# Residual Inhibition

by Clyde Witchard

## In brief

**What is residual inhibition?**

Residual inhibition is a temporary quieting of tinnitus that can happen after listening to the right type of sound. A sound that can cause this to happen is called here a *quieting sound* or a *trigger sound*. Although the effect of a single quieting sound is temporary, for many people tinnitus can be kept reduced for much longer periods, by simply keeping a repeating form of the quieting sound playing.

**Who discovered it?**

The effect was first discovered by a doctor called James Spalding in 1903. But it was not until the 1970s and 1980s that research began in earnest.

**How many people experience residual inhibition?**

The largest study to date (involving over one thousand people) found that more than 8 out of 10 people with tinnitus can experience at least some residual inhibition. Other studies, using various different methods, agree that most people with tinnitus can experience the effect.

**How many experience complete silence?**

The largest study found that just over half of people tested experienced a period of complete tinnitus silence after a quieting sound. Other studies, under different conditions, have found a variety of results.

**How long does residual inhibition last?**

The effect of a single quieting sound varies greatly from person to person. For most people the period of reduced tinnitus lasts somewhere from a few seconds up to a few minutes afterwards, although for some people it lasts much longer. Residual inhibition can be greatly extended by playing a repeating form of the sound. With this approach, for people who respond, their tinnitus can be kept suppressed for as long as the repeating sound plays.

**What research has been done on residual inhibition as a therapy?**

Three separate at-home studies have investigated residual inhibition treatment. These studies asked patients to listen to a single period of quieting sound, typically one or two times a day, rather than a repeating type of sound. One study reported that about half of the participants wished to carry on using the system after the trial. Another reported that 90% intended to use the therapy further, and the study's authors commented, "While Residual Inhibition Therapy does not appear to be an outright cure for tinnitus, its implications are positive." The third study did not ask their patients about future intended use, but it did report that 90% were able to achieve periods of reduced tinnitus after listening to the provided sounds.

**Are there any long-term benefits from prolonged use?**

Several published studies have reported on the effects of sustained or repeated listening to the sorts of sounds that can cause residual inhibition. These studies include reports, from well-known tinnitus researchers, of some patients who experienced longer-term improvement in their residual inhibition response. There are even reports of tinnitus going into complete remission for some patients – but these cases appear to be rare. It has been put forward that there may be more chance of longer-term improvement for people whose tinnitus has started recently, in line with evidence from animal studies.

**Is it safe?**

Yes. As with any type of sound, sound levels should be kept at general safe listening levels. After listening to a quieting sound, some studies have found that a small proportion of people experience a temporary increase (rather than a decrease) in the level of their tinnitus. Two studies have investigated how long such increases last for. They found that, for their few patients who experienced an increase, it lasted for less than a minute.

**What types of sound trigger the deepest and longest residual inhibition?**

Quieting sounds should be tuned to the pitch range of the person's tinnitus to have the most effect. That pitch range is generally found within the pitch range of their hearing loss. To be effective, the quieting sound must have sufficient volume and duration, and it should be played into the ear (or ears) in which the tinnitus is present. Three different types of sound waveform have seen the most investigation. In various studies all three have been found to be effective. Short and repeated forms of quieting sound have been found to be effective too.

**How does it work?**

It is not known for sure how residual inhibition works. However, four separate studies have used brain scanning techniques to investigate it. An emerging theory, supported by a number of researchers, is as follows. People with tinnitus normally have at least some hearing loss. This causes the brain to be starved of normal nerve signals from the ear, in the pitch range of the hearing loss. Groups of brain cells (neurons) react to this "missing input" by repeatedly misfiring, in a synchronized way. This ongoing abnormal activity is thought to be a major mechanism for tinnitus. If a quieting sound is played, this temporarily restarts the normal nerve signals to the brain in the pitch range of the hearing loss. In turn, this disrupts the abnormal activity of the groups of neurons, and reduces (or even stops) the tinnitus. In many people, this temporary return to normal activity (and reduced tinnitus) remains for a time after the quieting sound has stopped – the effect we call "residual inhibition".

*All of the above questions are answered more thoroughly (along with references) on page 4 onwards.*

**What is the HushTinnitus system?**

The HushTinnitus system ([www.hushtinnitus.com/ri](www.hushtinnitus.com/ri)) is a set of sound tracks that uses residual inhibition to provide immediate relief from tinnitus. Building on the independent research reviewed in this document, it brings together a wide range of optimizations. To cater for the great variability of tinnitus and residual inhibition, 67 sound tracks are custom-generated for each user. Each track is optimized to the user's own hearing.

In 2007, Larry Roberts (a leading tinnitus researcher, at McMaster University in Ontario) wrote that "residual inhibition can be a source of relief for tinnitus sufferers who have experienced a sound that otherwise has known only a life of its own. Vernon and Meikle (2003) have described instances where patients broke into tears at their first experience of silent ears after years of unremitting noise. Hence there is much current interest in optimizing residual inhibition for its possible clinical benefits." He also wrote that "considerable variability was present between subjects" and this "leads one to ask whether we can identify variables and procedures that may optimize residual inhibition for individual cases." The design of the HushTinnitus system pursued exactly these thoughts.

The HushTinnitus system uses a web-based application to gather relevant hearing estimates from the user. From this data it generates 67 custom sound tracks. These tracks are designed to cater for a wide range of responses, from people with strong residual inhibition responses (using "Type 1" sounds), through to people with less response (using "Type 2"), all the way through to an optimal masking sound ("Type 3a") that is suitable for people with no residual inhibition response at all. The main focus of the system is on sound tracks that repeat continuously and unobtrusively in the background. However, the Type 3 tracks can also be used for optimized one-shot stimulation, for people who get a useful longer residual inhibition response from this.

The system was designed to meet four main aims: to maximize effectiveness, to minimize obtrusiveness, to cater for variability, and to be easy to use. In terms of maximizing the depth and duration of residual inhibition, the design relies entirely on the results of independent mainstream published research, as reviewed in this document. Much of the design effort of the project was spent on minimizing the obtrusiveness of the sounds. The system was developed to be comfortable to use over extended periods of time – all day if required. A number of the methods and optimizations used by the HushTinnitus system are new. These include a new method for fading sound in and out in a less obtrusive way, and perceptually smoother forms of noise.

*A full outline of the HushTinnitus system starts on page 20.*

# In full

The following pages discuss the same questions as the preceding “In brief” section, but in greater detail. The discussion is split into two sections. The first section (starting on the next page) is a review of the published research to date into residual inhibition. The second section (starting on page 20) gives an outline of HushTinnitus, a new therapy system for tinnitus relief. The publications used are listed at the end of this document, and are referred to throughout the text with numbers in square brackets (e.g. “[1]”).

The whole document is written for the general reader. It keeps medical and technical terminology to a minimum. Where technical terms are necessary, they are explained.

# Research review

This review aims to cover all of the published research to date that has dealt directly with residual inhibition in a significant way. It focuses on a set of simple general questions (which are the headings for each section below). It is based on research from peer-reviewed academic publications, referred to throughout the text.

## What is residual inhibition?

If you play a specific pulse of sound to a person with tinnitus, in most cases you can reduce, or even silence, their tinnitus for a period of time after the pulse has stopped.[49][45][20][60][40][57] The effect of playing just a single such "trigger" pulse lasts anything from seconds or minutes, up to hours, or even a day or more in a small proportion of people.[49][20][17] The effect is known as *residual inhibition*.[13] The duration of the residual inhibition can be greatly extended by simply repeating the trigger pulse, with a period of silence following each pulse (i.e. pulse, silence, pulse, silence…).[9][24] In this way, continuous tinnitus suppression can be maintained for as long as the repeating sound is played.

It is worth covering a few points of terminology relating to residual inhibition. Firstly, the term *masking* benefits from some discussion. In the wider science of how we perceive sound (*psychoacoustics*), masking basically means "drowning out". So, for example a scientist might say, "Sound A masked sound B," rather than, "Sound A drowned out sound B." The term masking is also widely used in regard to tinnitus: if an external sound can drown out the tinnitus (i.e. make it imperceptible), then it is said to *mask* the tinnitus.[4][5][9] The sound or device that causes the tinnitus masking is sometimes called a *masker*.[14][20] Since the first studies, residual inhibition has been investigated alongside tinnitus masking: typically a burst of sound is played, during which masking of the tinnitus is often reported, and afterwards any continued suppression of the tinnitus (i.e. residual inhibition) is observed.[4][9] It has therefore become common terminology to refer to the sound that triggers residual inhibition as being "the masker".[14][23] However, when the focus is on using residual inhibition as a therapy, I feel that the term "masker" (as a general term for the trigger) is unhelpful, particularly for people new to the subject. The interest then is in the *after*-effects of the sound, not really the effects that happen *during* the sound. The problem is compounded by the fact that "tinnitus masking" (meaning here, the drowning out, or partial drowning out, of tinnitus with a continuous sound) is a widely known form of tinnitus relief in its own right. Therefore, I use the terms *quieting sound* or *trigger sound* in this document, to mean a sound for the purpose of triggering residual inhibition. ("Quieting sound" seems particularly helpful to people new to residual inhibition, as it is somewhat self-explanatory. "Trigger sound" seems useful in a more technical context: "trigger" clearly signifies a distinct event (i.e. the trigger sound) that causes a following action (i.e. the residual inhibition). In passing, some researchers have objected to the use of the term "masking" for any kind of tinnitus suppression, even when the suppressing sound is continuous.[46][47])

## Who discovered it?

The first known mention of the residual inhibition effect appears in a 1903 medical paper written by James Spalding, a doctor from Portland, Maine, who himself had tinnitus.[1] Spalding used to play various musical instruments to his patients to find the pitch of their tinnitus. He noticed that by playing prolonged notes on "an organ or reed pipe or violin string of the same pitch" he could reduce or remove his patients' tinnitus "for a certain length of time". He commented that this "seems to open a new field in the treatment of this obstinate affection."

The effect was then independently re-discovered a number of times. In 1928, Isaac Jones and Vern Knudsen, doctors from Los Angeles, noted the effect while trying to cure tinnitus with an unusual electrical device they had made themselves.[3] Their device simply played distorted "mains hum" to their patients through a speaker. Although none of their patients was cured by this, some did report a period of temporary relief or remission of their tinnitus after listening to the sound.

Three years later, Emanuel Josephson, a New York doctor, used an electronic valve sound generator to make pure tones (i.e. single frequencies), in order to find the pitch of his patients' tinnitus.[4] He then discovered that by playing this pure tone to them, their tinnitus could be "drowned out" and remained "lost for a variable period of time after stimulation has ceased."

In 1971, Harald Feldmann, a researcher at the University Clinic of Otolaryngology in Heidelberg, completed the first detailed investigation of tinnitus masking.[9] Studying around 200 patients, he again independently found the residual inhibition phenomenon – or simply "inhibition" as he termed both it and tinnitus masking, based on his theory that neural inhibition in the brain may be the basis of both mechanisms. Regarding the effect, Feldmann noted that "the duration of the stimulus itself and the duration of the testing as a whole can modify the result. In some subjects during the test-procedure the silent periods grew longer and longer, and finally the tinnitus vanished completely. By adapting the rate of on and off periods always an optimal mode can be found, so that the tinnitus is constantly suppressed by a minimum of external periodic stimuli. Prolonged stimulation may in some cases stop the tinnitus for quite a considerable period of time. There seems to be a neural mechanism able to block the spontaneous activity of the subjective* tinnitus. … If this mechanism could be trained or activated there might be a way to cure patients of their distressing tinnitus."

From the 1970s, the pace of tinnitus research accelerated. This included the contributions of Jack Vernon, a research psychologist working at Oregon Health and Science University. Vernon started the world's first clinic for treating chronic tinnitus, and he co-founded the American Tinnitus Association (ATA). In a 1975 paper, he was the first to coin the term "residual inhibition" – based on Feldmann's "inhibition", prefixed by "residual" to indicate its sustaining nature after removal of the masker sound.[13] Some researchers have commented that residual "suppression" would be a better term,[23][56] as it is more neutral to the possible mechanisms behind the effect; but the term residual "inhibition" is now widely used. Two years later, Vernon was the first to give a detailed proposal for noise masking as a therapy for tinnitus.[14] He commented, "What does residual inhibition include? This is an intriguing puzzle about which a great deal more work must be done. The existence of residual inhibition indicates that the masker may be interfering with whatever process it is that produced the tinnitus."

## How many people experience residual inhibition?

One of the largest studies of tinnitus patients, the Tinnitus Archive run by Oregon Health and Science University, tested residual inhibition in 1,451 people.[49] They found that 88% experienced at least some reduction of their tinnitus, after listening to 1 minute of broadband noise at a level moderately (+10 decibels) above the level that just masked their tinnitus. (*Broadband noise* is noise that spans a wide range of pitch.)

A separate study, by Marina Savastano of Padua University, tested 1,440 people.[45] 52% of this group showed at least some residual inhibition after listening to 1 minute of broadband noise. A likely reason for this figure being less than in the Oregon study is that the trigger sound seems to have been played quieter (by 10 decibels). (The importance of sound level is discussed more in a later section.)

* Tinnitus can be categorized as either "subjective", when it cannot be heard by anyone else; or "objective", when it can be heard by an examining clinician. Subjective tinnitus is far more common than objective tinnitus.

Using various forms of trigger sound, with various sound levels and durations, other studies have found a reduction of tinnitus in 78% (in over 200 people),[20] 69% (in 59 people),[60] 70% (in 46 people),[40] 63% (in 32 people),[23] and 90% (in 10 people) of people tested.[57]

## How many experience complete silence?

The Tinnitus Archive found that in 55% of patients their tinnitus disappeared completely for a period after the residual inhibition trigger.[49] The Savastano study found this to be 30%,[45] but again this lower figure may have been influenced by the quieter trigger sound. In other studies, 35% (in over 200 people),[20] 22% (in 59 people),[60] 54% (in 46 people),[40] and 20% (in 10 people) experienced total silence.[57]

## How long does residual inhibition last?

It has been found that there is a great deal of variation from one person to the next, in terms of how long tinnitus stays suppressed after a single trigger sound. Most studies of the effects of single trigger sounds have used quite long sounds (of a minute or more) to examine this. By arranging for the trigger sound to repeat at regular intervals, residual inhibition can be sustained for as long as the repeating sound is played, for those who have the necessary level of response. Studies have generally used particularly short trigger sounds to investigate this arrangement. This section looks first at studies of residual inhibition duration after a single trigger sound.

976 patients in the Tinnitus Archive listened to 1 minute of broadband noise, and then reported how long it took for their tinnitus to return to the normal level.[49] For 57%, this took less than 2 minutes; and for the remainder, the return of their tinnitus took varying amounts of time – for a few taking longer than 10 minutes. See below.

| | |
|---|---|
| Less than 2 minutes | 57% |
| 2 to 4 minutes | 23% |
| 4 to 6 minutes | 11% |
| 6 to 8 minutes | 4% |
| 8 to 10 minutes | 2% |
| More than 10 minutes | 3% |

Vernon described a typical patient response, after the trigger sound ends, as starting with a "silent phase" of 20 to 35 seconds (about which he commented, "seldom is it less but not infrequently it is greater"), followed by 30 to 35 seconds in which the tinnitus sound gradually recovers.[15] Three years later Vernon wrote, "In some cases residual inhibition can last a matter of hours after only a few minutes of masking at moderate intensities of 20 to 30 decibels sensation level."[17] He then described the way that tinnitus returns, referring to the shape of a graph showing how the intensity of the tinnitus changes over the time period of its return. "There are a number of different descriptions of the gradual return of tinnitus. For some it is a negatively accelerating curve†; for some it is a positively accelerating curve; to others it is a linear function; for some it is stair-stepped; for others it is scalloped; in rare cases it is all in one step; for a few the tinnitus appears and then disappears, repeating this process many times before gaining the normal level." Similarly, a paper by Tyler and colleagues shows five different graph shapes for the return of tinnitus described by their patients, and notes that there were other shapes besides.[28]

† A "negatively accelerating curve" is a curve that gets less steep as time increases. A "positively accelerating curve" is one that gets steeper as time increases. A "linear function" is simply a straight line on the graph.

Hazell and Wood agree that for most people the totally silent phase lasts for less than one minute, after one minute of broadband masker noise.[20] Regarding the use of longer periods of masker noise, they comment, “Commonly, patients wear a masker for 2 or 3 hours in the evening, removing it perhaps half an hour before going to bed and experiencing consequent relief from tinnitus during the important period while trying to get to sleep. In a small proportion (less than 4%) residual inhibition is very marked. Only 10 to 15 minutes of masker use is required to produce residual inhibition for the whole of the day. Sometimes residual inhibition can be achieved by using high levels of masking for short periods of time. On other occasions it is more effective to use the minimal masking level of the instrument for longer periods.”

