Evaluation of Different Forms of Compression in Digital Hearing Aids

by

Ivan Aguayo

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Bachelor of Science in Electrical Science and Engineering

and

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2001

Author ....
Department of Electrical Engineering and Computer Science
May 22, 2001

Certified by ....
Julie Greenberg
Research Scientist
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Evaluation of Different Forms of Compression in Digital Hearing Aids

by

Ivan Aguayo

Submitted to the Department of Electrical Engineering and Computer Science
on May 22, 2001, in partial fulfillment of the
requirements for the degrees of
Bachelor of Science in Electrical Science and Engineering
and
Master of Engineering in Electrical Engineering and Computer Science

Abstract

People with cochlear hearing loss have a reduced dynamic range of hearing, thus amplitude compression may provide adequate amplification of soft sounds without uncomfortable overamplification of loud sounds caused by conventional linear amplification. Although compression is conceptually straightforward, there are various design parameters that may affect the intelligibility of speech, the quality of sounds, and the perception of background noise. Four combinations of the Dual Front-End automatic gain control (AGC) system were implemented: (1) The Dual Front-End with a Hold Timer developed in Stone *et al.* [8], which aimed at reducing pumping effects while maintaining a relatively fast release; (2) The Dual Front-End with the SNR Estimator, investigated by Martin *et al.* [4], designed to provide a varying release time constant depending on the SNR level; (3) The Dual Front-End with both the Hold Timer and the SNR Estimator; (4) The Dual Front-End by itself, without the Hold Timer or the SNR Estimator. A fifth system, composed of linear amplification and compression limiting, was implemented to be used as a reference condition. A variety of stimuli consisting of speech at different levels and speech plus environmental sounds were processed by the five systems and presented over headphones to three hearing-impaired subjects. Subjects rated the processed stimuli for intelligibility and quality. While no clear differences were found among the four compression systems, there were some major differences between the Dual Front-End systems and the Linear system. The direction of these differences varied with subject, and to a lesser degree, with stimulus condition. In addition, compression systems generally performed better in stimuli conditions with low SNRs, indicating that compression may be useful for suppression of background noise.

Thesis Supervisor: Julie Greenberg
Title: Research Scientist

Acknowledgments

This work was supported by Grant No. NIH-2-R01-DC00117 from the National Institute of Deafness and Other Communicative Disorders, and by the Intel Foundation, which generously provided me with a scholarship.

First off, I want to thank the three subjects that participated in the experiment (PG, MG, and RG) for their extreme patience and for making this work possible. Many thanks out to Merry, for providing me with her subjects and being a fun person to work with. Everyone at the Sensory Communication Group deserve a gratitude for helping me along the way of this project. Michael "Q" for getting me started and teaching me the essentials in compression. My officemate Joe for putting up with me, providing me with his extended computer skills, and always having a great sense of humor. Jeanie and Panina for their friendliness and their willingness to help me out. A special thank you also goes to Professor Louis D. Braida for his friendliness, and the much respectful person that he is.

Many special thanks to my advisor Dr. Julie Greenberg for being a great person to work with. She has to be the greatest advisor one can have. Thank you so much for your patience, your shared knowledge, and for simply being a fun person to be around. You are a great role model.

Outside of the SCG laboratory I would like to thank all my friends at MIT who have made it a memorable 5 years. Thank you for your support and the great times we had together. You all know who you are. Finally, my greatest thank you goes to my parents and family. This page is not enough to express my gratitude to my parents for their unconditional support.

This Thesis is dedicated to my two grandmothers, Paula and Petra, and my cousin Chuy, who have passed away during my years at MIT.

Que Dios los tenga en la gloria!

Contents

List of Figures

List of Tables

Chapter 1

Background

1.1 Introduction

Hearing-impaired people have a limited dynamic range of hearing. A hearing-impaired person will not be able to hear a low-level sound that a normal-hearing person would, yet a loud sound may be perceived equally loud for the hearing-impaired as for the normal-hearing person. For this reason, linear amplification does not provide adequate amplification of weak sounds without intense sounds being overamplified and becoming uncomfortably loud. Such problem can be alleviated using amplitude *compression*, where weak sounds are amplified more than high-level sounds.

1.2 Compression

The idea of amplitude compression applies time-varying amplification dependent on the input level. Greater amplification is applied to lower- level sounds, while high-level sounds receive less amplification, and are often attenuated. Compression systems contain a *static input/output curve* (I/O curve) which describes the level of compression to be applied. Figure 1-1 shows an example of an I/O curve. Input levels below the compression threshold (Th) are linearly amplified; a constant gain of LG is applied, thus the output corresponding to the threshold is the threshold level plus LG. Compression is applied to input levels above the threshold, with a compression ration of CR. The slope of the compression curve in the I/O plot is thus 1/CR. To protect the user from receiving uncomfortable loud sounds, an upper-level threshold (Max) is determined so that all input levels beyond Max be presented at the highest comfortable level of the user. This is what is referred to as *compression limiting.*

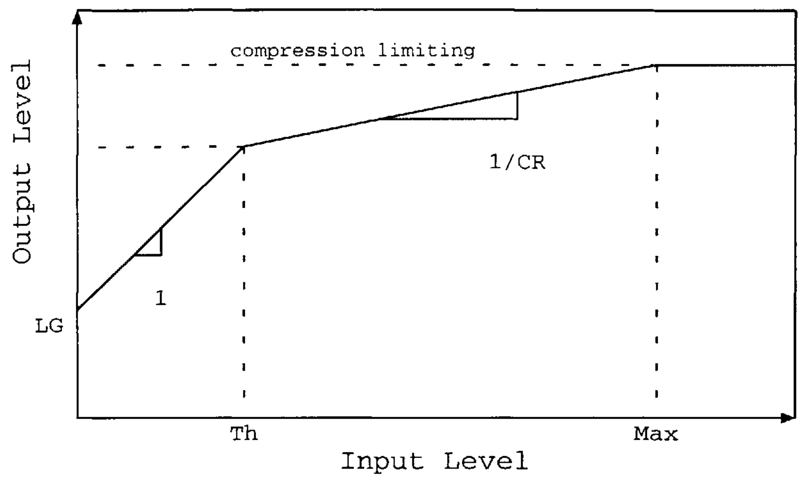

Figure 1-1: Example of a static I/O curve

1.2.1 Attack and Release Constants

An automatic gain control (AGC) compression system is dependent on the input level, which determines how much amplification or attenuation is needed. Compression includes attack and release time constants, which play a key role in the behavior of the AGC system. The temporal envelope of the signal must be calculated to serve as the input level value(s), which can be achieved by rectification followed by a low-pass filter. In calculating the temporal envelope, going from a given level to a higher-level sound, an *attack* is necessary to reach the higher-level value. The rate of the attack is determined by the attack constant, τ_a, where a larger value corresponds to a slower rate of increase. There is a similar occurrence when transitioning from a given level to a lower-level sound. In this case the system must *release* (or recover) to reach the lower-level value, determined by the release time constant, τ_r. Again, the larger the release value, the slower the decay rate.

1.2.2 Slow vs. Fast AGC Compression

There are two basic types of automatic gain control (AGC) compression systems, which are generally referred to as Slow-Acting Compression (Slow AGC) and Fast-Acting Compression (Fast AGC).[1] A Slow AGC compression system is composed of long attack and release time constants, producing a slowly-varying envelope over time, hence Slow AGC. One advantage of this system is that it prevents harmonic and intermodulation distortion due to its long recovery time [6]. Also, individual acoustic elements within speech may vary over a range of at least 40 dB, and the gain is desired to change little over such elements to prevent possible decrease in speech intelligibility [6]. Yet, a more significant advantage is reducing *pumping* or *breathing* effects. To better illustrate such phenomenon, suppose a speech signal contains a one second pause in between. During the speech before the pause, the gain may be low. At the moment of the pause, the gain will start increasing (i.e. envelope releasing), and then decreasing after one second when it detects speech again. If the release during the pause is fast, the gain will change significantly, causing an uprising "woosh" sound, referred as pumping. This effect is reduced with longer recovery times, as offered by the Slow AGC.

The Fast AGC system also offers its unique advantages. Contrary to the Slow AGC, the Fast AGC system is composed of fast (i.e. small values) attack and release time constants. When an intense sound transient is encountered, such as a door slam, a glass breaking or a sudden scream, protection to the user is necessary. Such system requires the gain to change rapidly in response to these transients, achieved with a fast enough attack. Also, immediately after cessation of the transient, the system should release relatively fast to avoid the aid going "dead" for a significant period of time. Since the system gain during the loud transient may be small, or even negative (i.e. attenuating), the time interval while the aid recovers to its original gain value may contain lower-level speech or important sounds that may be inaudible to the user.

[1]Fast-Acting Compression is also often referred as Syllabic Compression, since the gain changes over times comparable to the durations of individual syllables in speech.

1.3 Previous Studies

1.3.1 The Dual Front-End

Both the Slow AGC and the Fast AGC compression systems have their unique advantages and disadvantages, which makes it difficult to decide on which is the "best" method, or what aspects are more important and significant than the others. A better solution to these problems was developed by Moore and Glasberg [6], with a system referred to as the *Dual Front-End AGC*. The Dual Front-End AGC is composed of two control subsystems, a Slow AGC and a Fast AGC. Normal operation is determined by the Slow AGC, with slowly-varying gain over time. When intense sounds occur, the Fast AGC rapidly takes over and reduces the gain to protect the user. The Fast AGC can be triggered in two ways, depending on the design of the system. Either when the instantaneous input level goes above a fixed threshold level regardless of the slow control parameter (SCP) level, or when it goes above the SCP level by a fixed amount, generally 7 or 8 dB. Following cessation of the intense sound, the gain returns to its original value, determined by the Slow AGC.[2] The Dual Front-End provides outputs free of speech distortion because it is controlled by the Slow AGC most of the time, and provides protection to the user from intense transients. But like any system, it also contains a few disadvantages. Below are some of the Dual Front-End's advantages and disadvantages taken directly from Stone *et al.* [8].

Advantages of the Dual Front-End

- Speech is delivered at a comfortable level, regardless of the input level.
- The temporal envelope of speech is hardly distorted; envelope fluctuations at syllabic rates are preserved.
- The spectral pattern of sounds is not distorted.
- Harmonic and intermodulation distortion are minimal.

[2]The Fast AGC releases with a relatively fast time constant until it reaches its threshold of operation.

- Protection is provided from intense brief transients with little effect on the long-term gain.

Disadvantages of the Dual Front-End

- Loudness perception is not restored to "normal". Indeed, since the output level is held almost constant for input levels above the compression threshold, it may be difficult for the user to judge the strength of sound sources, for example, the volume setting on a television or radio.
- The system may not deal very effectively with situations where two voices alternate with markedly different levels, which is often the case when one of the voices is that of the aid wearer.
- When the user moves from a situation with high sound levels to one with lower levels, the gain takes a second or two to reach the value appropriate for the new situation.

A simplified diagram of the parallel implementation of the Dual Front-End is shown in Figure 1-2. The input level is simultaneously calculated by independent Slow AGC and Fast AGC control systems. The outputs of both systems, known as the *slow control parameter* (SCP) and the *fast control parameter* (FCP) for the Slow AGC and the Fast AGC, respectively, are compared at every instant. The maximum of the two control parameters is taken as the output of the Dual Front-End, which is used as the input level to calculate the gain used for compression. To ensure that the Slow AGC dominates the compression most of the time, x dB is added to the SCP, usually 7 or 8 dB. The Fast AGC will only then be triggered in sudden level increases. The FCP will rise much rapidly than the SCP and quickly surpass it by more than x dB, at which point the Fast AGC is activated.

The experiment conducted by Moore and Glasberg [6] showed significant benefit using Dual Front-End systems over linear amplification. In one version of the Dual Front-End, speech reception thresholds (SRTs) in noise were, on average, 4 dB better than for linear amplification.

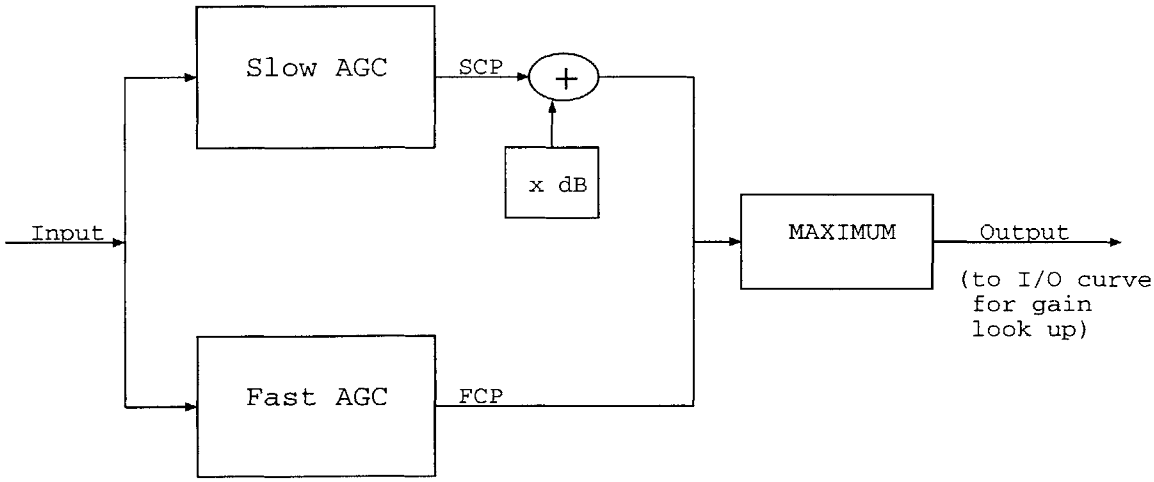

Figure 1-2: The Dual Front-End (parallel implementation).

1.3.2 The Hold Timer

A significant problem in fast-acting compression hearing-aids is the annoyance of pumping and breathing. In 1999, Stone *et al.* [8] developed a *Hold Timer* aimed at reducing pumping while maintaining a relative fast release constant. The basic idea of the Hold Timer is to prevent the system's gain from significantly changing during speech and during brief pauses in speech, or preventing the system from releasing too quickly. A counter is kept track of at every instant, where its value is incremented whenever the slow control parameter is increasing, and its value is decremented otherwise. The counter contains a minimum and a maximum value. When speech is presented after a moment of silence, the slow control parameter will rise, causing the counter to increase. If the control parameter keeps rising, the counter will reach its maximum and stay at the maximum value. When a pause in the speech is encountered, the control parameter will want to release, causing the counter to decrement. The rule is that the system gain is not allowed to increase (i.e. control parameter is not allowed to release) unless the counter value is at its minimum. Such rule will prevent the gain to change during the pause, until the counter reaches its minimum value. The duration of the gain hold clearly depends on the counter's rate of decrease, its minimum, and maximum values. If the level decrease detected in the input signal is a short pause in speech, then the hold timer will create a bridge, avoiding the unnecessary gain adjustment that might cause pumping. If on the other hand, the level decrease is lower-level speech from another talker or a new weaker environment, the system gain will be unchanged for the duration of the Hold Timer, but will contain a relative fast recovery time to quickly adapt to the new level speech or environment. Such fast recovery avoids long "dead" periods of the aid, found in systems with long recovery times.

In the experiment by Stone *et al.* [8], one of two AGC systems that contained the Hold Timer was slightly preferred by the subjects over three other AGC systems. An important factor to notice is that Stone *et al.* did not use a linear amplification system for reference. Such reasons provide motivation for further investigation of the Hold Timer.

1.3.3 SNR Estimator

Extreme variation in release time constants can produce either pumping for short release, or the aid going dead for a significant time when the release is slow. One may think that a mid-value of τ_r would be optimal, but it is sometimes desirable to have a short or long release. At times when release is due to short pauses in speech, a long recovery time is ideal since it will have little effect on the system's gain, thus avoiding pumping effects. On the other hand, when release is due to lower-level speech, such as another talker compared to the aid user's voice, a fast recovery time is necessary to prevent the aid from going dead and assure that the user will be able to hear everything said. Martin *et al.* [4] developed the idea of a time-varying release constant τ_r, which depends on a calculated signal to noise ratio (SNR). If the signal is detected to have an SNR level below a certain threshold, it may be assumed that release is due to background noise, thus the system will use a long recovery constant to maintain a relatively constant gain. If the SNR level is detected above the threshold, the system will switch to a short τ_r, to provide a fast recovery to the assumed low-level speech.

Martin *et al.* [4] estimated the SNR from the peaks and valleys of the signal's temporal envelope using an algorithm described by Festen *et al.* [3]. The results of the experiment in Martin *et al.* suggested that improved gain control can be achieved through the use of SNR estimates.

1.4 Goal of Thesis

The purpose of this research is to evaluate different combinations of compression systems, and determine any benefit in speech intelligibility and pleasantness under background noise from the Dual Front-End, the Hold Timer, and the SNR Estimator. The experiment developed by Stone *et al.* [8] evaluates various compression systems, including the Dual-LO system, but does not contain a reference linear system to compare. This experiment will allow for a similar Dual-LO to be compared with a linear amplification system. The experiment will also allow for comparison of the Noise Control System in Martin *et al* [4] to the Dual-LO. Finally, a combination using both the SNR Estimator and the Hold Timer will also be considered.

The experiment will be conducted using hearing-impaired subjects who have experience using hearing aids. Speech with a variety of SPL levels and background noises will be processed, and presented to the subjects monaurally over headphones in a sound proof booth. This laboratory experiment is different than the one in Stone *et al.* who had a wearable device that subjects wore for several weeks in their everyday acoustic environments.

Chapter 2

Implementation of Compression Systems

2.1 System Components

2.1.1 Pre-Emphasis

Prior to level calculation and compression, high-frequency emphasis was applied to the input signal in all five automatic gain control (AGC) systems. A rising gain of 3.3 dB per octave was applied between 500 Hz and 4 kHz. At and below 500 Hz the gain applied was 0 dB, and the gain at and above 4 kHz was 10 dB. The filter was designed in MatLab using a 65 point, linear phase, Remez Parks-McClellan optimal equiripple FIR filter. This emphasis stage, called pre-emphasis, served two purposes. One, it prevented the AGC systems from being dominated by low-frequency sounds such as car noise. Second, it prevented the AGC from being excessively activated by the user's own voice [5][8].

2.1.2 Calculating Temporal Envelope of Signal

Four of the five systems implemented used the Dual Front-End system. The fifth system, or the Linear System, made use only of the Fast AGC component to apply compression limiting. Both the Slow AGC and Fast AGC require the calculation of the signal's temporal envelope using different values of release τ_r, and attack τ_a time constants. The process to calculate the temporal envelope of a signal is shown in Figure 2-1. The signal is first rectified by taking its absolute value, then two independent low-pass filters are applied. One of the low-pass filters uses the attack time constant, τ_a, the other uses the release constant, τ_r. The maximum of the two low-passed signals is taken at every sample resulting in the desired temporal envelope, or slow control parameter. The low-pass filters used were first order infinite impulse response (IIR), described by the following difference equation, where F_s is the sampling rate and τ corresponds to the time constant used:

$$y[n] = \frac{1}{1 + \tau F_s} x[n] + \frac{\tau F_s}{1 + \tau F_s} y[n-1] \tag{2.1}$$

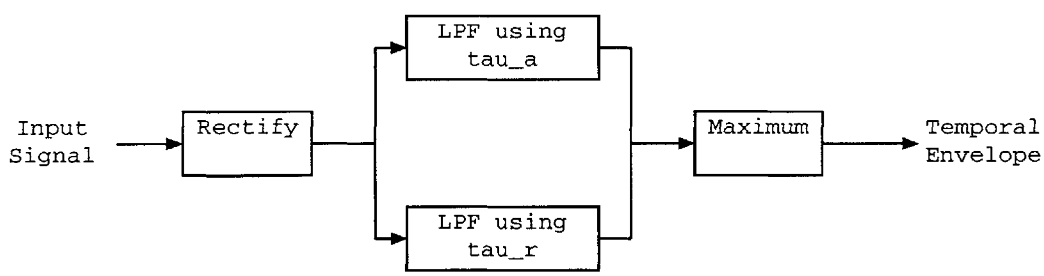

Figure 2-1: Block diagram of signal's temporal envelope calculation

The four compression systems using the Dual Front-End had time constants of $\tau_a = 350$ milliseconds and $\tau_r = 1$ second for the Slow AGC subsystem. Two of the systems, the SNR-Dual and the HT-SNR-Dual (see section 2.2), contained an additional shorter release constant of $\tau_r = 300$ milliseconds for use with the SNR Estimator. All five systems used the same time constant values for the Fast AGC subsystem, with an attack of $\tau_a = 5$ milliseconds and a release of $\tau_r = 75$ milliseconds.

2.1.3 Dual Front-End

The Dual Front-End component was the "brain" of four of the compression algorithms. The Dual Front-End was implemented using the series approach[1], where the output of the Slow AGC was the input to the Fast AGC [8], as shown in Figure 2-2. Compression was applied to the input signal by the Slow AGC, with slowly varying gain over time. The output signal of the Slow AGC was then fed to the Fast AGC system, where its function was to track those levels that exceeded a fixed value, and attenuate them accordingly.

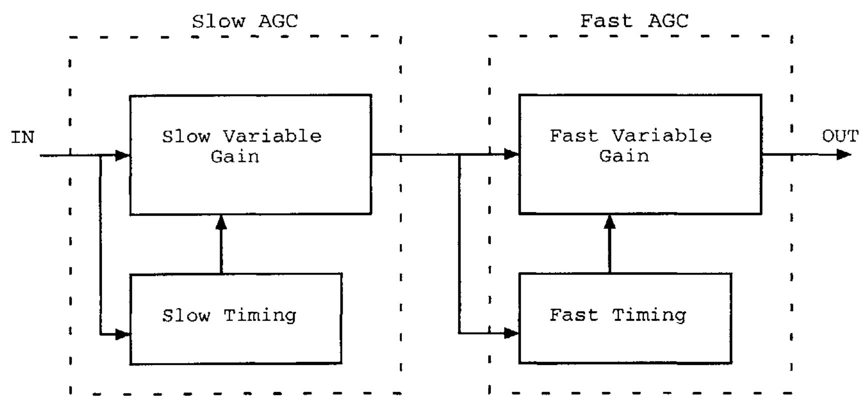

Figure 2-2: Series implementation of the Dual Front-End

Shown in Figure 2-3 is the static input/output compression curve of the Dual Front-End system. Input levels below the Slow AGC threshold, 55 dB SPL, were in the linear region and no compression was applied. The fixed linear gain was set to 0 dB at this stage of the processing, and was later applied using the Cambridge Formula and a fine tuning procedure (see Section 3.2.1). The fixed linear gain causes a vertical shift of the input/output curve, where the output level corresponding to the threshold is equal to the threshold plus the fixed linear gain. For this reason the output level in the figure was left unmarked. A compression ratio of 3 was applied for input levels above the threshold. The fast compressor was triggered for levels above a

[1]The Dual Front-End shown in figure 1-2 is implemented using the parallel approach, where the Slow AGC and the Fast AGC are computed in parallel [6].

fixed value, corresponding to 6 dB above the Slow AGC's output when its input was 95 dB SPL. Since the fixed linear gain was 0 dB at this stage, the fast compressor level was 74.3 dB SPL, calculated from the equation below, where FC is the fast compressor level, $IL_{threshold}$ is the threshold input level, IL_{max} is the maximum input level, HR is the 6 dB headroom, and CR is the compression ratio:

$$FC = IL_{threshold} + \frac{1}{CR}(IL_{max} - IL_{threshold}) + HR \tag{2.2}$$

$$FC = 55 + \tfrac{1}{3}(95 - 55) + 6$$
$$FC = 55 + 13.3 + 6$$
$$FC = 74.3 dBSPL$$

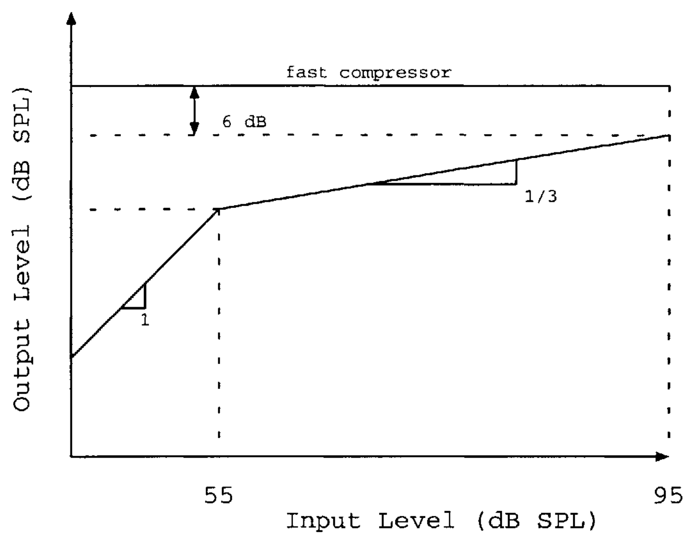

Figure 2-3: The static input/output curve.

