Do You ‘ear Wha’ I ‘ear?:

Lowering Voice Frequencies to Improve Hearing Assistance

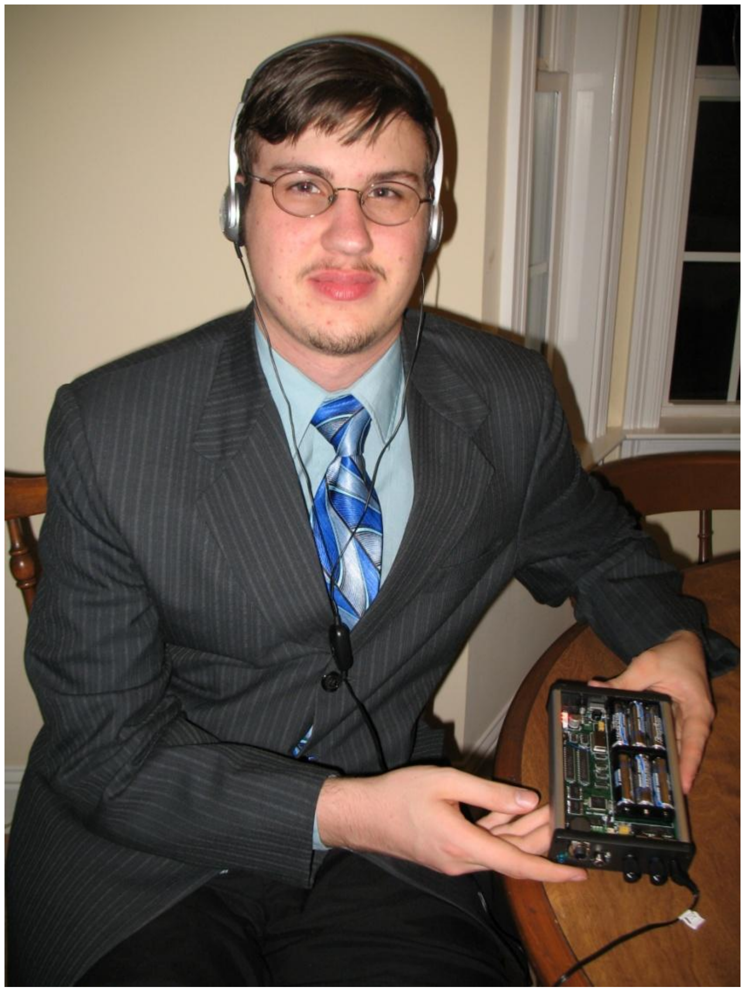

Nicholas M. Christensen, an Intel STS Finalist and four-year Intel International Science and Engineering Fair veteran, won 2nd place in the world in Computer Science for his Do You ‘hear Wha’ I ‘ear? project (© 2009) after winning 4th place for preliminary work in 2008. The project also took 2nd place in the Armed Forces Communications and Electronics Association high school science competition is 2008. He is now pursuing a patent for his revolutionary algorithm for hearing assistance technology. This year he is attending the University of Alabama at Huntsville, majoring in physics and computer science. He hopes to eventually do high-level research with quantum computing and/or possibly work with NASA.

I hear like an 85-year-old man, but I am not alone. Twenty-five million Americans are already affected by hearing loss ("Hearing lost statistics"), and this staggering number is expected to double by 2050 (qtd. in Schmid), especially considering how many students are currently damaging their ears by the combination of loud music and earphones. What they do not realize is that sound has a physical force that damages the stereocilia, the delicate hair cells in the cochlea that pick up vibrations. Once broken, those cells do not regenerate. The vast majority of people can expect hearing damage as they age. Others, like me, have damage from ototoxins; life-saving drugs like the ones that saved my life as a premature infant can cause unfortunate hearing impairment. That is the personal problem that led to my two-year science project, Do You 'ear Wha' I 'ear?, which explores the revolutionary concept of lowering sounds in pitch rather than simply making them louder. Current hearing aid technology is still based on increasing the volume; however, I know from personal experience that hearing aids really do not work well.

In studying audiology, I found that the stereocilia actually respond to different frequencies at different locations: those near the ear canal pick up the higher frequencies while those deeper inside the cochlea pick up lower frequencies as seen in Figure A below.