Olsen and colleagues tested 46 people in their clinic, measuring residual inhibition after they had listened to 2 minutes of sweeping-pitch pure tones.[40] The tones were individually tuned around the tinnitus pitch of each patient. Different individuals experienced quieting of their tinnitus for periods ranging from 10 seconds to almost 6 minutes. 23 people in the study group (all of whom experienced total silence during residual inhibition) went on to try the sweeping-pitch pure tones at home. Of these, 2 people reported experiencing complete residual inhibition (i.e. total silence) for several hours after listening to the 2-minute trigger sound.

The Roberts study also saw great variation between different people, in terms of the duration of their residual inhibition.[60] The in-clinic measurements were limited to 45 seconds by a limit in their software, but at home patients reported up to 20 minutes residual inhibition.

Sockalingam and colleagues played various types of sounds (including pure tones, warbling-pitch tones, narrowband noise and broadband noise) to 10 participants to test their residual inhibition.[57] In their clinic, they played these sounds for 5 minutes, at a sound level moderately above the tinnitus level (like the Oregon study). The participants experienced an average of 7 minutes residual inhibition in response to a pure tone tuned to the pitch of their tinnitus. One participant experienced residual inhibition for over 45 minutes after listening to this type of sound.

All the residual inhibition timings given above were measured after the patient listened to just a *single* pulse of sound. By repeatedly alternating a pulse period, then a silent period (i.e. pulse, silence, pulse, silence…), the total duration of residual inhibition can be extended. Of course, the residual inhibition can then last for as long as the repeating sound is played. The first experiments with repeating pulses were carried out by Harald Feldmann in 1971.[9] He used particularly short pulses of narrowband noise, each pulse lasting just half a second. In reference to just one patient, he reported periods of total silence lasting up to 1.8 seconds after each pulse, depending on the sound level. Feldmann was therefore able to create continuous tinnitus silence by repeating these pulses at an appropriate rate. In a later paper, he reported using repetition of pulses as short as one tenth of a second to sustain residual inhibition.[24]

## What research has been done on residual inhibition as a therapy?

Much of the published work has focused on studying the characteristics of the residual inhibition effect in a fairly “pure research” way, or sometimes as a side point in a paper that is mainly concerned with other matters. There are, however, some studies that have focused on present-day application of residual inhibition as a therapy.

The 1996 Olsen study, using sweeping pure tones matched to the tinnitus pitch, was the first to have an at-home field test.[40] 23 people participated, all of whom had been pre-selected on the grounds of experiencing a period of complete residual inhibition after listening to the 2-minute trigger sound. They tried the system over a period varying between 11 and 85 days, listening to the trigger sound between 1 and 5 times per day, e.g. before going to sleep, waking during the night, reading, before watching TV, studying or eating. At the end of the field test, 11 of this group of 23 said they would like to continue using the system.

In 2003, the Canadian Hearing Society ran an at-home study completed by 31 people.[44] The study used pure tones, warbling-pitch tones and narrowband noise, all tuned to the tinnitus pitch. The participants spent a week on each sound type, listening for 15-minute sessions once or twice per day, typically with one session before going to bed.[66] After the 3 weeks of at-home testing, the warble tone was found to be the most effective type. Comparing Tinnitus Handicap Questionnaire (THQ) scores before and after showed average improvements of 12% (social, emotional and behavioral), 8% (tinnitus and hearing) and 7% (outlook). A questionnaire on tinnitus severity, annoyance and effect-on-life improved by 5% on average, and there was also a 5% average reduction in the amount of time that people were aware of their tinnitus. 31% of the participants reported their tinnitus was better after the 3-week study, 41% unchanged, and 19% worse. 90% reported "that they would not want to return the CD at the conclusion of the study and were interested in using it on their own in the future". The report concluded, "While Residual Inhibition Therapy does not appear to be an outright cure for tinnitus, its implications are positive. It may be utilized as a more economical alternative with relative limited financial risk for clients wanting to explore a treatment for their tinnitus, but are not prepared for the subsequent expense of Tinnitus Retraining Therapy."

The 2007 Sockalingam study also had a 4-day at-home trial, completed by 9 people, using similar trigger sounds.[57] The focus of the study was to find the most effective type of trigger sound, including looking at the effect of pitch. The main conclusion was that sounds matched to the tinnitus pitch were more effective: more patients experienced residual inhibition, and the duration of tinnitus suppression was longer.

Another at-home study, run by Maren Struve and colleagues at the University of Heidelberg, commenced in 2007,[58] but the team has yet to publish results. This study took a combination-treatment approach, in which 30-minute trigger sound sessions were combined with the drug pregabalin, aimed at reducing over-activity of neurons in the brain. The treatment also used a psychological method aimed at "the extinction of negative emotional responses to the tinnitus". The method aims to achieve "extension in time of residual inhibition", by "employing imagery of the auditory stimulation" – for example, "participants are instructed to deliberately prolong or deepen the residual inhibition state focusing on the auditory signal and subsequent imagery."

## Are there any long-term benefits from prolonged use?

The short-term characteristics of residual inhibition have now been quite well established, by several independent research efforts. Therapeutically, short-term relief can be experienced for a period after listening to a single pulse of appropriate sound; or, for much longer periods by listening to a repeating sound trigger of a suitable type. For those who are responsive to residual inhibition, these are effective and immediate forms of relief in themselves, bringing a reduction of tinnitus that is available by few other methods. But can any longer-term tinnitus reduction be gained from sustained listening to these sorts of sounds?

The published research contains numerous reports of people who have experienced longer-term improvement, including accounts from some of the most prominent tinnitus researchers. However, in some reports only a small proportion of people under treatment were noted by the authors as seeing longer-term improvement. In some reports the proportions are not stated. As ever in medical research, results need to be considered against natural recovery rates (i.e. without any intervention) and any possible influence of the placebo effect.

Feldmann wrote in his original 1971 paper of patients for whom, during the test procedure their silent periods grew longer and longer after each trigger pulse.[9] In his 1983 paper he described "fatigue of tinnitus" in some individuals, in which again, the duration of residual inhibition increased with each successive trigger pulse, up to some level.[24] However, he also described "fatigue of masking" in other patients, in which ever longer trigger pulses were needed to maintain a residual inhibition response.

In 1978, Vernon and Schleuning reported on 82 patients who had tried ear-level tinnitus maskers (these are masking devices that are worn like a hearing aid).[15] They wrote, "Two patients are of particular interest. Both utilized the tinnitus masker and claim now to be cured. That is, they have returned their maskers after 6½ month's use in 1 case and 9 months in the other. … We do not claim that masking tinnitus is a curative procedure but clearly this is one aspect of masking which will bear additional attention. Both of these patients had had a long standing tinnitus, 10 years in one and 6 in the other. Both displayed residual inhibition which gradually increased as they continued using the masker. It is routine that patients using the masker on a daily basis experience 30 to 40 minutes or more of residual inhibition upon removal of the masker." Vernon and Meikle re-affirmed in 1981 that "patients often experience residual inhibition (temporary suppression of tinnitus upon cessation of masking) which may accumulate with sustained use of masking, in some cases becoming permanent".[18] The same year Vernon wrote in another paper, "Some patients experience residual inhibition which is excessively long after only very brief periods of masking. One of the most extreme cases … had tinnitus for 61 years and after 1 minute of masking in each ear the tinnitus completely disappeared for 8 weeks before it returned. Another patient with a 20-year history of tinnitus can mask for 15 minutes and obtain 12 hours of complete residual inhibition. One female patient … can mask from rising time until 10:00 a.m., after which she experiences complete residual inhibition until 3:00 p.m., after which she returns to masking until bedtime. Incidentally, this patient … only first began to experience it after about 4 months of masking. Another patient has a complete day of residual inhibition after 2 days of masking… And there are others who experienced very long periods of residual inhibition, for as long as several weeks. Some of these have gone on to find complete remission of the tinnitus."[17]

In the same year, Hazell and Wood wrote, "In a proportion of patients who experience residual inhibition, the effect becomes more marked as time goes by. Some patients who initially experience no residual inhibition may achieve this after regular use for, say, 1 month. Conversely, there is a small group who initially experience residual inhibition, but in whom the phenomenon disappears."[20]

The Roberts study in 2008 noted, "Follow-up reports volunteered by three subjects given maskers for home use described residual inhibition lasting longer (up to 20 minutes) after repeated masker use."[60] However, the main longer-term tests of this study (run over 2 to 3 weeks) revealed no significant changes. (Specifically, after listening to their 30-second trigger sounds "several times daily", the participants showed no change in tinnitus loudness, pitch range, or residual inhibition depth or duration. However, a limit in the testing software meant that residual inhibition durations longer than 45 seconds could not be measured.)

The 2003 Canadian Hearing Society study on 31 people, already described above, found a moderate average improvement in nearly all aspects measured, over the 3 weeks of the at-home residual inhibition therapy.[44]

Roberts' 2007 paper discusses a theory that treating patients early in the onset of their tinnitus may be important for long-term benefits to be seen.[56] Roberts suggests that we should assess whether continuous exposure to masker noise, tuned to the tinnitus pitch range, "can induce a lasting residual inhibition, particularly in new tinnitus cases". He cites two earlier studies [51][52] in which cats were exposed to damaging levels of noise. After this, one group of cats was kept in a quiet environment. Another group was kept in an "enriched acoustic environment" in which the cats were continuously played a sound, effectively like masking noise, that matched the pitch range of their expected hearing loss. After a number of weeks, the cats in the quiet environment showed classic symptoms: noise-induced hearing loss, distortion of their brains' physical "pitch map" (a *plastic* change in the auditory cortex), and neural signs of tinnitus (over-active groups of neurons, in the auditory cortex, firing synchronously – called *hypersynchrony*). The cats that had been kept in the enriched acoustic environment, however, showed much less hearing loss, no map distortion and no neural signs of tinnitus. Therefore, there is strong evidence that simply by playing pitch-matched sound to these cats, in the early weeks of recovery, the development of tinnitus was stopped; and, perhaps more remarkably, hearing loss was considerably reduced too. Given that the hearing systems of cats and humans are similar in a number of ways, there may be possibilities for human sound therapies in the early stages after noise damage. Roberts also comments that magnetic stimulation of certain brain regions more effectively triggers a residual-inhibition-like effect in more recent cases of tinnitus (2-4 years) than in longer-standing cases (9-15 years). He comments, "To be effective, therapies … may have to intervene before functional‡ plastic changes lead to structural ones."[56]

## Is it safe?

Based on the publications studied for this review, the answer is "yes". Specifically, there are no persistent adverse affects reported in these publications, from trigger or masker sounds used for residual inhibition. Clearly, as with most treatments or products, there are conditions for safe and comfortable use, such as keeping sound levels at general safe listening levels; but such general safety considerations would be expected with the use of any audio system.

Several studies have reported that a small proportion of people experience an increase in the loudness of their tinnitus immediately after listening to a trigger sound.[15][18][26][28][49][55][60][67] This was termed "residual facilitation" in a paper by Jack Vernon and Mary Meikle,[18] and "residual excitation" in a paper by William Sedley and colleagues.[67] There are no reports that this persisted in the long term, and most publications explicitly note that the effect was temporary.[15][18][26][28][55][67]

The Tinnitus Archive examined residual inhibition in 1,451 people with tinnitus, and reported "tinnitus exacerbation in a few cases".[49] In 1978 Vernon and Schleuning reported on 158 patients of the same clinic, noting, "In one case the tinnitus was temporarily elevated."[15] Later, Vernon and Meikle wrote, "We have also observed occasionally a brief 'residual facilitation' after the cessation of masking."[18] The Roberts 2008 study reports, "5 out of 59 subjects (8.5%) reported small increases in tinnitus after the masking stimuli."[60] In 1984 Tyler and colleagues found that one person in their ten-person study experienced temporary tinnitus exacerbation.[26][28] They wrote that her "tinnitus returned immediately after the masker but was louder than normal."[26] The Tyler team measured the time for her tinnitus to return to its normal level as being between 4 and 12 seconds. A 2007 study using a proprietary phase-shifting waveform reported, "Of those patients who showed intensity increases, only one patient's intensity continued to increase with each treatment. This patient elected to stop treatment after the experimental month. He was evaluated 1 month later, and his intensity had decreased to baseline levels."[55] A study in 2012 investigated temporary tinnitus exacerbation in 4 patients, who represented 13.3% of the 30 patients who had enrolled into the study.[67] The researchers played these patients noise spanning various ranges of pitch. They then measured the average (mean) time for their tinnitus to return to the normal level as being (for each individual) 5.0, 6.9, 20.6 and 36.3 seconds.

The 2003 Canadian Hearing Society (CHS) study, completed by 31 people, reported, "When asked subjectively about their tinnitus, 31% reported it was better, 41% reported it was the same and 19% reported their tinnitus was worse after having undergone residual inhibition therapy for a period of 3 weeks."[44] (Presumably the missing 9% did not answer this question, for some reason.) 19% (reporting "worse") is a notably higher figure than the figures reported in the other studies. However, it could be argued that this is not too surprising, if considered as follows. The CHS study differs from the others in that it did not look at tinnitus loudness *immediately* after listening to a trigger sound, but rather it looked at longer-term underlying change over 3 weeks. It

‡ "Functional plastic changes" refers to changes in brain *activity patterns*, whereas "structural plastic changes" refers to the changing of brain *structures*. In both cases, this is as a result of some other longer-term change (hearing loss, in this case). Structural plastic changes can be seen on many levels, from large-scale "regional" changes (which can even cause measurable change in the size of regions of the brain), down to changes in the detailed ways that individual neurons interconnect and influence each other.

seems reasonable to suppose that if there was any actual longer-term tinnitus change over the 3 weeks, it would (for most individuals) be smaller than the change observed *immediately* after a trigger sound. Therefore, it could be argued that its measurement would be more affected by other random factors. (Asking people to report their subjective assessment of any subtle change would be expected to yield less-accurate answers than asking the same for a large change.) For example, it is known that for many people their tinnitus is noticeably variable (often from day-to-day), so this is one possible random factor. Also, a question like "Is your tinnitus better, the same, or worse?" requires a subjective judgment, which must be prone to some amount of judgment error, and even some memory error over a period of 3 weeks – another possible random factor. Considering these two random factors together, if you simply asked people with long-term tinnitus (who are *not* under some treatment) how their tinnitus had changed over 3 weeks, you might expect some roughly symmetrical result like 15% / 70% / 15% (better / same / worse). (The specific numbers here are just examples, to illustrate the point.) Looking instead at the actual CHS study, we might then expect these symmetrical figures to be swayed by any genuine effects of the treatment and by any placebo effect (possibly significant, as the study did not use a placebo control). So, by this argument, the figure reported for "worse" could be significantly influenced by random factors (in addition to any influence from genuine effects or placebo effects). This might explain, at least in part, why the figure for "worse" is higher in the CHS study than in the other studies.

## What types of sound trigger the deepest and longest residual inhibition?

Over the years it has been discovered that various characteristics of the trigger sound have a significant effect on the depth and duration of residual inhibition. This needs to be borne in mind when comparing the results of the various studies, as differences in results may simply be due to differences in the characteristics of the trigger sounds used.

Roberts and colleagues showed that both the depth and duration of residual inhibition are increased when the trigger sound has the same pitch (frequency) range as the patient's tinnitus.[53][56][60] The Sockalingam study independently agrees with this, although the earlier Roberts work is not referenced.[57] A 2012 study, led by William Sedley of Newcastle University, reported comparable results.[67] A 1987 study on 117 people found that residual inhibition was longer if the tinnitus pitch range was present in the trigger sound.[31] A 1979 paper reported on one patient whose residual inhibition duration was more than quadrupled when a 30-second pure tone was played with a pitch close to that of his tinnitus, compared to pure tones at other pitches but with the same duration and *sensation level* (explained later).[16] A 1981 paper by Vernon and Meikle noted briefly, "The results from the Tinnitus Clinic suggest that the most effective way to produce residual inhibition is to provide masking that closely resembles the tinnitus."[18] The Terry study found that the effect of pitch was "complex", and that the most effective trigger was usually lower in pitch than the tinnitus.[23] Roberts suggests that this may simply be because hearing damage is generally less at lower frequencies, or because of uncertainties in finding the pitch of the tinnitus.[56] The 1984 Tyler study reported that there were no "obvious" differences for their 10 patients when the pitch of a pure tone trigger sound was varied.[28] However, it is likely that the method used by this study significantly contributed to the lack of an observed effect: whereas the Roberts study maintained the same perceived loudness at all the pitches tested, the Tyler study did not.[§] The original Feldmann paper comments briefly that the pitch of the trigger sound affects residual inhibition "to a less degree" than its sound level or duration,[9] although no details are given.

[§] The Tyler study seems to have maintained the same *sound pressure level*, for a particular patient, as the pitch was varied. This does not take account of differences in perceived loudness as pitch is varied. Even for a person with healthy hearing, at a fixed sound pressure level there is considerable variation in perceived loudness as pitch is varied. Among people with hearing loss, which includes most people with tinnitus, these differences vary greatly.