Slow AGC of Dual Front-End

A detailed implementation of the Slow AGC subsystem is shown in Figure 2-4. The input signal was first pre-emphasized, as described in section 2.1.1. The out signal was then buffered into 4 millisecond frames, with no overlap. The root-mean-square (RMS) was then calculated for every frame, creating a new signal, called the *RMS* *signal* of the Slow AGC. The temporal envelope of the RMS signal was computed, as described in the previous section, using 350 milliseconds and 1 second for the attack and release time constant, respectively. The temporal envelope of the RMS signal, known as the slow control parameter (SCP), was used as the input to the static input/output slow compressor curve, where a new signal was computed by applying the appropriate gain to every frame of the buffered signal. The resulting frame signal was then fed into the Fast AGC for further processing.

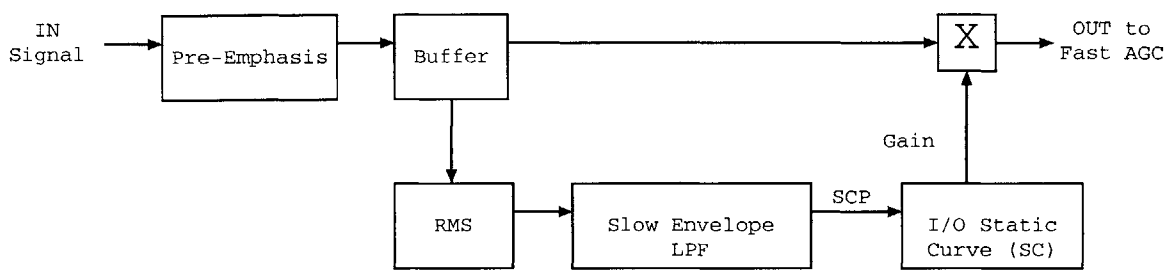

Figure 2-4: Slow AGC subsystem of the Dual Front-End.

Fast AGC of Dual Front-End

The Fast AGC component took the output signal of the Slow AGC as its input, with the exception of the Linear system.[2] The processing of the subsystem was similar to that of the Slow AGC. The input took two paths, one to apply compression limiting, and the other to include a small delay. In the latter, the input frame signal was reconstructed, a 4 millisecond delay was introduced, and once again buffered into frames. Such delay was sufficient to compensate for the attack time when calculating the fast control parameter (FCP). No delay was necessary for the Slow AGC since the Dual Front-End was implemented in series. The Slow AGC was meant to provide a slowly-varying gain, while the Fast AGC tracked high-level sounds that the Slow AGC was "too slow" to appropriately attenuate, caused at moments of attack. On the second path of the Fast AGC, the RMS was calculated and its temporal envelope was

[2]In the Linear system the input signal was simply pre-emphasized, buffered into frames, and passed to the Fast AGC.

computed using an attack of τ_a=5 milliseconds and a release of τ_r=75 milliseconds. The resulting FCP was fed as the input to the fast compressor, attenuating levels above FC = 74.3 dB SPL to FC. The output of the fast compressor was the gain signal to be applied to the buffered signal. Before multiplication of the gain and the buffered signal, a 4 millisecond zero padding was applied at the end of the gain signal to assure equal length vectors. Finally, the multiplied frame signal was reconstructed to a time signal.

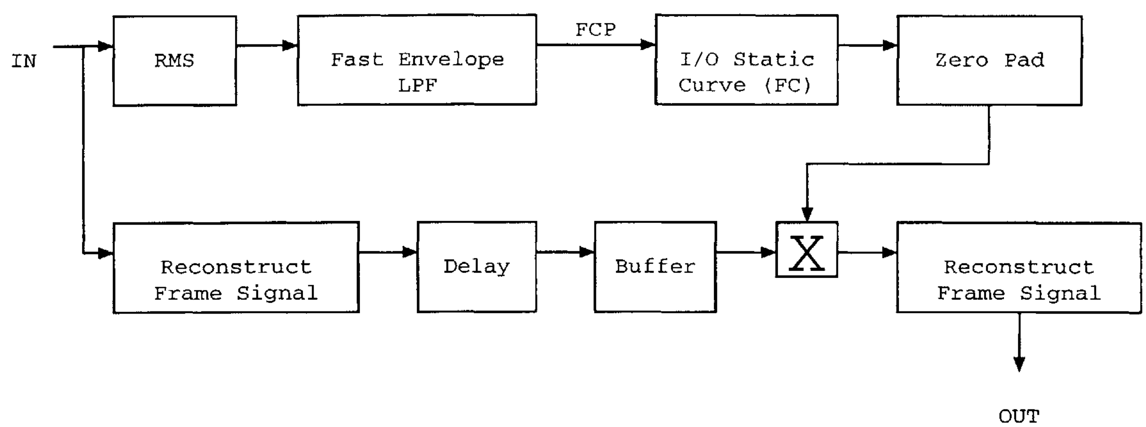

Figure 2-5: Fast AGC subsystem of the Dual Front-End.

2.1.4 Hold Timer/Counter (HT)

The Hold Timer (HT) proposed by Stone *et al.* [8] is a feature integrated into the Dual Front-End AGC system to reduce the effects of pumping without the long recovery times. The HT consisted of a counter with a minimum and maximum value of 0 and 600, respectively. The basic idea was to increment the HT counter whenever the slow control parameter (SCP) increased, and decrement the HT counter otherwise. The SCP was only allowed to decrease when the HT counter value was zero. In theory, the HT is supposed to hold the SCP at a constant level, thus keeping the system's gain constant, for 600 milliseconds.

The gain of the Slow AGC was always calculated directly from the value of the SCP using the static input/output compression curve, thus the SCP was modified by

Table 2.1: Three cases to determine the Hold Timer and the Slow Control Parameter

Case	Action
SCP < CFRMS **AND** CFRMS $\leq$ SCP + 12dB	HT = HT + 8 SCP = exp. average of CFRMS with previous SCP
CFRMS > SCP + 12dB	HT held constant SCP = exp. average of previous (SCP + 12dB) with previous SCP
CFRMS $\leq$ SCP	HT = HT - 4 If HT = 0, SCP = exp. average of CFRMS with previous SCP Else, SCP held constant.

the HT before compression was applied. In every frame, the value of the HT counter and of the modified slow control parameter was computed in one of the following three ways, as summarized in Table 2.1, depending on the relationship between the SCP and the current frame RMS (CFRMS) value.

The first scenario occurred when the SCP was meant to increase without a loud transient being present. This happened when the current frame value of the RMS signal (CFRMS) was greater than the current SCP value, but not by more than 12 dB. Under this condition, the HT was incremented by 8, corresponding to a rate of 2 units per millisecond since each frame was 4 milliseconds. An additional check was placed to make sure the HT did not increase above 600. The next value of the SCP was calculated by exponential averaging the next CFRMS value with the current SCP value, as described in the following difference equation:

$$SCP[n] = (\frac{1}{1+\tau_a F_s})CFRMS[n] + (\frac{\tau_a F_s}{1+\tau_a F_s})SCP[n-1] \tag{2.3}$$

The second case occurred when the CFRMS exceeded the value of the SCP by more than 12 dB, as expected in a loud transient. In such case the HT was not allowed to increase while the transient was present, thus it was held constant. The SCP was calculated by exponential averaging the previous SCP value plus 12 dB with the previous SCP value, as shown in difference equation 2.4:

$$SCP[n] = (\frac{1}{1+\tau_a F_s})(SCP[n-1] + 12dB) + (\frac{\tau_a F_s}{1+\tau_a F_s})SCP[n-1] \tag{2.4}$$

The last scenario consisted of the HT's discrimination between holding constant and releasing the SCP. If the CFRMS was less than or equal to the SCP, then the HT was decremented by 4, or a rate of one unit per millisecond. An additional check was incorporated to guarantee that the HT did not reach values below zero. The SCP was released only when the HT counter value was zero, with exponential averaging of the CFRMS and the previous SCP value using the release time constant, as described by the difference equation 2.5:

$$SCP[n] = (\frac{1}{1+\tau_r F_s})CFRMS[n] + (\frac{\tau_r F_s}{1+\tau_r F_s})SCP[n-1] \tag{2.5}$$

In this scenario, the value of the SCP was held constant for any frame with a non-zero value of the HT counter. Assuming the system started to release when the HT was at its maximum, the SCP was held constant for a maximum of 600 milliseconds (i.e. 150 frames). Since the rate of increase was twice the rate of decrease of the HT, a bias was created to maintain the HT at a maximum level, thus holding for the full 600 milliseconds in pauses between speech.

2.1.5 Signal to Noise Ratio (SNR) Estimator

Much research has been undertaken to develop systems that accurately estimate signal to noise ratio (SNR) levels, but little success has been found. Since developing such SNR techniques was not the focus of this research, an ideal SNR Estimator was used in two of the five AGC systems; the SNR-Dual and the HT-SNR-Dual. As shown in Figure 2-6, the ideal SNR Estimator took two input signals, a clean speech signal and a noise signal.[3] RMS signals were calculated for each, using a buffering of 4 millisecond

[3]For interference conditions of clean speech only, the noise input was made of zeros, thus providing a high SNR when comparing the two input signals.

frames. The two RMS signals were subtracted, resulting in a highly accurate SNR signal. It should be noted that implementing such an ideal SNR Estimator is only possible in computer simulation. Separate clean speech and noise signals are not typically available in a practical system.

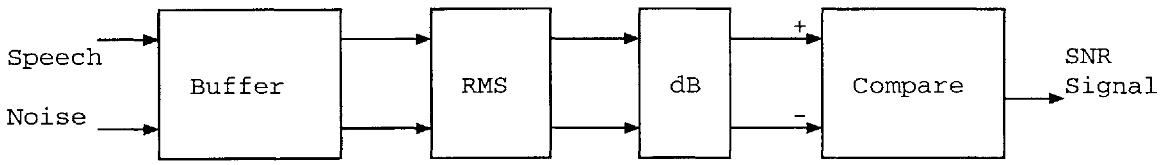

Figure 2-6: SNR Estimator.

The "golden" SNR signal was used to determine which of two release time constants would be used at every frame of the AGC's recovery. If at a frame of release the SNR level was below 0 dB, it was assumed the release was due to low-level background noise in the absence of speech (i.e. pause in speech). In this case, the long recovery constant of one second was provided to the AGC to prevent the system's gain from increasing significantly. If on the other hand, the SNR level was 0 dB or greater at a moment of release, it was assumed to be caused by lower-level speech, such as the transition from the aid user's voice to another talker nearby. Under such circumstance, a fast recovery of 350 milliseconds was provided to the AGC, at an aim to quickly increase and adjust the system's gain.

2.2 Compression Systems

Four AGC compression systems were implemented using different combinations of the components described in section 2.1, plus a fifth linear amplification system used for reference. All implementation was conducted in the MatLab programming language, and the source code can be found in Appendix A. A sampling rate of 11,025 Hz was used for all five AGC systems.

2.2.1 Linear System

The *Linear* system was composed of fixed linear amplification and was designed to be used as a reference algorithm. Input signals were fed straight to the Fast AGC subsystem of the Dual Front-End for limiting compression. Any levels above the fast compressor (FC) level, 74.3 dB SPL, were attenuated accordingly down to the FC level. The fixed, subject-dependent overall gain was later applied using a subject fitting procedure (see subject fitting, section 3.2.1).

2.2.2 Dual Front-End System

The first AGC compression system was composed of the *Dual Front-End* by itself, with the Hold Timer (HT) and the SNR Estimator components deactivated. Input signals to the Dual Front-End were first processed by the Slow AGC, as in figure 2-4, followed by the Fast AGC, implemented as shown in figure 2-5.

An artificial noise signal was created in order to demonstrate the effects of the Dual Front-End. The signal was six seconds in duration, and contained a one second pause and one half of a second loud transient in it. The RMS signal of the noise, shown in figure 2-7(a), provides a view of the temporal envelope. It was calculated by buffering the signal into 4 millisecond frames, and taking the root-mean-square of each frame. The superimposed dashed plot corresponds to the calculated slow control parameter (SCP) used as the input signal to the input/output slow compression curve. The RMS signal was re-calculated after the gain from the Slow AGC was applied, resulting in the solid line plot of figure 2-7(b). Note that the Slow AGC did not provide enough attenuation for the loud transient occurring at 4 seconds. The fast control parameter (FCP) shown as the dashed plot in figure 2-7(b) was fast enough to track the level above the fast compressor threshold, and attenuated the signal accordingly. The applied attenuation of the Fast AGC is shown in figure 2-7(c), which ensured that no part of the signal was presented above the fast compressor threshold.

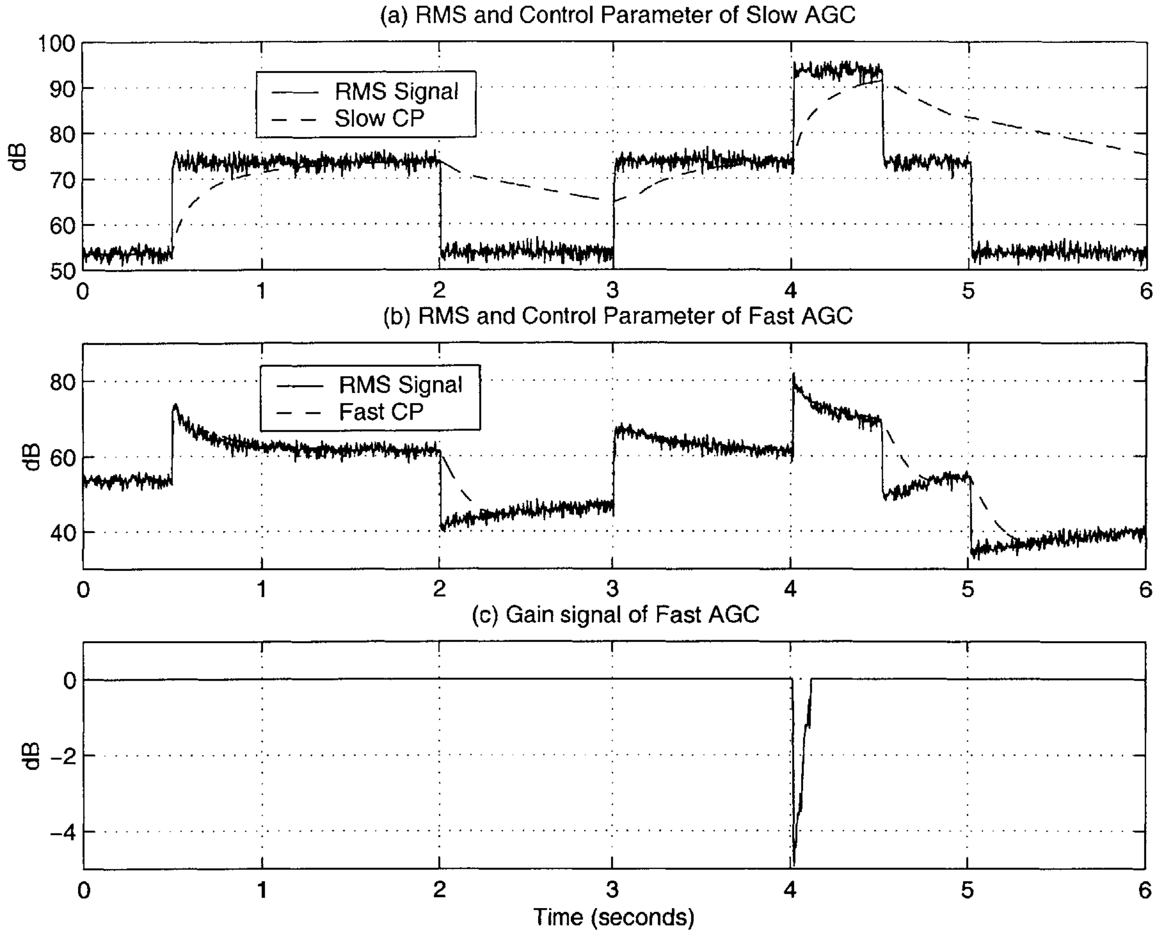

Figure 2-7: Effects of Slow and Fast AGC on a signal.

2.2.3 HT-Dual System

The *HT-Dual* AGC compression system was composed of the Dual Front-End with the Hold Timer. The behavior of the HT-Dual was similar to the Dual Front-End, with the exception that release was dependent on the value of the counter from the HT.

A seven-second duration artificial noise signal was created to show the effects of the HT-Dual system, and how it compares to the Dual Front-End. The noise's RMS signal is shown as the dash-dotted plot in figure 2-8(a). First note the superimposed dashed plot which corresponds to the SCP of the Dual Front-End. It attacked at 1 second, then released starting at 3 seconds when it detected a level drop. It continued to attack then release as expected, at 4 and 5.5 seconds, respectively. In contrast, the solid plot corresponds to the SCP calculated from the HT-Dual system and the HT counter, shown in figures 2-8(a) and 2-8(b), respectively. At time: (0) The SCP was flat and at a low-level. The HT counter gradually increased with fluctuations due to the random values of the RMS signal.[4] Recall that the increase rate was twice the decay rate in the HT counter; (1) The SCP started to attack, detecting a level increase. By this point, the HT counter was at its maximum, and it stayed at a maximum, given the SCP was attacking; (3) A level decrease was detected, but the SCP did not release since the HT counter was not at its minimum. The HT counter decreased at a constant rate of one value per millisecond; (3.6) The HT counter reached its minimum of zero, and the SCP began to release after being held constant for 600 milliseconds; (4-7) The process was repeated. The HT counter recharged to its maximum level while the SCP attacked to reach the new higher-level sound. The HT once again prevented the SCP from recovering at the 5.5 second time mark, until the the HT counter reached its minimum value 600 milliseconds later.

The above scenario showed an ideal case of the effects of the HT. The HT-Dual held the SCP constant for the full intended 600 milliseconds between pauses in the noise signal. However, in the presence of real speech, the HT-Dual system did not necessarily provide the full 600 millisecond hold for all pauses. Level variations due to the acoustic elements in speech produce frequent attacks and releases which in turn have a direct effect on the HT counter. Figure 2-9 shows plots of the same variables as in figure 2-8, but calculated using speech rather than an artificial noise signal. The speech was made of two sentences, separated by a pause of about 1400 milliseconds. Notice the level variations of about 30 dB between syllables in the RMS signal. This caused the HT counter to fluctuate significantly, preventing the HT from being at its maximum at the moment the pause was reached. As a consequence, the HT was not able to provide the full hold for the pause at the 3-second mark, but only about 300 milliseconds, or half the expected pause. At the end of the second sentence (at about

[4]The counter values increased and decreased somewhat randomly, causing an overall slower rate of increase than expected.

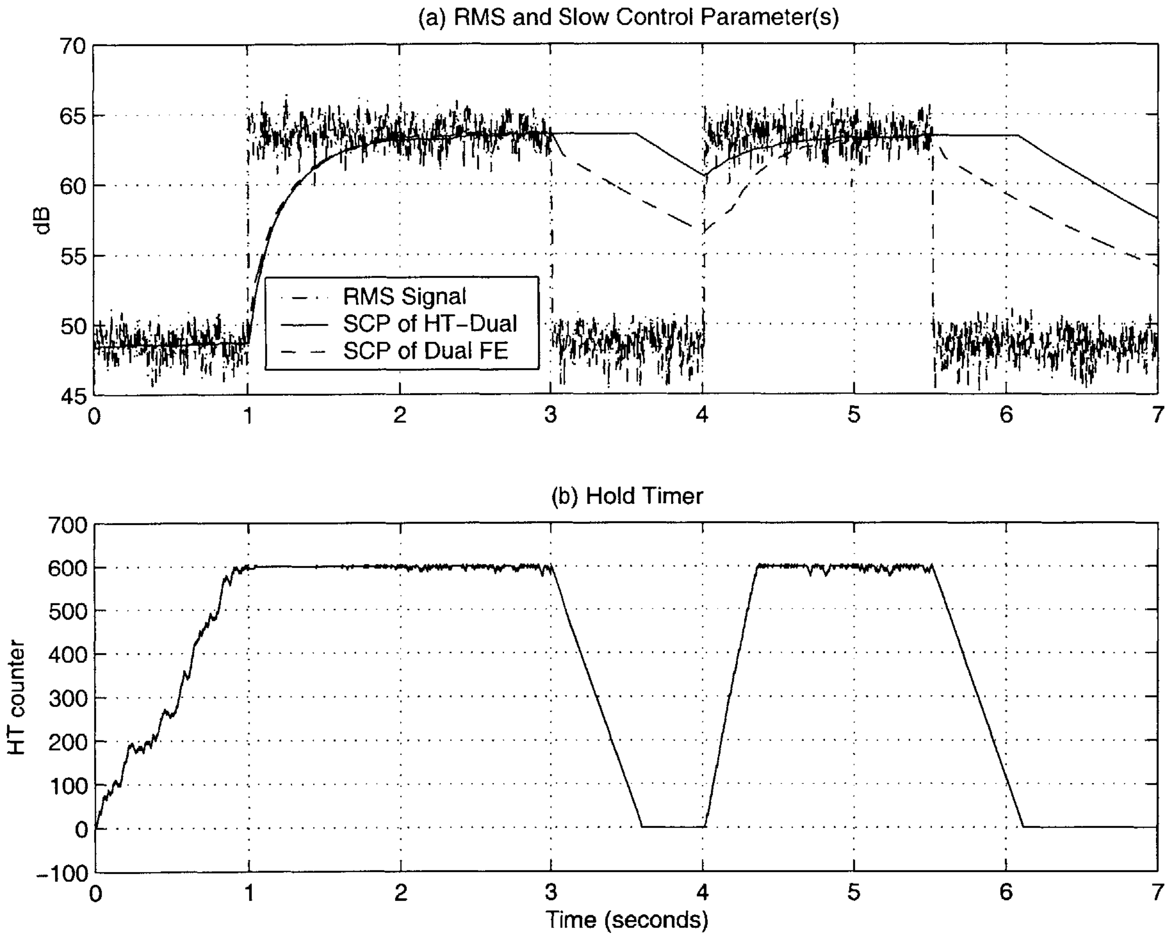

Figure 2-8: RMS, SCPs, and HT counter of a noise signal under HT-Dual.

the 6 seconds), the HT counter was at its maximum, and it did provide the full hold for the upcoming pause. Clearly, the hold time provided to a pause in speech by the HT depended on the acoustic characteristics of the speech. This "random hold" of pauses was the major disadvantage found in this implementation of the HT.