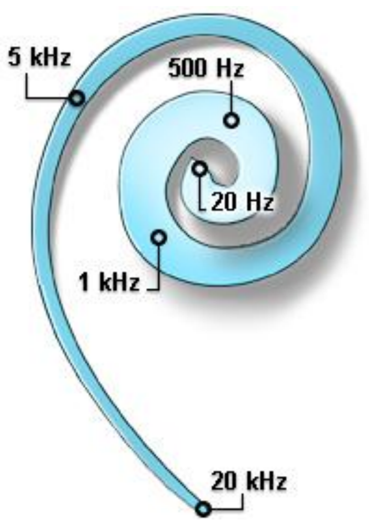

Figure A. The cochlea picks up high frequencies (20 kilohertz) at the base and low frequencies (20 hertz) further inside (Grau).

Because the area associated with high frequencies is more likely to get damaged, people with hearing loss tend to experience the most loss in high-frequency sounds, such as *s*, *sh*, *th*, and *f*. The most common type of hearing loss is presbyacusis, which is defined as high-frequency loss associated with aging (Gates and Mills). Therefore, a hearing-impaired individual may actually have more difficulty **distinguishing** similar-sounding words, such as *math*, *mass* and *mash*, rather than **hearing** them. In fact, he may be able to hear low frequencies normally. My basic hypothesis was that "lower, not louder" may be a more effective way to develop hearing technology.

For the first phase of my project, I wrote a computer program to omit sections from the sound waves of recorded voices, and then stretch the wavelength, making them lower in frequency. What is important to understand is that

$$f = \frac{\lambda}{v}$$

Frequency (*f*) is the number of wave cycles per unit of time, typically a second, which humans perceive as pitch. Frequency is inversely proportional with wavelength (λ), meaning that as one increases, the other decreases, given that the speed of the wave (*v*) remains the same. Figure B below illustrates a 25% decrease in frequency.

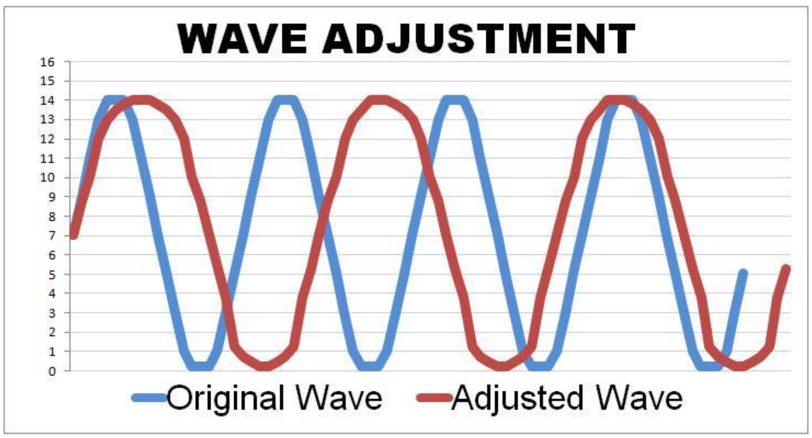

Figure B. The blue wave represents a sound wave with four cycles per unit of time, whereas the red wave has been stretched to include only three cycles in the same time period.

In order for a computer to understand a continuous wave, it must be broken into discrete parts through sampling: plotting the wave at points in time by a series of numbers called samples, each of which describes the wave at a different point in time. This is similar to the concept of finding the area under a curve in calculus, using Riemann sums. Figure C below compares a low and high sampling rate.

