DESIGN, VALIDATION, AND IN-FLIGHT EVALUATION OF AN AUDITORY ATTITUDE INDICATOR BASED ON PILOT-SELECTED MUSIC

Douglas S. Brungart, Brian D. Simpson

Air Force Research Laboratory, WPAFB, OH
douglas.brungart@wpafb.af.mil

ABSTRACT

Although all cockpits are currently equipped with visual displays that provide accurate information about the attitude of the aircraft, spatial disorientation continues to be one of the leading causes of aviation accidents. In this paper, we describe the design of an audio display that modifies the acoustic properties of an arbitrary audio input signal (i.e., pilot-selected music) to provide the pilot with supplementary information about the current attitude of the aircraft. Details are provided about how and why the cues were selected, and how they were implemented in a real-time audio system in the aircraft. Results are also provided from laboratory and flight tests that were used to evaluate the performance of the system.

1. INTRODUCTION

One of the leading causes of aviation accidents is spatial disorientation. Spatial disorientation occurs when pilots receive conflicting or misleading information from their visual, vestibular, and proprioceptive sensory systems that causes them to become confused about their physical orientation relative to the earth. This conflicting information is often so compelling that it causes pilots to question or even ignore their flight instrumentation and choose instead to fly on the basis of their own intuition about the true orientation of the aircraft. Sadly, such decisions often lead to tragic consequences.

An obvious way to counter the effects of spatial disorientation is to provide pilots with an additional source of information about aircraft orientation that does not depend exclusively on the visual "artificial horizon" indicator currently used in a typical aircraft cockpit. Ideally, this redundant information should be presented to a non-visual modality to counter the visual-vestibular illusions, such as the oculogravic illusion and the oculogyric illusions, that can impair a person's ability to interpret visual information while under linear or rotational acceleration. Consequently, a number of researchers have proposed the use of a rudimentary artificial auditory horizon to provide pilots with information about the pitch and roll of the aircraft through the manipulation of an auditory signal presented to the pilot through stereo headphones.

Such an auditory artificial horizon was first described by DeFlorez in 1936 [1]. DeFlorez used a continuous tone signal presented over headphones to provide two simple orientation cues to the pilot. The rate of turn of the aircraft was indicated by increasing the level of the tone in one ear and delaying the phase of the tone in the opposite ear, thus changing the apparent left-right location of the tone. The airspeed of the aircraft was indicated by increasing or decreasing the pitch of the tone. These audio cues were somewhat successful, and the investigators reported that it was possible to fly an aircraft in a stable manner for more than 40 minutes while blindfolded solely on the basis of these cues. However, they also noted that these tone-based cues were very fatiguing to the ear, and suggested that a better alternative might be to base the cues on a broadcast radio signal that might be more appealing for the pilot to listen to for long periods of time. They suggested using interaural amplitude cues to manipulate the apparent left-right position of the broadcast radio signal to indicate the rate of turn of the aircraft, and using low-frequency amplitude modulations in the radio signal to convey information about the rate of climb or dive of the aircraft. No specific details were given about the implementation of the proposed system, however, and there is no evidence that it was ever actually tested in flight.

From 1944 to 1945, T.W. Forbes at Harvard University conducted a number of experiments further exploring the use of an auditory attitude indicator in flight [2]. The most successful configuration involved a "three-in-one" sound source that used 1) a repetitive left-right sweeping sound (presumably generated by periodically changing the interaural level difference of the signal) to indicate rate of turn; 2) a variation in the pitch of the tone to indicate bank angle; and 3) a variation in the interruption rate of the tone (causing a "putt-putt" sound) to indicate the airspeed of the aircraft. This display was not tested in flight, but it was shown that it could be used to maintain a level flight pattern in a ground based "Link" trainer.

In 1990, the USAF Aerospace School of Medicine described a more elaborate auditory attitude indicator called an Auditory Orientation Instrument (AOI) [3]. The AOI provided acoustic representations of three critical flight parameters: 1) Airspeed, which was represented by the frequency of a square-wave signal that increased with increasing velocity; 2) Bank Angle, which was indicated by a left-right intensity panning of the sound; and 3) Vertical Velocity, which was indicated by amplitude modulating the envelope of the square wave, with repeated crescendos indicating an increase in altitude and repeated decrescendos indicating a decrease in altitude. These audio cues were found to increase the pilot's ability to maintain a steady airspeed, altitude, and bank angle when no visual cues were present, but not up to the level of performance achieved when visual cues were available.

More recently, Grohn, Lokki and Takala [4] have discussed the use of an auditory attitude indicator for maneuvering through a virtual environment with a 6-degree of freedom flight model. This attitude indicator was based on a 3-D audio display that used Head-Related Transfer Functions (HRTFs) to manipulate the apparent locations of sounds presented to the listener over headphones [5]. In order to determine where to place the virtual sound in their attitude display, Grohn, Lokki and Takala relied on what they called a "ball on a plate" metaphor. In this metaphor, the apparent direction of the sound source was determined by the downward direction a ball would roll if it were located on a plate with the same attitude orientation as the operator's vehicle. Thus one would assume that a plane tilting upwards would result in the perception of a sound source (in this case a pink noise pulsed at a 2.4-Hz rate) directly behind and below the pilot, while a plane rolling to the right would result in the perception of a sound source directly to the right of the pilot. Grohn, Lokki and Takala also reporting using three additional cues to provide the operator with information about the amount of tilt in the aircraft attitude: 1) a gain cue, where the level of the pulsed pink noise increased with the amount of tilt, and the pulsed noise was inaudible when the operator was level; 2) a pitch cue, where a narrow-band noise was added to the stimulus, with the center frequency of the noise varying from 50 Hz to 2 kHz as the amount of tilt increased; and 3) a rate cue, where the pulse rate of noise increased from 0.7 Hz when the operator was level to 8 Hz when the operator was fully tilted. These three conditions were tested in a virtual flight task, and all three were found to result in significantly lower pitch and roll errors than those obtained in a visual-only control condition with the same visual cues but no auditory horizon cue. No difference in performance was found between the three audio conditions, but the "gain cue" condition was preferred because it was the only one where the pulsed noise sound disappeared when the operator was in a level orientation. The subjects considered the other conditions to be "annoying" because they resulted in pulsed noise sounds even when the operator was flying straight and level.