The sound level (volume) of the trigger has long been recognized as being important for effective residual inhibition, certainly since the measurements of Feldmann in 1971.[9] There is good agreement between several studies that, for those who exhibit residual inhibition, it can be reliably produced when the trigger sound level is moderately (around 10 decibels) above the level that just masks the tinnitus (the "minimum masking level").[17][20][23][24][28] Additionally, the Terry study investigated louder trigger sounds (at 20 decibels above the minimum masking level), and found that this more than trebled the duration of the residual inhibition.[23] The study looked at quieter trigger sounds too (at exactly the minimum masking level), but this yielded no residual inhibition response in their patients. The Tyler study also examined the importance of sound level.[28] They looked at a group of 10 patients, most of whom had little residual inhibition response. In those patients who had more response, increased sound level did increase the duration of their residual inhibition, but not to the marked extent seen in the Terry study. (One reason for the difference may be the way that the patients were selected for these studies. The Tyler study used a random sample of 10 people with tinnitus. However, the Terry study took 32 people with tinnitus and selected from them, for these tests, the 10 who had the strongest residual inhibition.) The Roberts study investigated trigger sound level relative to the estimated loudness of the tinnitus, rather than relative to the minimum masking level.[60] The team looked at a group of patients who had showed a "good" residual inhibition response in their study. For these patients, the trigger sound level had been 13.3 decibels (on average) louder than their tinnitus (for trigger frequencies above 4,000 hertz).

The duration of the trigger pulse is also important. The Terry study found that although longer trigger pulses do produce longer residual inhibition, the two are *not* proportionally in scale.[23] For example, if you double the length of the trigger, the residual inhibition duration is generally *less* than doubled. (In mathematical terms, the duration of the residual inhibition is linearly related (approximately) to the logarithm of the trigger length.) The Tyler study found a roughly similar trend among those patients with a residual inhibition response.[28] While these findings may seem like "bad news" for the therapeutic use of residual inhibition, they are potential "good news" for suppressing tinnitus with repeated trigger pulses: if you reduce the trigger pulse length, it may have *less* effect on residual inhibition duration than you might suppose. Although the Terry team found that their 10 study subjects did not experience much residual inhibition for trigger pulses of less than 10 seconds, in other studies individuals have achieved residual inhibition with much shorter pulses. For example, Feldmann induced residual inhibition with trigger pulses of between just 100 and 500 milliseconds,[9][24] and the Tyler study reported on residual inhibition after trigger sounds with varying duration as short as 1 second.[26][28] The findings of Roberts and others suggest that the trigger sound in the Terry study (a fairly broadband noise, not matched to the patients' tinnitus pitch) was not optimal. This may have reduced its effectiveness, perhaps quashing any residual inhibition response for pulses below 10 seconds.

Some studies have looked at which ear (left or right) is most responsive to trigger sounds. For people who have tinnitus in just one ear, the general agreement is that the trigger sound should be played into that same ear for the most effective residual inhibition.[15][17][23] The early work of Feldmann shows an opposite result for one patient.[9] This patient experienced longer residual inhibition, and at lower sound pressure levels, from trigger sounds played to his non-tinnitus ear. However, this might be explained by the much higher level of hearing loss in his tinnitus ear. Masking has been found to be most effective if different noise signals are simultaneously played into each ear.[35]

Regarding the best waveform type for causing residual inhibition, two types have seen the most study: pure tones and noise. The noise used in these studies spanned various ranges of pitch (i.e. various *bandwidths*). Several papers report on comparisons between the types, with some comparing warbling-pitch tones too.[18][23][40][44][57][60] However, there is disagreement among these papers on the most effective type:

- The Roberts study, on 90 people, reports that pure tones and very narrowband noise (described as "ringing") were much less effective than wider-band noise (described as "hissing").[60] All 90 participants were tested with the "Tinnitus Tester" program, in which a test trigger sound was matched to both the pitch and bandwidth (pure tone, "ringing" or "hissing") of their tinnitus. The results seem somewhat surprising, as the thrust of the paper (the theory and all the other results) points to the most effective residual inhibition being triggered by the closest match to the person's own tinnitus sound. Therefore, someone with pure tone tinnitus, for example, would be expected to respond best to a correctly pitch-matched pure tone trigger sound, since all of the available sound power would be concentrated at the most effective pitch (according to the theory) for the deepest and longest residual inhibition. However, this was not seen in practice. (We could speculate, though, that more pitch mismatch error might occur when trying to match a narrower bandwidth of tinnitus with a narrower bandwidth of trigger sound. This could be a factor influencing these results, perhaps hiding the true effectiveness of an accurately pitch-matched arrangement.)

- Two studies, on a total of 41 people, found the opposite, namely that pure tones were more effective at inducing residual inhibition than narrowband noise.[44][57] Again, the sound was tuned to each patient's own tinnitus pitch. In addition, these studies looked at warbling-pitch trigger sounds. The Banks team found warbling to be the overall most effective trigger type. The Sockalingam team also found that warbling produced residual inhibition in more people than the other types, but that pure tones gave the longest average residual inhibition.
- Likewise, the Olsen study found that "a pure tone was more effective in inhibiting tinnitus than a narrow-band noise signal with the same loudness", in both cases tuned to the patient's tinnitus pitch.[40] However, this was observed only during an "informal study" that acted as a precursor to their main study.
- Vernon and Meikle reported a similar observation in 1981.[18] In a paper focused on tinnitus masking therapy, they noted briefly that "pure-tone masking, which is rarely acceptable to the patient, usually induces more residual inhibition than does masking with a narrow band of noise centered at the same frequency."
- The Terry team, studying 18 people, found an even split between the effectiveness of pure tones (giving most residual inhibition to 9), and noise (of two bandwidths, giving most residual inhibition to the other 9).[23] Again, a tuning process was used for each patient, to find the most effective pitch for all these trigger sounds.

Why is there not more agreement between these studies on the most effective sound type? Perhaps the reason might lie in differences in the methods of the experiments. For example, there are differences in the noise bandwidths used (and in some papers the bandwidths are not stated), sound levels, tinnitus pitch tuning methods, and various other experimental details.

Aside from pure tones, warble/sweep tones and noise, several other "more specific" waveform types have been studied over the years, to explore their residual inhibition properties. This includes ultrasonic sound, at various frequencies, and some very specific types of audible waveform. However, they tend to sit somewhat aside the mainstream research into residual inhibition, and to date none has conclusively been proven to be any more effective, so they are not covered by this review.

A quite different type of trigger sound is discussed in a paper by Henry and colleagues.[47] The Henry team studied a number of different types of continuous masking sound, to find which gave the most relief from tinnitus. They found that the most relief was gained from types of sound in which the sound level briefly, but regularly, deviated above the norm (by 5 to 15 decibels) in the mid to upper frequency region. The paper discusses that these deviations, which were just tenths or even hundredths of a second long, are much like Feldmann's short pulses for inducing residual inhibition. However, whereas Feldmann's pulses were abrupt sounding, because they were simply hard on-off bursts of noise, the sound used in the study was a more subtle mix of a continuous background noise with moderate additional deviations. The paper suggests that, due to the effective tinnitus suppression (and low annoyance) properties of such sounds, they could be the “background sounds of choice for promoting tinnitus habituation in Tinnitus Retraining Therapy patients”.

Sound is not the only way to trigger residual inhibition, however. People with cochlear implants, who also have tinnitus, usually find that their tinnitus is suppressed simply by switching on the implant (which then stimulates nerves in the inner ear).[50] When the implant is switched off, many people find that they then experience residual inhibition. Similarly, other methods of electrical stimulation can also trigger residual inhibition. One method involves passing a needle electrode through the eardrum, under local anesthetic, and placing it onto the promontory of the cochlea (this is termed electrical promontory stimulation, EPS). Residual inhibition from this method was first noted in 1993.[37] Another method is to apply electrical stimulation to the skin in the region of the ear (transcutaneous electrical stimulation). It was first reported in the 1980s that tinnitus could remain suppressed for a time after this type of stimulation.[19][29][33] An effect much like residual inhibition can also be triggered by magnetic stimulation of regions of the brain involved with tinnitus (transcranial magnetic stimulation, TMS),[56] or by electrical stimulation through scalp electrodes (transcranial direct current stimulation, tDCS),[68] or by surgically implanting electrodes into the brain,[65] placed either on the auditory cortex (auditory cortex stimulation) or deeper in the brain (deep brain stimulation, DBS). This review, however, focuses on sound as the trigger for residual inhibition.

## How does it work?

The biological mechanism that produces residual inhibition is not yet fully understood. However, four separate brain-imaging studies have specifically investigated residual inhibition.[38][50][59][67] These contributed new evidence, and ensuing theories. Additionally, several other arguments have been put forward since the 1970s. Some of the recent papers [59][60][67] share a number of common ideas, built upon prior theories for the mechanism of tinnitus. The theory in these papers is based on the idea that tinnitus is due to the abnormal, spontaneous and synchronized over-activity of groups of neurons in the brain. In turn, this is caused by a lack of normal nerve signal input from the ear in the pitch region of the patient’s hearing loss. During a trigger sound, the sound is sufficiently loud in the pitch region of hearing loss that normal nerve signal input is temporarily restarted. This reduces or stops the abnormal neural over-activity in the brain, and hence stops the perception of tinnitus. When the trigger sound stops, the normal brain activity patterns persist for a while (for reasons speculated, but not fully known), keeping the tinnitus suppressed – the effect we call residual inhibition. This section necessarily contains rather more medical and technical terms.

In 1995, Rumyana Kristeva-Feige and colleagues, at the Universities of Freiburg and Münster in Germany, used magnetic brain imaging (MEG) in a series of tests on a person with tinnitus of both ears, and four people who had normal hearing as control subjects.[38] MEG produces a view of electric currents in the brain, so it can be used to observe neural activity. The tests were carried out over five different days, and consistent results were repeatedly seen. These results revealed that the brain of the person with tinnitus produced unusual low-frequency activity during residual inhibition (which was triggered after playing a period of masker sound into the right ear). This activity was not seen when his tinnitus was present, nor was it seen at any time in the four people who had normal hearing. The unusual low-frequency activity was seen over quite a wide area (it was detected by all 37 sensors of the MEG system, which were placed over the left temporal plane). This suggested that neurons in a relatively large region (of the cerebral cortex / thalamus) may be producing this activity “in concert” (neural synchronization). The research team concluded that this unusual activity indicated that residual inhibition is unlikely to simply be the normal tinnitus mechanism temporarily ceasing, but rather that it somehow involves an “extraordinary and unusual process” – at least, in the one individual tested.

Ten years later, a large team from Japan, led by Yasuhiro Osaki from Osaka University Graduate School of Medicine, used another brain imaging method, positron emission tomography (PET), to study the mechanism of residual inhibition.[50] They studied three patients, all with cochlear implants in their left ears, and all having tinnitus of both ears. Six people with normal hearing were used as control subjects. Before the measurements, all of the subjects were injected with a radioactive tracer to enable the PET scan to show regional blood flow within the brain, another indicator of neural activity.[39] Unlike the MEG system that was used above, this PET scanner was able to show the whole brain, in 3D, and in much finer detail. The study found that a region in the right temporal cortex was distinctly more active during residual inhibition than when the patients could hear their tinnitus. A region in the right cerebellum was distinctly more active when their tinnitus returned. The normal-hearing group showed no such differences. Regarding theories behind the mechanism of residual inhibition and tinnitus, the paper advances the following thoughts:

- The temporal cortex activity seen during residual inhibition includes activity in a region related to the processing of tonal sounds and voices. This also has links to memory retrieval (for the recognition of these sounds). Therefore, "…our study would imply that tinnitus and its suppression are related to some memory functions".
- In this particular study, right-sided areas of the brain saw a greater difference in activity than left-sided areas. This should be expected, since all the patients were given masking sounds via their left ears only (through their cochlear implants), and it is known that the left ear routes primarily to the right half of the brain (and vice versa). However, all the patients reported that, during residual inhibition, their tinnitus was suppressed on *both* sides.
- Pitch information has previously been found to be predominantly processed on the right side. Further, the paper references earlier reports of right-sided activity being seen in studies on tinnitus and unpleasant sounds.
- A region of the cerebellum was more active when the tinnitus returned (after the period of residual inhibition). The paper links the cerebellum to earlier reports on tinnitus and unpleasant sounds, tinnitus annoyance, and the storage of the pitch and duration of sounds. Therefore, "Activation of the cerebellum during tinnitus perception might be related to attention and some memory functions."

The Osaki paper does not reference the earlier Kristeva-Feige study. Some comparative notes on the two papers are made here:

- In the earlier study, the residual inhibition trigger was applied to the right ear, and in the later study to the left ear. This is an important difference in terms of the sides of the brain that might be expected to see the greatest change in activity. If we account for this left/right difference, then there are a number of similarities in the observations made in the two studies.
- In both studies, during residual inhibition extra activity (compared to tinnitus periods) was seen in the temporal cortex area opposite to the ear that was stimulated.
- In both studies, that extra activity ceased when residual inhibition stopped.
- In both studies, the activity was seen over an extended region (specifically, the adjacent Brodmann areas 21 and 38 in the PET study) – although the implied area was perhaps larger in the one individual used in the MEG study.
- The MEG study reported that the extra activity seen in the patient was not seen in the control group with normal hearing. However, from the published results of the PET study, it may not be possible to compare regional brain activity between the patient and control groups. (The headline results of the PET study report *changes* in regional blood flow within the *same* group. The two groups may have had different baseline measured blood flow rates. Therefore, it may not be valid to infer differences in neural activity levels from the reported differences in blood flow rates between the two groups. Also, the regional blood flow rates shown in the paper may have been separately processed (normalized, or otherwise offset/scaled) for the patient and control groups – the text does not seem definitively clear on this.)

- The PET study was the first to see any brain activity "switch off" during residual inhibition, in part of the cerebellum. However, this is not inconsistent with the MEG study, as no MEG sensors were placed over the cerebellum area, so no observation was made there.
- A clear difference is that the second study used patients with cochlear implants. The stimulus, rather than just being a burst of white noise (as it was for the control patients), also involved switching the cochlear implant on at the start of the noise burst, and off at the end of the noise burst.
- An obvious and fundamental difference in the methods of the two studies is that one (MEG) measured electrical neural activity, whereas the other (PET) measured regional blood flow. However, both measures do relate to neural activity.[39]

In 2008, Nina Kahlbrock from Konstanz University and Nathan Weisz from INSERM completed a whole-brain MEG imaging study on 8 people who showed residual inhibition.[59] Their main finding was that brain activity in the delta frequency band (i.e. 1.5 to 4 hertz) was decreased during residual inhibition, in temporal regions. They put forward a theory that residual inhibition is caused by a temporary return of normal activity in the brain's auditory system. (The following quotation is loosely "translated" into more everyday English, with less medical terminology, in a footnote below.) "The current results suggest that changes of tinnitus intensity induced by residual inhibition are mediated by alterations in the pathological patterns of spontaneous brain activity, specifically a reduction of delta activity. Delta activity is a characteristic oscillatory activity generated by deafferented/deprived neuronal networks. This implies that residual inhibition effects might reflect the transient reestablishment of balance between excitatory and inhibitory neuronal assemblies, via reafferentation, that have been perturbed (in most tinnitus individuals) by hearing damage."** This is somewhat in contrast with the earlier single-patient MEG study of Kristeva-Feige and colleagues, in which *increased* activity was seen in the 2 to 8 hertz frequency band; and those authors put forward that residual inhibition involves some unusual extra brain activity, not seen in their control subjects. The earlier MEG study is cited by Kahlbrock and Weisz, but the differences are not discussed. The authors also cite the earlier PET study by Osaki and colleagues, and suggest that perhaps the relative increase in temporal cortex activity seen during residual inhibition was simply a return to normal healthy levels of activity (i.e. that these brain regions are abnormally under-active when tinnitus is present). "Both brain regions [Brodmann areas 21 and 38] are putatively involved in auditory processing and the result could be interpreted as a transient decrease of hypoactivation or reafferentation."