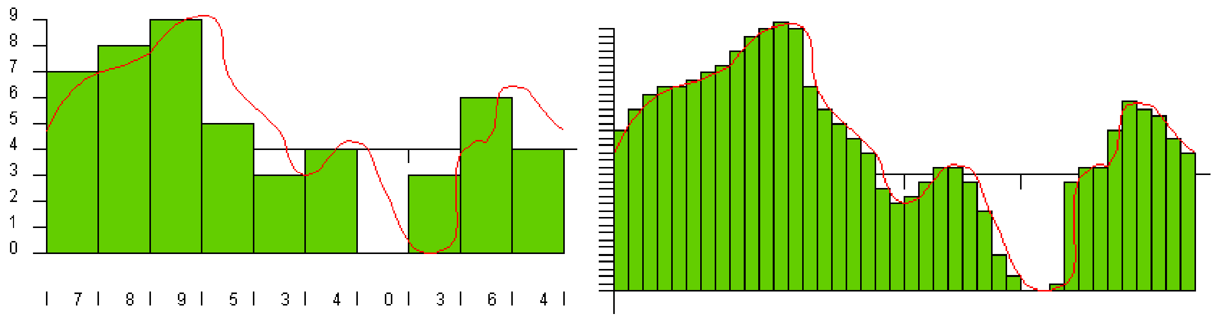

Figure C. The left graph has fewer samples per unit of time compared with the right graph. More samples means that the digital representation of the wave is better fitted to the wave itself ("Why does").

Obviously, a higher sampling rate is preferred for closer approximation. The Nyquist-Shannon sampling theorem states that "For lossless digitization, the sampling rate should be *at least twice* the maximum frequency responses" (Marshall). This is because if a point were selected for sampling at the same location for every cycle, it would give the false impression that the wave is a flat line. The standard sampling rate for audio files is 44,100 samples per second (Brain). Stretching the waves allows them to fit back into the same time frame as the original computer .wav files.

I started learning computer programming by using Pascal, a dated programming language but one that is easy for a beginner. One of the first concepts a programmer must realize is that computers are not intelligent; in fact, arithmetically, they can only do a few operations: addition, subtraction and comparison. In order to multiply, they have to add over and over. To divide, they must subtract repeatedly. The comparison capability is simple but an important tool; they can compare two numbers to check if they are equal (=), greater than (>) or less than (<). That is how the processor can follow the digitized curve of the wave: two sampling numbers are compared. If the second one is larger, the curve is rising; if lower, the curve is falling. As I needed to count wave cycles for my program, I divided them into quadrants, illustrated in Figure D, as they fell below zero, rose, rose above zero, then began to fall again:

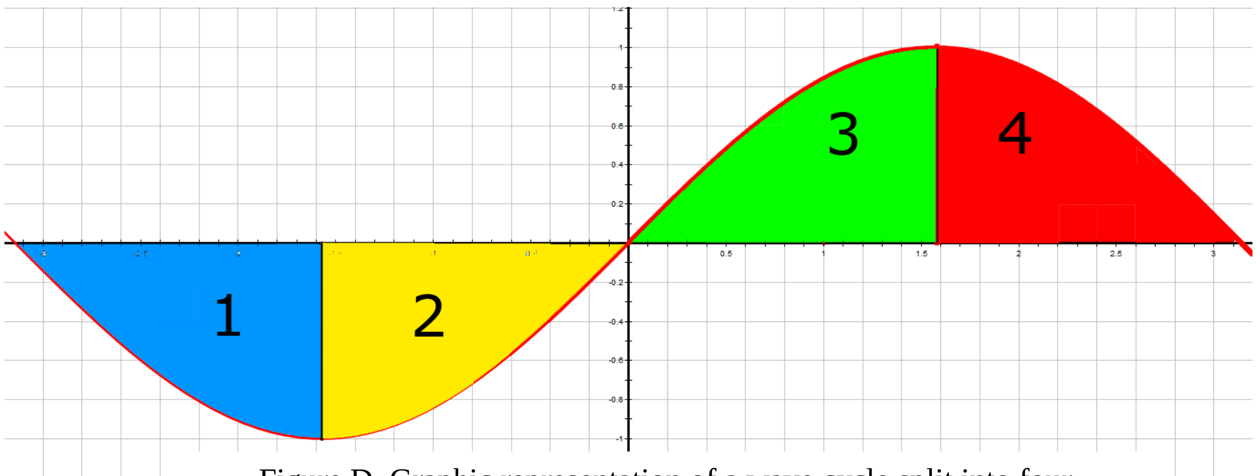

Figure D. Graphic representation of a wave cycle split into four quadrants, each defined by rising or falling, above or below zero.