On the other hand, the design of the HT affected the SCP not only in pauses of speech, but within speech as well. Even though the HT counter fluctuated significantly within speech, it was, most of the time, at a non-zero value. This prevented release of the SCP in the valleys found between syllables, as desired. Looking at the first sentence in figure 2-9(a), it can be noted that the SCP of the HT-Dual AGC is nearly flat, providing a constant gain by the Slow AGC. In contrast, the SCP of the

Dual Front-End contains undesired level variations within the speech of up to 5 dB.

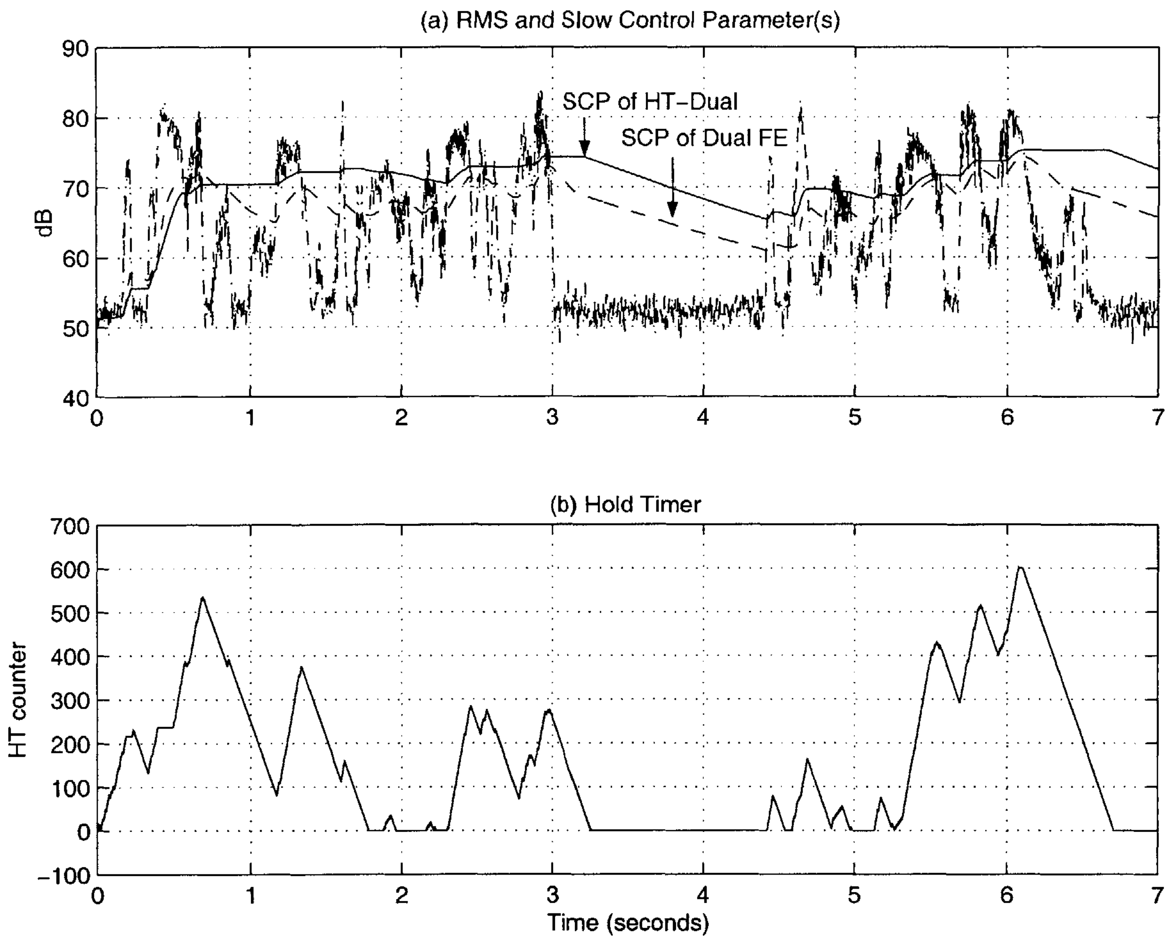

Figure 2-9: RMS, SCPs, and HT counter of a speech signal under HT-Dual.

2.2.4 SNR-Dual System

A third combination AGC system, named the *SNR-Dual*, was composed of the Dual Front-End plus the SNR Estimator. The included Estimator feature made use of the calculated SNR level at moments of release, and provided either fast or slow recovery if the lower-level sound was speech or noise, respectively. If the lower-level sound had an SNR below 0 dB, the SNR-Dual behaved exactly as the Dual Front-End, providing a release time τ_r of one second. Otherwise, a fast release of 300 milliseconds was used.

An artificial signal, similar to the noise signal created in the HT-Dual section, was used as a speech signal to demonstrate the effect of the SNR-Dual. Shown in figure 2-10(a) is the calculated RMS, which for demonstration purposes, it is assumed two talkers are present. The higher-level speech comes from the aid user, while the lower-level speech (about 60 dB) corresponds to the other converser. Very low-level background noise is assumed to be present so that the SNR level is always well above 0 dB. Superimposed in the figure are the solid and dashed plots corresponding to the SCPs of the SNR-Dual and the Dual Front-End system, respectively. At moments of release (i.e. 3 and 5.5 second marks), the SNR-Dual AGC assumes the presence of speech and provides the faster recovery. The system as a result adapts relatively quick to the low-level speech and provides adequate gain, as shown in figure 2-10(b).

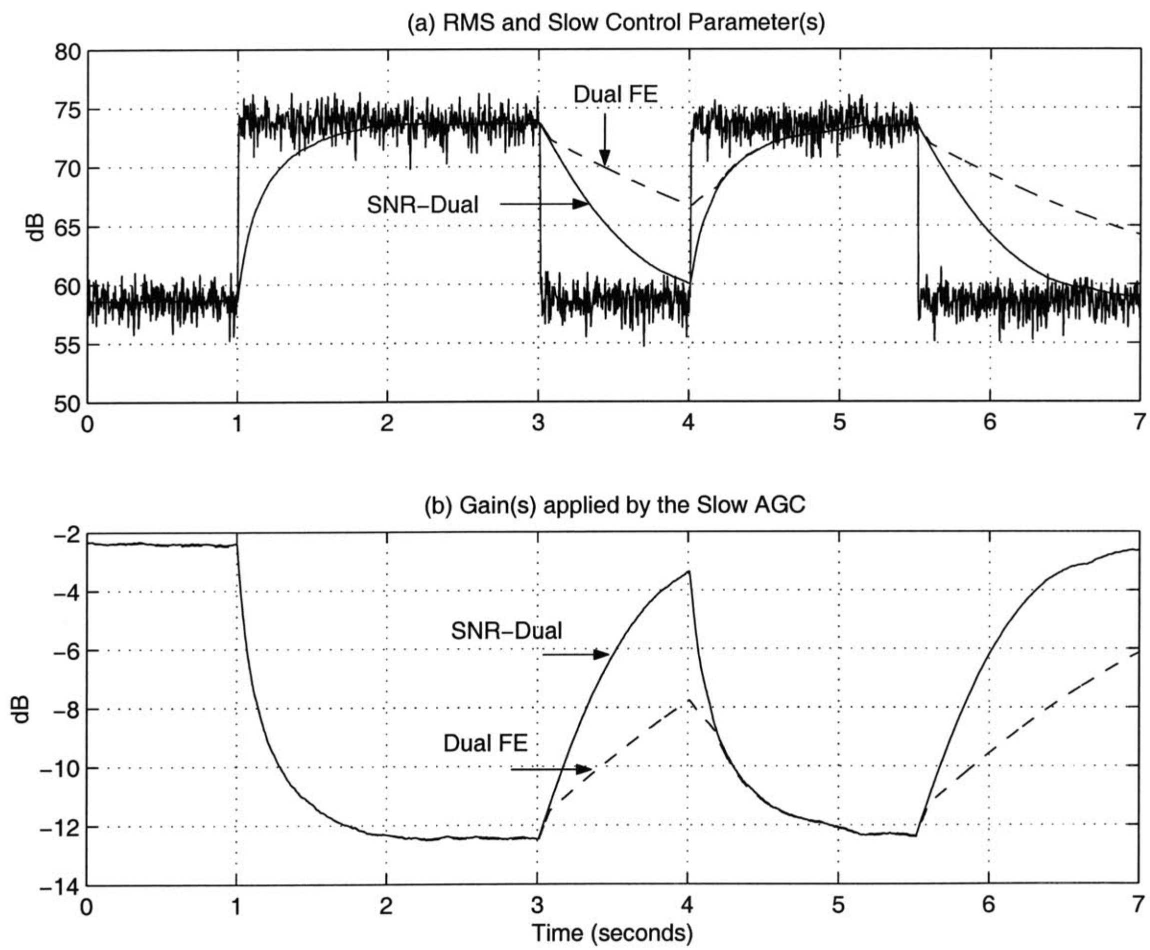

Figure 2-10: RMS, SCPs, and Slow AGC gains of a signal under SNR-Dual.

2.2.5 HT-SNR-Dual System

The final compression system was given the name of the *HT-SNR-Dual* AGC. Revealed by its name, the HT-SNR-Dual was a combination of the Dual Front-End, the Hold Timer, and the SNR Estimator. Figure 2-11 shows the same artificial noise used with the SNR-Dual to be used as a speech signal for demonstration of the HT-SNR-Dual AGC. The SNR level of the entire signal is assumed to be above the threshold. The SCP of the HT-SNR-Dual, shown as the solid plot, detects a decrease in level at the 3 seconds. Since the HT counter is at its maximum it prevents the SCP from releasing, thus holds it constant for 600 milliseconds. By the time the HT counter reaches its minimum value, the SCP is released at a relatively fast rate since the SNR level is above 0 dB. Clearly, under this situation where two talkers are presented at different levels, the SNR Estimator and the HT contradict each other. At the moments of release the HT holds the gain constant while the SNR Estimator changes the gain significantly. On the other hand, it may be argued that a maximum hold of 600 milliseconds followed by release at a fast rate is still an overall faster recovery than the use of a longer release constant. From the figure it is shown that it takes less than a second for the SCP of the HT-SNR-Dual to surpass the SCP of the Dual Front-End under slow parameters of τ_s = 300 and 1000 milliseconds, respectively.

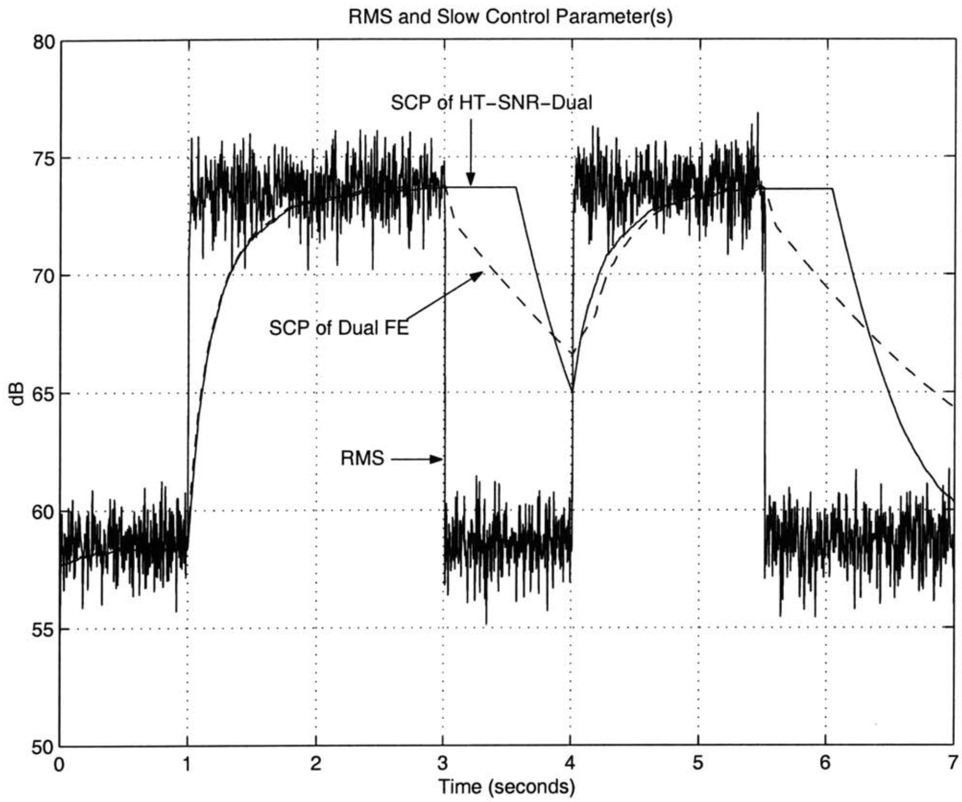

Figure 2-11: RMS and SCPs of a signal under the HT-SNR-Dual.

Chapter 3

Experimental Methods

The stimuli presented to the subjects were processed through various stages before reaching their ears. The stimuli were generated as described in section 3.1. The five systems (four AGC compression algorithms plus the linear reference system) were implemented in MatLab. All stimuli were pre-processed by the five algorithms and saved on CDROMs. During the actual experiment, subject-dependent gains were applied to the pre-processed stimuli using the MatLab programming language in a PC, as described in section 3.2.1. Finally, the analog audio output of the computer was passed through a transformer, followed by an attenuator, then to the earphones used, described in section 3.3.

All sounds were digital recordings, originally sampled at 44.1 kHz and downsampled to 11.025 kHz before any processing was preformed. MatLab was used to present stimuli at a sampling rate of 11.025 kHz via the computer's sound card.

3.1 Interference Conditions - Stimuli

The key components of the AGC compression systems included the Hold Timer, designed to reduce pumping effects; the SNR Estimator, used for adaptive release depending on whether low-level sound was speech or background noise; and the Dual Front-End which provided protection from loud transients and preserved speech intelligibility by maintaining a slowly-varying gain over time. Twenty interference conditions were carefully designed to activate all the various system components and provide information whether such component(s) showed any benefit. The twenty interference conditions are summarized in Table 3.1. All speech used was taken from *The American Presidents* CD, by Kunhardt *et al.*, and narrated by Richard Neustadt. Clean speech segments ranging from 15 to 25 seconds in duration were extracted from the CD, making sure that each sentence began and ended on sentence boundaries and that no music or sound effects were included.

Interference conditions 1-3, denoted by $I_1 - I_3$, were 20 second passages composed of speech with background multi-talker babble at different signal-to-noise ratio (SNR) levels. The speech level for all three conditions was set to 65 dB SPL, and the babble level was digitally adjusted to provide the appropriate SNR.

Conditions $I_4 - I_8$ were composed of clean speech alone presented at different intensity levels. The speech segments contained pauses, ranging from 0.5 seconds to 1.5 seconds. Such pauses were intended to test the effect of the Hold Timer. Speech passages ranged from 15 to 20 seconds each.

The next set of interference conditions, $I_9 - I_{10}$, presented speech alone at alternating intensity levels. The original speech utterance was digitally modified so that the first four seconds were presented at a low level, denoted by L. The sound level then rose at a logarithmic constant rate for half of a second, reaching a high level, H. Another four seconds were presented at H, followed by a 0.5 second constant decrease, reaching L. This semi-square wave period of 9 seconds was replicated and applied to the rest of the speech utterance, which ranged from 20 to 25 seconds in duration. Condition I_9 had a low level L of 50 dB SPL, and an H value of 65 dB SPL, thus alternating between 50 and 65 dB SPL. Condition I_{10} contained values of $L = 65$ dB SPL and $H = 80$ dB SPL.

Table 3.1: Description of Interference Conditions

Condition	Description
1.	Speech @ 65 dB SPL plus Babble: SNR = 0 dB
2.	Speech @ 65 dB SPL plus Babble: SNR = -6 dB
3.	Speech @ 65 dB SPL plus Babble: SNR = -12 dB
4.	Speech alone with pauses @ 50 dB SPL
5.	Speech alone with pauses @ 58 dB SPL
6.	Speech alone with pauses @ 65 dB SPL
7.	Speech alone with pauses @ 72 dB SPL
8.	Speech alone with pauses @ 80 dB SPL
9.	Speech alone with alternating levels: 50-65 dB SPL
10.	Speech alone with alternating levels: 65-80 dB SPL
11.	Speech plus Environmental Noise (City Noises)
12.	Speech plus Environmental Noise (Hair Dryer)
13.	Speech plus Environmental Noise (Power Drill)
14.	Speech plus Environmental Noise (Water Running in Sink)
15.	Speech plus Environmental Noise (Vacuum)
16.	Speech plus Environmental Noise (Air Conditioner)
17.	Speech plus Environmental Transients (Nails into Wood)
18.	Speech plus Environmental Transients (Glass Breaking)
19.	Speech plus Environmental Transients (Wood Dropping)
20.	Speech plus Environmental Transients (Dog Barking)

The following set of interferences presented speech with semi-steady-state environmental background noises. For all conditions, $I_{11} - I_{16}$, both speech and the background noise were set to 65 dB SPL each. Preliminary listening of these passages at +0 dB SNR was undertaken to make sure that speech intelligibility was adequately affected, without completely masking out the speech. Passage duration ranged from 20 to 25 seconds.

The last set of interference conditions, $I_{17} - I_{20}$, presented speech at 65 dB SPL plus loud environmental transient sounds. Passages ranged from 15 to 20 seconds in duration. The environmental sounds were digitally adjusted so that their peak levels were 90 dB SPL, high enough to activate the Fast AGC in the Dual Front-End system. The environmental steady-state and transient sounds used for background noise in conditions $I_{11} - I_{20}$ were taken from *Living Sound Effects*, volumes 1-4, produced by Records Bainbridge.

3.2 Software Processing

All the software processing, including the AGC compression systems were implemented in the MatLab programming language. The stimuli presented to the subject were processed through two main components; compression then subject fitting, as shown in Figure 3-1.

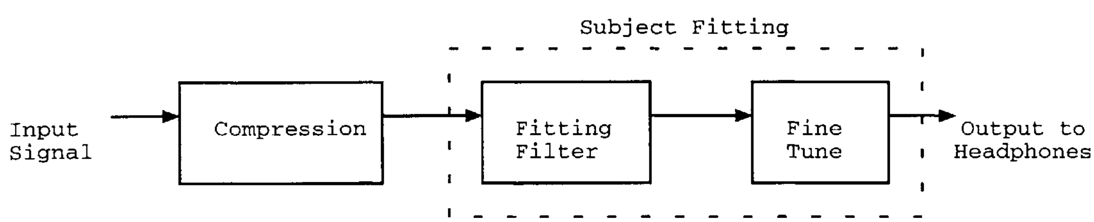

Figure 3-1: Three stage process of stimuli.

The compression box corresponded to the pre-processing of the stimuli by the five AGC compression systems. All subjects received equal compression, thus all stimuli were processed once by each system, and saved. The subject fitting procedure,

Table 3.2: Values of the intercept (INT) in the Cambridge Formula.

Frequency (kHz)	0.125	0.25	0.5	0.75	1.0	1.5	2.0	3.0	4.0	5.0
INT	-11	-10	-8	-6	0	-1	1	-1	0	1

composed of a frequency-dependent filter and a flat fine tuning gain, was the only difference between subjects. Its purpose was to apply adequate frequency-dependent gain to compensate for the subject's hearing loss.

3.2.1 Subject Fitting

Hearing-loss is frequency-dependent, thus it is essential to apply frequency-dependent gain in the attempt to compensate for the hearing loss. This is what is generally referred to as subject fitting. There were two parts of the fitting procedure, an initial fitting based on the audiogram of the subject, and an adaptive fine-tuning procedure. A single fitting procedure was done for every subject, and applied in all five AGC compression systems.

The initial fitting using the audiogram was achieved using the Cambridge Formula, which was developed using the loudness model proposed by Moore and Glasberg [7]. The formula to calculate the frequency-dependent initial gain (IG) is shown in Equation 3.1, where HL is the absolute threshold in dB HL and INT is a frequency-dependent intercept. The INT values for given audiometric frequencies are shown in Table 3.2.

$$IG = HL * 0.48 + INT \tag{3.1}$$

The computed IG gain values at the audiometric frequencies were linearly interpolated to produce a value at every 100 Hz, evenly spaced out. The interpolated data was then used to create a 65-point linear-phase finite impulse response (FIR) filter whose magnitude response was a weighted least-squares approximation [to the data].

The second part of the fitting procedure was the fine tuning, conducted exactly as the experiment by Stone *et al.* [8]. The goal of the fine tuning procedure was to adjust the overall wideband gain so that input speech to the aid at 85 dB SPL was presented to the subject as their "highest comfortable level". Ten clean speech segments at 85 dB SPL, of 15 to 20 seconds in duration each, were processed by the HT-Dual algorithm and the subject's fitting filter determined in the first part of the fitting procedure. The stimuli were presented to the subject one at a time. After each presentation the subject rated the loudness, and then the overall gain was adaptively modified for the next speech segment depending on the subject's response. The following menu was presented to the subject on the computer screen in front of him/her:

"Please judge the loudness of the sentence you just heard as compared to the HIGHEST volume that you would be comfortable listening to for a long time:

7. Far too loud
6. Much louder than I like
5. Somewhat louder than I like
4. The highest volume I like
3. Somewhat softer than the highest volume I like
2. Definitely softer than the highest volume I like
1. Far too soft"

The initial fine tuning gain was set to 0 dB, thus the first speech segment of this procedure was presented at 85 dB SPL with compression plus the gain applied from the subject's fitting filter. The overall gain, or fine tuning gain, was adjusted according to the following rules:

If 7 was pressed, the fine tune gain was decreased by 4 dB.
If 6 was pressed, the fine tune gain was decreased by 2 dB.
If 5 was pressed, the fine tune gain was decreased by 1 dB.
If 4 was pressed, the fine tune gain was left unchanged.
If 3 was pressed, the fine tune gain was increased by 1 dB.
If 2 was pressed, the fine tune gain was increased by 2 dB.

If 1 was pressed, the fine tune gain was increased by 4 dB.

This process was repeated, using a randomly selected speech segment each time, until the subject responded 4 (the target response) twice in succession.

3.3 Hardware Setup

Stimuli was presented to the subject monaurally over a set of Telephonics TDH-39P earphones in a soundproof booth. The analog audio output of the computer fed into a set of transformers and attenuators for isolation, and to provide the desired SPL at the output of the earphones. The computer's keyboard was inside the booth, but the computer and monitor were placed outside the booth to avoid fan noise. The booth contained a window were the subject was able to see the monitor and type his/her ratings. Sound level calibration was made before subjects were tested using a General Radio Company, type 1565-A sound level meter. Various frequency sine waves, at different levels each, were measured directly from the speaker of the earphones to verify the correct output levels. The sound level meter was set to the C_s weighting characteristic position, which had a nearly flat frequency response at 0 dB. Sound level verification was also done at the beginning of every session using an oscilloscope and listening to the tone, simultaneously.

3.4 Experimental Design

For each of the twenty interference conditions presented in Section 3.1, five stimuli were made, denoted by $S_1 - S_5$. Each combination signal I_jS_k[1] contained a different speech segment, totaling to 100 different stimuli. Each I_jS_k signal was pre-processed by the five compression systems, producing 500 stimuli to be presented to each subject, denoted by $A_iI_jS_k$, where A_i is one of the five algorithms. The signals were split into five sets of 100 passages, where one set was presented for a given session.

[1] The value I_j denotes the 20 interference conditions, and S_k denotes the five stimuli per condition.

Table 3.3: Pre-processed stimuli for each of the five sets

Set	1	2	3	4	5
	A_1S_1	A_1S_2	A_1S_3	A_1S_4	A_1S_5
Five signals	A_2S_2	A_1S_3	A_1S_4	A_1S_5	A_1S_1
per Interference	A_3S_3	A_1S_4	A_1S_5	A_1S_1	A_1S_2
	A_4S_4	A_1S_5	A_1S_1	A_1S_2	A_1S_3
	A_5S_5	A_1S_1	A_1S_2	A_1S_3	A_1S_4

Each set contained five stimuli for each of the 20 interference conditions. For a given interference condition, the five stimuli were pre-processed by a different compression system and contained different speech segments. Table 3.3 shows the stimuli for the five sets. Each set of 100 pre-processed passages was recorded into a CDROM, to allow for random order of set presentation among subjects.

