustment towards better distinction results in a change of compression ratio. The third dimension, sound quality, is automatically optimized for each change of the essential fitting dimensions within the fitting software. Thus, when the balance between audibility and distinction has been optimized for a particular user, sound quality is always automatically adjusted to maintain optimal clarity.

Sound quality can be characterized in more detail in terms of sound clarity and hearing comfort. On occasion, a patient may need a different balance between clarity and comfort. In this case, it is possible to further adjust the sound quality with 4 predefined settings to balance optimum clarity and personal hearing comfort. An adjustment to increase hearing comfort results in a shift of the upper cut-off frequency towards the maximum audible bandwidth.

Figure 8 illustrates the fitting procedure for SoundRecover2 with fine-tuning of audibility and distinction and subsequent fine-tuning of sound quality. This method is implemented with two perceptual sliders within the fitting software.




Figure 8: The SoundRecover2 fitting procedure. The basis of the fitting is the default setting as derived from pre-calculation. Fine-tuning allows the optimization of the audibility of high frequency sounds and the distinction of the lowered high frequency components. The resulting settings include an automatically optimized sound quality to achieve maximum sound clarity. The sound quality can be further fine-tuned towards more hearing comfort in four predefined steps.

Outcomes

This fitting procedure, based on perceptual dimensions, allows the hearing care professional to adjust SoundRecover2 in an intuitive, understandable and user-friendly way. Reports from patients during frequency lowering fitting relate much better to perceptual dimensions than to technical parameters or the "weaker-stronger" control of the original SoundRecover.

The default settings derived from the pre-calculation set an excellent starting point for the fitting by providing good audibility, together with sufficient distinction of compressed sounds and acceptable sound quality.

Conclusions

The new SoundRecover2 frequency lowering algorithm is specifically designed to allow operation with lower cut-off frequencies and weaker compression ratios in order to expand the already existing user base of SoundRecover. This ambitious goal was accomplished by introducing an additional cut-off frequency and by adaptively adjusting the starting point of compression to the energy content of the input signal. By instantaneously switching between the lower and upper cut-off frequencies, high frequency components can be compressed into a larger frequency range with a weaker compression ratio, while low frequency components can be protected and are left unaltered.

A preliminary external study showed improved word recognition in quiet and improved recognition of plurals, with no detriment in the detection or recognition of consonants. Long-term users of the original SoundRecover feature were able to switch to SoundRecover2 without a long period of acclimatization (Wolfe et al. 2016). In an internal study at Phonak headquarters with a group of subjects with profound hearing loss, the Phoneme Perception Test revealed significantly better detection thresholds for 3 out of 4 tested stimuli and a significantly better recognition threshold for 1 out of 4 high frequency stimuli.

Together with the new signal processing, a new fitting concept has been developed which allows easier management of the difficult trade-off between audiological benefits and sound quality in frequency lowering systems. The resulting perceptual fitting is based on the "trade-off triangle" which depicts the interrelated perceptual dimensions: audibility, distinction and sound quality. The fitting pre-calculation offers an optimal starting point for good audibility of high frequency sounds, enables sufficient distinction between compressed sounds and delivers an acceptable overall sound quality. Fine-tuning to match users' individual needs can easily be carried out with the use of perceptual sliders in the Target fitting software.

References

- McDermott, H. (2010). SoundRecover - The importance of wide perceptual bandwidth. Phonak Background Story.
- Wolfe, J., Duke, M., Schafer, E., Rehmann, J., Jha, S., John, A., Jones, C. (2016). Preliminary evaluation of a novel non-linear frequency compression scheme for use in children. Submitted to International Journal of Audiology.
- Schmitt, N., Winkler, A., Boretzki, M., Holube, I. (2016). A phoneme perception test method for high frequency hearing aid fitting. Journal of the American Academy of Audiology.

Authors

Julia Rehmann is an Audiological Engineer at Sonova Headquarters in Switzerland.

Siddhartha Jha is a DSP Engineer in the Science & Technology department at Sonova Headquarters in Switzerland.

Silvia Allegro Baumann is a Senior DSP Engineer in the Science & Technology department at Sonova Headquarters in Switzerland.