Once the wave cycles were counted, they could be processed for omissions, and the .wav files were adjusted to a lower frequency, I tested 120 subjects for word recognition when listening to the lowered words containing voiceless phonemes (*s*, *sh*, *th*, and *f*). The results were encouraging, with overall 6% improvement and individual improvements as high as 65%. I personally have a 40% bilateral hearing loss. Although I did not include my own test results in the overall findings, when I took the test, I had a 72% increase in understandability, from seven misses at normal range to no misses in the 25% lowered range (Christensen). The findings supported my hypothesis that lowering frequency could be a helpful tool in hearing assistance technology.

Because the results of my first year's testing were so promising, during the second year I worked with an engineer to develop a customized circuit board, using a 40MHz microprocessor to process at five million instructions per second. This allows the program to bypass the sampling step because it is not recorded in a digital file but is automatically converted back to an analog wave.

Before I could work with programming, I had to understand that computers use the 0s and 1s of base-two mathematics, known as binary, in which each positional value multiplies by two as it moves left. In other words, reading right to left, the first value equals 1, the second equals 2, the third equals 4, the fourth equals 8, and so on. Using the method shown below, the binary number 01010001100 is equivalent to the decimal number 652.

Binary 0 1 0 1 0 0 0 1 1 0 0

Decimal $0\times2^{10} + 1\times2^{9} + 0\times2^{8} + 1\times2^{7} + 0\times2^{6} + 0\times2^{5} + 0\times2^{4} + 1\times2^{3} + 1\times2^{2} + 0\times2^{1} + 0\times2^{0} = 652$

Processors approximate the values not by an electric circuit that is on for one and off for zero but by higher and lower voltages; for silicon, the charges start to go through at 0.3 volts with the current rising considerably at 0.7 volts. At low voltage, the transistor acts like a closed gate that is it is turned off, while at a higher voltage it is open or on. These two states correspond to the two digits in binary: 0 or 1 (“Gates”; Kuphaldt). Figure E below is a graphic representation showing how the transistors work to perform basic operations using electricity:

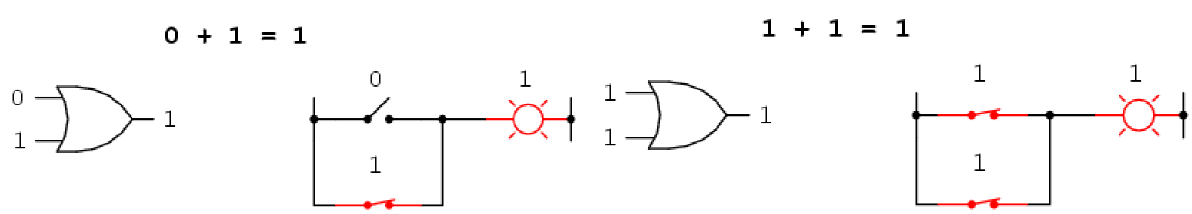

Figure E. Two possible combinations of electrical circuitry in a gate, showing the gates as opened or closed (Kuphaldt).

In working with an eight-bit processor, it is also important to understand the concept of low bytes and high bytes. A sixteen-bit value has to be split, needing extra lines of code. The first eight places from right to left are designated as the low byte and are handled separately from the high byte, the second eight, as in the computer code in Figure F, dealing with different bytes.

incf LOW variable, f
btfsc STATUS, Z
    incf HIGH variable, f

Figure F. The command incf tells the processor to add one to the variable. The command btfsc checks the variable to see if it is equal to zero. If it is zero, then proceed to add one to the high byte; if it is not zero, then skip the next instruction (Predko 390-392).

The sound waves are sampled by an analog-to-digital converter (ADC) so that the computer has binary numbers to work with. In my project I also had to have a digital-to-analog converter (DAC) to convert those numbers back into sound. The general plan in block diagram in Figure G below shows the primary elements for my customized EarMeNow circuit board. Besides the ADC and the DAC, it contained as well as a microphone input, operational amplifiers, buffer RAM, clocks, a battery power supply, and a headphone output.