All of these auditory horizon systems have reported some degree of success in terms of their ability to provide useful attitude information to a pilot, but their drawbacks have thus far prevented them from gaining even limited acceptance in the aviation community. In order to understand this failure, it is perhaps useful to look at the attributes that should be present in an ideal auditory horizon indicator and discuss how well these earlier systems have achieved those design goals. Based on this analysis, it will be clear that some prior systems have achieved some of these objectives, but that no previously described system has adequately addressed all of the requirements necessary to make an auditory attitude indicator practical for everyday use in actual aircraft.

1.1. Objective 1: The attitude indicator should provide information about the pitch and roll of the aircraft

Most spatial disorientation accidents are the result of "vestibular illusions" that cause the pilot to "feel" that the plane is in a different orientation than it actually is. For example, the oculogravic illusion is one common vestibular illusion that results when there is a change in linear acceleration; acceleration produces a pitch up illusion while deceleration produces a pitch down illusion. This can result in a fatal accident when, after a night take off and while still accelerating, the pilot falsely senses an excessive pitch angle and compensates with a (unnecessary) pitch down stick input resulting in impact with the ground.

The "leans" is another common vestibular illusion that occurs when the pilot is in a prolonged turn and the vestibular organs adapt to the point that they register the angle of bank used during the turn as being "vertical." When the plane rolls to wings level to terminate the turn, the pilot may perceive this rotation as a bank and turn in the opposite direction. This can cause pilots to lean in an attempt to assume what they think is a vertical posture. The leans can also occur when the pilot performs very slow roll to the left that does not stimulate the vestibular apparatus and then rolls rapidly to the right to level flight. Such a maneuver can generate the false impression that the plane has only rolled to the right.

Since the goal of the auditory horizon display is to provide the pilot with an additional tool to combat these vestibular illusions, it makes sense for the display to focus on the attitude information that is most likely to be compromised by vestibular disturbances, specifically the pitch and roll of the aircraft. It is no coincidence that these are exactly the parameters that are indicated by the critical visual "artificial horizon" display that is prominently displayed at the center of every aircraft cockpit and is considered essential for avoiding spatial disorientation in instrument flight.

Surprisingly, neither of these two parameters was displayed by the system described by DeFlorez, which provided information only about the turn rate and airspeed of the aircraft. All of the other auditory attitude indicators described above provided some means of indicating bank angle or roll, but only one [4] has provided any means to determine changes in the pitch of the aircraft. These systems have instead focused on providing indications of airspeed or vertical velocity. While these parameters are undeniably useful for flying an aircraft, the vestibular system is not particularly sensitive to either of these parameters and thus the probability of a sensory conflict that causes the pilot's "seat of the pants" feeling about of the orientation of the aircraft to contradict the readings on these instruments is not as high as it is for the pitch angle of the aircraft.

1.2. Objective 2: The indicator should have an intuitive "anchor point" for a straight and level flight

Since the primary goal of a spatially disoriented pilot is to restore the aircraft to a "straight-and-level" flight path, a well-designed auditory horizon system should provide a very intuitive way for the pilot to know when straight and level flight has been achieved. For example, the display described by Lyons et al. [3] used left-right amplitude panning to indicate bank angle, which means that the pilot could effectively level the bank of the aircraft simply by trying to "center" the sound in the middle of the head. This contrasts with the system described by Forbes [2], which used the pitch of a tone to indicate bank angle, meaning the pilot would have to memorize an arbitrary frequency to determine if the plane were "straight and level" (a difficult task for all but the small percentage of the population with "perfect pitch"). Of the systems that have provided an auditory cue for pitch angle (all variations of the system described by [4]), only the "gain" configuration provided a clear indication when straight and level flight had been achieved. In that system, the sound was turned off when the aircraft attitude was set to a level pitch angle, thus giving a clear indicator of a "straight and level" flight path.

1.3. Objective 3: The indicator should easily distinguish between pitch up and pitch down orientations

In recovering an aircraft from an unusual attitude, it is obviously extremely important to know whether the aircraft is currently pitched downward or pitched upward. As discussed previously, the only prior system that provides direct information about aircraft pitch is the system by Grohn et al. [4], which provided a pitch cue by rendering a virtual sound source at the position that a ball on a plate would roll if it were tilted with the same attitude as the aircraft. However, prior research has shown that it is very difficult to distinguish between sound sources located at the same lateral angle in front and behind the head. Such sources produce nearly identical binaural cues, and thus are said to fall on the "cone of confusion" with respect to the listener (e.g. see Wallach [6]). Thus one would guess that the listener might easily confuse upward and downward pitch angles in the such a system

1.4. Objective 4: The attitude indicator should be based on a sound source that can be tolerated by the listener for long periods of time, such as the length of a cross-country flight.