In 2012 a research group from University College London and Newcastle University, led by William Sedley, published another whole-brain MEG study on 17 patients with chronic tinnitus.[67] Of these 17 patients, 14 exhibited residual inhibition, and 4 exhibited the reverse response, i.e. temporary tinnitus exacerbation. The paper terms this "residual excitation". (One patient experienced both residual inhibition and excitation, in reaction to different sounds.) The study reported that, "Our striking and consistent finding was that in residual inhibition, auditory cortex gamma power positively correlates with tinnitus intensity, and in residual excitation it shows the opposite correlation." In other words, activity in the auditory cortex, in the gamma frequency band (which covers 30 to 150 hertz), decreased during *both* residual inhibition and residual excitation. The team also looked at lower frequency activity in the auditory cortex (in the delta band, from 1.5 to 4 hertz; and the theta band, from 4 to 8 hertz). Again, this activity consistently decreased during residual inhibition. This aspect of their results agrees totally with the Kahlbrock and Weisz MEG study, but not with the original Kristeva-Feige MEG study. During residual excitation,

** "Our results suggest that changes in tinnitus intensity during residual inhibition are due to changes in abnormal patterns of unprompted brain activity; specifically, a reduction of activity in the delta frequency band. Activity in this band is typical of the repetitive signaling generated by "input-starved" ("deafferented") neuronal networks. This implies that residual inhibition effects might reflect a temporary return to the normal balance of activity between groups of neurons that increase the firing-rate in others, and groups of neurons that do the opposite. This is caused by the temporary restoration of neuronal input ("reafferentation") from the masker sound; input that is normally missing, due to the hearing damage present in most people with tinnitus."

however, no significant change in auditory cortex lower frequency activity was seen. The study also looked at brain activity in areas other than the auditory cortex, and found that individual patients did indeed have regions in which activity consistently changed as their tinnitus level changed (i.e. during residual inhibition or excitation). However, these regions were not consistently located from patient to patient, even allowing for many inter-patient differences, such as left- versus right-sided tinnitus. Regarding a theory for residual inhibition, the authors propose that, "Residual inhibition is achieved by a transient and partial normalization of the deafferentation of auditory thalamus that leads to the generation of tinnitus." (This agrees with the theory of Kahlbrock and Weisz, with the added detail of the normalization occurring in the auditory thalamus – this following from an earlier "thalamocortical dysrhythmia model" cited by the Sedley team.) The authors also present a theory to explain why the auditory cortex gamma activity decreases during both residual inhibition and excitation. They propose that auditory cortex gamma oscillations act to suppress the perception of tinnitus (whereas previous research had proposed that they did the opposite, i.e. that they were a driving force behind tinnitus perception). The authors present quite a detailed speculative model, based on neural mutual inhibition, to explain the mechanisms of both residual inhibition and excitation, and the role of auditory cortex gamma activity. They highlight the possibility of drug treatments that might increase this gamma activity, and hence reduce or stop tinnitus.

Aside from the brain-imaging studies, the following theories regarding residual inhibition have been put forward over the years:

- In 1971 Harald Feldmann proposed a theory that tinnitus masking and residual inhibition occurred in the brain rather than the ear. Specifically, he suggested that this may be due to neural inhibition.[9] He based his argument on evidence from the unusual frequency characteristics of tinnitus masking, the ability to mask tinnitus in one ear with a relatively low-level masking sound at the other, and the existence of what we now call residual inhibition.
- Mark Terry and colleagues from UWIST in Wales proposed a theory in 1983 for residual inhibition based on an unusual form of temporary threshold shift (TTS) that they had discovered during residual inhibition.[23] TTS is a temporary hearing loss, which normally occurs after exposure to loud sound. However, the UWIST team noticed that a small TTS occurred during residual inhibition, in the frequency region of the tinnitus. They suggested that residual inhibition may occur because "the tinnitus 'signal' drops below the temporarily raised threshold."
- In 1987, Juergen Tonndorf from Columbia University put forward a theory for residual inhibition based on the neuroscience of pain.[32] It was based on a particular hypothesis in which the mechanism of subjective tinnitus (for many patients) was assumed to be in the cochlea, in the inner ear. (As an aside, the location of the mechanism of subjective tinnitus has been a subject of intense debate among researchers. However, since Tonndorf wrote his paper, evidence and argument has grown for the mechanism (typically) being based in the brain, even if the tinnitus started as a result of cochlear damage.) Tonndorf outlines the situation for most cochlear damage associated with tinnitus: the stronger inner hair cells (with large diameter nerve fibers) remain relatively undamaged, while the delicate outer hair cells (with small diameter nerve fibers) are destroyed. This leaves the small diameter nerve fibers input-starved ("deafferented"), which is known to increase spontaneous activity in the nerve. This extra activity is the supposed source of the ongoing tinnitus. Making an analogy with a neural theory regarding chronic intractable pain, Tonndorf then suggests that, "Acoustic masking with its relatively short 'residual inhibition' (typically measuring in minutes) might mechanically re-activate the large diameter, inner-hair-cell fibers in largely the same manner as the large-diameter pain fibers are temporarily re-activated by scratching or by vibratory stimulation." Then, by a "gate control theory" for pain (first proposed in 1965), the activity of the large diameter, inner-hair-cell fibers acts to shut-off (for a time) the aberrant signaling from the small diameter, outer-hair-cell fibers. Thus perception of the tinnitus is stopped for a period of time. (Similarly, but more briefly, in 1981 Jack Vernon and Mary

Meikle speculated that the mechanism of residual inhibition may be related to mechanisms that suppress pain for a period of time after electrical stimulation.[18])

- Pawel Jastreboff and Jonathan Hazell, the pioneers of tinnitus retraining therapy, suggested in a 2004 book on the subject that residual inhibition can "be easily explained by the rebound phenomenon. The rebound phenomenon is well recognized in neurophysiology. If the activity of a neuron, as the result of sound stimulation, is increased, cessation of the signal frequently results in activity decreasing below the previous level of spontaneous activity occurring before stimulation. If stimulation was causing inhibition of neuronal activity, then switching off the sound results in an enhancement of spontaneous activity for some time. After a while, the neuronal activity returns to the pre-stimulus level."[46]
- In a number of papers between 2004 and 2008, Larry Roberts and colleagues published observations and arguments supporting a particular theory for the mechanism of tinnitus, along with ideas for the mechanism underlying residual inhibition.[48][53][56][60] They observed that the frequency region of hearing loss corresponded to the range of pitch of people's tinnitus, and also to the frequency region that stimulates the deepest and longest residual inhibition. Based on this, and previous research, they argued that the primary auditory cortex (the seat of auditory perception) is starved of nerve signal inputs in the frequency region of hearing-loss, and groups of neurons associated with that region then start to spontaneously "self-fire" in synchronous fashion (hypersynchrony). The team suggested a number of more detailed mechanisms by which that might occur, one of which is the breakdown of a system of feed-forward neural inhibition that normally keeps neurons switched off in frequency regions corresponding to silence.[60] If this mechanism fails, the perception of silence could be broken at those frequencies, causing tinnitus. When a masking sound is played at these frequencies, this may inject new states of excitation or inhibition into the region, thus disrupting (or at least segregating) the over-active synchronous neurons responsible for tinnitus. Then, "residual inhibition could reflect a temporary adaptation of neurons involved in synchronous activity," or a "rebalancing" of the inputs to those neurons, or "other mechanisms that subside" over the same durations as residual inhibition.

# An outline of HushTinnitus, a new sound therapy system for tinnitus relief

(This section assumes basic familiarity with residual inhibition. Please see the previous section, starting on page 4, for an introduction.)

## 1 - Introduction

The HushTinnitus system ([www.hushtinnitus.com/ri](www.hushtinnitus.com/ri)) is a set of sound tracks that uses residual inhibition to provide relief from tinnitus. The sound tracks are custom-generated for each user, with the aim of maximizing the effectiveness (i.e. depth and duration) of their residual inhibition. The system primarily focuses on playing repeating trigger sounds, of various available types. For those users who have the necessary level of residual inhibition response, this allows them to stay in a state of residual inhibition, i.e. to keep their tinnitus suppressed. This can be maintained for as long as the track is played – all day, if desired. (The tracks are designed to be played in “single track repeat” mode on the user’s player, so the playing time is effectively unlimited.) For those who have less residual inhibition response, or even none at all, custom-optimized continuous masking sounds are provided.

Most of the design effort of the project was spent on minimizing the obtrusiveness of the sounds. A number of the methods and optimizations used are new to tinnitus sound therapy, and they are outlined below.

The main aims of the design were to:

- Maximize the effectiveness of the sounds (i.e. the depth and duration of residual inhibition)
- Minimize the obtrusiveness of the sounds (with the particular aim of being less obtrusive than other alternatives, such as broadband noise masking)
- Cater for variability (both person-to-person variability, and variation over time within an individual)
- Be easy to use (both during set-up, and in use)

In terms of identifying the most effective trigger sounds, the work rests entirely on prior independent published research (as reviewed on page 12). In terms of minimizing obtrusiveness, the design uses a number of techniques built on various established perceptual (psychoacoustic) models, and on signal analysis. These techniques include a minimally obtrusive method for fading sound in and out (optimal fading), statistically and perceptually smoother forms of noise, and psychoacoustically optimized warble tones. Variability is addressed both by the custom-generation of sound tracks, based on simple hearing measurements, and also by the provision of a range of sound track types. Ease-of-use features are discussed as they arise in the document.

The following sections outline how the design addresses the design aims.

## 2 - Catering for variability

Since early studies, the person-to-person variability of tinnitus and residual inhibition has been widely commented on.[9][28][45][49][67] One area of variability is the duration of residual inhibition after the trigger sound has ceased (p. 7), if indeed it occurs at all (p. 6). In response to this, the HushTinnitus system provides three categories of basic sound “structure”, called Type 1, 2 and 3 (detailed below). For each user, the system custom-generates 67 sound tracks. Sixty-three of these are of Type 1, one is Type 2, and three are Type 3. The types are summarized as follows:

- **Type 1.** Each track of this type contains an optimally-faded pulse of trigger sound, followed by a longer period of silence. (As already noted, all the tracks should be played in single-track-repeat mode, so this simply keeps repeating: pulse, silence, pulse, silence…)
- **Type 2.** This track comprises 15 minutes of continuous sound. (Again, it is designed to keep repeating.) It contains many ultra-short trigger sounds per second. These are subliminally mixed into a continuous noise signal. (This is discussed in detail later.)
- **Type 3.** Each track of this type contains 15 minutes of continuous sound (again, designed to repeat continuously). These are continuous masker sounds, of various types, discussed in more detail below.

As noted, of the 67 sound tracks provided to the user, 63 are of Type 1. The reason for having so many Type 1 tracks is to cater for variability: both from person-to-person, and variation over time within an individual. The Type 1 tracks include different combinations of trigger-sound duration and silent-period duration. They also include different trigger-sound waveform types. These are termed Type 1a, 1b and 1c; and are discussed further later. (The same waveform types are also used in the continuous Type 3 tracks, and are termed Type 3a, 3b and 3c.) All 67 tracks in the system are listed below.

| Album name | Track name | Waveform (simple description) |
|---|---|---|
| Type 1 with 2-second quieting sounds | 01 Type 1a then 4 seconds silence | Noise |
| | 02 Type 1a then 8 seconds silence | |
| | 03 Type 1a then 15 seconds silence | |
| | 04 Type 1a then 30 seconds silence | |
| | 05 Type 1a then 60 seconds silence | |
| | 06 Type 1a then 120 seconds silence | |
| | 07 Type 1a then 240 seconds silence | |
| | 08 Type 1b then 4 seconds silence | Pure tone |
| | 09 Type 1b then 8 seconds silence | |
| | 10 Type 1b then 15 seconds silence | |
| | 11 Type 1b then 30 seconds silence | |
| | 12 Type 1b then 60 seconds silence | |
| | 13 Type 1b then 120 seconds silence | |
| | 14 Type 1b then 240 seconds silence | |
| | 15 Type 1c then 4 seconds silence | Warble tone |
| | 16 Type 1c then 8 seconds silence | |
| | 17 Type 1c then 15 seconds silence | |
| | 18 Type 1c then 30 seconds silence | |
| | 19 Type 1c then 60 seconds silence | |
| | 20 Type 1c then 120 seconds silence | |
| | 21 Type 1c then 240 seconds silence | |
| Type 1 with 4-second quieting sounds | 22 Type 1a then 4 seconds silence | Noise |
| | 23 Type 1a then 8 seconds silence | |
| | 24 Type 1a then 15 seconds silence | |
| | 25 Type 1a then 30 seconds silence | |
| | 26 Type 1a then 60 seconds silence | |
| | 27 Type 1a then 120 seconds silence | |
| | 28 Type 1a then 240 seconds silence | |
| | 29 Type 1b then 4 seconds silence | Pure tone |
| | 30 Type 1b then 8 seconds silence | |
| | 31 Type 1b then 15 seconds silence | |
| | 32 Type 1b then 30 seconds silence | |

| | | |
|---|---|---|
| | 33 Type 1b then 60 seconds silence | |
| | 34 Type 1b then 120 seconds silence | |
| | 35 Type 1b then 240 seconds silence | |
| | 36 Type 1c then 4 seconds silence | Warble tone |
| | 37 Type 1c then 8 seconds silence | |
| | 38 Type 1c then 15 seconds silence | |
| | 39 Type 1c then 30 seconds silence | |
| | 40 Type 1c then 60 seconds silence | |
| | 41 Type 1c then 120 seconds silence | |
| | 42 Type 1c then 240 seconds silence | |
| Type 1 with 8-second quieting sounds | 43 Type 1a then 4 seconds silence | Noise |
| | 44 Type 1a then 8 seconds silence | |
| | 45 Type 1a then 15 seconds silence | |
| | 46 Type 1a then 30 seconds silence | |
| | 47 Type 1a then 60 seconds silence | |
| | 48 Type 1a then 120 seconds silence | |
| | 49 Type 1a then 240 seconds silence | |
| | 50 Type 1b then 4 seconds silence | Pure tone |
| | 51 Type 1b then 8 seconds silence | |
| | 52 Type 1b then 15 seconds silence | |
| | 53 Type 1b then 30 seconds silence | |
| | 54 Type 1b then 60 seconds silence | |
| | 55 Type 1b then 120 seconds silence | |
| | 56 Type 1b then 240 seconds silence | |
| | 57 Type 1c then 4 seconds silence | Warble tone |
| | 58 Type 1c then 8 seconds silence | |
| | 59 Type 1c then 15 seconds silence | |
| | 60 Type 1c then 30 seconds silence | |
| | 61 Type 1c then 60 seconds silence | |
| | 62 Type 1c then 120 seconds silence | |
| | 63 Type 1c then 240 seconds silence | |
| Type 2 | 64 Type 2 | Rapid subliminal triggers (continuous) |
| Type 3 | 65 Type 3a | Noise (continuous) |
| | 66 Type 3b | Pure tone (continuous) |
| | 67 Type 3c | Warble tone (continuous) |

The overall idea is that the Type 1 tracks are generally found to be less obtrusive than the Type 2 track. Likewise, the Type 2 track is generally found to be less obtrusive than the Type 3 tracks. (In fact, this extends to the waveform types too, so in increasing order of likely obtrusiveness we have: 1a, 1b, 1c, 2, 3a, 3b, 3c.[††]) Of course, we want the user to find the least obtrusive sound that works for them. Therefore, a simple "fitting" procedure is given in the supplied User Guide, for the user to follow when they first get their HushTinnitus tracks. It works very simply like this:

- The user tries Type 1a. If they find it works (i.e. if it effectively suppresses their tinnitus), they make a note of it. It is likely that they will find Type 1a to be the least obtrusive type; but of course preferences vary from person to person, so even if it suppresses their tinnitus well, they carry on to try the other types.

[††] This is just intended as an approximate ranking of obtrusiveness. However, there is some research basis for it. Waveform types a (noise), b (pure tone) and c (warble tone) were found least annoying, in that order, in a 1986 study into preference for different types of masker sound.[30] Type 1 is placed ahead of Type 2 on the grounds that it is generally silent for most of the time – although preference may vary. Type 2 is placed ahead of Type 3 on the basis that it contains recurring trigger sound features of a type found in the least annoying sounds identified by a 2004 study [47] (discussed further later).

- Then the same is done for (in order): 1b, 1c, 2, 3a, 3b, 3c.
- At the end, the user has a list of types that work for them, in order of likely obtrusiveness. The user can review this, and make a final decision on which type they would prefer to use.

> *As an example* of general usage of the system, I will describe my own experience with it, in the early weeks of my tinnitus (when it was present every day, and much more intensely than ever happens now). I found that all of the track types in the system could keep my tinnitus silenced. Therefore, I mainly used Type 1a, on account of its low obtrusiveness. From the selection of Type 1a tracks, I generally found that 2-second trigger sounds worked well for me, but occasionally I used 4-second trigger sounds. In terms of the duration of the silent period (between the triggers), I found that the longest period I could use (while still keeping my tinnitus silent) varied between 30 seconds and 4 minutes. (As the weeks progressed, this more commonly varied between 2 and 4 minutes.) Most of the variation happened overnight, with slower variation (if any) through the day. I accommodated for any slow daytime variation by occasionally increasing or decreasing the duration of the silent period. (In practice, this simply involved pressing the "next track" button or "previous track" button on my MP3 player.) In this way, I kept my tinnitus silent throughout the day. After getting accustomed to the system, I found that I was generally completely unaware of the occasional soft, gently-faded trigger sounds in the Type 1a tracks. Using miniature open-ear headphones, I was no more aware of the system than wearing a pair of glasses. Through the day, I preferred the total silence in the Type 1 sounds to the continuous Type 2 or Type 3a sounds. However, to fall asleep, I sometimes used Type 2. I found the constant low-level sound helpful for sleep. For this purpose, I preferred Type 2 over Type 3a, as I could play it at a lower volume level, while still keeping my tinnitus silenced.

To summarize, the general idea behind providing so many sound tracks is to offer a range of sound time structures, waveform types and other controllable features to address variability in residual inhibition duration and preferred stimulus type (preferred in terms of both its effectiveness and its obtrusiveness). In addition, by simply ordering the sound types in an approximate order of likely obtrusiveness, some guidance is given during the user fitting procedure to help identify the likely least obtrusive fit.