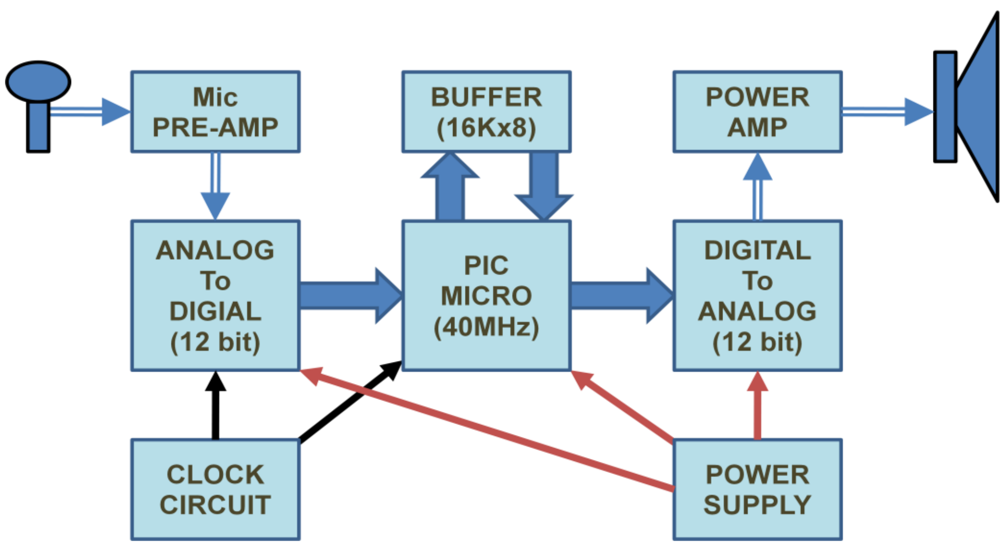

Figure G. This block diagram shows the main components and connections needed for my circuit board, including the PIC18F8722 microprocessor.

My EarMeNow hearing assistance device works in real time because the clocks are synchronized with the same percentage of time adjustment as of frequency adjustment. This means that as the digital-to-analog converter outputs the results, it also expands the wave and lowers the pitch for the listener. The microprocessor works so fast that the time lapse between input and output is only 50 milliseconds, which most people cannot perceive.

In order to test the device, I recorded four voices (male high, male low, female high, female low) saying similar-sounding words with voiceless phonemes, such as ***s***ick, ***st***ick, ***th***ick, ***ch***ick and la***sh***ing, la***st***ing, and la***tch***ing. For different adjustment levels from 12.5%-25% lower in frequency, I chose twelve words, three from each speaker, making sure that the words varied with voiceless phonemes at beginning, middle, and end. I then randomly listed the twelve words and saved them into five separate .wav files, each in a different order.

One hundred subjects were tested, ranging in ages from "19 & under" to "80 & over." The largest percentage (82 people—82%) were in the "over 40" category because they were the most likely to have some hearing loss. There were 29 with documented loss, though not necessarily high-frequency loss. (Some had tinnitus or injured eardrums, for instance.) Many people, however, have loss that is undocumented because hearing tests are not as commonly given as eye tests. Figure H below divides the subjects into age groups by percentage.

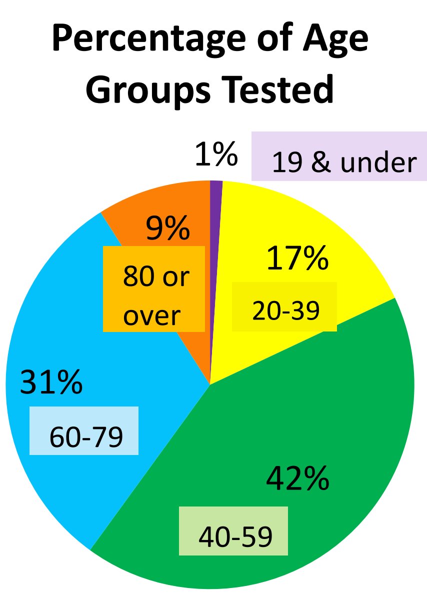

Figure H. The majority of subjects tested were above 40 because they were most likely to have hearing loss due to aging, but one test subject who was 19 also had hearing loss.