Because the auditory indicator only provides a secondary way for a pilot to access information that is already displayed on the visual displays of a standard cockpit, there is no reason to believe that any pilot would want to use an auditory horizon cue unless that cue is based on sound that is pleasant to listen to and not annoying in any way. Indeed, this point was emphasized in every previous report that has described an auditory attitude indicator. DeFlorez [1] noted that the use of a tone-based attitude indicator "is somewhat fatiguing to the ears and its use as a signal seems inadvisable." Forbes [2] noted the importance of producing an audio cue that did not "capture" the attention of the pilot to the exclusion of other information, including radio traffic. Lyons et al. [3] noted that the sounds they used (which were based on an amplitude-modulated square wave) interfered with performance in secondary tasks, and opined that "by tailoring acoustic signals to meet criterion for auditory comfort as well as for discrimination, such perceptual 'overloading' might be avoided. And Grohn et al. [4] noted that their subjects preferred the "gain" condition of the their display because it was the only one where the pulsed-noise cue was turned off in straight and level flight. The subjects considered the other conditions, which required them to listen to pulsed noise for the full duration of their experiment, to be "annoying."

These results seem to suggest that an auditory attitude indicator is unlikely to gain widespread acceptance in the aviation community unless it is based on a signal the pilot actually wants to hear, rather than an artificial signal the pilot has to listen to in order to obtain the useful attitude information. Even a system that only produces sound when the aircraft attitude deviates from straight-and-level flight path is unlikely to be satisfactory, because it will be unable to distinguish between intentional maneuvers (takeoffs, landing, and turns) and unintentional maneuvers (such as slow onset rolls), and thus will likely be producing unwanted sound even when the pilot is engaged in normal flight procedures.

An obvious solution to this problem is to design the auditory attitude indicator in such a way that it can be applied to any arbitrary audio signal that the pilot might select to listen to for extended periods of time on a cross-country flight. Most US drivers choose to listen to music, sports, talk radio, or some other form of auditory entertainment on their car stereos on extended drives, and to a lesser but growing extent many private pilots are now listening to prerercorded music stored on a personal music player such as the Apple IPod when they are making long cross-country flights. If the auditory attitude indicator could be superimposed on top of this kind of user-selected auditory entertainment, pilots might view it as a benefit rather than an annoyance and the likelihood that it would actually be used in practice would be dramatically higher than for previous displays that have been based on tones, noises, square-waves, or other abstract, aesthetically unpleasant sounds. To this point, only DeFlorez [1] has actively proposed the use of an auditory attitude display based on a user selected signal rather than an abstract sound.

As stated earlier, we believe that an auditory attitude indicator should focus on replicating the functionality of the "artificial horizon" and thus should provide aircraft pitch and roll information rather than turn angle and airspeed information. Clearly the left-right panning described by DeFlorez [1] could easily be adapted to present bank information rather than pitch information. However, there is no clear way to use "low-frequency amplitude beats" to indicate pitch information. The biggest problem is that it is a unidimensional parameter that does not provide a way to distinguish between upward and downward pitch angles. It also has the potential to be quite annoying, as it would introduce temporal distortions into the signal that might severely distort the sound of the music.

In this paper, we describe the design, validation, and in-flight verification of an auditory artificial horizon display consitent with DeFlorez's suggestion of adding aircraft orientation cues to an arbitrary audio entertainment signal such as broadcast radio. The system uses a variation of the left-right amplitude panning that has previously been suggested as a means to indicate the bank angle of the aircraft [3], but it improves upon that implementation by adding interaural time delays information and by using a more intuitive mapping of the left right amplitude pan to the bank angle of the aircraft. The display also uses a novel method of indicating aircraft pitch that is based on two types of auditory processing: 1) manipulating the interaural correlation of the stimulus, which effects the apparent width of the auditory image; and 2) introducing a "repetition pitch" into the stimulus, which produces the illusion of a pitch signal even when the source material does not contain significant energy at the fundamental frequency of the apparent pitch. This combination of cues provides a robust and intuitive way to add aircraft pitch information to an arbitrary audio stimulus, and it does so in a way that is unlikely to significantly impair the auditory comfort of a pilot over long periods of continuous listening. Thus it is hoped that the display might improve aviation safety by decreasing the likelihood of slow-onset spatial disorientation in cross-country flights.

2. AUDITORY DISPLAY DESIGN

Figure 1 shows the basic block diagram of the auditory artificial horizon system. The core of the system is the Audio Processor Unit, shown in the center of the figure. The Audio Processor Unit is designed to add attitude information to any arbitrary audio input selected by the user. In the block diagram, this input is assumed to come from an external system (labeled Audio Source) that allows the content of the signal to be selected by the user. This source could be an external CD Player, a satellite radio system, an MP3 player, or any other device for playing back sound.