Although the focus of the system is on sound tracks that repeat continuously "in the background", it should be noted that the Type 3b and Type 3c tracks are not really intended to be used this way. These tracks have been rated as the most obtrusive: Type 3b is a continuous pure tone, and Type 3c is a continuous warble tone. It is unlikely that people would want to listen to these for very long periods. Rather, they are intended to be used in a "single shot" manner, for those who can get useful longer residual inhibition durations from this kind of stimulation. These tracks are intended to be played in a single burst, say for a minute or a few minutes, at a slightly louder volume, in order to trigger a longer single period of residual inhibition. (This might be useful before going to bed, for instance.) Type 3a can be used in this way too.

Another major form of variability is person-to-person variation in hearing thresholds, at different pitches. (In other words, inter-person differences in the audiogram.) The pitch region of hearing loss (as seen on the audiogram) has been shown to correspond to the pitch region of tinnitus, and to the trigger sound pitch that gives the most effective residual inhibition.[60] This form of person-to-person variability is handled by the HushTinnitus system using approximate measurements of various hearing thresholds of each user. All of the sound tracks are then generated accordingly. This is discussed in the next section.

## 3 - Maximizing effectiveness

One of the design aims of the HushTinnitus system was to maximize the effectiveness, i.e. the depth and duration, of the residual inhibition caused by each trigger sound. In this regard, the design is guided by the findings of mainstream published research, as reviewed on page 12. This section summarizes the system features influenced by those findings.

The Roberts team,[53][56][60] and others,[16][57] found that residual inhibition was most effective when the trigger sound was pitch-matched to the tinnitus pitch region. Another team found supporting evidence, that residual inhibition was more effective if the tinnitus pitch region was present in a more broadband trigger sound.[31] The generation of all of the HushTinnitus sound tracks follows these principals.

There are a number of different ways that tinnitus pitch can be estimated. Some procedures ask the user to tune (or select) an external pure tone, or note, to match the pitch of their tinnitus.[1][4][5][8][9][23][25][49] Other methods ask the user to numerically score several differently-pitched external sounds in terms of "likeness" to their own tinnitus.[43][60][64] In many cases, people with tinnitus report that their tinnitus is not like a single pure tone, but rather that it is noise-like (for example, a "hissing" type sound).[49] In these instances, methods that ask users to match their tinnitus to just a single pure tone have shortcomings, as noise-like sounds occupy a *range* of pitches (frequencies). Various studies have looked at the different kinds of tinnitus sounds that people hear. They found that 70%‡‡ (in 1,625 people),[49] 73% (in 200 people),[8] 47% (in 90 people),[60] and 53% (in 17 people) [67] described their tinnitus as sounding like something *other* than just a pure tone. Indeed, it has been argued that even if people describe their tinnitus as a pure tone, it is in fact a narrow range of pitches.[6][43][60] Also, repeated testing of individuals, with single tone matching methods, has shown that many do not give consistent estimates of their tinnitus pitch from one test to the next.[23][40][64] Clearly, a single tone matching procedure is not an ideal general method for estimating tinnitus pitch.

For these reasons, pure tone matching is not used by the HushTinnitus system. Instead, the system exploits another finding of the Roberts team, namely that the pitch region of tinnitus corresponds to the pitch region of hearing loss.[60][64] (In fact, there is a prior history of similar brief observations of this correspondence. In papers in 1928 [2] and 1940,[5] Edmund Fowler commented on it. He wrote, "When tinnitus occurs with sharp dips or narrow deep troughs in the audiogram curve I have observed that its frequency occurs within the limits of the slopes."[5] In 1931, Emanuel Josephson wrote, "The frequency of the tinnitus is generally at the pitch, or pitches, where the auditory curve drops off sharply."[4] Feldmann made a similar observation in 1971.[9])

Hearing loss is usually depicted on an *audiogram*. This is a graph that shows the quietest sound level a person can hear (their *threshold of hearing*), in each ear, at a number of different pitches (frequencies). The levels are shown relative to those of a typical young person with healthy hearing (using a standardized unit of measure: *decibels Hearing Level (dB HL)*). Hearing loss is seen as increased threshold levels (*threshold shift*) on this graph. Above a certain pitch (the *audiometric edge*), many adults' thresholds start to rise significantly. This pitch region, above the audiometric edge, is the predominant region of hearing loss in most adults. It occurs in most people due to age (*presbycusis*) or exposure to excessive sound levels (*noise induced hearing loss*). Some people also have "notches" in their audiogram: these are isolated pitch regions of hearing loss that occur below the audiometric edge. (Other forms of threshold shift are seen on audiograms, too, as a result of other types of hearing loss. For example, Ménière's disease, and conductive hearing loss, can both have distinctive audiogram characteristics. However, discussion of these is beyond the scope of this document.)

‡‡ The Tinnitus Archive asked 1,625 patients to select one or more descriptions of their tinnitus: 24.5% selected just "ringing", 4.4% just "clear tone", and 1.0% just "whistle". These descriptions (and no others) have been interpreted here as representing a "pure tone".

The HushTinnitus system estimates the user's audiometric edge, for each ear. This is then used by computations that tune the pitch of various features of the HushTinnitus sound tracks. The tuning process is guided by the findings of Roberts et al, and also by obtrusiveness considerations (p. 36). All of the measurements, and the generation of the sound tracks, are handled by a web-based application, on the HushTinnitus website. To find the audiometric edge, the system plays a quiet tone, through headphones, to each ear in turn. The tone starts at a high pitch (above the audiometric edge for most adults) and slowly descends in pitch. The user indicates when they can first hear the tone, by clicking a button or pressing return, and the system uses that as the estimate of the audiometric edge for that ear. The tone is played at a nominally fixed hearing level (i.e. at a nominally fixed value in dB HL) throughout its pitch descent. Therefore, the audiometric edge estimate represents the pitch at which the user's threshold crosses this level. The pitch descends in much smaller steps than used in a standard audiogram test, giving greater precision to the estimate. Also, the descending tone "beeps" at a regular rate, to help the user distinguish it better, especially against any tinnitus sounds. (The beeping is "optimally faded" to make sure that the user cannot hear any clicks, or "spectral leakage" effects, before the descending tone actually reaches their audiometric edge. Optimal fading is discussed more on page 30.) One advantage of the descending tone method is that it is quick, and simpler for the user than either a pure tone matching method or a likeness scoring method. All of the HushTinnitus threshold tests start with sound that is imperceptible (because the frequency is too high, or the sound is too quiet), and then move slowly to perceptibility. At this point the user presses a button, and that test is then terminated. The whole testing process is therefore kept as quiet and gentle as possible. This is important because it is a self-test procedure, and users may have altered loudness perception or comfort levels (including loudness recruitment, softness imperception or hyperacusis).

After estimating the audiometric edge, the HushTinnitus system then moves on to test for audiometric notches. It has been found that, like the hearing loss region above the audiometric edge, notched regions can also correspond to regions of tinnitus pitch, and to the trigger sound pitch region that gives the most effective residual inhibition.[43][56][60] The system estimates, for each ear, the user's thresholds at the specific frequencies that have been most commonly reported for notches: 6,000 hertz, 4,000 hertz and 2,000 hertz (which is sometimes called the "Carhart notch").[62][42][63][7][10][11] Frequencies are skipped if they are above the user's audiometric edge, to avoid frustrating the user with a needless test that they may not be able to hear. At each frequency tested, a test tone starts very quietly (below audibility for most people), and the sound level is slowly increased. When the user first hears the test tone, they click a button or press return, and the system uses that as the estimate for their threshold level (for that particular frequency and ear). Again, optimally-faded "beeping" is used, for the same reasons as above.

It should be noted that these measurements are *not* intended to be equivalent to a full calibrated audiometric test (for example, as used for general diagnostic purposes in a general hearing test). The system delivers the test sounds through the user's own headphones (connected to their computer, tablet or smartphone), and this part of the system is not formally calibrated. Rather, the system is looking only for certain specific large features (such as the audiometric edge, or substantial notches) in the user's threshold curves. These large and distinct features can be identified robustly in the presence of relatively large amounts of calibration error and typical equipment response variation. The system does use a calibration technique to counter the effect of differences in the overall system loudness (known technically as the system *gain*), due to different system volume settings etc. It does this by finding the user's hearing thresholds at 500 hertz, a frequency that is relatively robust to the most common forms of hearing loss.[§§] The 500

[§§] Ménière's disease is an exception, and people with Ménière's therefore should not use the HushTinnitus system. (Ménière's disease causes a spinning type of dizziness and hearing loss, often accompanied by low-pitched tinnitus.) A Ménière's audiogram can have a "reverse slope": the thresholds at higher frequencies can be relatively unaffected, but there can be significant low-frequency hearing loss, which may include 500 hertz.

hertz threshold of the user's best (i.e. lowest threshold) ear is used to calibrate all the other measurements. (More technical detail on this is available by completing the online free preview, and then clicking the "TECHNICAL" button at the end of the hearing check: the results of your individual measurements are given, along with technical explanatory notes.)

Another issue is identifying the most effective waveform type, or types, for the trigger sounds. The research on this was reviewed on page 13. Three general waveform types have seen the most investigation: noise, pure tones, and warble tones. Each of these waveform types has been investigated by a number of different research teams, and some studies have compared their relative effectiveness.[18][23][40][44][57][60] These studies did not agree on a single most effective waveform type. However, even if a common winner (averaged across participants) had emerged, there is evidence of significant person-to-person variation.[23][44][57] For these reasons, the HushTinnitus system includes all three waveform types. Of course, there are many important details and characteristics that need to be decided for each of these waveform types, particularly for the noise and warble tone types. These are discussed in the next section.

Sound level (volume) plays an important part in effectiveness too (p. 13). Of course, the sound level is under the control of the user in the HushTinnitus system. The general finding from research is that, to achieve effective residual inhibition, the trigger sound level should be at least 10 decibels louder than the minimum masking level (p. 13). It has been found that the minimum masking level is generally less than 10 decibels louder than the estimated loudness of the tinnitus.[36] This would imply that the trigger sound should usually be set to a level at least somewhere between 10 to 20 decibels louder than the user's tinnitus. (This agrees well with a 13.3 decibel figure found by the Roberts team[60]). Increasing the sound level increases the duration of the residual inhibition.[23][28] However, it also increases the obtrusiveness of the trigger sound,[30] and of course it must not exceed general safe listening levels. So there is a balance to be struck. Clearly, the user typically does not have any way to measure the various levels. Rather, the HushTinnitus User Guide aims to give general sensible guidance in setting the volume level, bearing these various factors in mind.

The effect of trigger sound duration was reviewed on page 13. Residual inhibition effectiveness varies greatly from person to person (p. 7), so some people require a longer trigger sound to achieve a particular period of residual inhibition, if they can achieve it at all. To cater for this wide spread of responsiveness, the HushTinnitus system provides a range of different sound tracks, as outlined in the previous section (2 - Catering for variability). Within the Type 1 tracks, different trigger and silence durations are provided. Type 1 uses trigger sounds that are a few seconds in length, followed by longer periods of silence. It is designed for people whose residual inhibition response lasts throughout the longer period of silence. This allows these people to experience tinnitus suppression *and* stimulus silence for most of the time. For others, Type 2 (discussed next) can be tried. This provides many triggers per second, but the individual triggers are perceptually "hidden" with a smoother-sounding continuous noise. Finally, for people who respond to neither Type 1 nor Type 2 (including people with no residual inhibition response at all), Type 3 provides custom-optimized continuous masking.

The HushTinnitus Type 2 waveform is distinctly different from the noise, pure tones and warble tones used in Type 1 and 3. Type 2 is a form of continuous noise – although, to listen to, it is distinctly different in tone quality (timbre) to Type 3a. It consists of two types of noise, mixed (added) together. One type of noise has been designed, in terms of its statistical properties, to have distinct, sudden, and relatively large changes in level (i.e. large, sudden changes in *instantaneous amplitude*), many times per second. Its purpose is to provide a rapid stream of recurring residual inhibition triggers, in the general manner of Feldmann's work,[9][24] and later observations by Henry et al.[47] I call this the *trigger noise*. Like all the HushTinnitus trigger sounds, to maximize effectiveness the trigger noise is pitch-matched to the user's region of hearing loss, as discussed above. If played on its own, the trigger noise has quite a rough sounding quality, and the largest level changes are distinctly perceptible as "pops" or "crackles". To make it much less harsh for the listener, it is mixed with another form of noise. This other noise has been designed to have the opposite statistical properties, i.e. to be stable and deviate very little from a defined range of level (amplitude). It spans a wider pitch range, too. I call it the *base noise*. The base noise has a much smoother sounding quality, with no discernible pops or crackles. The two types of noise are mixed together in a ratio that ensures that the base noise hides the perception of any pops and crackles from the trigger noise, but *only just*. I term this ratio the *subliminal mixing point*, as the individual trigger features in the trigger noise are delivered subliminally (i.e. they are not directly perceptible).

In the pursuit of maximizing effectiveness, the design of Type 2 uses a number of findings reported in a paper by James Henry and colleagues, from the US Department of Veterans Affairs.[47] The Henry team looked at various types of continuous sound for suppressing tinnitus. They asked 21 patients to rate the annoyance of their tinnitus before listening to the suppression sound. Then, during the suppression sound, they asked the patients to rate the annoyance of the combination of the sound and their tinnitus (if any was still perceptible). All of the sound types were found to reduce annoyance, but some were substantially more effective in reducing annoyance than others. The design of Type 2 builds upon a number of characteristics of the most effective sounds in this study. Firstly, the most effective sounds in the study were relatively broadband, but with a frequency spectrum that increased quite steeply with increasing pitch over much of the audio frequency band. The base noise used in the HushTinnitus Type 2 sound has been engineered to have this type of frequency spectrum. Secondly, the two most effective sounds in the Henry study had notable “dynamic” content: recurring deviations of 5 to 15 decibels above the average, lasting hundredths or tenths of a second, in the mid to upper frequency region. There are marked similarities with the Type 2 trigger noise in this regard. In particular, one of its statistical properties*** was designed to give deviations similar those in the Henry study. As already described, during the design of the Type 2 sound, the ratio of the trigger noise to the base noise was adjusted until these individual deviations were found to be *only just* imperceptible. The spectrum of the resultant combined sound was then examined. It was found – perhaps surprisingly – that the trigger noise, which spans the estimated region of the user’s hearing loss, was quite visible, standing well above the adjacent base noise. Specifically, at the user’s audiometric edge frequency, the addition of the trigger noise to the base noise raises the average level (the *power spectral density*) by around 12 decibels, with the deviating features standing an additional 10 to 15 decibels above that (i.e. at 22 to 27 decibels above the original base noise level). These figures are in good accord with the most effective suppression sounds in the Henry study. Further, whereas each sound type in the Henry study was fixed (i.e. the same for all users), the trigger noise in Type 2 is tuned to the estimated region of hearing loss for each user. Direct evidence and theory suggest that this is the optimal pitch region for triggering the most effective residual inhibition.[59][60][67]

The HushTinnitus Type 3 sound tracks are simply continuous versions of the Type 1 sound tracks. (So, Type 3a is continuous noise, Type 3b is a continuous pure tone, and Type 3c is a continuous warble tone.) As already discussed, Type 3 can be used in a “one shot” manner, e.g. by playing it for a few minutes at a slightly louder level, to induce a longer single period of residual inhibition. Alternatively, Type 3 can be used for extended periods for continuous masking, by playing it “in the background” at a quieter level, to mask (or partly mask) the tinnitus. (In this case, as noted earlier, it is only really Type 3a (continuous noise) that is intended for this purpose.) Like Type 1, the Type 3 tracks are all tuned to the user’s estimated pitch region of hearing loss. So, for one-shot use, they are optimal for triggering residual inhibition, as already argued for Type 1. However, when used as a continuous masker, what is the effectiveness of such sound? (Put another way, how unobtrusive is the overall result?) For people who have no residual inhibition response (even to the rapid recurring triggers in Type 2), masking is the next available option for suppressing tinnitus, so this is an important question.

It seems intuitive to link masking (i.e. tinnitus suppression during the external sound) with residual inhibition (i.e. tinnitus suppression that continues after the external sound), and to regard them as

*** Technically, the *probability density function* of its amplitude. This is discussed in more detail later.

being two aspects of the same underlying mechanism. Indeed, since early studies, they have been observed together,[4][9] and theorized as being a single extended response from the same mechanism.[9][32][56][59][60][67] Following this, we might then expect an external sound that gives optimally effective residual inhibition to also be optimally effective at masking. There is direct evidence to support this. A number of different studies [9][22][27] have revealed that pure tones or narrowband noise can mask tinnitus at a lower *sensation level*††† when they are played in the pitch region of hearing loss.[56] For most people, the predominant pitch region of hearing loss is in the high-pitched region above the audiometric edge. Roberts commented, "Almost all participants in these studies report masking when presented with sounds in this region (94% in the study of Mitchell, 1983)."[56] Similarly, a 1981 paper reported that the lowest minimum masking levels are, in most cases, found in the region of the tinnitus pitch.[18] A 1986 study tested a range of different masker sounds on a group of tinnitus patients.[30] The masker sound judged by the group to give the "least annoyance" was noise that contained only high-pitched content. (As might be expected, most of these tinnitus patients had significant high-frequency hearing loss.) By contrast, the noise type giving the "most annoyance" contained only low-pitched content. A 2004 study found similar results, namely that masker sounds with predominantly high-pitched content, and little low-pitched content, were generally the most effective at reducing the annoyance of tinnitus.[47] A 1981 paper presented comparable findings.[18] The 2004 study noted that "many tinnitus patients have tinnitus that seems to be suppressed most effectively by band-pass noise in the frequency region that encompasses the tinnitus."