Presentation of stimuli to the subject was structurally randomized. For each session, a different set was selected. The set order was different for each subject. Within a session, the order of the 20 interference conditions were randomized [by MatLab]. For a given interference, the five corresponding passages A_iS_k of Table 3.3 were presented together, but also in a random order.

3.5 Rating of Intelligibility and Quality

At the beginning of each experimental session[2], the subject was presented with the instructions for rating intelligibility and quality of the speech passages he/she was about to listen to. The rating instructions are found in Appendix B. For each passage presented, the subject was asked to subjectively rate the intelligibility of speech

[2]For the first session, rating began after the subject was done with the fine tuning procedure.

and the overall quality of the sound. Ratings of intelligibility were between 0 and 10, where 10 meant the subject was able to understand everything that was said, regardless of the annoyance of background noise. The following intelligibility rating menu was presented to the subject in front of the computer screen after listening to every speech passage:

"Intelligibility Rating	Percent of Words Understood
10	ALL
9	
8	80%
7	
6	60%
5	
4	40%
3	
2	20%
1	
0	NONE

Rate the INTELLIGIBILITY of the speech passage you just heard, then hit ENTER:"

Ratings of sound quality were also from 0 to 10, where 10 indicated that the overall sound quality was extremely pleasant; when the speech was clear and the background noise was easily ignored. A rating of 0 indicated that the overall sound quality was extremely annoying; when the background noise was extremely distracting and unpleasant, or the speech was very distorted, much too loud, or much too soft. The following quality menu was presented after every passage, following the intelligibility rating menu:

"Quality Rating	Overall Quality
10	

9	Very Pleasant
8	
7	Somewhat Pleasant
6	
5	Neutral
4	
3	Rather Annoying
2	
1	Very Annoying
0	

Rate the QUALITY of the sound you just heard, then hit ENTER:"

3.6 Hearing-Impaired Subjects

Three hearing-impaired subjects (PG, MG, RG), with a minimum of five years experience wearing hearing aids, participated in the experiment. The hearing loss (HL) levels of the subjects are summarized in Table 3.4. The three subjects had an almost bilaterally symmetric sensorineural hearing loss, with little difference between their two ears. The "better" ear was chosen for all subjects. Subject PG had a flat moderately-severe hearing loss, and MG had a sloping moderate to severe HL. Subject RG had a moderate HL in the mid-range frequencies, recovering to normal in the low frequencies and at 4 kHz, and sloping to moderate at frequencies above 4 kHz. After submission of the subject's HL levels to the Cambridge Formula (equation 3.1), the resulting gain values at the audiometric frequencies were interpolated to create the subjects' fitting filters, shown in figure 3-2. As for the fine tuning procedure, subject PG reached his most comfortable level (MCL) with 19 dB, subject MG with 4 dB, and RG with 20 dB of fine tune gain.

Table 3.4: Hearing Loss Levels of Subjects

	Hearing Loss Level (dB)		
Frequency (Hz)	PG	MG	RG
250	65	45	15
500	65	35	20
1000	70	50	45
1500	70	55	45
2000	65	55	45
3000	70	60	35
4000	70	65	20
Ear	R	R	R
Sex	M	M	F
Age	61	56	76
Fine Tune Gain (dB)	19	4	20

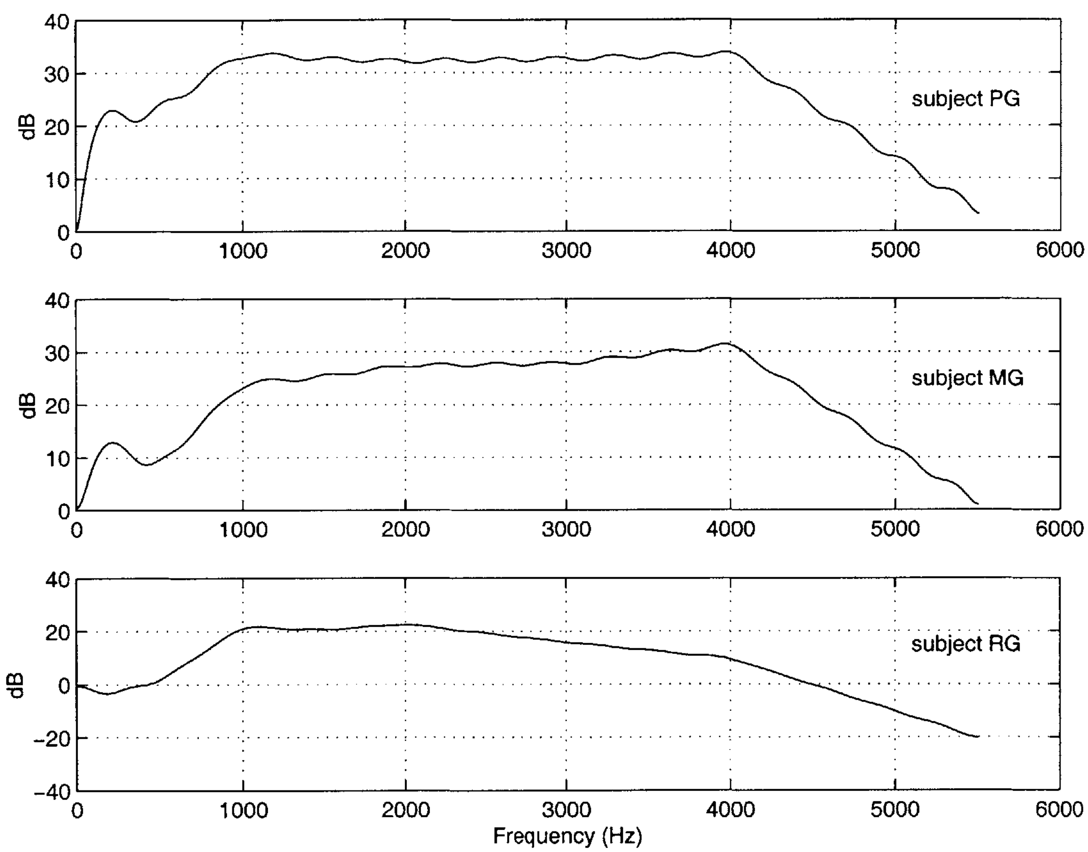

Figure 3-2: Fitting filters of hearing-impaired subjects

Chapter 4

Experimental Results

4.1 Average Intelligibility and Quality

The raw data of intelligibility and quality ratings from the three hearing-impaired subjects (PG, MG, RG) are contained in Appendix C. The mean rating value was calculated, for both intelligibility and quality, for every five passages presented per interference condition, per algorithm, as described in the following equation:

$$\overline{A_i I_j} = \frac{1}{5}\sum_{k=1}^{5} A_i I_j S_k, \qquad\qquad i = 1, ..., 5, j = 1, ..., 20 \qquad (4.1)$$

where $\overline{A_i I_j}$ is the average rating of algorithm A_i and interference condition I_j, and $A_i I_j S_k$ is the individual rating given to segment S_k of algorithm A_i, interference I_j. The average rating values are plotted in figures (4-1) - (4-3). The error bars correspond to the minimum and maximum individual ratings of the five segments. The left and right column plots show the intelligibility and quality ratings, respectively. The five plots per rating were organized into groups with relevant interference conditions; (1) speech plus multi-talker babble at different SNR levels, (2) speech alone at different SPL levels, with pauses, (3) speech alone with alternating SPL levels, (4) speech plus environmental steady-state sounds, and (5) speech plus environmental loud transient sounds.

These plots show that in general, both ratings increase with increases in SNR or in speech level. For speech plus babble, $I_1 - I_3$, the three subjects show higher ratings of intelligibility and better quality for higher SNR levels in the five systems, as expected. Conditions $I_4 - I_8$, speech alone with pauses, reveal that higher levels of SPL increase intelligibility and improve quality for subject MG, in four of the five systems. Subjects PG and RG, on the other hand, show no noticeable difference in intelligibility for the different SPL levels. The difference between increase in intelligibility as SPL increases for MG, and equal intelligibility for PG and RG may be due to the significant gap of fine tune gains between the subjects. While MG had 4 dB of fine tune gain, PG was provided with 19 dB, and RG with 20 dB of gain. The smaller gain of MG is the only clear indicator that shows the effect of intelligibility difference at different SPL levels, as compared to PG and RG.

At the end of the first experimental session, subject RG expressed that the "wood dropping" condition, I_{19}, was uncomfortably loud. An adaptive procedure was implemented to present RG the same passages at lower levels. A 7 dB attenuation was enough to present condition I_{19} at a comfortable level for the following four experimental sessions. It was first intended to null the data of RG for this condition, but after seeing that the variation between minimum and maximum values for all five AGC systems was small, it was decided to incorporate and include the data.

Figure 4-1: Mean values for subject PG

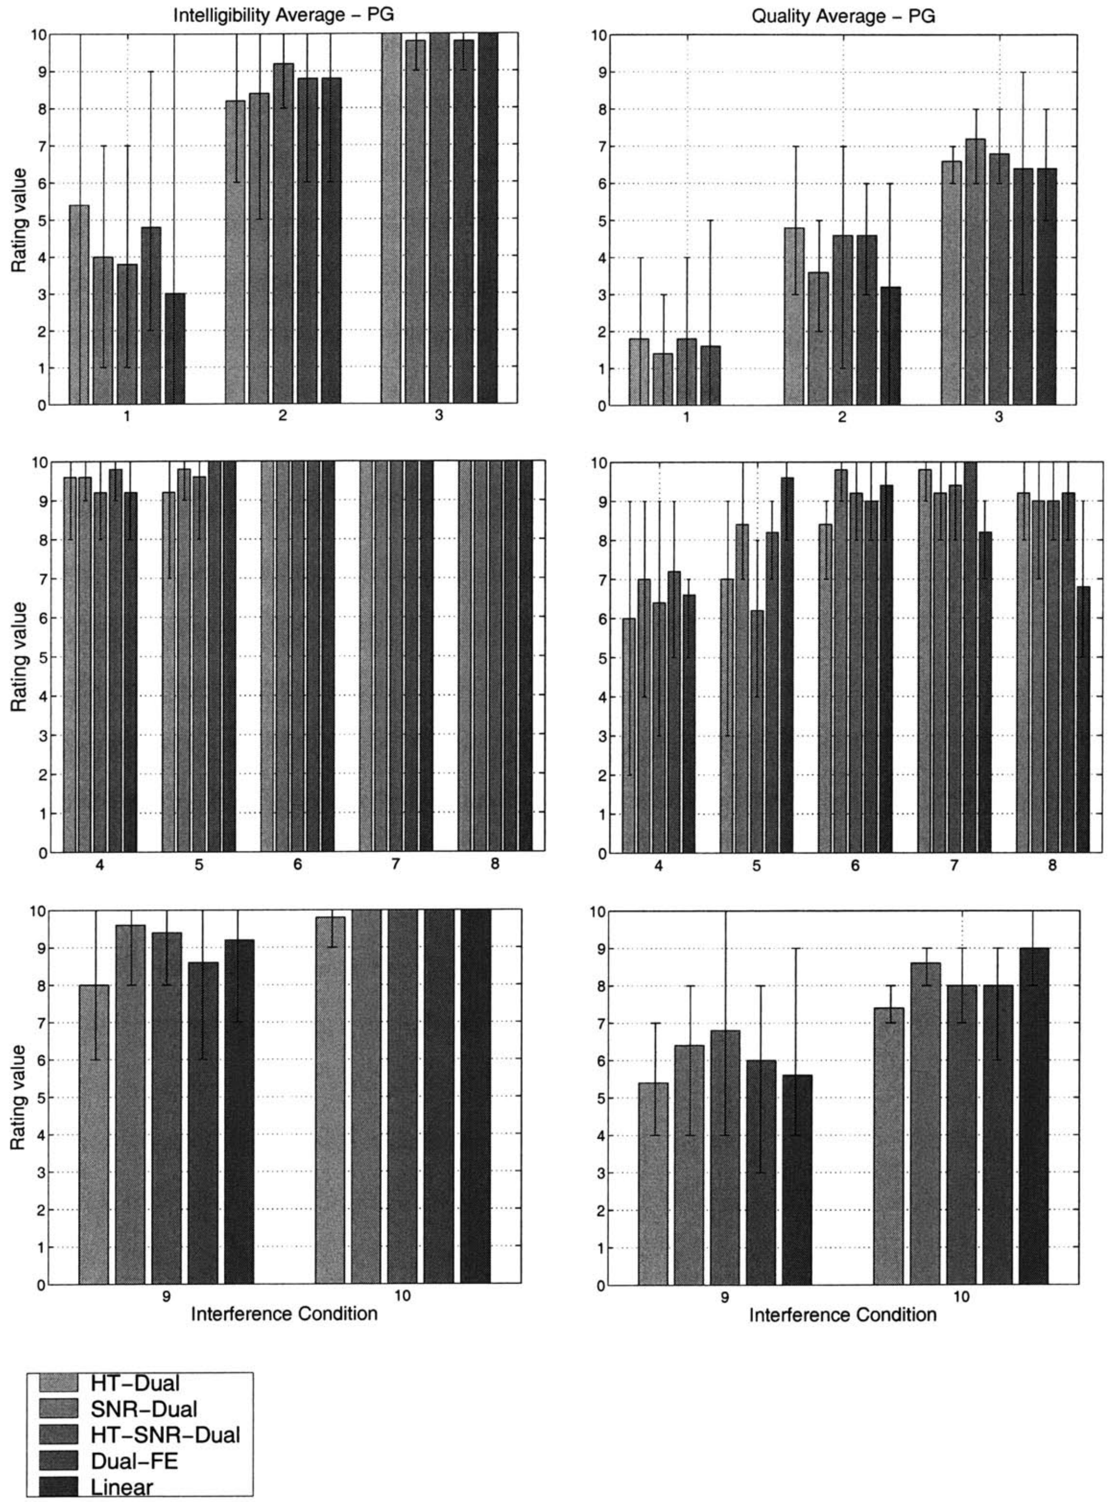

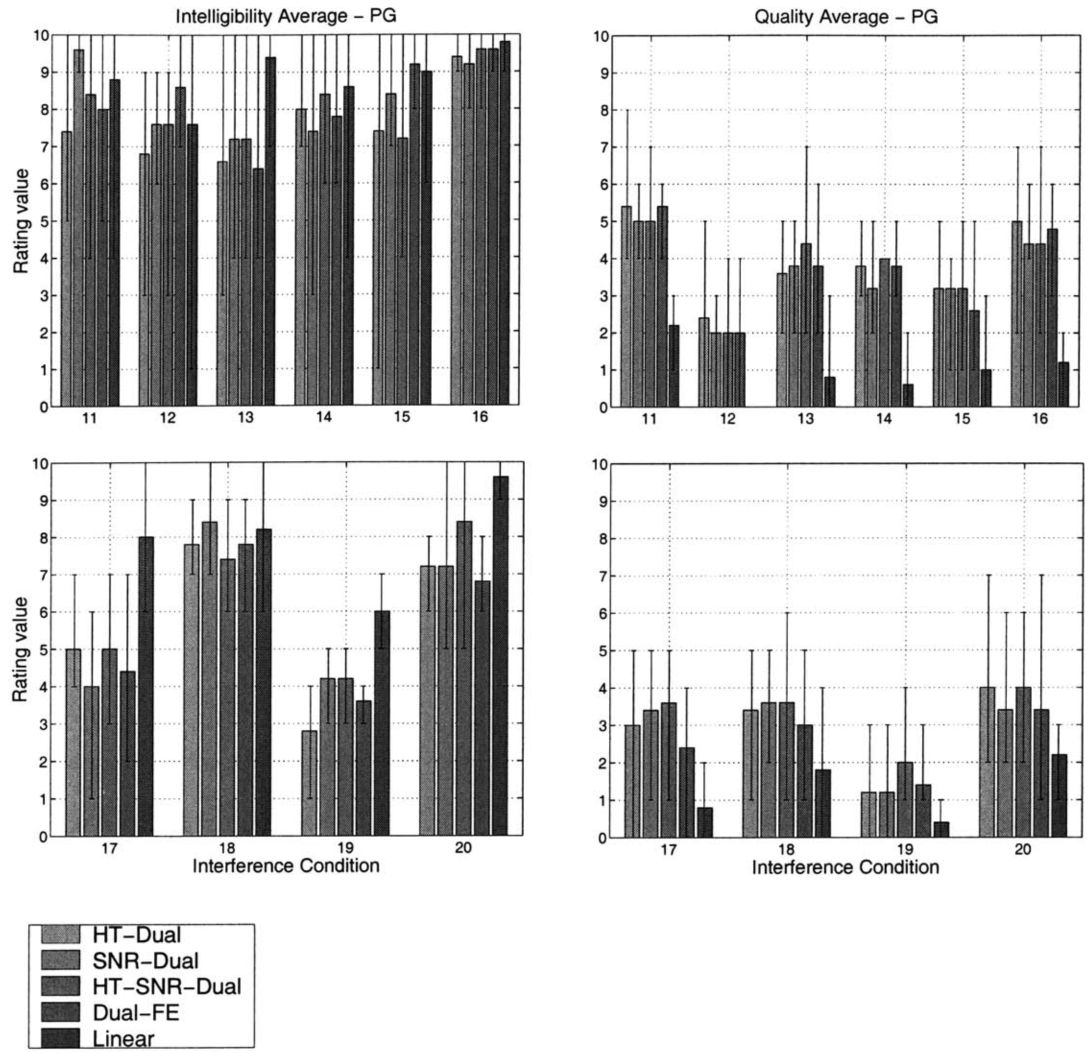
Intelligibility Average - PG
Quality Average - PG
Rating value
Interference Condition
HT-Dual
SNR-Dual
HT-SNR-Dual
Dual-FE
Linear

Figure 4-2: Mean values for subject MG

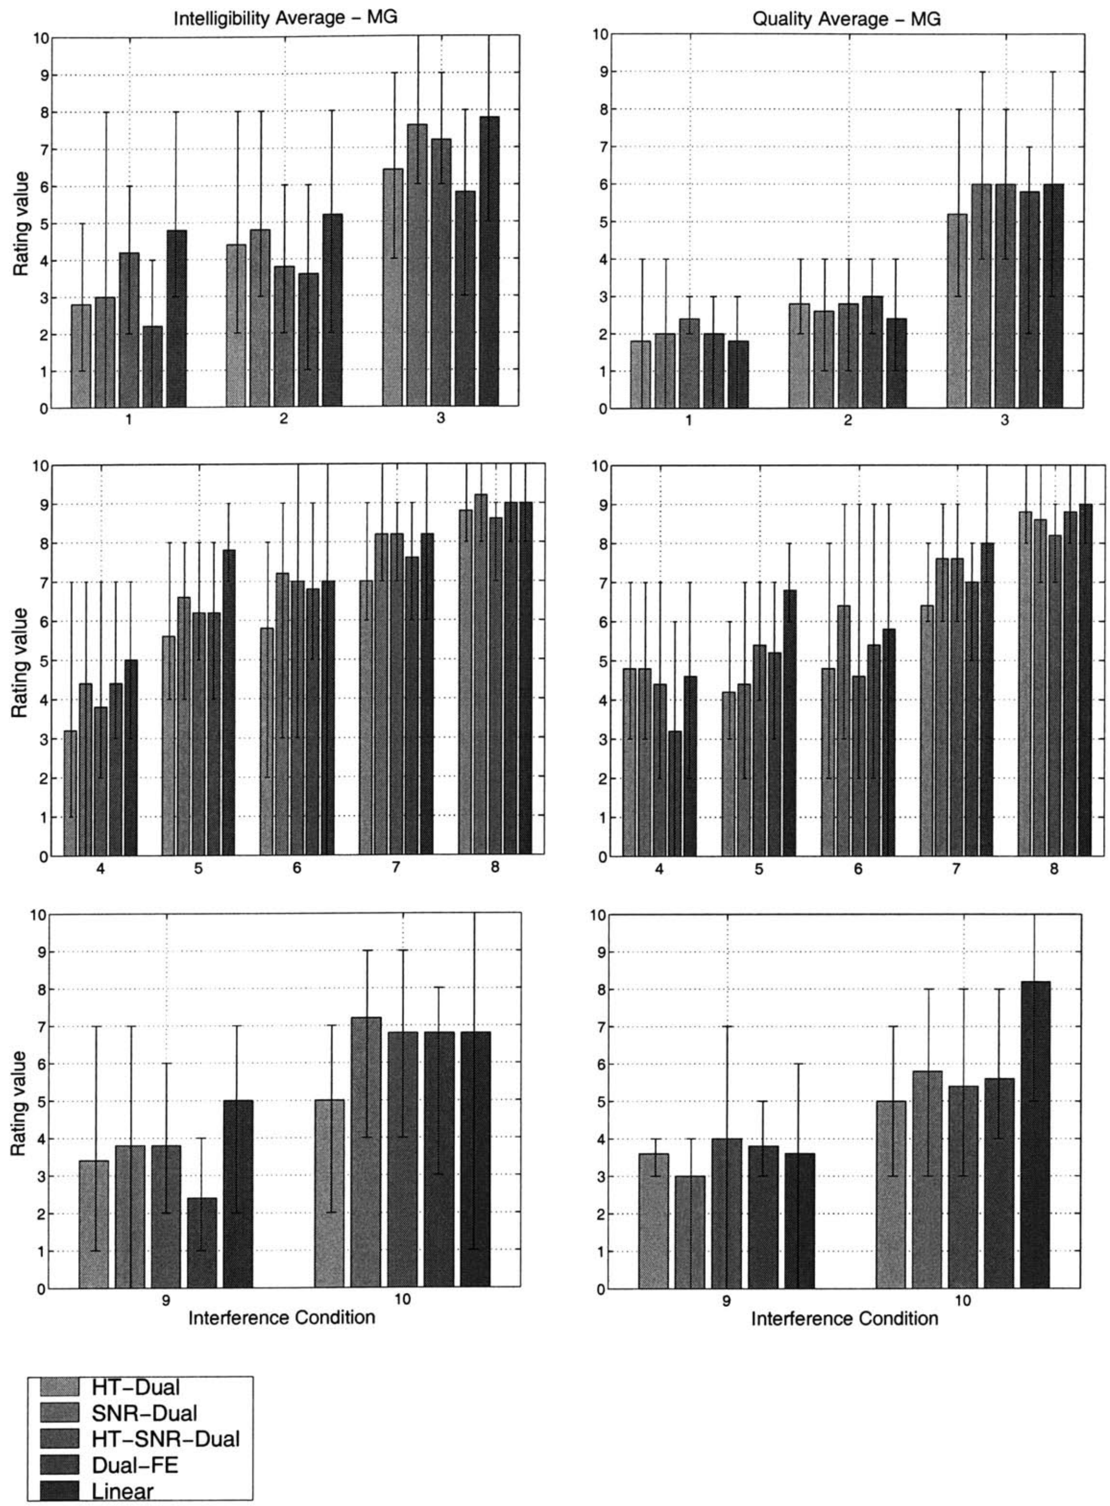

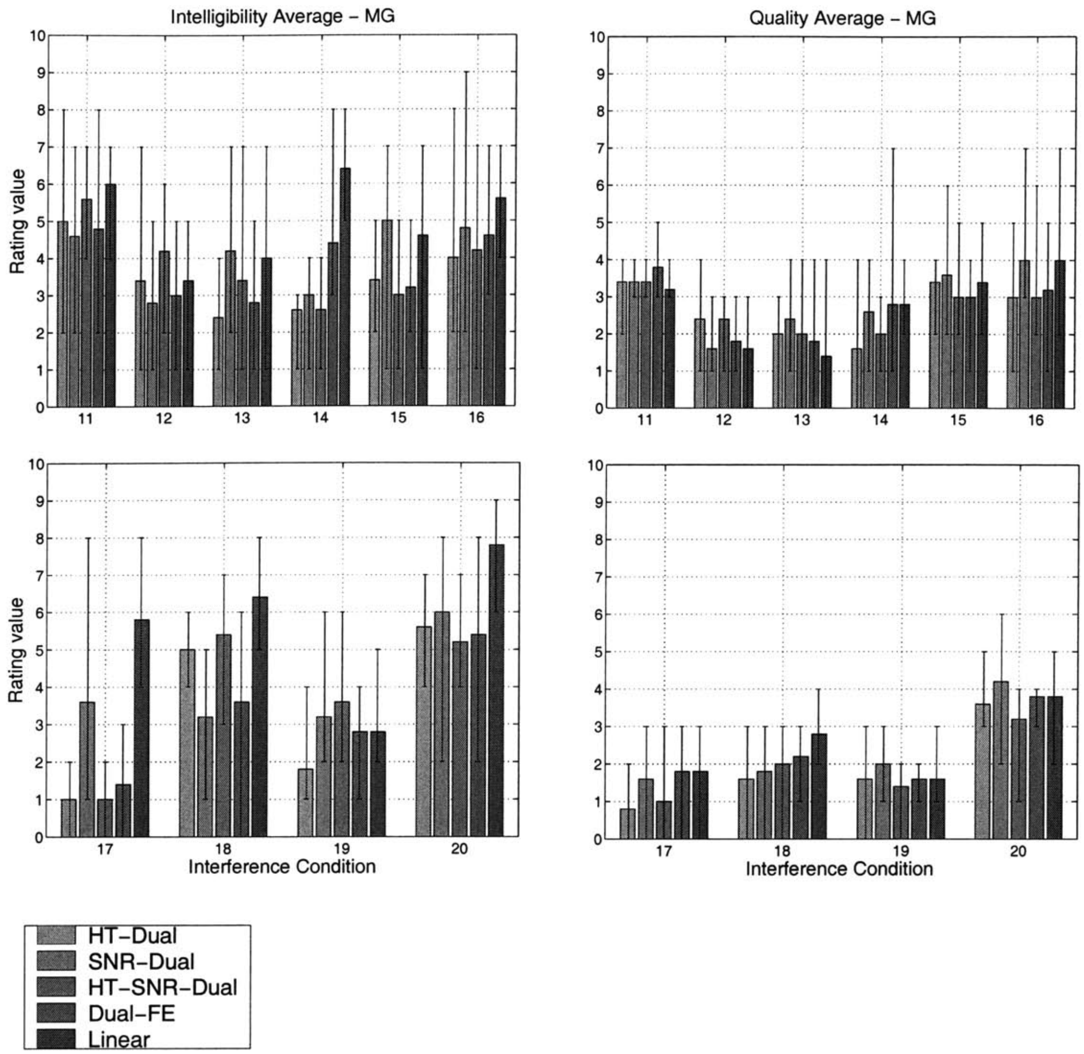
Intelligibility Average – MG
Quality Average – MG
Rating value
Interference Condition
HT–Dual
SNR–Dual
HT–SNR–Dual
Dual–FE
Linear