If the audio source is in stereo, the left and right outputs are first mixed together, and then processed by two seperate selectable 128 point finite impulse response filters (labeled $h_L[n]$ and $h_R[n]$ in the figure). The coefficients of these filters are selected from a lookup table that is driven by the pitch and roll orientations of the aircraft. This information is obtained from the external Aircraft Avionics Systems via a real-time interface such as an RS-232 serial interface. In order to maintain aircraft radio communications, it is necessary for the system to mix the stereo output of the audio horizon display with the audio output of the Aircraft Intercom System. This could be a simple mixing device, or it could be a more

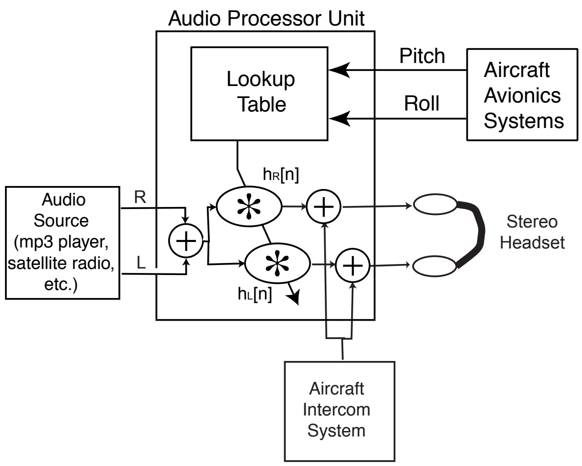

Figure 1: Overall block diagram of the attitude indicator system

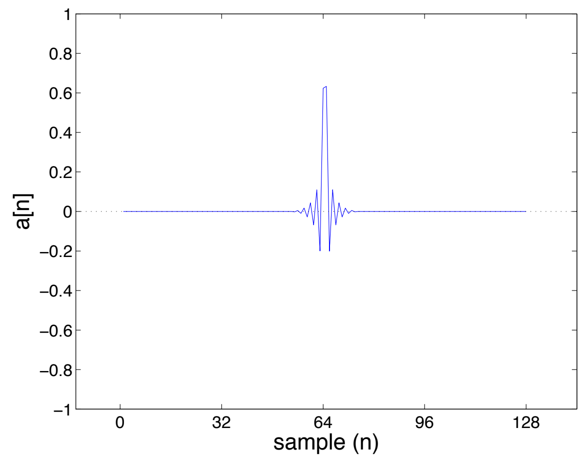

Figure 2: Allpass filter a[n] used as basis for creation of AITFs.

sophisticated system that attenuates or mutes the audio attitude indication signal whenever voice activity is present on the intercom system (thus eliminating the possibility that the attitude indication signal might interfere with normal radio communications in the aircraft). The resulting mixed signal is then played to the pilot via a stereo aviation headset.

Note that the core of the audio processing unit can easily be implemented in software on a standard general purpose personal computer, such as a windows-based PC. The current system has been implemented as a modification of the Sound LAB (SLAB) open-source software package developed by NASA to produce a real-time spatialized audio display [7]. In order to acheive this, the SLAB software is designed to process an audio signal with left and right 128-point FIR filters that are switched in real time in response to the azimuth and elevation of the listener's head. In order to modify the software to be an attitude indicator, we simply changed the FIR filters from the set of azimuth and elevation dependent Head Related Transfer Functions (HRTFs) normally used for 3D audio displays to a set of pitch and roll dependent Attitude Indication Transfer Functions (AITFs) designed to provide an auditory horizon cue to the pilot of the aircraft. In these modified AITFs, we desired a system state where aircraft azimuth would have no impact on the rendered sound, but changes in pitch and roll would consistently modify the audio signal. In SLAB, this can be accomplished by setting the source location directly above the listener, setting the listener's head position to match the azimuth, elevation, and roll of the aircraft, and setting the AITF for each azimuth and elevation location in the HRTF to account for the corresponding unique pitch and roll value of the aircraft. The next section describes the procedures used to generate the AITFs.

2.1. Generation of the Attitude Indication Transfer Functions

2.1.1. Normalized Pitch and Roll Values

One important design parameter in an auditory attitude indicator is the functional relationship between the physical change in the aircraft attitude and the size of the resulting changes in the auditory cues present in the auditory horizon indicator. Although some acrobatic aircraft can operate at virtually any pitch and roll values, most general purpose aircraft are restricted to a limited range of pitch and roll for normal safe operation. However, the size of the range can vary significantly across different aircraft types, and perhaps different aircraft missions (e.g., cropdusting versus passenger ferrying). Thus, in describing the algorithm, we will refer to aircraft attitude only in terms of the non-dimensional normalized roll value β and the non-dimensional normalized pitch value ρ. Under this convention, β is the normalized roll of the aircraft, with a -1 indicating the maximum acceptable roll angle to the left and +1 indicating the maximum acceptable roll angle to the right. Similarly, ρ is the normalized pitch of the aircraft, with -1 indicating the maximum acceptable downward pitch and +1 indicating the maximum acceptable upward pitch.

2.1.2. Baseline all-pass transfer functions

The Attitude Indication Transfer Functions (AITFs) were created by modifying the 128 point symmetric all-pass impulse function a[n] shown in Figure 2. This baseline filter was converted into the frequency domain representation A[ω_k] with an Discrete Fourier Transform (DFT), and then the individual coefficients of A[ω_k] were modified with three different cues associated with the upward or downward pitch of the aircraft:

2.1.3. Low or High Frequency Emphasis

The first cue added to the AITFs was a simple filtering function that emphasized the high frequencies of the AITFs associated with downward aircraft pitch and the low frequencies of the AITFs associated with upward aircraft pitch. This filtering was designed to produce filter slopes ranging from 0 dB per octave for a level aircraft to 6 dB per octave for an aircraft at the maximum safe upward or downward pitch. This cue was implemented by multiplying each coefficient of A[ω_k] by $10^{(log_2 k*6*\rho)/20}$.