Finally, on the subject of effectiveness, the earlier review (p. 12) noted that noise masking has been found to be more effective if a different noise signal is played to each of the left and right ears.[35] In line with the above discussion, it has been speculated that this may apply to residual inhibition too.[56] Guided by this, all of the HushTinnitus noise sound tracks (including Type 2) present a different signal to each ear, as do the warble tone sound tracks. In practice, however, it is common for users to inadvertently swap the left and right sides of their headphones or earphones (or even speakers). Therefore, although the HushTinnitus noise/warble tracks have a different signal for each ear, the two signals have the same characteristics (in terms of sound level, pitch-tuned features, etc.). This is true even when the hearing measurements for the two ears are different. A single "worst-case composite" of the hearing loss in *both* of the user's ears is used to tune the relevant features of the sound tracks. Users with tinnitus in both ears can play the tracks in stereo to both ears, whereas users with tinnitus in just one ear may choose to play the sound to just the affected side, e.g. by wearing an earphone on that side only. (This would be in keeping with research findings on the most effective ear to use.[15][17][20][23][28])

## 4 - Minimizing obtrusiveness

Most of the work during the HushTinnitus system design was spent on efforts to minimize the obtrusiveness of the trigger sounds. Some aspects of this have already been covered in the previous sections. This includes the design of the Type 2 sound track, which uses trigger features that are subliminally hidden in a broader bandwidth, and statistically smoother, noise type (p. 26). Minimally obtrusive continuous masking (with Type 3a) was discussed on page 27. The general ordering of the sound track types, in terms of likely obtrusiveness, was covered on page 22. This section deals with the remaining methods used in the design of the HushTinnitus sounds.

The HushTinnitus Type 1 and Type 2 sound tracks use repeating trigger sounds to maintain the state of residual inhibition. Type 1 has a particularly simple structure: alternating periods of trigger sound and silence, that repeat for as long as they are played. Since this is such a simple (and seemingly obvious) way to extend a person's residual inhibition, for long periods if wanted, we might wonder why it is not a commonly tried tinnitus suppression therapy. Indeed, the basic idea

††† Sensation level takes into account the person's threshold of hearing. (Therefore, it takes any hearing loss into account.) For example, if a pure tone is played at a level of "5 decibels Sensation Level" (written "5 dB SL"), that means it is 5 decibels above that particular person's threshold of hearing at that frequency.

is far from new. In 1971 Harald Feldmann published his now-classic study of tinnitus masking and residual inhibition, in which the technique was described for the first time.[9] This paper is widely referenced by later publications, so it is well known among researchers. Since, for those people who respond to this stimulation, their tinnitus is suppressed during the playing of the sound, this would seem to leave just one possible auditory objection: the "obtrusiveness" of the repeating trigger sound.

A few research papers do contain references to the obtrusiveness of certain types of repeating trigger sound, but generally just briefly. In 1977 Vernon wrote, "Another approach which has not been tested would be to utilize residual inhibition in the form of a pulsed masker; that is, to depend on residual inhibition to keep the tinnitus continually suppressed while at the same time arrange the pulse rate so as to disrupt speech intelligibility as little as possible. Even if this approach could be made to work, the drawbacks are so numerous as to indicate its use in only the very severe case."[14] Vernon did not give any further details on the "numerous" drawbacks he was thinking of, but we could imagine (from the context of research at the time) that they might include:

- The abrupt nature of starting and stopping a sound (compared to a continuous masking sound). The prior research by Feldmann had simply instantly switched the trigger sound between being completely off, to being completely on, to being completely off again, etc.[9] This is a particularly attention-grabbing form of transition, as discussed later.
- Masking problems, particularly with speech. The wearable tinnitus maskers available at the time (*ear-level* maskers) produced a predominantly low-pitched noise,[18] covering pitch regions important for speech intelligibility.
- Perhaps Vernon was thinking of annoyance from rapidly occurring pulses of noise, regularly demanding auditory attention. (Feldmann's work had used very short and rapidly occurring pulses.[9]) From later work [30] we would expect annoyance to be particularly high from the low-pitched masker noise available at the time.

More positively, in 1981 Hazell and Wood commented that some patients may benefit from "some kind of intermittent masking noise rather than a continuous one, and it is hoped that the next generation of maskers will include pulsed tones."[20] In 1983 Terry and colleagues wrote, "However, for residual inhibition to have clinical usefulness, some assessment needs to be made, both of the relief provided by residual inhibition and of any incapacity produced by the masker. Such factors are central to the issue of whether residual inhibition is useful as a clinical tool. The design of a tinnitus masker may therefore have to take into consideration both effects when the masker is present (i.e. the masker should be less annoying than the tinnitus, and if possible it should not mask speech) and the after-effects of the masker (i.e. residual inhibition, temporary threshold shifts)."[23]

In 1986, Terry and Jones completed a study that specifically investigated the relative annoyance of a number of different masker sounds.[30] This included a pure tone (at 2,000 hertz) that was alternately switched instantly on and then off, either every half-second, or every quarter-second. These "interrupted tones" were judged by a group of tinnitus patients to be more annoying than constant noise (centered at either a high pitch, or a medium pitch) or a constant pure tone; but less annoying than constant noise centered at a low pitch, or various other ongoing sounds that were tested. Again, there are some features that we might immediately suspect may have increased the annoyance of the interrupted tones: the abrupt instant transition between on and off, the rapid switching rate, the use of a tone frequency (2,000 hertz) placed in a sensitive part of the human range (that is also important for speech intelligibility), and perhaps the use of a tone rather than a band of noise centered at a high pitch (the latter being found in the same study to be less annoying).

The design the HushTinnitus system proceeded with the aim of significantly reducing the obtrusiveness of these early and experimental discontinuous sounds. Clearly, minimizing obtrusiveness is particularly important for users who want to use the sound tracks for long periods at a time, such as for all-day use. As work progressed, it became clear that there were a number of areas where optimization was possible – not just for discontinuous trigger sounds, but in some areas for continuous masking sound too.

(In order to describe the methods used here – some of which may be new to tinnitus sound therapy – a little more technical detail is sometimes used than in the rest of the document. However, within each of the following sections, the opening description has been kept reasonably non-technical, so the reader may just skip to the next section to avoid the more technical material.)

**Optimal fading**

The HushTinnitus Type 1 sounds contain alternate periods of trigger sound and silence. What is the best way to transition, or *fade*, from one to the other? Is there some way that we can optimize this fading, using *objective* (i.e. scientifically quantified) measures? Of course, the aim is to find the least obtrusive form of fading, which is ultimately a “subjective” (i.e. perceptual) matter, judged by the listener. However, there is an established body of scientific study that relates the physical (and easily measurable) properties of sound to the way we perceive them; a subject called *psychoacoustics*. Hugo Fastl, a prominent researcher in the field of psychoacoustic sound engineering, has proposed that when designing a new sound, the fading (also called *gating*) should be one of the first aspects considered – as it is of “utmost importance”.[54]

In the Type 1 sounds, the transition between trigger sound periods and silent periods is gradual, not instant. During this transition period, the sound level (technically, the *gain*) changes gradually, following a particular curve (of gain against time). This curve has been selected to meet certain objective (directly measurable) criteria. Those criteria have in turn been derived by considering some established psychoacoustic effects, with the aim of minimizing the “obtrusiveness” (or perceptual *salience*) of the transition. I have called this type of optimized transition “optimal fading”, and the method is outlined as follows.

The 1986 Terry and Jones paper comments that “if the energy distribution in time is uneven, as in the interrupted tones, then the annoyance rating is higher than that of the steady tones. It might be expected that sounds that were repeatedly switched on and off would be more annoying since in general perceptual systems react strongly to change in stimulation.”[30] The HushTinnitus optimal fading method considers, firstly, the change in perceived loudness with time (during the fading). In line with the comments of Terry and Jones, a limit is put on the rate of change of loudness, with the aim of minimizing the strength of the perceptual reaction to this aspect of the transition. Loudness models are well known in psychoacoustics, and such a model is used here to define a measure for loudness, and to put a limit upon its rate of change. (It is worth noting that an instant off-on or on-off transition is the worst possible violation of such a limit. In other words, the experimental discontinuous sounds of Feldmann and Terry/Jones were as *non*-optimal in this regard as it is possible to get. It is also noteworthy that a simple “straight line” fade (technically, a constant rate of change of amplitude with time) performs poorly in the model: it has a particularly high rate of change of loudness around the lower part of the line.)

Secondly, the HushTinnitus optimal fading method considers how the fading curve affects the frequency spectrum. In the field of signal processing, it is well known that when a fading curve (or *window function* as it is generally known in this field) is applied to a signal, it spreads the spectral energy of the signal into other parts of the spectrum. This is known as *spectral leakage*. So, even if the spectral energy of a continuous signal is confined to a certain range of frequencies, applying fading (or *windowing*) spreads some of the energy into other frequencies. This is exactly the situation we have here with the Type 1 sounds: we start with a continuous signal whose spectrum is well constrained to a wanted band of frequencies (or even just a single frequency in the case of a pure tone), and then we apply fading. It is mathematically unavoidable that this will spread the spectral energy to other frequencies, to at least some extent. This is undesirable from the point of view of minimizing obtrusiveness. For example, the spectral leakage can descend from the target frequency band above the audiometric edge, down into more sensitive (and hence more salient) frequency regions of hearing.

However, with the right choice of fading curve, the spectral leakage effect can be minimized. With this in mind, the HushTinnitus optimal fading method sets another type of limit, defining the allowed spectral leakage. It is worth noting that *instantly* switching the trigger sound on and off creates a high level of spectral leakage. So, again, the early experimental sounds of Feldmann and Terry/Jones perform poorly in terms of this aspect of obtrusiveness. A "straight line fade" over a period of time gives lower levels of spectral leakage, but is still a long way short of the performance of an optimized fading curve over the same period.

So, the HushTinnitus optimal fading curve is produced to meet both of these limits: a temporal limit (maximum rate of change of loudness with time), and a spectral limit (an allowed spectral leakage profile). It is worth noting immediately that there is a very simple solution: if we make the fading slow enough then we can always satisfy both criteria. This works for any shape of fading curve, as long as the rate of change is always kept below a certain limit. However, this simple solution is not suitable for the present problem. The Type 1 trigger sound pulses are only a few seconds long. Since the sound level has a significant influence on the effectiveness of residual inhibition,[23][28] and since we wish to maximize the effectiveness of each trigger sound pulse, we want to keep the recorded signal at (or near) its maximum level for as much of the trigger pulse period as possible. This consideration pushes us to shorten the fading time. In turn, that pushes us to not simply slow the fading (with some arbitrary curve), but rather to produce an optimally shaped curve that gives better performance in the time available. The use of an optimally shaped fading curve allows us to simultaneously: (*a*) set relatively strong temporal and spectral limits for minimizing obtrusiveness; and (b) keep the fading time relatively short, to maximize effectiveness.

Optimal fading is applied to all of the HushTinnitus Type 1 trigger sounds, for the reasons discussed. A briefer version of the optimal fading curve is also used at the start and end of the "continuous" Type 2 and Type 3 sound tracks. These sound tracks are 15 minutes long, and are designed to repeat continuously (using single-track-repeat mode). However, most MP3 players insert a very short gap of silence when they reach the end of the track and go back to the start. Without any fading, this can sound like an obtrusive "click", every 15 minutes, in an otherwise continuous sound. (The obtrusiveness is in line with the temporal and spectral arguments just given.) Therefore, fast optimal fading is used, to minimize the perceptual salience of the gap.

**Smoother noise**

As a technical term, "noise" can have various meanings. One of its meanings is to denote a *random* ongoing signal. We are interested specifically in a random ongoing *sound* signal. If we look at the waveform of such a noise sound, it changes from one moment to the next, in an essentially unpredictable way. People generally perceive this sort of sound as having a "hissing" quality. (As an aside, the word noise can have distinctly different meanings in general usage. Dictionary definitions of noise (as used generally) contain descriptions like "loud", "unpleasant" and "undesired". The word noise, used in our present technical context of tinnitus sound therapy, implies none of these. Indeed, the right type of noise (specifically, noise centered at a high pitch) was found in one study to give the lowest annoyance out of 9 different types of sound.[30])

Although noise signals are random, they can still be characterized in various ways. Perhaps the most widely discussed characteristic is the frequency *spectrum*‡‡‡ of the noise. The frequency

‡‡‡ Strictly, as a technical detail, we are talking about the *energy spectrum* or *power spectrum* (also called the *power spectral density*) here. In wider technical use, the term "spectrum" can refer to a closely related (but slightly different) concept: the Fourier transform. In addition to giving the "strength" (the amplitude) of the signal at each frequency, the Fourier transform gives information on the "relative timing" (the *phase*) at each frequency. Two different signals can have exactly the same energy spectrum or power spectrum, but they will always have a different Fourier transform. For the sake of brief terminology, however, the word "spectrum" is used in this section to mean energy or power spectrum.

spectrum (or simply, the spectrum) is a graph that shows how much energy the signal has at different frequencies. Noise with any shape of spectrum can be created. Some shapes have been given names, following an analogy with the spectrum of visible light. *White* noise has equal energy at all frequencies, so its spectrum is flat. The spectrum of *pink* noise is a tilted line, with more energy at lower frequencies. Other types of spectrally shaped noise include brown, blue, violet and grey. It is common to use a signal-processing function called a *filter* to shape the spectrum of noise. A filter can be used to confine the noise to just a certain band of frequencies.

So does the spectrum of the noise *completely* describe it? The answer turns out to be *no*.[§§§] There is another characteristic called the *probability density function* (PDF). The PDF is quite a straightforward concept, and it is a fundamental property of any random signal (or, indeed, any randomly varying quantity). The PDF is a graph that shows how the signal is distributed over the range of values it varies over. For example, if a signal varies between a value of –10 and +10, spending an equal amount of time (on average) around any value, then its PDF graph will simply be a flat line (from –10 to +10). This type of distribution is called a *uniform* distribution. However, if the signal spends more time in one range of values than another, then the PDF will not be flat: the graph will be raised in the range in which the signal spends most time, showing the proportion of time spent there. In complex systems, including many found in nature, a "bell shaped" PDF very often occurs. This is called a *Gaussian* (or *normal*) distribution. The PDF of a Gaussian distribution has a central bell shaped hump, but the opening of the bell flares out at the bottom into two "tails". (In theory, these tails extend out infinitely; but of course, in practical reality something always limits the random value to a finite range.)

A quantity related to the PDF is the peak to average power ratio (PAPR). This is a single number, describing a signal, that simply gives the ratio of the maximum (peak) power in the signal to the average power, averaged across the whole signal. The PAPR of a signal can be derived from the PDF of the signal. The PAPR value can give an idea of the variability (i.e. the "dynamic" nature) of a signal, in the sense that it is a measure of how much its power varies. For example, if a signal has recurring brief but large deviations away from the norm (so it has quite a high variability, in this sense), then it will have a high PAPR value. On the other hand, if a signal varies with a more even distribution within a given range (so it has less variability, in this sense), then it will have a lower PAPR value.

It transpires that if a signal has a white spectrum, then it can have *any* PDF: Gaussian, uniform, or indeed any shape imaginable. As regards electronic systems, analog circuits nearly always produce noise with a Gaussian PDF. In digital circuits, or computer programming languages, it is easy to generate noise signals with a uniform PDF; but with a little more complexity, absolutely any PDF can be produced. The designer of a digital white noise generator therefore has total freedom over the PDF that is used. If the signal spectrum is not white, e.g. if it is filtered to contain just a certain band of frequencies, then there are certain constraints on the PDF shapes that are possible. However, there is still a great deal of freedom, especially if the bandwidth is fairly broad.

So we have established that noise (i.e. random) signals can be mathematically described by two types of graph: the spectrum and the PDF. We have also seen that there is great flexibility in terms of being able to engineer different shapes for both of these graphs. But what do different graph shapes actually *sound* like? In other words, what are the perceptual consequences of changing these graphs? Regarding the effect of different spectrums, that question has been relatively well explored by the research reviewed in this document. Unsurprisingly, noise whose spectral energy is primarily at high frequencies sounds high pitched, and noise with spectral energy mainly at low frequencies sounds low pitched. In terms of obtrusiveness, noise that is confined to a frequency band centered at a high pitch has been found to be generally less

[§§§] Actually, the Fourier transform of a signal (the meaning that some documents ascribe to "the spectrum") *does* completely and uniquely describe that signal. However, the energy or power spectrum (the meaning of "the spectrum" here) does not.

obtrusive.[30] Noise in the same pitch range as tinnitus has been shown to be the most effective at achieving both residual inhibition and tinnitus masking (see pages 12 and 27).