Figure 4-3: Mean values for subject RG

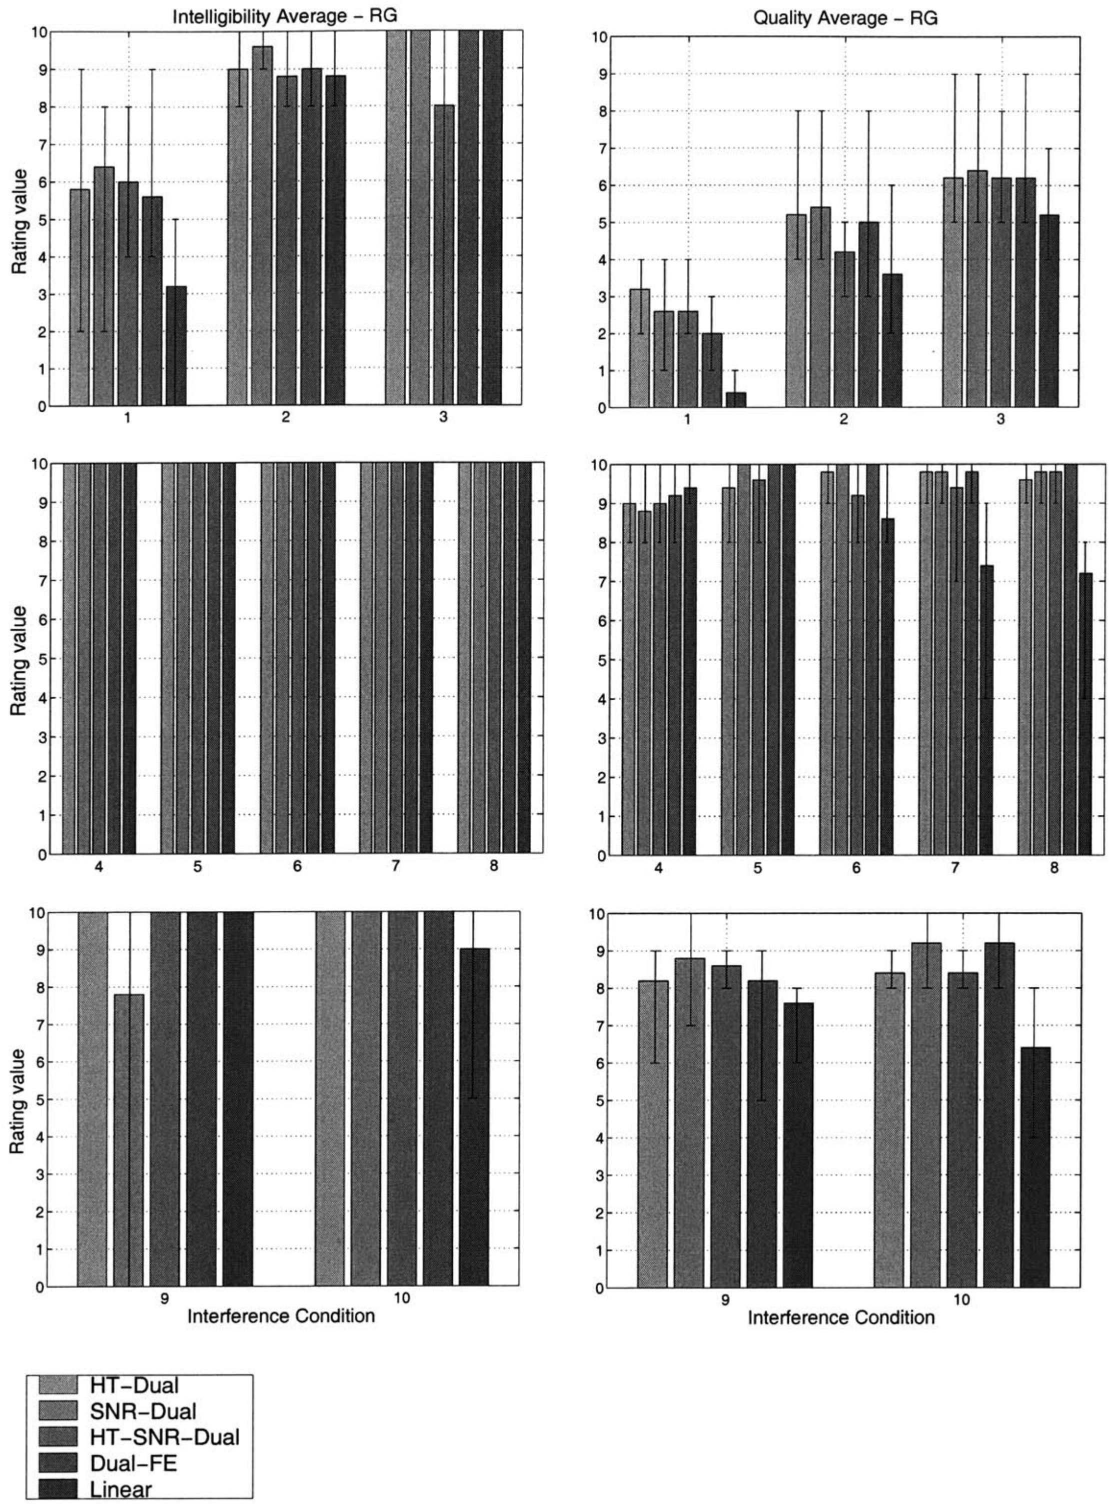

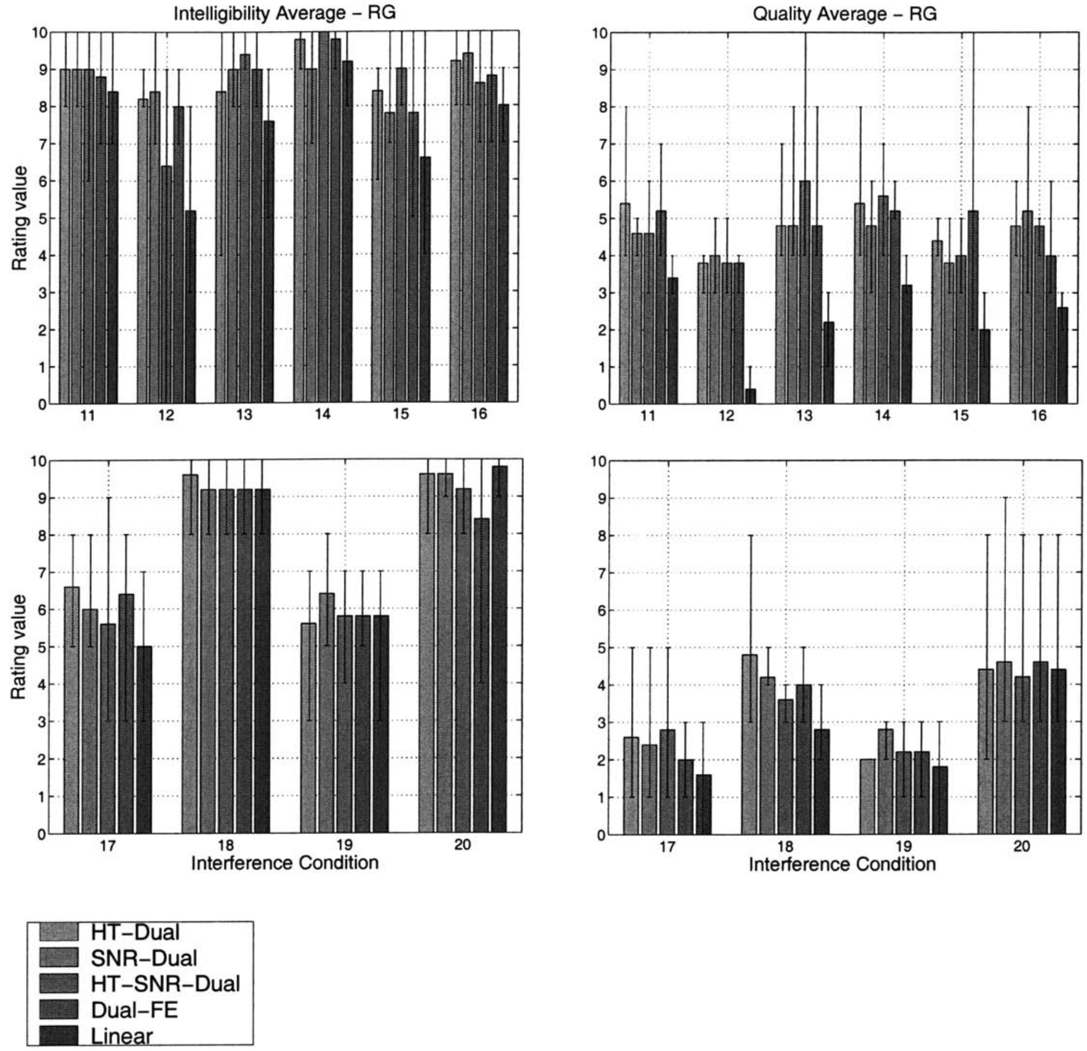

Intelligibility Average - RG
Quality Average - RG
Rating value
Interference Condition
HT-Dual
SNR-Dual
HT-SNR-Dual
Dual-FE
Linear

4.2 Difference Relative to the Linear System

A comparison relative to the Linear system was performed for the four AGC compression systems. The average of the five segments S_k per condition, per algorithm was previously calculated from equation 4.1, resulting in values of $\overline{A_i I_j}$. To provide an adequate comparison, the mean value of the Linear system $\overline{A_5 I_j}$ was subtracted from the averages of the other four AGC systems. This normalization provides a clear view of the level of increased or decreased performance of the four compression systems relative to the reference system. The average differences relative to the Linear system, $\Delta A_i I_j$, were calculated as follows:

$$
\begin{aligned}
\overline{A_i I_j} &= \frac{1}{5} \sum_{k=1}^{5} A_i I_j S_k, && i = 1, ..., 5, j = 1, ..., 20 \\
\Delta A_i I_j &= \frac{1}{5} \sum_{k=1}^{5} A_i I_j S_k - \frac{1}{5} \sum_{k=1}^{5} A_5 I_j S_k \\
\Delta A_i I_j &= \overline{A_i I_j} - \overline{A_5 I_j}, && i = 1, ..., 4, j = 1, ..., 20
\end{aligned}
\tag{4.2}
$$

Figures (4-4) - (4-6) contain the values of $\Delta A_i I_j$ for the three subjects. Again, the left and right column plots correspond to the intelligibility and quality values of $\Delta A_i I_j$, respectively. Negative values indicate a worse performance, and positive rating values correspond to better performance relative to the Linear system.

Figure 4-4: Average difference, relative to Linear, for subject PG

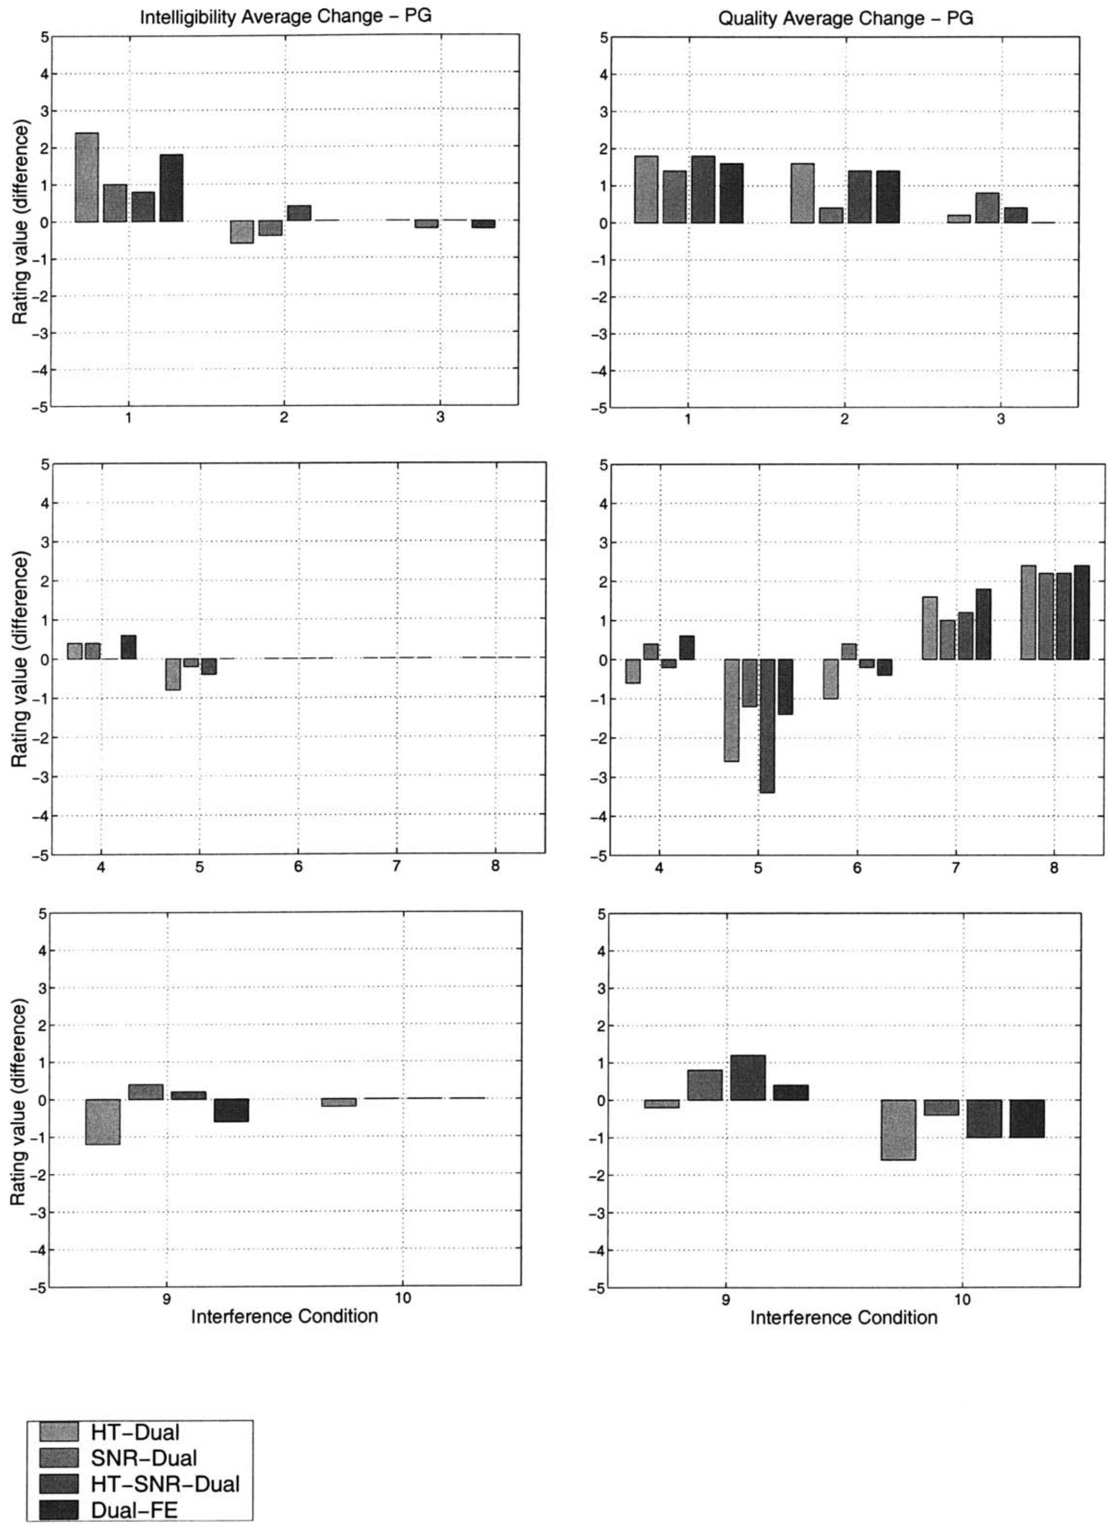

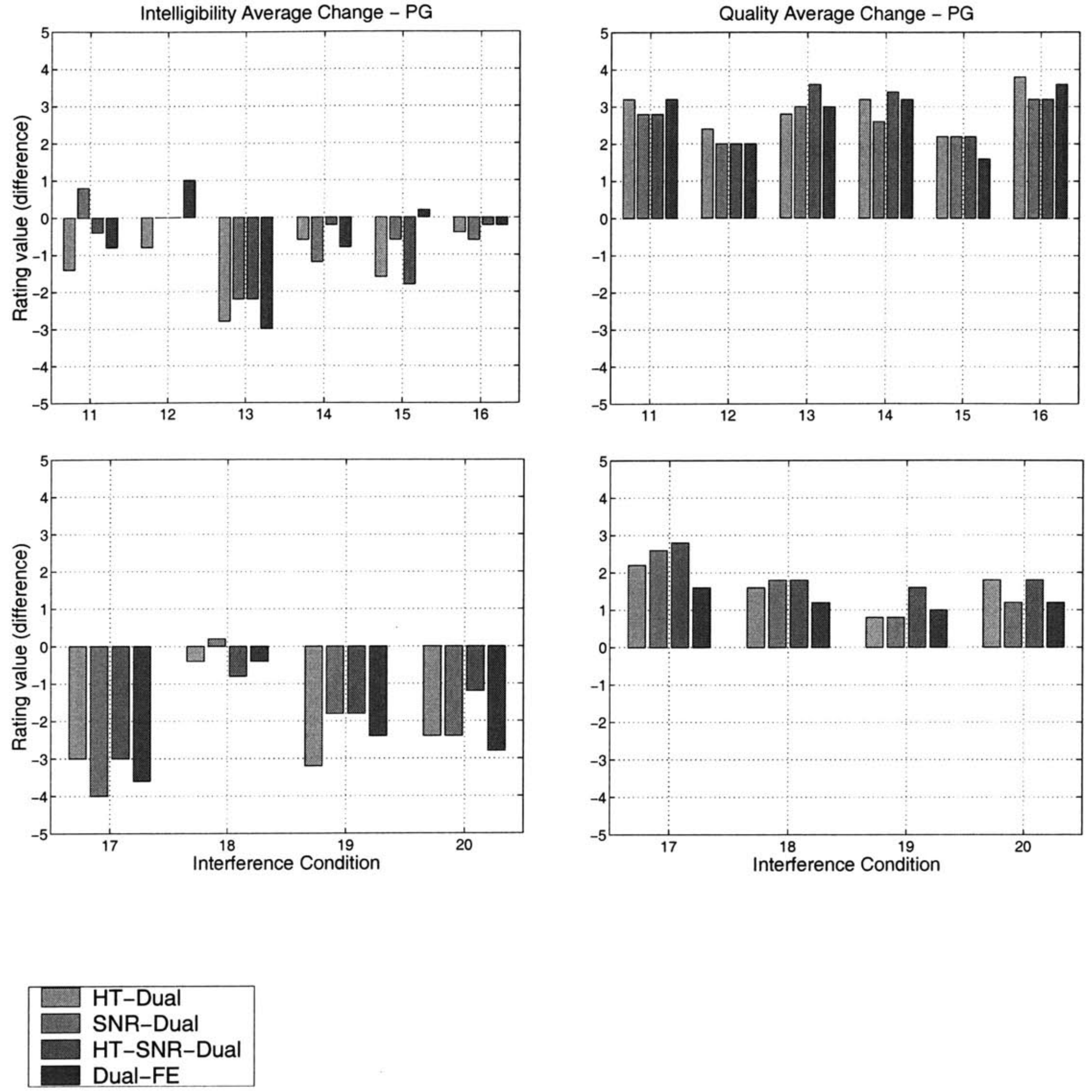
Intelligibility Average Change - PG
Quality Average Change - PG
Rating value (difference)
Interference Condition
HT-Dual
SNR-Dual
HT-SNR-Dual
Dual-FE