2.1.4. Repetition Pitch

The second cue added to the AITFs was a periodic spectral emphasis that was designed to create an apparent pitch in the audio signal. This emphasis increased the magnitude of every 3rd coefficient of A[ω_k] in the pitch-down condition and every 7th coefficient of A[ω_k] in the pitch-up condition by a factor of $1 + 9|\rho|$.

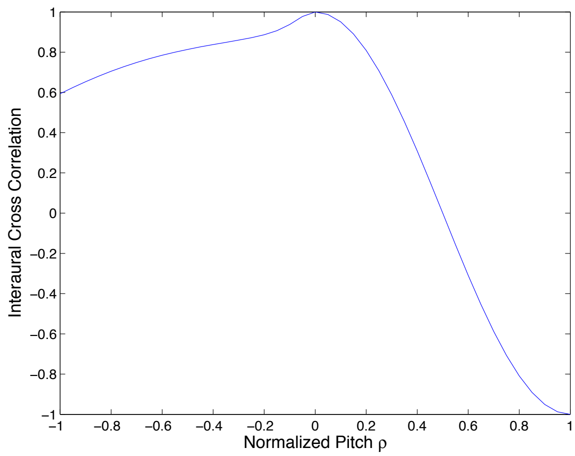

Figure 3: Interaural correlation coefficient for the AITF as a function of normalized pitch ρ.

This creates a linear ramp in amplitude that increases very rapidly a low pitch values, and asymptotes to a maximum amplitude boost of 20 dB at the maximum pitch value of the aircraft.

Note that the advantage of using a periodic spectral emphasis, in addition to a simple low or high pass filter, is that it generates a "repetition pitch" that produced the illusion of a consistent pitch value at the repetition rate (known as the "missing fundamental") even if the processed audio signal only contains energy in a limited range of frequencies that does not include the fundamental harmonic frequency. In this implementation, the boosting of every 3rd coefficient creates a repetition pitch of approximately 1000 Hz in the pitch-up condition, and the boosting of every 7th coefficient creates a repetition pitch of approximately 2400 Hz in the pitch-down condition.

2.1.5. Interaural Decorrelation

The final cue added to the AITFs was an interaural difference cue designed to decorrelate the left and right ear signals for aircraft attitudes above and below the horizontal plane. In order to further differentiate between the audio cues present in the signal in pitch-up and pitch-down attitudes, a different method of decorrelation was used for each condition. In the pitch-up condition, the decorrelation was implemented by introducing a rapidly-changing frequency-dependent interaural phase difference into the signal. In the pitch-down condition, the decorrelation was instead implemented by introducing a rapidly-changing frequency-dependent interaural level difference into the signal. Each technique is described below.

• *Pitch Up: Interaural Phase Decorrelation*

The interaural phase decorrelation was implemented by multiplying the DFT coefficients of the left ear AITF H(ω_k) by $e^{0.5\pi|\rho|(-1)^k j}$ and the DFT coefficients of the right ear AITF H(ω_k) by $e^{-0.5\pi|\rho|(-1)^k j}$. This manipulation resulted in interaural correlation coefficients ranging from -1 for pitch values equal to p_{max} to +1 for pitch values of 0. The perceptual effect of this decorrelation cue was an systematic change in the apparent width of the auditory stimulus, from a compact, punctate source for a pitch of 0° to a diffuse, broad sound source for a pitch of $\pm 10°$ [8].

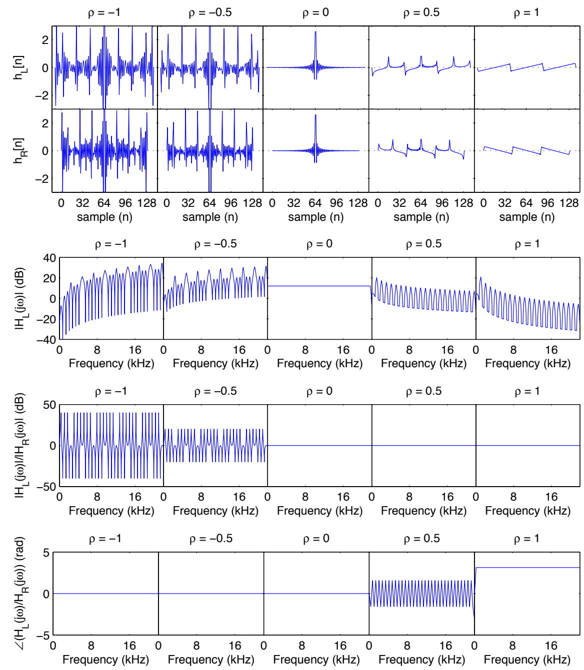

Figure 4: Time and frequency domain characteristics of AITFs as a function of normalized aircraft pitch ρ. See text for details.

• *Pitch Down: Interaural Level Decorrelation*

The interaural level decorrelation was implemented by multiplying the DFT coefficients of the left ear AITF H(ω_k) by $10^{|\rho|(-1)^k}$ and the DFT coefficients of the right ear AITF H(ω_k) by $10^{-|\rho|(-1)^k}$. Note, however, that these manipulations were not applied to the coefficients where the overall intensity was enhanced to produce the repetition pitch cue in the stimulus (every 7th coefficient in the DFT).

Because interaural correlation is more sensitive to phase differences than to level differences, this manipulation did not result in as large a change in interaural correlation coefficient as the phase-based decorrelation cue. As shown in Figure 3, the interaural correlation only decreased to approximately 0.6 when the pitch was set at the largest negative value. However, the perceptual effect of this interaural level decorrelation is similar to that experienced for interaural phase decorrelation: a systematic broadening in the apparent width of the stimulus with decreasing aircraft pitch [8]. The ILD manipulation also results in an apparent high-pitched tonal component, which is why the level decorrelation was paired with the higher repetition pitch employed in the pitch-down condition.