However, the perceptual effects of different PDFs seem not to have previously been researched. It has just been noted that different PDFs can be associated with different noise signal dynamics, and that the PAPR value is one measure of the variability of changes in signal level over time. In terms of perception, do signals with different PDFs (and particularly, with different PAPR values) *sound* different? It can readily be demonstrated that sounds with the same power spectrum and average power, but with different PDFs, *can* indeed sound distinctly (and even dramatically) different. For example, for a white noise generator we can design a PDF that keeps the signal within a small low-level range (just a few decibels sensation level) for most of the time, and that very occasionally allows it to jump up to a much higher value (say 80 decibels sensation level), but that assigns a tiny (if any) probability of occurrence to the values in between. The waveform of this sound looks like low-level smooth noise (around a few dB SL) for most of the time, but occasionally it has very large spikes (at 80 dB SL), occurring at random intervals. Large transients like this are well known to sound like “clicks”. Such a sound, unsurprisingly, is found to be perceived as a generally steady and smooth quiet noise that is occasionally interrupted by very loud and distinct clicks. This is an extreme example, with a very large PAPR value, just to demonstrate that varying the PDF of a noise signal can have a large perceptual effect. Clearly, an infinite variety of more moderate examples (with more moderate perceptual effects) can be created, by appropriate manipulation of the PDF.

To get a first quick (and informal) idea of the kind of perceptual effects that occur with different PDFs, I carried out various listening tests on myself. By changing the PDF, to produce ever larger PAPR values, I found that I perceived the sound to be increasingly “rough”, “gritty”, “crackly” or having large distinct “clicks”. Also, as the PAPR increased, I judged the sound to be ever more obtrusive. For all of these tests, the power spectrum and average power were kept the same.

Of course, single-person tests are not ideal. To improve upon this, and to explore a known form of obtrusiveness, we can look to a well-established psychoacoustic concept called *roughness*.[54] In psychoacoustics, roughness is a perceptual effect associated with moderately rapid sound variations. It occurs for variation rates between about 15 hertz and 300 hertz.[34] Various models have been devised that can estimate the level of perceived roughness of a sound. These models have been developed to match the responses of groups of real people, in various studies. Roughness can be regarded as a form of obtrusiveness, as it is generally linked to a reduction in “sensory pleasantness and the quality of noises”.[41]

For a noise signal, as the PAPR is decreased these psychoacoustic measures of roughness are also generally seen to decrease. This provides further weight for the argument that lower PAPR is associated with lower obtrusiveness. In a different line of research, although Terry and Jones did not directly explore the effect of noise PDF or PAPR, they did find a general result that if the energy distribution in time is more even, then the annoyance level is lower.[30] Again, this would suggest that a lower PAPR should generally be less obtrusive.

All of these considerations suggest that we should use a low PAPR form of noise. For example, uniform noise has a PAPR value of 4.8 decibels. Gaussian noise, however, has a much larger PAPR.**** Given a choice of these two noise types, uniform noise would therefore seem to be preferable. However, with other forms of PDF, it is possible to do better still. The HushTinnitus Type 1a and 3a sound tracks use noise with PAPR significantly lower than that of filtered uniform noise, and very much lower than that of Gaussian noise. This noise type is also used for the base noise of the Type 2 sound track (p. 26). However, the trigger noise of Type 2 exploits the opposite approach. It is uses a specifically designed high PAPR PDF, which has relatively large recurring

**** In theory, an unending Gaussian signal has an *infinite* PAPR. In practice, of course, the audio equipment limits the maximum PAPR that can be achieved. Shorter listening times also reduce the probability of observing a large PAPR.

deviations, for the purpose of triggering residual inhibition. This is then perceptually hidden in the low PAPR base noise, as discussed earlier.

### Optimal warble tones

The HushTinnitus Type 1c and Type 3c sound tracks use "warble tones". These tracks contain a pure tone whose frequency (pitch) is varied in a repetitive fashion. Such frequency variation is known technically as *frequency modulation* (FM), and so these tones are also known as FM tones. ("Modulation" is simply a technical term for "variation".) In terms of basic theory, the method is the same as that used in FM radio transmission, and some of the theory that was developed for FM radio is touched on here. This section looks at how the HushTinnitus warble tones were optimized to minimize obtrusiveness.

There are four design decisions that need to be made in order to specify how the frequency varies in a repetitive warble tone (also called a *sweep tone*): the lowest frequency of the sweep, the highest frequency of the sweep, the shape of the curve (of frequency against time) swept out between these extremes (called the *modulation waveform*), and how often the sweep curve repeats (called the *modulation frequency*). A few papers have previously reported on the use of warble tones for triggering residual inhibition.[40][44][57] However, there is incomplete information in these papers regarding the four specifications just mentioned. The first paper [40] does state that the frequency was swept over a range that included the estimated tinnitus pitch, and it includes a software screen-shot showing a 1.2 hertz modulation frequency. It does not state the modulation waveform used, however. The other two papers [44][57] again state that the frequency sweep range was matched to the tinnitus pitch, but they do not give the modulation waveform or the modulation frequency. The 1986 study by Terry and Jones looked at two warble tone sounds for the purpose of tinnitus masking, rather than residual inhibition.[30] The upper and lower frequencies of the sweep are given, and the modulation frequency is given as 2 hertz. However, again the modulation waveform is not stated.

In the wider research field of hearing perception in general, i.e. psychoacoustics, again studies have been done that are highly relevant to the present task of minimizing the obtrusiveness of an FM tone. This research looks at the temporal (time varying) nature of an FM audio signal, and how that relates to perceived effects.

The subject of *roughness* (as discussed in the previous section) is revisited below, this time in the context of FM tones; but first another psychoacoustic concept is discussed: *fluctuation strength*. In psychoacoustics, fluctuation is a perceptual effect associated with sound variations, but variations that are slower than those that cause roughness. The sensation of fluctuation is strongest for variations at around 4 hertz, and it occurs in response to variation in either frequency or sound level (amplitude).[54] The sensation of fluctuation plays a crucial role in how the human brain decodes speech.[54] Fluent speech contains around 4 syllables (i.e. major variations of frequency or amplitude) per second, matching the 4 hertz peak for perceived fluctuation strength. Fastl wrote, "As can be expected in nature, the human speech organ produces speech sounds with dominant envelope fluctuations at such a rate for which the human hearing system is most sensitive. In sound engineering, fluctuation strength is indispensable for the creation of warning signals. Actually, most warning signals used in daily life show large values of fluctuation strength."[54] Indeed, many emergency alarms and sirens use FM (i.e. pitch variation) at a rate of a few hertz, exploiting human sensitivity to this particular frequency of variation. Clearly, in order to be unobtrusive, the HushTinnitus modulation frequency should be kept well away from this attention-grabbing region. (Of note, the FM tones of Olsen (at 1.2 hertz) [40] and Terry/Jones (at 2 hertz) [30] both had a relatively high fluctuation strength, with consequent high obtrusiveness.)

To minimize the sensation of fluctuation, in principal we can clearly either increase or decrease the modulation frequency away from the peak sensitivity at 4 hertz. Substantially decreasing the frequency is not really a suitable option for the HushTinnitus Type 1c sounds, however, as they can be very short. (Taking the fading periods into account, some of these sounds are too short to properly represent the signal at very low frequencies.) Increasing the frequency is therefore the option to pursue here. As modulation frequency is increased upward from 4 hertz, the sensation of fluctuation decreases; but the sensation of *roughness* starts to increase. As noted in the previous section, roughness is also a form of obtrusiveness, as it is associated with making sounds "unpleasant". For frequency modulated tones, roughness reaches a peak at a modulation frequency between 40 and 80 hertz.[21][61]

As modulation frequency is increased further, roughness diminishes. So, is there anything to stop us from simply setting the modulation frequency to a very large value, say several thousand hertz? It transpires that there is, and the issue is with some mathematically inevitable features of the frequency spectrum of the signal. Firstly, there are features called *sidebands*. Sidebands are essentially the same kind of effect as spectral leakage, discussed on page 30. Spectrum analysis of an FM signal shows that (perhaps surprisingly) the spectrum is *not* constrained within the limits of the highest and lowest frequency of the sweep.†††† Rather, it spills out either side of that range, into the sidebands. These sidebands can be problematic for the present purposes: they spread the sound's energy away from the desired frequency region, potentially reducing the effectiveness of the trigger sound; and they can potentially spill down into sensitive pitch regions of hearing, increasing the obtrusiveness of the sound. Fortunately, the sidebands can be controlled to a large extent. The choice of modulation waveform (i.e. wave shape) is one factor, with a pure tone (also called a *sine wave*) causing the least sideband energy, for a given modulation frequency. Therefore, this is what the HushTinnitus system uses. The modulation frequency is critical too. In 1922, an engineer and communications theorist called John Renshaw Carson published the first work investigating this. He found that, although FM sidebands are theoretically infinitely wide, nearly all of their energy is actually contained within a defined frequency range. His formula for this is known as Carson's bandwidth rule. It shows that larger modulation frequencies cause wider sidebands, spreading more of the signal's energy outside of the sweep range. The choice of modulation frequency in the HushTinnitus system was guided by Carson's bandwidth rule: it is placed high enough to minimize perceived roughness and fluctuation, but low enough to minimize the negative effects of sideband energy.

Aside from sidebands, spectrum analysis also shows another unwanted effect of high modulation frequencies. The spectrum of an FM tone is not a smooth continuous curve, but rather it is composed of a series of isolated and extremely narrow "spikes". That is, the energy of the sound exists only at a set of distinct and separate frequencies. As the modulation frequency increases, these separate frequencies in the spectrum become more widely spaced apart. In the extreme, at several thousand hertz or more, we can end up with few (if any) of these frequencies existing in the target frequency region. This is obviously to be avoided, and so this frequency-spacing effect is another consideration in setting the modulation frequency.

In terms of the four design decisions introduced at the start of this section, two (modulation waveform and modulation frequency) have so far been discussed. The remaining two, the choice of the upper and lower frequency limits of the sweep, are also clearly of great importance. The considerations here, in order to maximize effectiveness (see pages 12 and 27) and minimize obtrusiveness (see page 36), are essentially the same as for the upper and lower frequency limits of the noise based trigger sounds (i.e. Type 1a, Type 3a, and the trigger noise of Type 2). Therefore, the warble tone frequency limits are set to be the same as the frequency limits for

†††† You might ask, how can this be? After all, have we not confined the frequency exactly, by the very way it is generated; i.e. by *specifying* it to vary between an upper and lower frequency limit? The answer turns out to be in the "mathematical small-print". Mathematically, there are actually two types of "frequency". "Instantaneous frequency" is allowed to vary with time. Within this section, up to this point the word "frequency" has actually meant "instantaneous frequency" – a concept very useful for describing FM signals. However, in spectrum analysis (which in turn is based on *Fourier analysis*), the concept of frequency does *not* vary with time. The method inherently analyses any time-varying signal as being composed of a number of *fixed* frequencies. Mathematically, the two types of frequency are quite different: the description of a signal using one type can be quite unlike the description of the same signal using the other type.

these noise sounds. This frequency placement is in general accordance with the methods of studies that have observed residual inhibition in response to warble tones.[40][44][57]

So, the design process for the HushTinnitus warble tones has minimized obtrusiveness due to perceived fluctuation, roughness, and also the negative spectral effects of excessive modulation frequency. However, there is one final problem to address regarding the warble tones. If we follow the principals just outlined, we find that some combinations of the four design decisions result in a sound that is perceived as distinctly discordant and harsh. Spectrum analysis shows that the spectrum of these sounds does not conform to the pattern of normal tonal (or *harmonic*) sounds. Other combinations, however, lead to a sound that is perceived as a tone, with no discordant quality. Spectrum analysis reveals that these sounds do conform to the normal pattern of harmonic sound. These types of effect were first observed, and mathematically analyzed, in the late 1960s and 1970s by a composer, theorist and Stanford professor called John Chowning. In 1973 he published a paper containing the mathematical conditions necessary for producing harmonic (as opposed to inharmonic) sound by frequency modulation.[12] (Chowning's work went on to become the basis of the first digital music synthesizers, which became widely used in the 1980s.) The HushTinnitus warble tones are generated according to Chowning's conditions, thereby mathematically guaranteeing that the sounds are harmonic.

**Optimal tuning**

As already discussed on page 28, all of the HushTinnitus sound tracks are tuned to the estimated tinnitus frequency region, with the aim of maximizing effectiveness. In addition to effectiveness, the tuning process also needs to be designed to minimize obtrusiveness. In particular, there are some specific considerations for the sound types that occupy a range (a *band*) of frequencies, namely the noise and warble sounds.

As discussed on page 24, most adults have less hearing sensitivity at high frequencies than that of young people with healthy hearing. The main frequency region of hearing loss is above the *audiometric edge*. The audiometric edge is the highest frequency at which a person's hearing sensitivity is considered to be "normal". However, the audiometric edge is not a sudden "cliff" in the *threshold of hearing* on the *audiogram* (see page 24 for definitions of these terms). As frequency increases, a typical audiogram shows that, above some frequency, the threshold of hearing gradually starts to rise; and thereafter it continues to rise as the frequency is increased further. At some point, the threshold crosses from what is generally regarded as normal hearing, into a region regarded as having "hearing loss"; and this frequency point is taken to be the audiometric edge.[60][64] So, the region of hearing loss is really a region of "partial hearing", with least impairment just above the audiometric edge. As discussed earlier (p. 24), tinnitus usually occurs in a bandwidth of frequency in the region of hearing loss. Trigger or masker sounds played in that frequency region (which still has partial hearing) have been found to be the most effective (pages 12 and 24).

It is clear from the research that the trigger or masker frequency band should cover the tinnitus frequency region. However, what is the effect of *excess bandwidth*, i.e. trigger or masker frequency range that overshoots either the lower or upper limit of the tinnitus frequency span? In terms of effectiveness, the emerging theories [59][60][67] might suggest that excess bandwidth, in itself, should not detract from the mechanisms of residual inhibition or masking, as "reafferentation" (i.e. re-establishment of neural input, p. 17) is still occurring in the relevant frequency region. However, there is an aspect of spectrum analysis that would seem relevant. If a sound (in this case a band of noise or warble tone) is kept at a constant overall power while its bandwidth is increased, then the power at any one frequency drops. This is because the available power is now more "thinly spread" over a greater frequency range. In technical terms, for each doubling in bandwidth, the power (strictly, the *power spectral density*, or *PSD*) at a particular frequency within the band drops by 3 decibels. So, for example, we could start by playing a bandwidth of noise that exactly covers a person's tinnitus frequency range. Then, if we double the bandwidth, we could compensate for the drop in PSD by simply increasing the overall sound power by 3 decibels. (3 decibels is a relatively small change, in terms of perceived loudness, or movement of a volume control on an audio player.) In this way, the signal power across the tinnitus frequency range would be kept the same. We might therefore expect that the residual inhibition and masking effectiveness would stay the same, at least approximately. Aside from theory, there is direct evidence that increasing the sound level of a trigger sound increases its effectiveness (p. 13), and therefore we might expect that increases in sound level could be used to correct for reduction in effectiveness due to other factors.

So there would appear to be an argument that, from the point of view of effectiveness, even quite large excess bandwidth can be accommodated for by relatively moderate increases in overall sound power. Of course, as ever we must keep sound levels at safe and comfortable limits (and within limits of the audio player), but the fact that only a moderate power increase is involved makes it a practically useable option in many cases.

Looking next at obtrusiveness, for the predominant region of hearing loss above the audiometric edge, the excess bandwidth can be divided into *upper excess* (i.e. sound frequencies above the tinnitus range) and *lower excess* (i.e. sound frequencies below the tinnitus range). Upper excess is much less of a problem, as the user's rising hearing threshold in this region approaches, and then exceeds, the sound level. (In other words, the user can hardly hear this excess; or perhaps not even hear it at all.) Lower excess is far more of an issue, as the reverse is true: the excess descends down into the user's region of good, and typically ever more sensitive, hearing. So, although a relatively small amount of lower excess would seem to have little impact on effectiveness (or required overall sound power increase), it can have a noticeable effect on obtrusiveness. This latter effect was directly observed by Terry and Jones.[30] In one experiment, they studied the effect of "high pass noise" (noise with only high frequencies present). As the lower frequency limit of the noise was reduced, their patients reported that it was increasingly annoying. For these reasons, the design of the HushTinnitus system pays much more attention to the exact location of the lower frequency limit of the sound, keeping it close to the estimated audiometric edge. (As noted earlier, the HushTinnitus system's audiometric edge measurement is determined with much smaller frequency steps than in a standard audiogram test, enabling greater precision here.)