Figure 4-5: Average difference, relative to Linear, for subject MG

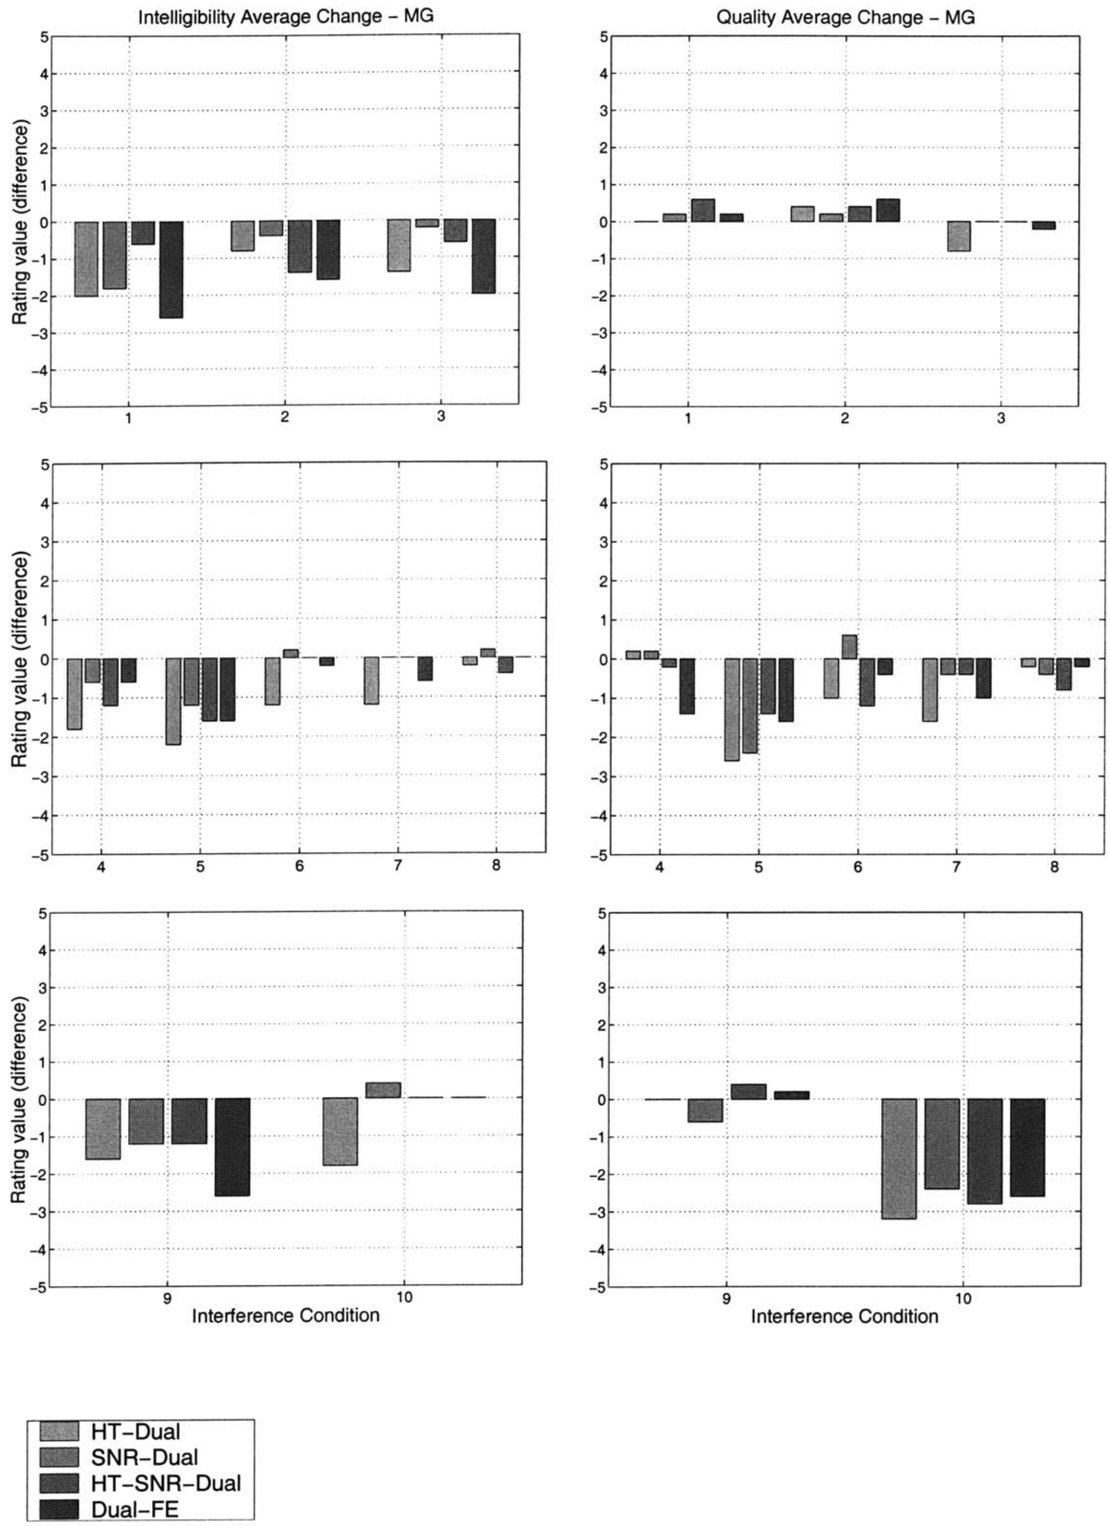

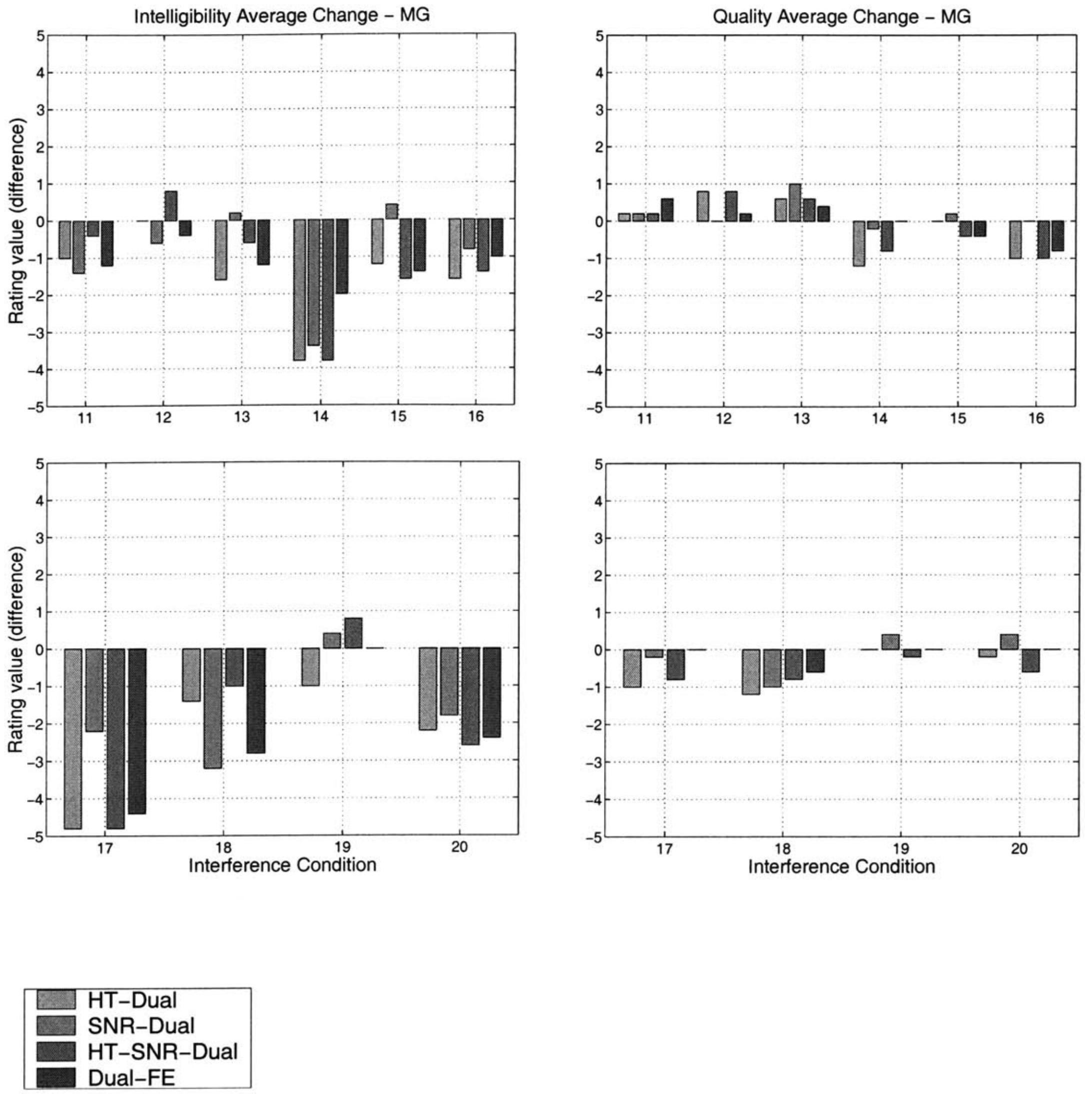
Intelligibility Average Change - MG
Quality Average Change - MG
Rating value (difference)
Rating value (difference)
11
12
13
14
15
16
17
18
19
20
Interference Condition
Interference Condition
HT-Dual
SNR-Dual
HT-SNR-Dual
Dual-FE

Figure 4-6: Average difference, relative to Linear, for subject RG

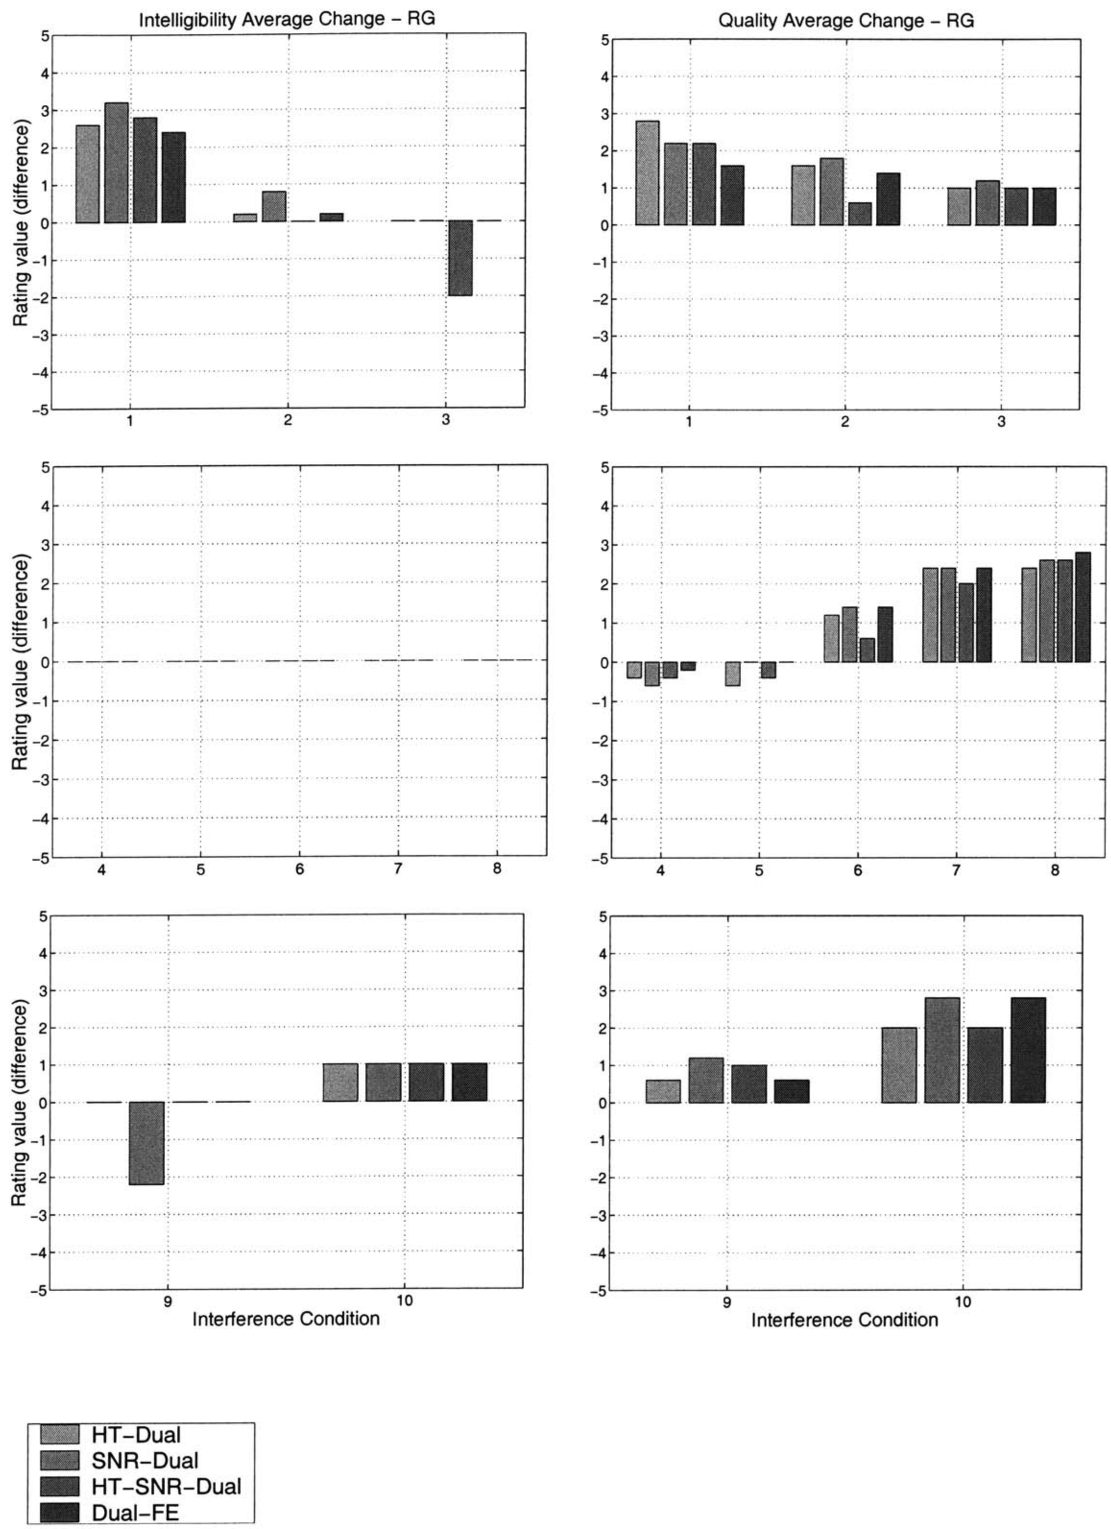

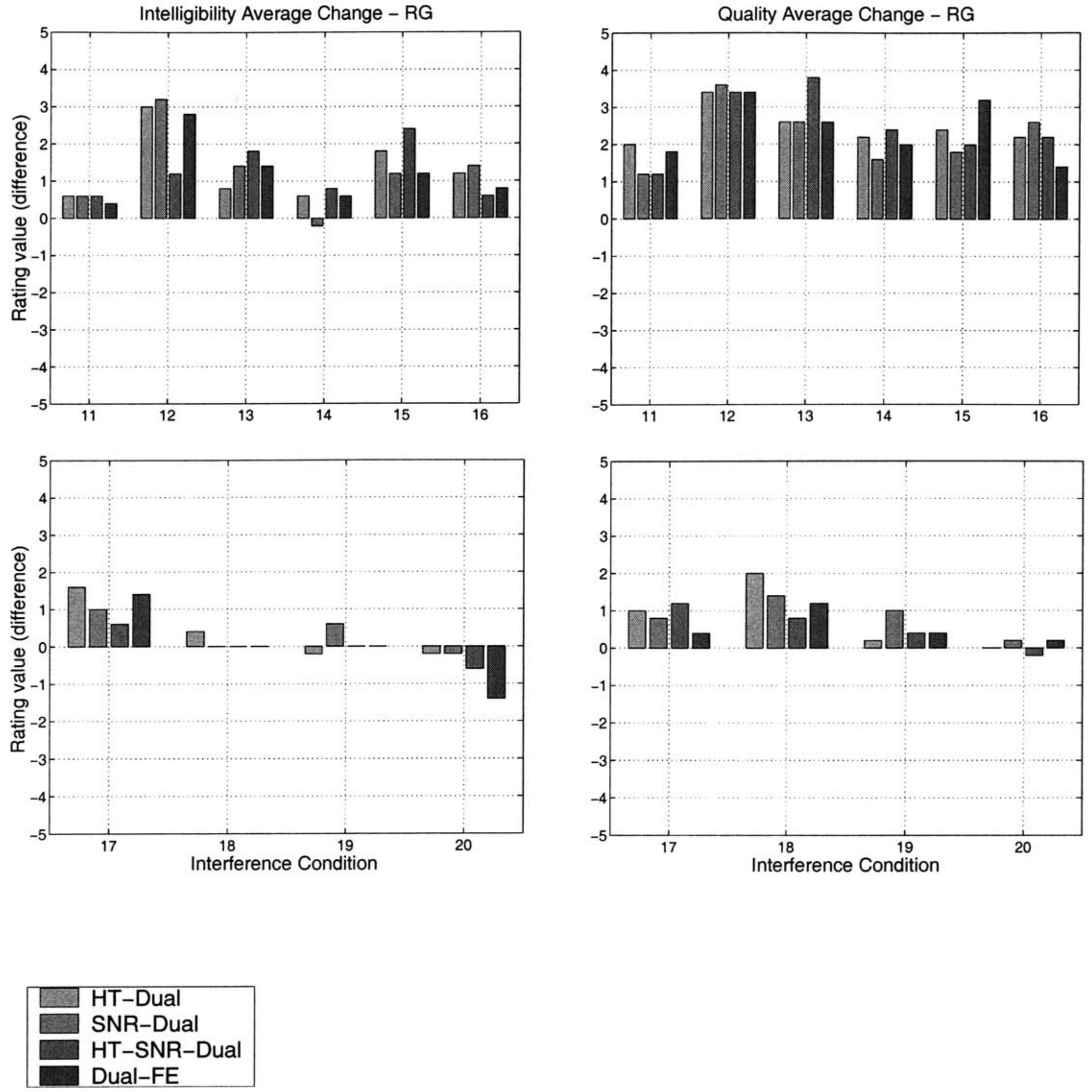

Intelligibility Average Change - RG
Quality Average Change - RG
Rating value (difference)
Interference Condition
HT-Dual
SNR-Dual
HT-SNR-Dual
Dual-FE

4.3 Relative Average Group Difference

A final step of combining the data into relevant interference condition groups was undertaken. The average difference, relative to the Linear system, for algorithm A_i and interference condition group G_g was calculated as A_iG_g using the following equations:

$$A_iG_1 = \frac{1}{3}\sum_{j=1}^{3} \Delta A_i I_j, \qquad i = 1, ..., 4 \tag{4.3}$$

$$A_iG_2 = \frac{1}{5}\sum_{j=4}^{8} \Delta A_i I_j, \qquad i = 1, ..., 4 \tag{4.4}$$

$$A_iG_3 = \frac{1}{2}\sum_{j=9}^{10} \Delta A_i I_j, \qquad i = 1, ..., 4 \tag{4.5}$$

$$A_iG_4 = \frac{1}{6}\sum_{j=11}^{16} \Delta A_i I_j, \qquad i = 1, ..., 4 \tag{4.6}$$

$$A_iG_5 = \frac{1}{4}\sum_{j=17}^{20} \Delta A_i I_j, \qquad i = 1, ..., 4 \tag{4.7}$$

where $\Delta A_i I_j$ was previously calculated as the average difference for algorithm A_i, interference I_j, in equation 4.2. The description of the five groups G_g are summarized in table 4.1 below. The calculated average group differences A_iG_g are plotted in figure 4-7, and are discussed in detail in the following chapter.

Table 4.1: Description of five interference condition groups

Interference Condition Group	Description
G_1: S + B	Speech plus multi-talker babble ($I_1 - I_3$)
G_2: S w P	Speech alone with pauses ($I_4 - I_8$)
G_3: SAL	Speech alone at alternating levels ($I_9 - I_{10}$)
G_4: S + SSE	Speech plus env. steady-state sounds ($I_{11} - I_{16}$)
G_5: S + T	Speech plus env. loud transients ($I_{17} - I_{20}$)

Figure 4-7: Average group interference condition difference

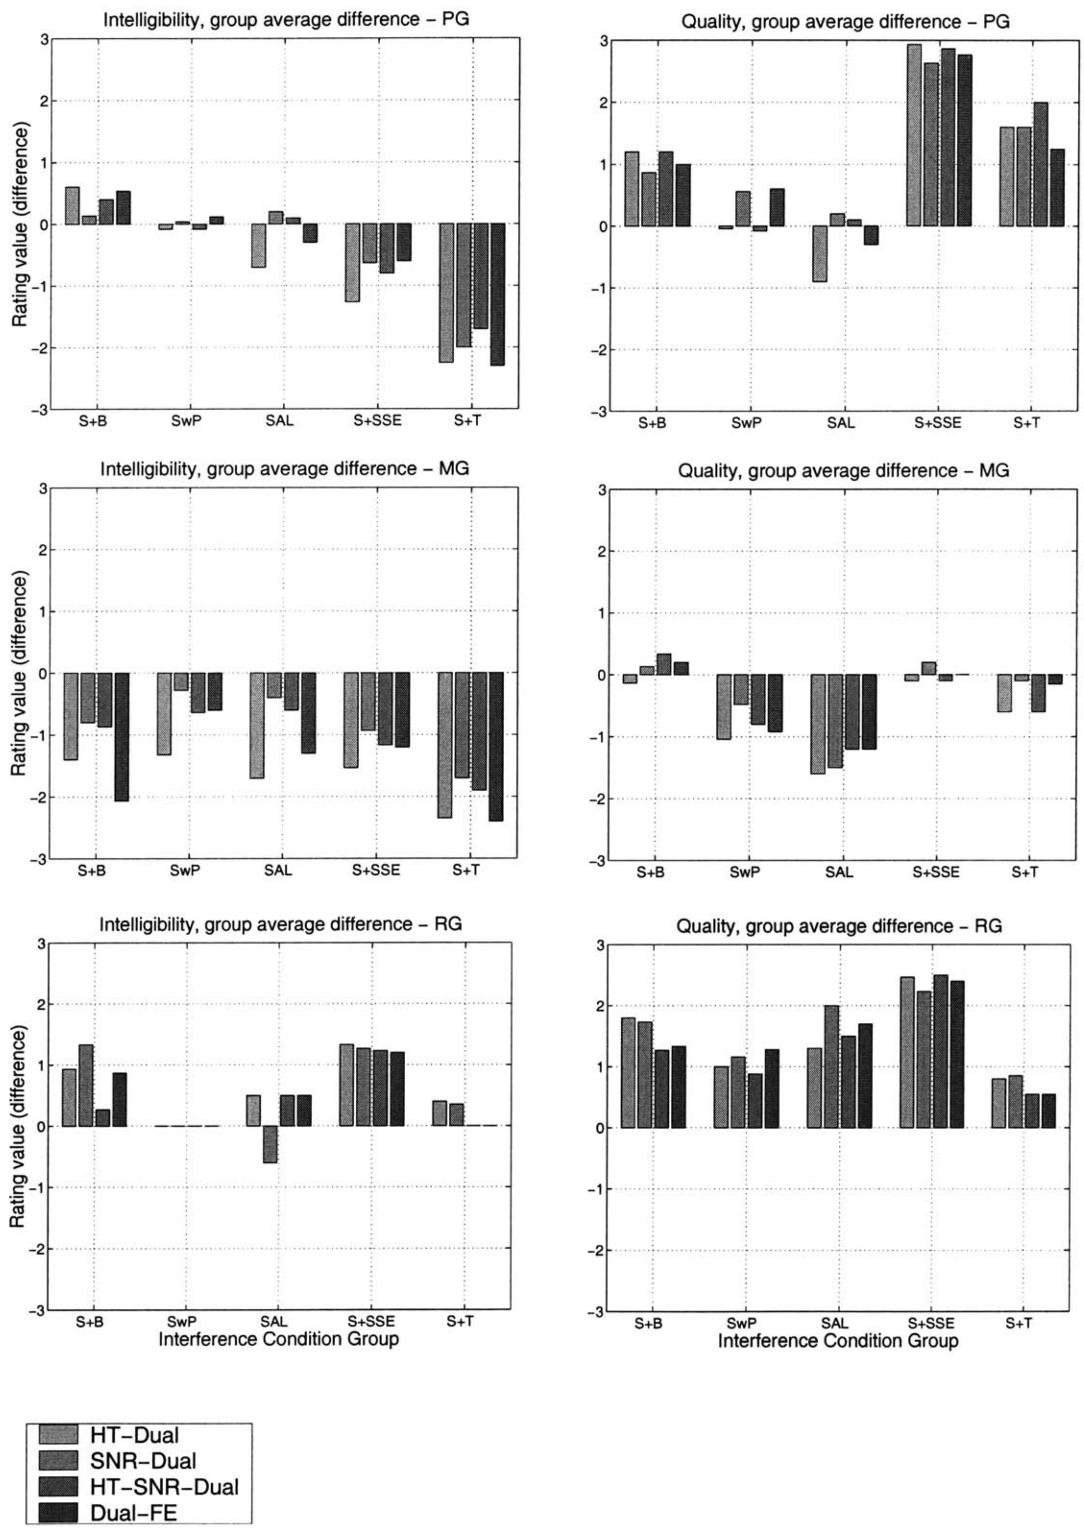

Chapter 5

Discussion

5.1 Four Compression AGCs vs. Linear System

No clear pattern was found for any of the AGC compression systems that indicates a particular preference for all three subjects. The subjects typically had distinctive ratings for both intelligibility and quality of the four compression AGCs compared to the Linear system, but a few similarities are found between subjects, especially between PG and RG.