2.1.6. Filter Generation

Once the coefficients of the allpass filter were modified with the three transformations outlined above, they were forced to be complex-conjugate-symmetric and converted back into the time domain with the inverse DFT. They were then equalized in level. This was achieved by convolving each filter with a 1024-pt filter designed to match the frequency content of a sample of pop music, and determining the proper scaling value to equalize the RMS output en-

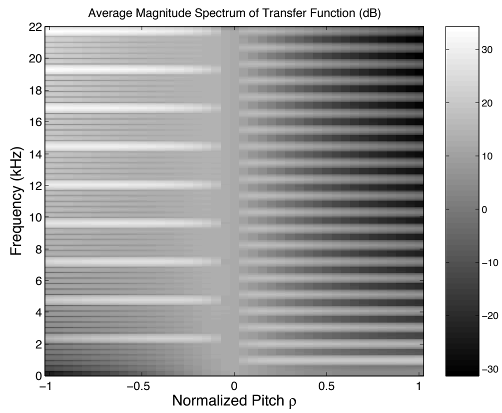

Figure 5: Overall magnitude spectrum of the AITF as a function of the normalized pitch ρ

ergy for this music waveform. Figure 4 shows some examples of the resulting AITFs for an implementation with a maximum pitch of 10°. The top two rows of the figure show the time domain impulse responses $h_l[n]$ and $h_r[n]$ for the left and right ear AITFs. Note that, at extreme positive pitches, the right ear transfer functions are simply inverted versions of the left ear transfer functions. This produces an interaural cross-correlation value of -1. Also note that the positive pitch AITFs appear much lower in magnitude than the negative pitch AITFs. This occurs because they have a low-frequency emphasis and thus have been adjusted to compensate for the relatively greater amount of energy that typically occurs in the low-frequency range of pop music.

The 3rd row of Figure 4 shows the magnitude of the frequency response for the left ear transfer functions. Note the periodic ripples related to the repetition pitch added to the AITFs, with a greater spacing between the ripples for the higher frequency repetition pitch that occurs at positive attitudes.

The 4th row of Figure 4 shows the interaural level difference between the left and right ear filters at each frequency. At negative pitch values, there is a rapidly alternating interaural level difference of up to $\pm$ 40 dB. However, at positive pitch values, where interaural correlation was manipulated by changing the interaural phase only, there is little or no interaural level difference at any frequency.

Finally, the last row of Figure 4 shows the interaural phase difference between the left and right ears. As expected, these graphs show little or no interaural phase difference at negative pitch values (where interaural correlation was manipulated with a rapidly varying ILD), a rapidly varying interaural phase of $\pm\pi$ radians when the pitch is equal to half its maximum value (and interaural correlation is zero), and a constant ILD value of π radians when the pitch is equal to its maximum value (and the correlation coefficient is -1).

Figure 5 shows the the magnitude of the frequency response for the AITFs as a function of the pitch of the aircraft, averaged across the two ears to eliminate the rapidly varying interaural level cues at downward pitch angles. Note the increasingly prominent "ripples" resulting from the introdution of the repetition pitch cue at non-zero attitudes. Also note the general low-frequency emphasis for positive attitude values, and the high-frequency emphasis for negative attitude values.

2.1.7. Roll Cues: Interaural Level Differences

The final component of the AITFs was the interaual level difference cue that was added to provide an indication of the roll value of the aircraft. This level difference cue was implemented by attenuating the AITF in the ear in the direction of the roll by $48sin(\frac{\beta\pi}{3})$ dB, and then normalizing the left and right ear AITFs to have a constant total power independent of the bank angle of the aircraft (i.e. so $\Sigma_k h_L(n)^2 + \Sigma_k h_R(n)^2 = C$, where C is a constant for all possible pitch and roll values of the aircraft).

2.1.8. Subjective Impression

The net perceptual effect of the AITF is an audio signal that systematically changes with the attitude of the aircraft in a way that provides a subtle cue about to the direction to maneuver to restore straight and level flight. When the aircraft is pitched up, the AITFs produce a spatially diffuse sound with a relatively low pitch characteristic indicating the aircraft needs to be maneuvered down in pitch. When the aircraft is pitched down, the AITFs produce a diffuse sound with a relatively high pitch characteristic indicating that the the pilot needs to pull the aircraft up. When the plane banks left, the sound shifts to the right, and and when the plane banks right, the sound shifts to the right. When the aircraft is straight and level, the listener hears a compact, centered sound with no alterations to its original spectral characteristics.

3. LABORATORY VALIDATION

3.1. Methods

Nine paid volunteer listeners with normal hearing participated in the laboratory validation. In this experiment, listeners who were sitting in front of a PC while wearing headphones first familiarized themselves with the auditory cues by listening to sounds while using a mouse to adjust the simulated pitch and roll of the aircraft. Then they were asked to listen to a sequence of sounds and, for each sound, to identify which direction the aircraft should be manuevered in to restore it to level flight. This was done in a five-alternative forced choice task where they were presented with a stimulus that was proessed with the appropriate AITFs to simuluate either a straight and level flight condition (0° pitch and 0° roll), one of ten pitch values ranging from -10° (ρ=-1) to +10° (ρ=+1), or one of ten roll values ranging from -30° (β=-1) to +30° (β=+1). The listeners then had to chose whether they should roll left, roll right, pitch up, pitch down, or make no change in order to restore the aircraft to a straight and level attitude. The same task was performed with two kinds of sounds: white noise, and music from a CD that was selected from the personal collection of that listener. It was also performed with two durations: unlimited, where the listener could listen as long as desired before making a selection, and limited, where the sound was turned off after two seconds. Over the course of the experiment, each listener responded to a total of 5 trials with each combination of stimulus type, stimulus duration, simulated pitch or simulated roll.