Test results found that people with documented hearing loss missed an average of 35% at the normal speech frequency, but dropped to 22%, 24%, 27%, and 25% at lower frequencies as shown by the charts below in Figure I. There was definitely a hearing advantage (up to 13%) gained by lowering the frequency. The optimal frequency adjustment was adjustment 7 or 14.6% lower than normal. In two cases, people with certain kinds of documented hearing problems, such as scarring or ear drum damage, had no benefit. However, they missed none and had no apparent high-frequency loss, so according to the hypothesis, it is not expected for them to benefit.

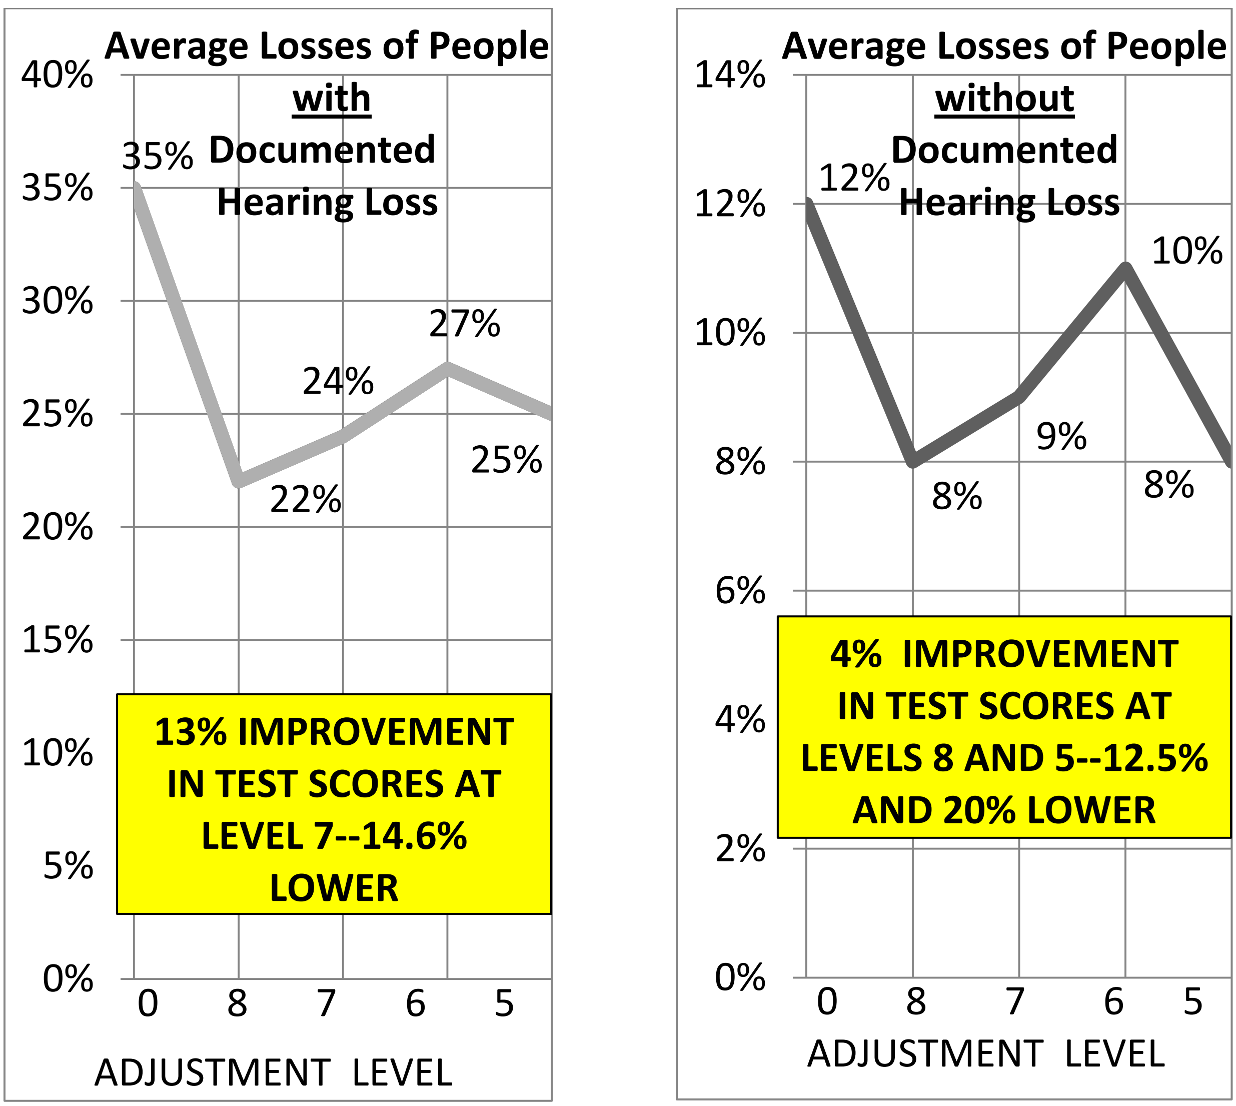

Figure I. The average losses of people with and without documented hearing loss.

People with possible (but not documented) hearing loss missed, on average, 12% at normal and dropped to 8%, 9%, 10%, and 8% at lower frequencies as seen in Figure I. They, too, gained an advantage (up to 4%). However, the optimal frequency adjustment for those with no hearing loss is adjustment 8 or 12.5% lower than normal, as well as adjustment 5 or 20% lower.