Regarding the placement of the upper frequency limit of the sound, this has just been argued to be less critical, in terms of obtrusiveness. However, it is still desirable to prevent the excess growing too far, in the interests of moderating the overall sound power increase required to maintain the effectiveness. With this in mind, the HushTinnitus system design was guided by data from a 2011 study that reports the individual audiograms of 67 people with tinnitus, along with statistical data on the slope (steepness) of the thresholds and tinnitus pitch.[64] Tinnitus audiograms were studied in other papers too.[9][14][15][52][60] Considering this data, a suitable upper frequency limit was chosen for the HushTinnitus sounds, dependent on the audiometric edge frequency. The study data suggests that this limit should provide coverage of the audible "partial hearing" band for most people with tinnitus, given likely playback sound levels (see page 26 and [49]). (Audiometric notches are handled too, but are beyond the scope of this document.)

## 5 - Summary

Roberts wrote, “Although residual inhibition duration induced by current methods is typically brief, residual inhibition can be a source of relief for tinnitus sufferers who have experienced a sound that otherwise has known only a life of its own. Vernon and Meikle (2003) have described instances where patients broke into tears at their first experience of silent ears after years of unremitting noise. Hence there is much current interest in optimizing residual inhibition for its possible clinical benefits.”[56] He also wrote that “considerable variability was present between subjects” and this “leads one to ask whether we can identify variables and procedures that may optimize residual inhibition for individual cases.”[53] The design of the HushTinnitus system followed exactly these thoughts. It pursued 4 design aims, namely to:

- maximize effectiveness,
- minimize obtrusiveness,
- cater for variability, and
- be easy to use.

The HushTinnitus system uses a web-based application to gather relevant audiometric estimates from the user. From this data it generates 67 custom sound tracks. These tracks are designed to cater for a wide range of responses, from people with strong residual inhibition responses (e.g. using Type 1 sounds), through to people with less response (e.g. using Type 2), all the way through to an optimal masking sound (Type 3a) that is suitable for people with no residual inhibition response at all. The main focus of the system is on sound tracks that repeat continuously and unobtrusively in the background. However, the Type 3 tracks can also be used for optimized one-shot stimulation, for people who get a useful longer residual inhibition response from this.

Data on human perceptual responses, regarding either the effectiveness or obtrusiveness of sounds, was sourced through a thorough review of mainstream peer-reviewed research. More specifically, this review used independent publications in the fields of both tinnitus research and psychoacoustics.

In terms of effectiveness, the aim was to identify waveform types, and other relevant sound parameters, that have been found by multiple independent teams to be the most effective. The key research finding here is that trigger sounds and masker sounds are most effective when they match the pitch range of the tinnitus; which in turn is to be found in a pitch range of hearing loss. In terms of waveform types, tuned noise, pure tones and warble tones have all been observed to be the most effective, in different studies. Also, some studies show person-to-person variability in terms of the most effective type. Therefore the HushTinnitus system provides all of these waveform types.

For the purpose of minimizing obtrusiveness, psychoacoustic models were used throughout the optimization work. The models were taken from independent studies on significant numbers of people. Although the models were taken from prior studies, the design of the HushTinnitus system applied them in some seemingly new ways. These included the development of optimal fading, low-roughness (low-PAPR) noise, and psychoacoustically optimized warble tones.

29 May 2013 (Version 1.1)

# References

(in chronological order)


[1] Spalding J.A. (1903). Tinnitus, with a plea for its more accurate musical notation. *Archives of Otology, 32(4)*, 263-272. Link: www.residualinhibition.com/spalding.pdf

[2] Fowler E.P. (1928). Marked deafened areas in normal ears. *Archives of Otolaryngology, 8(2)*, 151-155.

[3] Jones I.H., Knudsen V.O. (1928). Certain aspects of tinnitus, particularly treatment. *Laryngoscope, 38(9)*, 597-611.

[4] Josephson E.M. (1931). A method of measurement of tinnitus aurium. *Archives of Otolaryngology, 14(3)*, 282–283.

[5] Fowler E.P. (1940). Head noises: significance, measurement and importance in diagnosis and treatment. *Archives of Otolaryngology, 32(5)*, 903-914.

[6] Fowler E.P. (1944). Head noises in normal and in disordered ears: significance, measurement, differentiation and treatment. *Archives of Otolaryngology, 39(6)*, 498–503.

[7] Carhart R. (1950). Clinical application of bone conduction audiometry. *Archives of Otolaryngology, 51(6)*, 798-808.

[8] Reed G.F. (1960). An audiometric study of two hundred cases of subjective tinnitus. *A.M.A. Archives of Otolaryngology, 71(1)*, 84-94.

[9] Feldmann H. (1971). Homolateral and contralateral masking of tinnitus by noise-bands and by pure tones. *International Journal of Audiology, 10(3)*, 138-144, May-June 1971.

[10] Carhart R. (1971). Effects of stapes fixation on bone-conduction response. In Ventry I.M., Chailkin J.B., Dixon R.F. (Eds.), *Hearing measurement: A book of readings*, 116-129. Appleton-Century-Crofts, New York.

[11] Tonndorf J. (1971). Animal experiments in bone conduction: clinical conclusions. In Ventry I.M., Chailkin J.B., Dixon R.F. (Eds.), *Hearing measurement: A book of readings*, 130-141. Appleton-Century-Crofts, New York.

[12] Chowning J.M. (1973). The synthesis of complex audio spectra by means of frequency modulation. *Journal of the Audio Engineering Society, 21(7)*, 526-534, September 1973.

[13] Vernon J. (1975). Tinnitus. *National Hearing Aid Journal, 28*, 13 and 82.

[14] Vernon J. (1977). Attempts to relieve tinnitus. *Journal of the American Audiology Society, 2(4)*, 124-131.

[15] Vernon J., Schleuning A. (1978). Tinnitus: a new management. *Laryngoscope, 88(3)*, 413–419.

[16] Bailey Q. (1979). Audiological aspects of tinnitus. *Australian Journal of Audiology, 1*, 19-23.

[17] Vernon J.A. (1981). Some observations on residual inhibition. In Paparella M.M., Meyerhoff W.L. (Eds.), *Sensorineural hearing loss, vertigo and tinnitus (Ear clinics international), 1*, 138–144. Lippincott Williams & Wilkins, Baltimore.

[18] Vernon J.A., Meikle M.B. (1981). Tinnitus masking: unresolved problems. In Evered D., Lawrenson G. (Eds.), *Ciba Foundation Symposium 85 - Tinnitus*, John Wiley & Sons, Chichester, UK, 239-262.

[19] Chouard C.H., Meyer B., Maridat D. (1981). Transcutaneous electrotherapy for severe tinnitus. *Acta Oto-laryngologica, 91(5-6)*, 415-422, May-June 1981.

[20] Hazell J.W.P., Wood S. (1981). Tinnitus masking – a significant contribution to tinnitus management. *British Journal of Audiology, 15(4)*, 223-230, November 1981.

[21] Kemp S. (1982). Roughness of frequency-modulated tones. *Acta Acustica united with Acustica, 50(2)*, 126-133, February 1982.

[22] Mitchell C. (1983). The masking of tinnitus with pure tones. *Audiology, 22(1)*, 73–87.

[23] Terry A.M.P., Jones D.M., Davis B.R., Slater R. (1983). Parametric studies of tinnitus masking and residual inhibition. *British Journal of Audiology, 17(4)*, 245–256, November 1983.

[24] Feldmann H., (1983). Time patterns and related parameters in masking of tinnitus. *Acta Oto-laryngologica, 95(5-6)*, 594-598.

[25] Vernon J., Fenwick J. (1984). Identification of tinnitus: a plea for standardization. *The Journal of Laryngology & Otology, 98, Supplement S9 (Proceedings of the II International Tinnitus Seminar, New York, 10-11 June 1983)*, 45-53.
[26] Tyler R.S., Babin R.W., Niebuhr D.P. (1984). Some observations on the masking and post-masking effects of tinnitus. *The Journal of Laryngology & Otology, 98, Supplement S9 (Proceedings of the II International Tinnitus Seminar, New York, 10-11 June 1983)*, 150-156.
[27] Tyler R.S., Conrad-Armes D. (1984). Masking of tinnitus compared to masking of pure tones. *Journal of Speech and Hearing Research, 27(1)*, 106-111, March 1984.
[28] Tyler R.S., Conrad-Armes D., Smith P.A. (1984). Postmasking effects of sensorineural tinnitus: a preliminary investigation. *Journal of Speech and Hearing Research, 27(3)*, 466-474, September 1984.
[29] Engelberg M., Bauer W. (1985). Transcutaneous electrical stimulation for tinnitus. *Laryngoscope, 95(10)*, 1167-1173, October 1985.
[30] Terry A.M.P., Jones D.M. (1986). Preference for potential tinnitus maskers: results from annoyance ratings. *British Journal of Audiology, 20(4)*, 277-297, November 1986.
[31] Kitajima K., Kitahara M., Kodama A. (1987). Can tinnitus be masked by band erased filtered masker? Masking tinnitus with sounds not covering the tinnitus frequency. *The American Journal of Otology, 8(3)*, 203-206, May 1987.
[32] Tonndorf J. (1987). The analogy between tinnitus and pain: a suggestion for a physiological basis of chronic tinnitus. *Hearing Research, 28(2-3)*, 271-275, 1987.
[33] Kuk F.K., Tyler R.S., Rustad N., Harker L.A., Tye-Murray N. (1989). Alternating current at the eardrum for tinnitus reduction. *Journal of Speech and Hearing Research, 32(2)*, 393-400, June 1989.
[34] Zwicker E., Fastl H. (1990). *Psychoacoustics: facts and models.* Springer-Verlag, Berlin.
[35] Johnson R.M., Hughes F.M. (1992). Diotic versus dichotic masking of tinnitus. In Aran J.M., Dauman R. (Eds.), *Tinnitus 91: Proceedings of the Fourth International Tinnitus Seminar*, Kugler Publications, Amsterdam, 387–390.
[36] Mitchell C.R., Vernon J.A., Creedon T.A. (1993). Measuring tinnitus parameters: loudness, pitch, and maskability. *Journal of the American Academy of Audiology, 4(3)*, 139-151, May 1993.
[37] Okusa M., Shiraishi T., Kubo T., Matsunaga T. (1993). Tinnitus suppression by electrical promontory stimulation in sensorineural deaf patients. *Acta Oto-laryngologica Supplementum, 501*, 54-58.
[38] Kristeva-Feige R., Feige B., Kowalik Z., Ross B., Feldmann H., Elbert T., Hoke M. (1995). Neuromagnetic activity during residual inhibition in tinnitus. *Journal of Audiological Medicine, 4(3)*, 135-142.
[39] Jueptner M., Weiller C. (1995). Review: does measurement of regional cerebral blood flow reflect synaptic activity? Implications for PET and fMRI. *NeuroImage, 2(2)*, 148-156, June 1995.
[40] Olsen S.Ø., Nielsen L.H., Osterhammel P.A.A., Rasmussen A.N., Ludvigsen C., Westermann S. (1996). Experiments with sweeping pure tones for the inhibition of tinnitus. *Journal of Audiological Medicine, 5(1)*, 27-37.
[41] Daniel P., Weber R. (1997). Psychoacoustical roughness: implementation of an optimized model. *Acta Acustica united with Acustica, 83(1)*, 113-123, January-February 1997.
[42] McBride D., Williams S. (2001). Audiometric notch as a sign of noise induced hearing loss. *Occupational and Environmental Medicine, 58(1)*, 46-51, January 2001.
[43] Noreña A., Micheyl C., Chéry-Croze S., Collet L. (2002). Psychoacoustic characterization of the tinnitus spectrum: implications for the underlying mechanisms of tinnitus. *Audiology & neuro-otology, 7(6)*, 358-369, November-December 2002.
[44] Banks R., Atrach M. (2003). Effectiveness of Residual Inhibition Therapy (RIT). Published by the Canadian Hearing Society, Toronto.
[45] Savastano M. (2004). Characteristics of tinnitus: investigation of over 1400 patients. *The Journal of Otolaryngology, 33(4)*, 248-253, July 2004.
[46] Jastreboff P.J., Hazell J.W.P. (2004). *Tinnitus Retraining Therapy: Implementing the Neurophysiological Model*. Cambridge University Press.

[47] Henry J.A., Rheinsburg B., Zaugg T. (2004). Comparison of custom sounds for achieving tinnitus relief. *Journal of the American Academy of Audiology, 15*, 585-598.
[48] Eggermont J.J., Roberts L.E. (2004). The neuroscience of tinnitus. *Trends in Neurosciences, 27(11)*, 676-682, November 2004.
[49] Meikle M.B., Creedon T.A., Griest S.E. (2004). *Tinnitus Archive, second edition*. Oregon Hearing Research Center, Department of Otolaryngology, Oregon Health & Science University.
[50] Osaki Y., Nishimura H., Takasawa M., Imaizumi M., Kawashima T., Iwaki T., Oku N., Hashikawa K., Doi K., Nishimura T., Hatazawa J., Kubo T. (2005). Neural mechanism of residual inhibition of tinnitus in cochlear implant users. *Neuroreport (Lippincott Williams & Wilkins), 16(15)*, 1625-1628, 17 October 2005.
[51] Noreña A.J., Eggermont J.J. (2005). Enriched acoustic environment after noise trauma reduces hearing loss and prevents cortical map reoganization. *The Journal of Neuroscience, 25(3)*, 699-705, 19 January 2005.
[52] Noreña A.J., Eggermont J.J. (2006). Enriched acoustic environment after noise trauma abolishes neural signs of tinnitus. *Neuroreport, 17(6)*, 559-563, 24 April 2006.
[53] Roberts L.E., Moffat G., Baumann M., Bosnyak D.J. (2006). Residual inhibition functions in relation to tinnitus spectra and auditory threshold shift. *Acta Oto-Laryngologica, 126*, 27-33.
[54] Fastl H. (2006). Psychoacoustic basis of sound quality evaluation and sound engineering. In Eberhardsteiner J., Mang H.A., Waubke H. (Eds.), *Proceedings of the Thirteenth International Congress on Sound and Vibration (ICSV13), Vienna, 2-6 July 2006.*
[55] Lipman R.I., Lipman S.P. (2007). Phase-shift treatment for predominant tone tinnitus. *Otolaryngology – Head and Neck Surgery, 136(5)*, 763-768, May 2007.
[56] Roberts L.E. (2007). Residual inhibition. *Progress in Brain Research, Tinnitus: Pathophysiology and Treatment, Elsevier, 166*, 487-495, 22 October 2007.
[57] Sockalingam R., Dunphy L., Nam K., Gulliver M. (2007). Effectiveness of frequency-matched masking and residual inhibition in tinnitus therapy: A preliminary study. *Audiological Medicine, 5(2)*, 92-102.
[58] Struve M., Diesch E., Flor H. (2007). Extinction training for tinnitus. *Progress in Brain Research, Tinnitus: Pathophysiology and Treatment, Elsevier, 166*, 461-466, 22 October 2007.
[59] Kahlbrock N, Weisz N. (2008). Transient reduction of tinnitus intensity is marked by concomitant reductions of delta band power. *BMC Biology, 6:4*, 16 January 2008.
[60] Roberts L.E., Moffat G., Baumann M., Ward L.M., Bosnyak D.J. (2008). Residual inhibition functions overlap tinnitus spectra and the region of auditory threshold shift. *Journal of the Association for Research in Otolaryngology (JARO), 9(4)*, 417-435, 20 August 2008.
[61] Van Egmond R., Lemmens P., Pappas T.N., de Ridder H. (2009). Roughness in sound and vision. *Proceedings of SPIE 7240, Human Vision and Electronic Imaging XIV, 72400B*, 10 February 2009.
[62] Osei-Lah V., Yeoh L.H. (2010). High frequency audiometric notch: an outpatient clinic survey. *International Journal of Audiology, 49(2)*, 95-98, February 2010.
[63] Wilson R.H. (2011). Some observations on the nature of the audiometric 4000 hz notch: data from 3430 veterans. *Journal of the American Academy of Audiology, 22(1)*, 23-33, January 2011.
[64] Sereda M., Hall D.A., Bosnyak D.J., Edmondson-Jones M., Roberts L.E., Adjamian P., Palmer A.R. (2011). Re-examining the relationship between audiometric profile and tinnitus pitch. *International Journal of Audiology, 50(5)*, 303-312, 9 March 2011.
[65] Vanneste S., De Ridder D. (2012). Noninvasive and invasive neuromodulation for the treatment of tinnitus: an overview. *Neuromodulation: Technology at the Neural Interface, 15(4)*, 350-360, 11 April 2012.
[66] Atrach M. (2012). *Private communication*, 19 June 2012.
[67] Sedley W., Teki S., Kumar S., Barnes G.R., Bamiou D.E., Griffiths T.D. (2012). Single-subject oscillatory gamma responses in tinnitus. *Brain: a journal of neurology, 135(10)*, 3089-3100, 13 September 2012.

[68] Song J.J., Vanneste S., Van de Heyning P., De Ridder D. (2012). Transcranial direct current stimulation in tinnitus patients: a systemic review and meta-analysis. *TheScientificWorldJournal, 2012*, Article ID 427941, 17 October 2012.