3.2. Results

The results of the experiment are shown in Figure 6. Three-way ANOVAs on the percentage of correct responses versus the independent variables of stimulus duration, stimulus type, and simulated roll or pitch revealed that stimulus duration did not have a

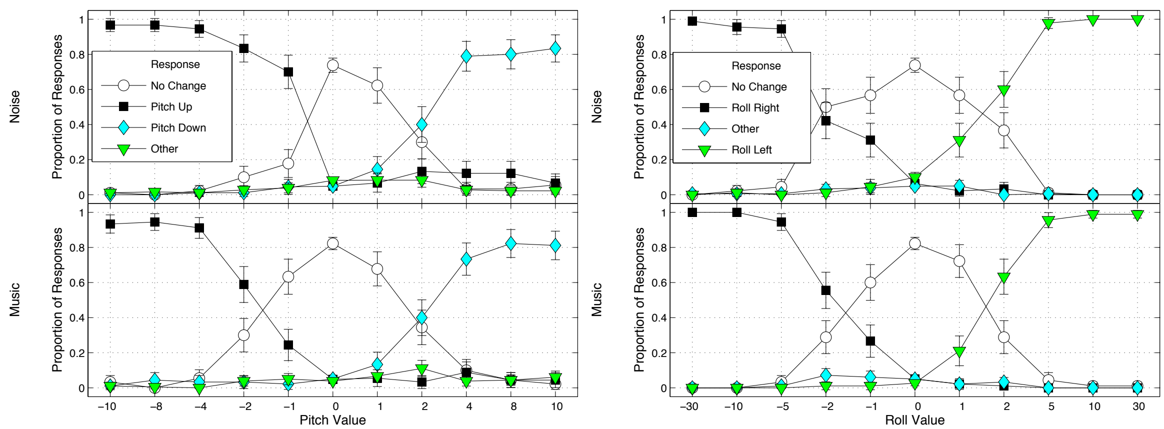

Figure 6: Laboratory Validation of Pitch and Roll Cues for a broadband noise stimulus and for a user-selected musical stimulus. The left panel shows the percentage of correct pitch identifications, and the right panel shows the percentage of correct roll identifications. The error bars show the 95% confidence intervals around each data point.

significant impact on performance in any condition of the experiment. Thus the 2s and unlimited stimulus duration conditions have been collapsed together in the figure. The left panel shows the distribution of subject responses in the Pitch condition as a function of the simulated pitch of the processed stimulus. The results show that listeners were clearly able to discern the simulated pitch of the signal, with correct up or down responses approaching 80% whenever the simulated pitch exceeded $\pm 4^\circ$, and the percentage of correct "no change" responses approaching 80% when the pitch was level. Comparing the noise and music stimulus conditions, the biggest difference is that the music produced a less robust pitch cue than the noise when the simulated pitch was slightly negative (-1° or -2°). It is also apparent in both stimulus conditions that the pitch cues were somewhat less robust when the pitch value was positive than when it was negative.

The right panel Figure 6 shows performance in the Roll condition of the experiment. In both the noise and music conditions, the listeners were able to correctly identify the roll in nearly 100% of the trials when the roll angle was greater than $\pm 5^\circ$. In contrast to the pitch condition, performance in the roll condition was slightly more sensitive in the Music condition than the Noise condition.

While it is true that pitch and roll detection was poor for pitch roll values, near 0°, it is also true that it is desirable for the sound source not to change too rapidly in the range of values where the plane attitude is likely to vary a lot due to turbulence, etc. during normal operation (otherwise the cue might become annoying). Thus, this linear mapping between aircraft attitude and the normalized pitch and roll values was deemed adequate for use in our initial in-flight evaluations of the technology.

4. IN-FLIGHT EVALUATION

After the display was determined to be acceptable in laboratory listening tests, its real-world performance was evaluated in a flight test that was conducted at the NASA Langley Research Center using general aviation aircraft. The full details of these flight tests are beyond the scope of this manuscript. However, a brief summary of the test and results are provided below.

The experiments were conducted with a Cirrus SR-22 aircraft (Figure 7) that was equipped with a laboratory computer running real-time custom audio software based on the NASA SLAB audio rendering software [7]. The input signal consisted of music (typically pop or rock music) that was selected by the individual test pilot prior to the flight. This music was processed according to the aircraft pitch and roll as outlined above, mixed with the aircraft intercom signal with an analog audio mixer, then presented to the subject pilot via stereo ANR headphones (DRE-6500).

A total of 16 pilots participated in this flight test. Prior to the flight, the specifics of the audio horizon display were described to the subject pilot and the pilot was given an opportunity to interact with the audio horizon display using a custom-built, PC-based flight simulator in which attitudinal changes in the aircraft were reflected in the display. Then the pilots were briefed on the safety procedures, familiarized with the aircraft controls, and taken into normal straight-and-level flight by an accompanying NASA safety pilot. The subject pilots were then given an opportunity to gain familiarity with the audio horizon display while navigating the aircraft to the test area.