Overall, the main point is that almost everyone who was tested seemed to benefit from lowered frequencies. Those with the most significant hearing loss increased their word recognition by the largest percentage, thus supporting my hypothesis. (See Figure J.)

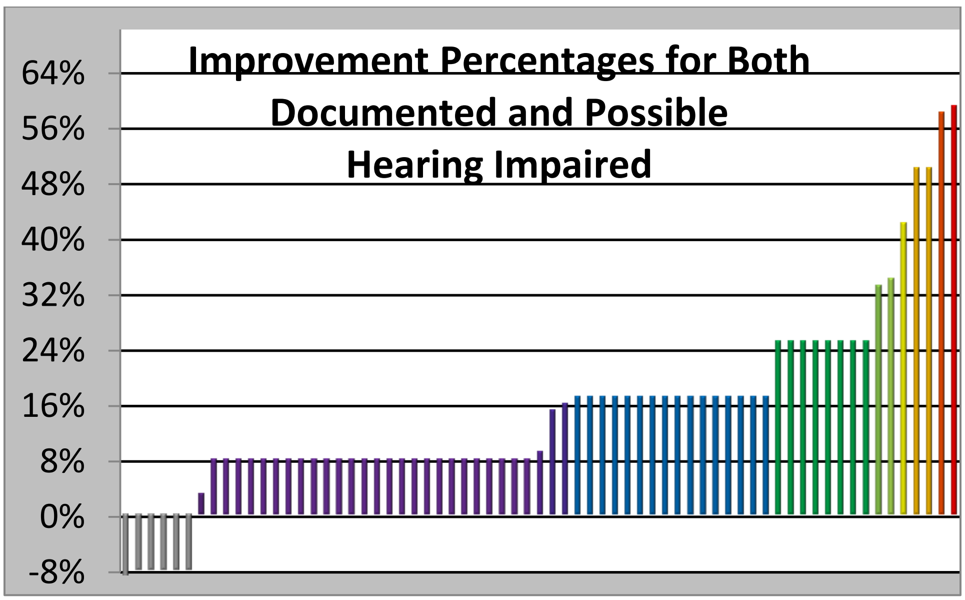

Figure J. Individual tests show that at least two subjects improved nearly 60% from lowering the frequency of voices in real time (see **red** columns). Even though many subjects had **no** documented hearing loss, they still improved around 8% (see **purple** columns), whereas those **with** documented hearing loss generally improved 25% or more in word recognition (see **blue**, **green**, and **orange** columns). Only a few people's scores actually worsened (see **grey** columns), probably because of the distortions due to the omission of information.

Another important point is that the higher-pitched female voice was, as expected, the voice that was most commonly missed during the tests. Interestingly, many subjects who claimed to have no known hearing loss actually misunderstood words with voiceless phonemes, especially those spoken by a female. Notice the percentages of misses of female voice as compared with male voices in Figure K below.