In the first condition group G_1, speech plus babble (S+B), a clear improvement in quality is found in the four compression systems relative to the Linear system for subjects PG and RG. However, no significant difference is shown for subject MG. For intelligibility, a similar effect is shown; slight benefit from the compression systems for PG and RG, but a clear decrease of intelligibility for MG. Interestingly, the intelligibility benefit for both PG and RG is due only to the speech plus babble at the lowest SNR of 0 dB. For I_2 and I_3, S+B at +6 dB and +12 dB SNR, respectively, no benefit is found for any of the algorithms. The relative quality ratings also decreases for PG and RG from I_1 (0 dB SNR) to I_3 (+12 dB SNR). This indicates that compression may have greater benefit (both intelligibility and quality) over linear amplification at lower SNR levels, as compared to higher SNR levels of speech.

To support the above statement of greater benefit from compression at low SNR levels of speech, consider interference group G_2, composed of clean speech at different

SPL levels with pauses (SwP). Stimuli in group G_2 contained high SNR levels, and clearly there is no difference in intelligibility in all five AGC systems for subjects PG and RG. A slight quality benefit is found for two compression systems under PG, and a more significant benefit for the four compression AGCs under RG. Regardless, the average quality benefit in the SwP group for PG and RG is less than the quality benefit in the S+B group. This supports the idea that more intelligibility and quality benefit is acquired from compression, relative to linear amplification, under low SNRs, compared to higher SNR levels of speech.

Even though two subjects show no clear difference in intelligibility among the five algorithms for all clean speech ranging from 50 dB SPL (I_4) to 80 dB SPL (I_8), all three subjects show that the effect of compression on sound quality depends on the SPL level. Clear quality benefit, relative to the Linear system, is shown for two subjects (PG and RG) in the four compression systems under conditions I_7 (speech at 72 dB SPL) and I_8 (80 dB SPL). As expected, compression provides adequate attenuation to present loud speech at a comfortable level, making the quality of sounds much more pleasant than with linear amplification. Compression did not improve quality (or intelligibility) for MG; it actually decreased intelligibility and quality for most of the interference conditions in group G_2. But the actual difference between compression ratings and Linear system ratings decreased as the SPL of speech increased. This indicates that either the compression systems performed better as the SPL went up, the Linear system's performance worsened as SPL increased, or both.

Interference condition group G_3, clean speech with alternating levels (SAL), provides mixed results for the three subjects. Compared to the Linear system, the four compression AGCs improve quality for RG, but degrade quality for MG. To add to the mixture, subject PG shows little difference in quality between 3 compression systems and the Linear AGC. The fourth compression AGC (HT-Dual) is almost one rating unit worst than the Linear system. Interestingly, most of the negative quality weight in the group average (relative to the Linear) for MG comes from interference condition I_{10}, speech alternating between 65 dB SPL and 80 dB SPL. For subject RG, condition I_{10} provides most of the beneficial weight in the average of group SAL.

The other condition in group SAL, I_9, is lower level speech alternating from 50 to 65 dB SPL, which is less affected by compression than I_{10}. This opposing difference between MG and RG indicates that one subject (RG) receives outstanding benefit from compression as compared to linear amplification, and the other (MG) is more "comfortable" with a Linear system than using one with compression. This can be supported by noticing that RG has positive differences for most interference groups (figure 4-7) for the four compression systems, while MG is the opposite, having negative differences for most interference groups, for the four compression systems, and for both intelligibility and quality[1].

The next interference group G_4, composed of speech plus steady-state environmental sounds (S+SSE), shows clear quality benefit for the compression systems. This condition group exhibits the most quality benefit of any of the other groups for subjects PG and RG, with an average of almost 3 and 2.5 units or rating difference, respectively, compared to the Linear system. Interference conditions I_{11} - I_{16} contained steady-state background noise with an SNR of 0 dB. Recall that I_1, speech plus babble at 0 dB SNR, provided most of the intelligibility and quality benefit for the compression systems in group G_1 for PG and RG. The proposition that compression is more beneficial at low SNRs is again upheld by the quality ratings in the interference group S+SSE. Comparing the group quality of S+SSE with the quality of group G_2 (SwP) also shows that speech at lower SNRs (i.e. group S+SSE) has a greater beneficial gap difference than speech at higher SNR levels (i.e. group SwP) between the compression systems and the Linear system, for subjects PG and RG. Even for MG, who tends to prefer the Linear system over any of the compression systems, the quality difference is reduced under group S+SSE, indicating that at least there is no negative effect on quality from the compression systems at low SNR levels compared to the Linear system.

As for intelligibility ratings in group G_4, there is a mixture of indications between subjects. Subject RG's result support the idea that any of the compression systems

[1]There is a relative greater negative effect in the compression systems for intelligibility than for quality, for subject MG.

improve intelligibility at low SNRs, relative to linear amplification. Comparing group S+SSE with group SwP for RG, a benefit of more than one rating unit is seen for the four compression AGCs in S+SSE, while zero difference is found in the five systems for the SwP interference group. The opposite effect is shown for PG, where compression has a negative effect on speech intelligibility. Apparently there is a trade-off for PG between intelligibility and quality for this condition group. Subject MG on the other hand, is consistent with his preference of linear amplification. On the average, he rates the intelligibility of compression systems about 1 rating value (except the HT-Dual which is about 1.5 units) below the Linear system's rating.

The fifth interference group G_5 tends to indicate that while quality is increased using compression, some intelligibility of speech is lost. Group G_5 was composed of speech plus loud transients (S+T), and compression was aimed at protecting the user from uncomfortably loud sounds. A clear trade-off is seen for PG, where the quality is significantly better using compression, but the intelligibility of speech is greatly worse than with linear amplification. This effect is consistent with the interference group S+SSE. A slight difference is found with RG, where quality under the four compression systems is also better, but there is minimal difference in speech intelligibility for all five systems. Looking at the third subject (MG), speech intelligibility is again significantly worse under any of the compression AGCs compared to the Linear. Little difference between the five systems is found, though, in the MG ratings of sound quality.

The substantial quality rating difference in group G_5 between MG and the other two subjects may be due to the fine tune gains provided. Recall that the fine tune gains for PG, MG, and RG were 19, 4, and 20 dB, respectively. Since subjects PG and RG preferred a much greater gain relative to that prescribed by the Cambridge Formula, they were presumably operating closer to their level of uncomfortable loudness. As a result, speech was generally more intelligible (see figures (4-1)-(4-3)), but at a greater danger of becoming uncomfortably loud. This in turn, caused PG and RG to prefer sound quality of compression over that of linear amplification. Subject MG had a much lower fine tune gain. It may be the case that speech sounds were not amplified sufficiently for this subject, so that any reduction in gain applied by compression had a negative effect on intelligibility. On the other hand, this lower gain did not operate the systems close to MG's uncomfortable level, thus compression did not improve sound quality relative to the Linear system.

Overall, the three subjects in the experiment had distinctive preferences over compression vs. linear amplification. Subject MG found all the compression systems to be worse than the Linear system for all the group conditions of speech intelligibility, and two (out of five) of the quality groups. At the other end is RG, who found compression to be much more beneficial, compared to the Linear AGC system, for all the sound quality groups and three of the five speech intelligibility groups. In the middle is subject PG. He finds significant compression benefit in three sound quality interference groups, but at the cost of giving up speech intelligibility in two of those interference groups.

5.2 Differences in Compression AGC Systems

Clear differences are found for most conditions when comparing the four compression AGCs with the Linear system, but not so much when comparing the compression systems among themselves. A few trends are apparent in the data of one subject (MG) in the comparison of the four compression AGC systems, but the differences are minimal.

Examining the interference condition group data of subject PG in figure 4-7 shows no clear differences among the four compression algorithms. There is a small negative difference for the HT-Dual relative to the other three compression AGCs, in the SAL and S+SSE intelligibility, and the SAL quality interference groups. A slight preference for the HT-SNR-Dual system is indicated from PG in the S+T group, for both intelligibility and quality. The difference is about 0.75 and 0.5 of a rating unit better than the Dual Front-End for intelligibility and quality, respectively. The only other slight difference for PG is found in the quality ratings of the SwP group. The HT-Dual and the HT-SNR-Dual algorithms are about 0.5 of a rating unit worse than the other compression AGCs, but looking at the average differences of the individual 5 interference conditions of group SwP (figure 4-4) the only substantial difference comes from only one condition (I_5). This may indicate that there is no clear difference in quality between the four compression algorithms in clean speech.

Subject RG also shows no clear preference among the four compression systems. At first sight of the intelligibility group ratings of RG (figure 4-7) it may seem that there is a significant loss from the HT-SNR-Dual and the SNR-Dual systems in the S+B and SAL interference groups, respectively, relative to the other compression AGCs. This is believed to be errors in data. First, the only negative intelligibility value (relative to the Linear system) found in group S+B is in condition I_3, with the HT-SNR-Dual rating difference of negative two (figure 4-6). Looking further into the intelligibility mean value of interference condition I_3 (figure 4-3), the minimum value of the 5 segments is zero, while the other four ratings are 10. All 24 intelligibility RG rating values (i.e. five systems, five segments per system, minus one corresponding to the rating in question) of condition I_3 are 10, except the single value under the HT-SNR-Dual which is zero. It is believed that RG meant to input a value of 10 and might have typed the digit 1 before she was prompt to enter her rating. In turn, only the digit 0 was received by the computer and saved. The same error is believed to have happened in the intelligibility rating of condition I_9 for the SNR-Dual system, thus producing a final negative effect in the SAL group of RG relative to the other compression AGC algorithms.

Subject MG preferred the Linear system over any of the compression systems for most interference condition groups, for both intelligibility and quality. Within the four compression systems, there is a noticeable difference in intelligibility ratings due to the use of the Hold Timer and the SNR Estimator. Looking at the intelligibility group plot for MG in figure 4-7, it can be noticed that the SNR-Dual is consistently rated higher than the other three algorithms, while the Dual Front-End and the HT-Dual system are generally rated lower. This may indicate that subject MG receives some benefit in speech intelligibility from the SNR Estimator, even though it is still not enough to outperform the Linear system. This idea of SNR Estimator benefit is supported by the performance of the HT-SNR-Dual system, which also incorporates the SNR Estimator. It performed slightly worse than the SNR-Dual, but the differences were minimal in all conditions, thus still significantly better than the HT-Dual and the Dual Front-End algorithms.

Even though consistent trends are seen for intelligibility for the SNR-Dual (and the HT-SNR-Dual) compared to the HT-Dual and the Dual Front-End, this is not the case for quality ratings of MG. The biggest quality rating difference among the four compression algorithms is under the speech plus loud transients group (S+T). The HT-Dual and the HT-SNR-Dual both performed an average of half (0.5) of a rating unit worse than both the SNR-Dual and the Dual Front-End. However, this may not be large enough to conclude that the Hold Timer has a negative effect on sound quality for this subject, in this condition group.

While there was a slight intelligibility preference for the SNR-Dual and the HT-SNR-Dual systems over the other two compression AGCs for MG, the overall ratings of the subjects show no substantial difference among the four compression AGC systems. These results show no evidence that the HT or the SNR Estimator have a negative or positive effect on speech intelligibility or quality when incorporated to the Dual Front-End system.

5.3 Summary and Relation to Other Work

The current study was primarily motivated by the work of Stone *et al.* [8], which evaluated four compression algorithms. Three of the four compression algorithms considered in that work included the Dual Front-End AGC system. Subjects were fitted with wearable devices that implemented one algorithm at a time and asked to use it in their everyday lives for 2-3 weeks, before returning for testing and to be fitted with the next algorithm. Objective intelligibility tests indicated no significant effect of the compression algorithms. APHAB results, a test developed by Cox and Alexander [1] where subjects rate how often they had problems in specific situations, and informal reports indicated that there was “a slight overall preference for the Dual-

LO system."[2] However, that study did not include any linear algorithm or reference condition.

The current study was entirely lab-based, so subjects did not get extensive exposure to the algorithms as in Stone *et al.* In addition, subjects did not have the opportunity to listen to the algorithms in the acoustic environments encountered in their everyday lives. However, the stimuli were generated from a carefully selected set of speech segments and background noises to be representative of typical listening situations and environmental sounds. Moreover, the design of the experimental conditions permitted a more direct comparison between algorithms under the same interference condition than is possible in a field trial.

The current study attempts to address two questions. First, does the version of the Dual Front-End AGC considered by Stone *et al* provide benefits in terms of speech intelligibility and/or sound quality relative to a linear reference condition? Second, which, if any, optional components of the Dual Front-End improve its performance? The optional components considered include the Hold Timer proposed by Stone *et al.* and the SNR detector proposed by Martin *et al.* [4].

The results of the current study can be summarized as follows:

- Substantial differences in both intelligibility and quality ratings are seen when comparing all four compression systems to the linear reference condition. However, the direction of these differences varies with subject, and to a lesser degree, with stimulus condition. Of the three hearing-impaired subjects in this study, one generally found compression beneficial to both intelligibility and quality, one generally found compression detrimental to both intelligibility and quality, and one found that compression sometimes improved quality, at the expense of intelligibility. This mixed result may indicate that there are hearing loss characteristics that affect the benefit of compression over linear amplification, or it may be related to the differences in fine tune gain selected by the subjects.

[2]The Dual-LO system considered by Stone *et al.* was essentially the same as the HT-Dual system in the current experiment, including the same compression ratio and time constants.

- Both the Hold Timer and the SNR Estimator perform as expected from an engineering point of view, but there are no substantial differences in intelligibility or quality ratings among the four compression systems. Thus, there is no evidence that the Hold Timer and SNR Estimator provide positive or negative effects on speech intelligibility or sound quality for the stimuli conditions included in this experiment.

- In general, the benefits of the compression systems were greater at low SNR stimuli conditions than for clean speech, indicating that the compression systems may be providing suppression of background noises.

5.4 Future Research

Future work should first include a more thorough analysis of the data collected in this study using analysis-of-variance to test for the statistical significance of the apparent effects discussed in sections 5.1 and 5.2. After such analysis has been completed, it may be useful to design a future investigation to explore the large subject-dependent differences seen in this experiment. This logically follows the current study since one subject clearly preferred compression, another subject preferred conventional linear amplification, and the third subject had mixed results.

It was earlier suggested that compression has greater benefit over linear amplification, for speech intelligibility and sound quality, under speech at low SNR levels then under speech with high SNRs (i.e. clean speech). To verify this idea, another suggested future investigation would include a study with two algorithms, a compression and a linear system. Conditions should include stimuli at different SNR levels of speech presented to subjects with both systems. It is necessary to include data specifically testing speech at different SNR levels to make a strong conclusion.

Appendix A

MatLab Source Code

```
function [out,sig] = htandsnr (speech, noise, Fs, alg);
% speech = input speech signal
% noise = input noise signal
% Fs = sampling frequency (Hz)
% alg = Algorithm to be used
%       Enter 1 for Hold Timer ON, SNR OFF (HT only, called HT-Dual)
%       Enter 2 for Hold Timer OFF, SNR ON (SNR only, called SNR-Dual)
%       Enter 3 for Hold Timer ON, SNR ON  (both HT and SNR, HTandSNR-Dual)
%       Enter 4 for both HT and SNR OFF (called Dual Front-End)
%       Enter 5 for linear system (called Linear)

% Ivan Aguayo, 2001

%check if valid algorithm
if alg>5 | alg<1
  disp ('Not a valid Algorithm!');
  return;
end;

%Call SNR Estimator (to combine speech and noise, and calculate SNR)
[snr, sig] = gold_snr (speech, noise, Fs);
%If alg=1 turn SNR off by making all snr values negative.
%This assures that release time will always be 1 sec.
if alg == 1
  snr = -1 .* ones(size(snr));
end;

%Pre-Emphasis (high frequency emphasis)
signal = emphasis (sig);

%Buffer signal into frames
frame = floor (Fs * .004);    %4ms frames
[buf_sig, remain] = buffer (signal, frame);
nframes = size (buf_sig, 2);

%Calculate RMS of frames
```

```
rms1 = rms (buf_sig);
rms1db = 20*log10(rms1);
initVal1 = mean (rms1(1:10));

%Time Constants, and Compression ThresholdsM̂
tau_sa = 0.350;             %attack of 350ms
tau_sr = 1.00;              %release of 1sec
tau_sr3 = 0.300;            %release of 300ms

%compression threshold for slow control (dB)
thresh = 55;
%highest input level (dB)
thresh2 = 95;

%For Algorithm without Hold Timer or SNR Estimator
%Exp. Averaging of SLOW ControlM̂

%Attack (350msec)
Sattack = exp_filt (rms1, tau_sa, Fs/frame, initVal1);

%Release (1sec)
Srelease = exp_filt (rms1, tau_sr, Fs/frame, initVal1);

%slow parameter (sp) is input level
sp = max(Sattack, Srelease); %slow parameter (linear)

% For Algorithms 1 and 3.
% HOLD TIMER (HT) and SLOW MEAN (SM) Calculation
if alg==1 | alg==3
  %Attack filter constants
  aA = (tau_sa * Fs/frame) / (1 + tau_sa * Fs/frame);
  bA = 1-aA;

  %Release filter constants (1sec)
  aR = (tau_sr * Fs/frame) / (1 + tau_sr * Fs/frame);
  bR = 1-aR;

  %Release filter constants (300msec, for use of SNR)
  aR3 = (tau_sr3 * Fs/frame) / (1 + tau_sr3 * Fs/frame);
  bR3 = 1-aR3;

  %Initialize
  HT = zeros(1, nframes);       %Hold Timer (HT)
  spHT = zeros(1, nframes);     %slow parameter using HT
  spHT(1) = sp(1);
  spHT(2) = sp(2);

  %Run loop, calculate every frame
  for n = 2:nframes

    %Case 1: RMS is greater than Slow Mean, but not by more than 12dB
    if spHT(n-1) < rms1(n-1) <= spHT(n-1) + 10^(12/20)
       %Update Hold Timer (increase)
```

```
    if HT(n-1) >= 600    %reached max value
      HT(n) = 600;
    else
      HT(n) = HT(n-1) + 8; %increment hold timer
    end;
    %Exp. ave. of current RMS with previous SM value
    spHT(n) = aA * spHT(n-1) + bA * rms1(n);
  end;              %end of Case 1

  %Case 2: RMS is greater than Slow Mean by at least 12dB
  if rms1(n-1) > spHT(n-1) * 10^(12/20)
    HT(n) = HT(n-1);      %keep constant
    %Exp. ave. of previous SM + 12 dB with previous SM value
    spHT(n) = aA * spHT(n-1) + bA * (spHT(n-1)*10^(12/20));
  end;              %end of Case 2

  %Case 3: RMS is below Slow Mean
  if rms1(n-1) <= spHT(n-1)     %current frame is lower than mean
    %Update Hold Timer (decrease)
    if HT(n-1) <= 0
      HT(n) = 0;
    else
      HT(n) = HT(n-1) - 4; %decrement hold timer
    end;

    if HT(n) <= 40 %threshold
      %Exp. ave. of current RMS with previous SM value
      %Decide on release constant depending on SNR
      if snr(n) >= 0
        spHT(n) = aR3 * spHT(n-1) + bR3 * rms1(n);
        state(n) = 3;
      else
        spHT(n) = aR * spHT(n-1) + bR * rms1(n);
      end;
    else
      %HT nonzero, hold SM constant
      spHT(n) = spHT(n-1);
    end;
  end;           %end Case 3
 end;            %end of for loop
end;   %end of slow parameter calc. for algorithms 1 or 3

% For Algorithm 2
% Gain calculation with Hold Timer OFF, SNR ON
if alg==2
 %Attack signal (350msec)
 Sattack = exp_filt (rms1, tau_sa, Fs/frame, initVal1);

 %Attack update constants
 aA = (tau_sa * Fs/frame) / (1 + tau_sa * Fs/frame);
 bA = 1-aA;

 %Release update constants
 aR = (tau_sr * Fs/frame) / (1 + tau_sr * Fs/frame);
```

```
  bR = 1-aR;

  %Release with 300ms constant (for use when SNR Estimator is on)
  aR3 = (tau_sr3 * Fs/frame) / (1 + tau_sr3 * Fs/frame);
  bR3 = 1-aR3;

  spSNR = zeros(1, nframes);  %slow parameter using SNR only
  spSNR(1) = Sattack(1);

  for n = 2:nframes
    if Sattack(n) >= Sattack(n-1)   %env. increasing
      spSNR(n) = aA * spSNR(n-1) + bA * rmsl(n);
    elseif snr(n) < 0               %env. decreasing and low SNR
      spSNR(n) = aR * spSNR(n-1) + bR * rmsl(n);
    else                            %env. decreasing and high SNR
      spSNR(n) = aR3 * spSNR(n-1) + bR3 * rmsl(n);
    end;
  end;        %end of for loop
end;          %end of slow gain calc. for algorithm 2

%I/O Curve for Slow AGC. Use compression ratio of 3.
%Get Slow Parameter depending on Algorithm chosen
if alg==1 | alg==3          %algorithms with HT
  slow_param = spHT;
elseif alg==2               %algorithm with SNR only
  slow_param = spSNR;
elseif alg==4               %dual front-end algorithm
  slow_param = sp;
else                        %linear (make dummy slow_param)
  slow_param = sp;
end;
Slow_Param = 20.*log10(slow_param);

%Gain below threshold (0 dB)
Sgain = zeros(1, nframes);

%Gain above threshold (in dB)
high = find (Slow_Param > thresh);
Sgain(high) = -2/3 * ( Slow_Param(high) - thresh );

%convert from dB to linear
slow_gain = 10 .^ (Sgain ./ 20);

%if algorithm is linear, make slow_gain values to 1
%may be modified later for better efficiency
if (alg == 5)
  slow_gain =  ones (size(slow_gain));
end;

%Apply slow variable gain
slow_out = zeros (size(buf_sig));
for j = 1:frame
  slow_out(j,:) = slow_gain .* buf_sig(j,:);
end;
```

```
Sout = slow_out(:);          %output of slow AGC

%Delay signal
delay = round (Fs * .004);  %4 ms delay
del_sig = [zeros(delay,1) ; Sout];
[del_slow_out, del_remain] = buffer (del_sig, frame);

%Parameter, RMS, and Exp. Averaging of Fast Control
%rms of slow control output
rms2 = rms (slow_out);
initVal2 = mean (rms2(1:10));

%Attack
tau_fa = .005;
Fattack = exp_filt (rms2, tau_fa, Fs/frame, initVal2);

%Release
tau_fr = .075;
Frelease = exp_filt (rms2, tau_fr, Fs/frame, initVal2);

fast_param = max (Fattack, Frelease);
Fast_Param = 20.*log10(fast_param); %convert to dB

%Fast control gain
%Find levels above upper threshold (in dB)
fixed_high = thresh + (thresh2-thresh)/3 + 6;
loud = find (Fast_Param > fixed_high);

%calculate fast gain (in dB)
Fgain = zeros(1, nframes);
Fgain(loud) = fixed_high - Fast_Param(loud);

%convert to linear
fast_gain = 10 .^ (Fgain ./ 20);

%pad gain vector to match size of signal
len = length(fast_gain);
pad = size(del_slow_out,2) - len;
fast_gain = [fast_gain  fast_gain(len).*ones(1,pad)];

%Apply fast variable gain to delayed buffered signal
fast_out = zeros (size(del_slow_out));
for j = 1:frame
   fast_out(j,:) = fast_gain .* del_slow_out(j,:);
end;
out = fast_out(:);
```

```
function [snr, sig] = gold_snr (speech, noise, Fs);
% speech = input speech waveform
% noise = input noise waveform
% Fs = sampling frequency(Hz).
% Ivan Aguayo, 2000.

%Make inputs same size
speech = speech(:);
noise = noise (:);
dif = diff( [length(speech) length(noise)] );
if dif > 0
   speech = [speech ; zeros(dif,1)];
elseif dif < 0
   noise = [noise ; zeros(-dif,1)];
end;

%added signals for optional output
sig = speech + noise;

%Buffer signals into frames and do RMS calculation
frame = floor (Fs * .004);  %4ms frames
[buf_speech, remain_speech] = buffer (speech, frame); %buffer speech sig
[buf_noise, remain_noise] = buffer (noise, frame); %buffer noise sig
nframes = size (buf_speech, 2);

speech_rms = rms (buf_speech);  %rms of speech sig
speech_rmsdb = 20*log10(speech_rms);

noise_rms = rms (buf_noise);      %rms of noise sig
noise_rmsdb = 20*log10(noise_rms);

%Lowpas Filter design
tau = 0.003;            %3ms time-constant
beta = tau*Fs / (1 + tau*Fs);
A = [1 -beta];
B = [(1-beta)];

%Apply filter to signal and noise to get envelope
speech_env = filter (B, A, abs(speech_rms));
noise_env = filter (B, A, abs(noise_rms));

%Convert linear envelopes to dB
speech_envdb = 20 .* log10(abs(speech_env + .00001));
noise_envdb = 20 .* log10(abs(noise_env + .00001));
snr = speech_envdb - noise_envdb;
```

```
function y = emphasis (x);
%Provides high-frequency emphasis to signal x
x = x(:);
N = 64;
F = [0, 0.1, 0.11, 0.79, 0.8, 1];
A = [1, 1, 1.05, 3.14, sqrt(10), sqrt(10)];
B = remez (N, F, A);
y = filter (B, 1, x);
```

```
function y = exp_filt (signal, tau, Fs, initial);
% y = exp_filt (signal, tau, Fs, initial)
%
% This function takes in a SIGNAL and low-pass
% filters it with a time constant of TAU (sec).
% Sampling frequency of Fs.
% Optional input INITIAL used as initial
% condition when filtering. Default is 0.
% If SIGNAL is a matrix, filter operates on the
% columns of SIGNAL.

% Ivan Aguayo, 2001

if nargin < 4
  initial = 0;
end;
beta = tau*Fs / (1 + tau*Fs);
A = [1 -beta];
B = 1-beta;
y = filter (B, A, signal, initial);
```

```
function val = rms (x);
% val = RMS (x);
% If input x is a vector frame of data, function
% RMS calculates VAL, its root-mean-square value
% of such vector.
% If input is a NxM matrix, RMS calculates root-mean
% square of every column. Output vector of 1xM

% Ivan Aguayo, 2001

y = sum (x .^ 2, 1);
val = sqrt (y);
```

Appendix B

Passage Rating Instructions

PLEASE READ ENTIRELY BEFORE CONTINUING

In this experiment, you will be presented with a series of speech passages. Some passages will be accompanied by various background noises. For each passage, you will be asked to provide two numerical ratings.

First, you will be asked to judge and rate the INTELLIGIBILITY of the speech on a scale of 0 to 10. By intelligibility, we mean the proportion of words that you understood. For example, if you understood everything that was said, you should give a rating of 10, for 100% intelligible. If you did not understand any of the words, you should give a rating of 0, for 0% intelligible. If you understood some, but not all of the words, select a value between 0 and 10, corresponding to the percentage of words understood.

Second, you will also be asked to judge and rate the sound QUALITY. This rating should be based on your overall impression of both the speech and the background noise. For example, a rating of 10 indicates that the overall sound quality is extremely pleasant. This is true when the speech is clear and the background noise is easily ignored. A rating of 0 indicates that the overall sound quality is extremely annoying. This can be true when the background noise is extremely distracting and unpleasant OR the speech is very distorted, much too loud, or much too soft. If the overall quality is somewhere in between, select a value between 0 and 10.

Each speech passage will be between 15-25 seconds long. Please listen to the entire passage, then enter your ratings and press ENTER.

Today's session consists of one practice plus 100 regular passages. You will be offered an opportunity to take a break after every 25 passages. Press any key when you are ready to start the experiment.

Appendix C

Subjects' Raw Data

Algorithm Number	Name of AGC System
A_1	HT-Dual
A_2	SNR-Dual
A_3	HT-SNR-Dual
A_4	Dual Front-End
A_5	Linear

Table C.1: Raw data of PG (I1 - I8)

Interference Condition	Intelligibility					Quality				
	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
1	10	6	5	3	3	4	2	1	0	0
	0	5	2	2	0	0	3	3	0	0
	3	1	7	4	1	0	0	4	3	0
	6	1	1	9	1	1	0	0	5	0
	8	7	4	6	10	4	2	1	0	0
2	8	10	9	6	6	5	3	1	4	3
	6	10	10	10	8	4	5	7	3	3
	8	5	10	10	10	5	2	5	6	0
	9	7	8	10	10	3	5	4	4	4
	10	10	9	8	10	7	3	6	6	6
3	10	10	10	10	10	6	8	6	7	6
	10	10	10	10	10	7	8	7	6	8
	10	9	10	10	10	6	6	8	7	5
	10	10	10	10	10	7	6	6	9	6
	10	10	10	9	10	7	8	7	3	7
4	10	10	9	9	10	4	8	7	7	5
	10	9	10	10	8	2	7	7	7	7
	8	10	8	10	10	7	4	6	8	7
	10	9	9	10	8	8	7	3	5	7
	10	10	10	10	10	9	9	9	9	7
5	10	10	10	10	10	8	9	8	8	10
	9	10	10	10	10	8	10	4	9	10
	7	10	10	10	10	3	7	7	7	8
	10	9	10	10	10	9	7	8	9	10
	10	10	8	10	10	7	9	4	8	10
6	10	10	10	10	10	9	10	10	8	10
	10	10	10	10	10	9	10	9	9	9
	10	10	10	10	10	7	9	9	9	8
	10	10	10	10	10	9	10	10	10	10
	10	10	10	10	10	8	10	8	9	10
7	10	10	10	10	10	10	8	10	10	9
	10	10	10	10	10	9	8	8	10	9
	10	10	10	10	10	10	10	9	10	9
	10	10	10	10	10	10	10	10	10	7
	10	10	10	10	10	10	10	10	10	7
8	10	10	10	10	10	9	7	8	10	7
	10	10	10	10	10	9	10	9	8	9
	10	10	10	10	10	10	10	9	8	5
	10	10	10	10	10	8	10	10	10	6
	10	10	10	10	10	10	8	9	10	7

Table C.2: Raw data of PG (I9 - I16)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
9	10	10	10	6	7	7	7	7	3	4
	6	10	10	10	9	4	4	7	6	6
	7	10	10	10	10	4	8	6	8	9
	9	8	9	9	10	6	5	10	8	5
	8	10	8	8	10	6	8	4	5	4
10	10	10	10	10	10	7	8	7	9	8
	9	10	10	10	10	7	8	7	7	10
	10	10	10	10	10	8	9	8	9	9
	10	10	10	10	10	7	9	9	6	9
	10	10	10	10	10	8	9	9	9	9
11	5	9	4	7	4	4	6	5	4	3
	8	10	10	5	10	5	5	7	6	1
	9	9	10	9	10	5	4	4	6	3
	5	10	9	10	10	5	4	5	5	3
	10	10	9	9	10	8	6	4	6	1
12	3	6	3	7	1	1	2	3	1	0
	6	7	8	9	9	0	3	4	3	0
	8	8	9	9	9	2	0	2	4	0
	8	9	9	10	10	4	2	0	2	0
	9	8	9	8	9	5	3	1	0	0
13	5	10	6	5	7	5	5	4	2	1
	3	6	8	4	10	5	5	7	2	0
	6	4	10	10	10	2	2	6	6	0
	9	9	4	9	10	3	3	3	5	0
	10	7	8	4	10	3	4	2	4	3
14	7	7	6	6	4	4	2	4	3	0
	7	10	10	6	10	3	3	4	5	2
	9	9	10	10	9	5	4	4	5	0
	10	8	8	10	10	4	5	4	3	1
	7	3	8	7	10	3	2	4	3	0
15	10	7	7	9	9	5	4	2	1	0
	1	7	4	8	6	0	4	5	2	0
	9	10	10	10	10	4	1	4	5	3
	7	8	8	10	10	2	3	1	4	1
	10	10	7	9	10	5	4	4	1	1
16	9	9	8	10	9	5	4	7	3	1
	9	10	10	9	10	2	4	3	6	2
	10	8	10	10	10	6	4	5	5	1
	9	9	10	10	10	7	4	2	5	0
	10	10	10	9	10	5	6	5	5	2

Table C.3: Raw data of PG (I17 - I20)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
17	7	3	7	2	6	4	5	4	0	2
	4	6	4	6	8	5	3	5	3	0
	4	1	6	3	8	0	4	3	4	1
	6	6	3	7	8	3	1	5	2	1
	4	4	5	4	10	3	4	1	3	0
18	8	10	9	6	8	5	5	4	1	2
	8	7	8	9	6	2	4	6	5	0
	7	7	6	9	10	1	2	4	4	2
	9	8	7	9	9	4	3	3	3	4
	7	10	7	6	8	5	4	1	2	1
19	4	4	3	4	7	1	1	4	1	0
	2	5	5	3	5	1	1	3	3	1
	3	3	5	3	6	0	0	1	1	1
	1	5	3	4	5	3	1	1	1	0
	4	4	5	4	7	1	3	1	1	0
20	7	7	9	8	9	4	4	6	1	2
	7	5	9	8	9	3	2	6	7	1
	8	7	10	6	10	2	3	4	4	3
	8	10	5	6	10	7	2	2	2	2
	6	7	9	6	10	4	6	2	3	3

Table C.4: Raw data of MG (I1 - I8)

Interference Condition	Intelligibility					Quality				
	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
1	3	8	5	3	8	2	4	3	1	0
	1	1	5	2	3	0	2	3	3	1
	2	0	2	4	4	0	0	2	3	3
	3	3	6	2	5	4	1	2	3	3
	5	3	3	0	4	3	3	2	0	2
2	2	5	5	3	5	2	3	4	2	2
	4	4	6	3	5	2	2	4	4	1
	3	4	2	6	8	2	3	1	4	4
	5	3	3	1	6	4	1	3	2	4
	8	8	3	5	2	4	4	2	3	1
3	6	10	9	6	5	5	9	8	6	3
	5	6	7	8	7	3	6	7	7	4
	4	7	6	7	8	3	5	7	7	8
	9	7	7	3	9	8	4	4	7	6
	8	8	7	5	10	7	6	4	2	9
4	1	7	2	4	4	7	7	4	3	2
	3	0	5	3	3	4	4	6	4	3
	2	5	7	7	4	3	4	7	6	7
	3	5	3	4	7	4	3	2	0	7
	7	5	2	4	7	6	6	3	3	4
5	6	8	6	7	9	4	7	4	6	8
	4	6	6	8	7	3	4	7	6	6
	5	4	6	4	8	4	2	5	4	6
	8	7	8	8	7	6	4	7	7	7
	5	8	5	4	8	4	5	4	3	7
6	6	9	7	5	10	6	9	5	2	9
	7	7	9	8	7	4	6	9	6	6
	6	9	3	9	8	4	8	3	9	5
	2	3	10	5	0	2	3	2	5	0
	8	8	6	7	10	8	6	4	5	9
7	6	9	9	8	8	6	9	9	8	8
	6	8	9	9	8	6	7	8	8	7
	7	7	8	9	9	6	6	8	8	8
	9	7	7	6	10	8	7	6	6	10
	7	10	8	6	6	6	9	7	5	7
8	9	9	9	9	8	9	8	9	9	8
	8	10	9	10	9	8	10	8	10	9
	8	8	9	9	9	8	7	8	9	9
	10	9	7	8	9	10	8	7	8	9
	9	10	9	9	10	9	10	9	8	10

Table C.5: Raw data of MG (I9 - I16)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
9	4	3	2	1	4	3	4	4	5	0
	2	4	6	4	2	3	4	7	4	4
	3	7	4	2	5	4	3	4	3	3
	7	0	5	3	7	4	0	0	3	6
	1	5	2	2	7	4	4	5	4	5
10	3	8	7	8	10	3	8	4	8	9
	7	4	8	7	1	5	3	7	4	10
	6	7	4	8	9	6	5	3	6	8
	2	9	6	3	9	4	8	5	4	9
	7	8	9	8	5	7	5	8	6	5
11	2	5	5	3	5	2	4	3	3	3
	4	5	6	2	7	3	4	4	3	3
	6	7	7	8	4	4	3	3	4	3
	5	2	6	6	7	4	3	4	5	4
	8	4	4	5	7	4	3	3	4	3
12	1	5	2	1	5	2	3	3	1	2
	4	1	6	2	4	1	1	3	2	2
	2	4	5	3	2	3	1	2	3	1
	3	2	5	5	5	2	1	1	2	3
	7	2	3	4	1	4	2	3	1	0
13	1	7	2	1	3	2	3	2	0	0
	2	3	4	2	1	0	2	4	3	0
	4	3	7	5	6	2	2	2	4	2
	1	2	1	2	7	3	1	0	1	4
	4	6	3	4	3	3	4	2	1	1
14	1	1	3	3	7	1	1	3	1	2
	3	2	1	5	5	0	2	3	3	3
	3	4	3	8	5	1	3	2	7	3
	3	4	4	3	7	2	4	1	1	4
	3	4	2	3	8	4	3	1	2	2
15	3	1	2	3	6	4	2	2	2	5
	5	3	1	2	6	4	4	3	1	3
	4	7	5	2	7	3	6	5	4	3
	3	7	4	4	1	2	4	3	4	2
	2	7	3	5	3	4	2	2	4	4
16	2	2	7	3	4	3	4	6	1	2
	4	2	4	7	6	3	3	3	5	4
	3	5	1	4	7	1	3	2	3	7
	8	6	4	4	7	5	3	2	4	3
	3	9	5	5	4	3	7	2	3	4

Table C.6: Raw data of MG (I17 - I20)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
17	0	2	2	2	7	0	2	3	2	0
	1	8	2	3	5	0	0	2	3	2
	0	1	0	2	5	0	0	0	2	3
	2	4	0	0	4	2	3	0	2	3
	2	3	1	0	8	2	3	0	0	1
18	6	2	4	6	6	0	2	3	3	2
	4	1	7	5	5	0	1	3	3	2
	4	4	6	3	6	2	0	2	3	3
	5	4	3	0	7	3	3	0	1	3
	6	5	7	4	8	3	3	2	1	4
19	2	4	6	3	3	2	2	2	1	1
	1	2	6	3	2	0	2	2	2	1
	1	2	2	4	2	1	1	1	2	2
	1	6	2	1	5	2	3	0	1	3
	4	2	2	3	2	3	2	2	2	1
20	7	8	4	6	8	4	4	4	4	4
	6	8	7	6	9	3	6	4	4	4
	4	2	4	8	7	3	2	1	4	4
	7	7	5	5	9	5	5	3	4	5
	4	5	6	2	6	3	4	4	3	2

Table C.7: Raw data of RG (I1 - I8)

Interference Condition	Intelligibility					Quality				
	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
1	9	8	7	9	4	4	3	2	3	0
	2	2	4	4	0	3	1	2	1	0
	4	7	4	5	3	2	4	2	2	0
	7	8	8	5	4	3	3	4	1	1
	7	7	7	5	5	4	2	3	3	1
2	8	10	8	9	8	4	4	5	4	2
	10	10	10	9	10	5	6	4	8	4
	9	9	8	8	8	4	5	3	3	6
	8	9	8	9	9	8	4	4	4	3
	10	10	10	10	9	5	8	5	6	3
3	10	10	10	10	10	5	5	8	6	6
	10	10	10	10	10	6	5	5	9	4
	10	10	0	10	10	6	7	5	5	7
	10	10	10	10	10	9	6	7	5	5
	10	10	10	10	10	5	9	6	6	4
4	10	10	10	10	10	9	8	10	8	9
	10	10	10	10	10	9	9	9	10	9
	10	10	10	10	10	8	8	9	10	10
	10	10	10	10	10	10	9	9	9	9
	10	10	10	10	10	9	10	8	9	10
5	10	10	10	10	10	9	10	8	10	10
	10	10	10	10	10	10	10	10	10	10
	10	10	10	10	10	10	10	10	10	10
	10	10	10	10	10	8	10	10	10	10
	10	10	10	10	10	10	10	10	10	10
6	10	10	10	10	10	10	10	10	10	10
	10	10	10	10	10	10	10	10	10	9
	10	10	10	10	10	10	10	9	10	8
	10	10	10	10	10	9	10	8	10	8
	10	10	10	10	10	10	10	9	10	8
7	10	10	10	10	10	9	10	7	9	8
	10	10	10	10	10	10	10	10	10	8
	10	10	10	10	10	10	10	10	10	4
	10	10	10	10	10	10	10	10	10	9
	10	10	10	10	10	10	9	10	10	8
8	10	10	10	10	10	10	10	10	10	8
	10	10	10	10	10	9	10	10	10	8
	10	10	10	10	10	10	10	10	10	8
	10	10	10	10	10	10	10	10	10	8
	10	10	10	10	10	9	9	9	10	4

Table C.8: Raw data of RG (I9 - I16)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
9	10	10	10	10	10	9	10	8	9	8
	10	10	10	10	10	6	9	9	5	8
	10	10	10	10	10	9	8	9	9	6
	10	0	10	10	10	8	10	8	9	8
	10	9	10	10	10	9	7	9	9	8
10	10	10	10	10	10	9	10	8	10	8
	10	10	10	10	10	8	9	8	9	8
	10	10	10	10	10	8	10	8	8	4
	10	10	10	10	5	8	9	9	9	5
	10	10	10	10	10	9	8	9	10	7
11	9	9	6	9	10	4	5	3	4	4
	9	8	10	7	7	5	4	5	5	3
	10	9	10	8	7	5	5	5	5	3
	8	9	9	10	9	8	4	6	5	3
	9	10	10	10	9	5	5	4	7	4
12	8	7	7	8	3	4	4	3	4	0
	8	8	9	8	7	4	4	3	3	1
	8	8	8	7	5	4	4	3	4	0
	8	10	8	8	3	3	5	5	4	0
	9	9	0	9	8	4	3	5	4	1
13	4	9	9	10	8	4	4	7	4	3
	10	10	10	9	8	5	4	5	8	2
	10	8	9	9	8	4	4	4	4	2
	9	9	9	9	5	7	4	4	4	1
	9	9	10	8	9	4	8	10	4	3
14	10	8	10	10	10	5	5	7	5	4
	10	10	10	10	10	4	5	7	5	3
	10	10	10	9	8	5	5	5	5	3
	10	10	10	10	8	8	6	4	5	3
	9	7	10	10	10	5	3	5	6	3
15	9	7	9	8	5	4	3	3	4	2
	6	7	9	7	4	4	4	3	2	1
	9	7	9	9	9	5	4	5	4	3
	9	8	8	5	5	5	3	4	10	1
	9	10	10	10	10	4	5	5	6	3
16	10	8	7	9	8	5	3	5	3	3
	8	10	9	7	8	4	5	4	3	2
	10	10	10	10	7	6	6	5	4	3
	9	9	8	8	8	5	4	5	4	2
	9	10	9	10	9	4	8	5	6	3

Table C.9: Raw data of RG (I17 - I20)

I_j	A_1	A_2	A_3	A_4	A_5	A_1	A_2	A_3	A_4	A_5
17	5	6	3	3	3	1	2	5	1	0
	7	8	9	8	5	2	3	4	3	1
	5	5	4	5	3	2	1	1	1	3
	8	6	7	8	7	5	1	2	2	2
	8	5	5	8	7	3	5	2	3	2
18	10	9	8	10	10	4	4	4	3	3
	8	9	9	8	9	3	4	3	4	2
	10	10	10	10	9	4	4	4	5	4
	10	10	9	9	10	8	4	4	4	3
	10	8	10	9	8	5	5	3	4	2
19	7	8	7	7	7	2	3	3	3	2
	6	6	7	5	7	2	2	2	3	3
	3	8	4	7	3	2	3	1	2	0
	5	5	7	5	7	2	3	3	1	2
	7	5	4	5	5	2	3	2	2	2
20	10	9	8	10	10	4	3	8	4	3
	10	10	9	8	10	4	4	3	8	3
	10	10	9	4	9	2	3	3	4	8
	8	10	10	10	10	8	4	4	4	5
	10	9	10	10	10	4	9	3	3	3

Bibliography

[1] Cox, RM, and Alexander, GC. *The abbreviated profile of hearing aid benefit.* Ear Hear. 1995; 16: 176-186.

[2] Dillon, Harvey. *Tutorial. Compression? Yes, But for Low or High Frequencies, for Low or High Intensities, and with What Response Times?* Ear and Hearing. 1996.

[3] Festen, JM, Van Dijkhuizen, JN, and Plomp, R. *Considerations on Adaptive Gain and Frequency Response in Hearing Aids.* Acta Otolaryngol (Stockh). 1990; Suppl. 469: 196-201.

[4] Martin, GR, Dix, AK, and Zurek, PM. *Comparison of Automatic Gain Control Algorithms.* Unpublished manuscript.

[5] Moore, BCJ. *Signal processing to compensate for reduced dynamic range.* Recent Developments in Hearing Instrument Technology, edited by Beilin, J and Jensen, GR. 1993.

[6] Moore, BCJ and Glasberg, BR. *A comparison of four methods of implementing automatic gain control (AGC) in hearing aids.* British Journal of Audiology. 1988; 22: 93-104.

[7] Moore, BCJ and Glasberg, BR. *A model of loudness perception applied to cochlear hearing loss.* Auditory Neuroscience. 1997; 3: 301-319.

[8] Stone, MA, Moore, BCJ, Alcantara, JI, and Glasberg, BR. *Comparison of different forms of compression using wearable digital hearing aids.* Journal of the Acoustical Society of America. 1999; 106: 3603-3619.

[9] Stone, MA, Moore, BCJ, Wojtczak, M, and Gudgin, E. *Effects of fast-acting high-frequency compression on the intelligibility of speech in steady and fluctuating background sounds.* British Journal of Audiology. 1997; 31: 257-273.