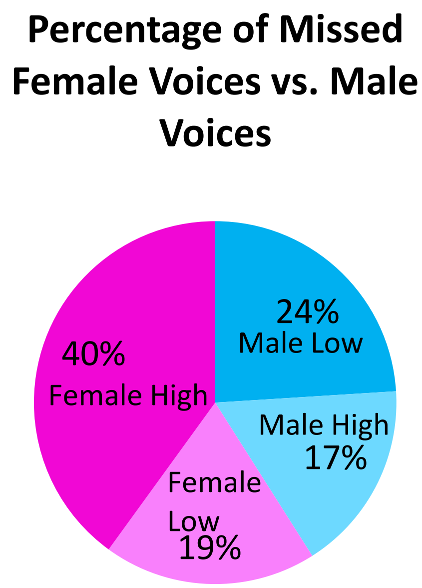

Figure K. The female high-pitched voice was missed the most often at 40% missed. The male low-pitched voice was missed the second most often at 24%, which may have been caused by lowering the frequency too much.

Studying the statistics of the test results is important. It was necessary to have a test group that is clearly defined, sufficient in number and representative of the population. The state and international science fairs impose a minimum of 30 tests for a statistical study. I gave each subject five test sections, for a total of 500. Professional research may deal with 1000s of test subjects in order to obtain more accurate results and note important trends.

Clearly, the results of my testing support my hypothesis that lowering frequencies is beneficial in improving hearing assistance technology. The obvious application is in the form of improved hearing aids, which in my own case would be helpful. However, because this program is very compact and can be run by an inexpensive eight-bit microprocessor, it could be easily integrated into many communication devices, such as radios, telephones, earphones and televisions. It could also be of value to the military and first responders who often rely on two-way radio communication because I found that even people without documented hearing loss benefited somewhat in word recognition. Combining audiology, math, electronics, and programming, I demonstrated that “lower, not louder” is the hearing assistance of the future.

Works Cited

Brain, Marshall. “How Analog and Digital Recording Works.” HowStuffWorks. 1998-2008 22 Sept 2008 <http://communication.howstuffworks.com/analog-digital3.htm>.

Christensen, Nicholas M. Do You ‘ear Wha’ I ‘ear?: Redigitizing Voice Signals into Lower Frequencies to Revolutionize Hearing Assistance Technology. Unpublished, but copyrighted 2008.

“Gates.” U.F.R. Mathématiques et Informatique Département d'Informatique. <http://dept-info.labri.fr/~strandh/Teaching/AMP/Common/Strandh-Tutorial/gates.html>.

Gates, George A., and John H. Mills. "Presbycusis." Lancet 366.9491 (24 Sep. 2005): 1111-1120. Biomedical Reference Collection: Comprehensive. EBSCO. South U, Montgomery, AL. 2 June 2007. <http://search.ebscohost.com/login.

Grau, Harold J. “Hearing.” biologyreference.com 7 July 2007 <www.biologyreference.com/Gr-Hi/Hearing.html>.

“Hearing lost statistics by the National Institute on Deafness and Other Communication Disorders (NIDCD).” NIDCD. National Institute on Deafness and Other Communication Disorders. 23 Jan 2007. 4 June 2007 <http://www.nidcd.nih.gov/health/statistics/ hearing.asp>.

Kuphaldt, Tony R. “Lessons in Electric Circuits – Volume IV.” University of Cambridge. 1 Jan 2004. 24 Aug 2008 <http://www.eng.cam.ac.uk/DesignOffice/mdp/electric_web/ Digital/DIGI_7.html>.

Marshall, Dave. “Nyquist’s Sampling Theorem.” Cardiff School of Computer Science. 2001. 21 Sept 2008 <http://www.cs.cf.ac.uk/Dave/Multimedia/node149.html>.

Predko, Myke. Programming and Customizing the PIC Microcontroller: 3rd Editon. New York: McGraw Hill, 2008.

Schmid, Randolph E. “Aging nation faces growing hearing loss.” Montgomery Advertiser 18 Feb. 2007:4A.

"Why does it say 1-bit Dual D/A converter on my CD player?" HowStuffWorks.com. 23 April 2001. 05 July 2010. <http://electronics.howstuffworks.com/question620.htm>.