Each pilot conducted a total of two 1-1.5 h flights with the aircraft. In the first portion of each flight, the pilots conducted an audio navigation task, which has been described elsewhere [9]. In the second leg of the flight, each pilot was blindfolded (Figure 7) and asked to conduct one of two different tasks measuring the effectiveness of audio attitude indicator. These two tasks are described below.

4.1. Change in Attitude Detection

The first task tested how well the pilots were able to detect a slow change in the pitch or roll of the aircraft, with or without the audio display. On each trial, the safety pilot slowly changed the attitude of the aircraft in one axis only (i.e., pitch or roll) at a rate of approximately 1°/sec. As soon as a change in attitude was recognized, the subject pressed a button on a response box and provided a verbal response identifying the axis and direction of change that had occurred (e.g., roll left). The verbal response and response time were recorded, and the safety pilot returned the aircraft to a level attitude before the start of the next trial. A total of 20 trials were run for each subject, 12 with an audio cue and, as a con-

Figure 7: Cirrus SR-22 aircraft used for the flight tests (left) and illustration of blindfolded pilot in flight task (right).

trol, 8 with no audio cue (i.e., the subject had only G-loading, or 'seat-of-the-pants', information on those trials).

The results of the experiment were very encouraging. The addition of the audio cues reduced the percentage of trials where the pilot incorrectly identified the direction of the change in aircraft pitch from roughly 18% to roughly 12%. The addition of the audio cues decreased the number of incorrect roll identifications even more dramatically, from 21% to less than 3%. Thus the results of the experiment provide strong evidence that an auditory horizon cue can increse the pilot's awareness of the attitude of an aircaft in flight.

4.2. Recovery from Unusual Attitude

The second task tested how well the pilots could use the audio cues to recover the aircraft from an unusual attitude. On each trial, the safety pilot maneuvered the aircraft into a displaced attitude, varying in both roll ($\pm 25^\circ$ angle of bank) and pitch ($\pm 10^\circ$ pitch). Once the desired attitude was achieved, the control of the aircraft was given to the subject pilot, who was required to recover the aircraft, in a smooth and controlled manner, to level flight using only the audio horizon. The subject announced when he believed the recovery was complete, at which point the response time and attitude of the aircraft were recorded, and the safety pilot once again took control of the aircraft. Ten such trials were completed for each of 16 subjects, for a total of 160 trials. If, during any recovery, the subject pilot maneuvered the aircraft beyond $\pm 45^\circ$ angle of bank or $\pm 20^\circ$ pitch, or the aircraft descended to an altitude less than 3000 ft MSL, the safety pilot took control of the aircraft and the trial was aborted.

Overall, the results of this experiment suggest that the audio cues were quite effective for allowing the pilots to maneuver an aircraft out of an unusual attitude and into a more stable straight-and-level flight pattern. Despite the fact that the pilots were completely blindfolded and had only audio cues to guide them, there were only 3 trials out of 160 where the pilot attempted to maneuvered the aircraft outside of the safe operating envelope of the aircraft. In more than 85% of the trials, the initial maneuver made by the pilot was in the correct direction to stablize the aircraft, and in 92% of the trials the ending attitude of the aircraft represented an improvement over the initial aircraft attitude. Thus, while clearly the audio attitude indicator is no substitute for visual instruments, there is strong evidence that the audio cue provided useful information about the attitude of the aircraft.

5. CONCLUSIONS

In this paper, we have described the design, implementation, and evaluation of an audio attitude indicator that can modify the acoustic characteristics of an arbitrary audio signal to provide real-time information about the orientation of the aircraft. In general, the results of the in-flight evaluations show that performance of the audio indicate was quite good. It reduced the number of errors in identifying slow changes in the attitude of the aircraft by more than 50%, and it allowed completely blindfolded pilots to maneuver the aircraft out of an unusual attitude and back into a straight and level flight pattern. The subjective feedback provided by the pilots was also quite positive. In terms of improving the display, the most common complaint about the display is that the pitch cue was relatively difficult to perceive. In part, we believe this was due to the fact that the normal operating pitch of the SR-22 aircraft in straight-and-level flight is roughly $+4^\circ$, and not the 0° point that was assumed to be the center point for generating pitch in the design of the pitch cue. We believe that even better performance could be obtained by shifting the center point of the pitch display to account for this disrepancy. Overall, however, we think the results of this experiment show very strong potential for the operational utility of audio attitude indicators of the type described here. These displays clearly can provide supplemental information about aircraft attitude, and they are likely to achieve a high degree of acceptability among pilots. We also believe that they could be implemented in a very affordable way. Considering the exceptionally high costs that can occur due to even a single incident of spatial disorientation during aircraft operations, we believe a strong case can be made for the adoption of audio attitude indicators in operational general aviation aircraft.

6. REFERENCES

[1] Luis DeFlorez, "True blind flight," *Journal of the Aeronautical Sciences*, vol. 3, pp. 168–170, 1936.

[2] T.W. Forbes, "Auditory signals for instrument flying," *Journal of the Aeronautical Sciences*, vol. 13, pp. 255–258, 1946.

[3] T. J. Lyons, K. K. Gillingham, D. C. Teas, W. R. Ercoline, and Oakley C., "The effects of acoustic orientation cues on instrument flight performance in a flight simulator," *Aviation, Space, and Environmental Medicine*, vol. 61, pp. 699–706, 1990.

[4] Matti Grohn, Tapio Lokki, and Tapio Takala, "An orientation experiment using auditory artificial horizon," in *Proceedings of the International Conference on Auditory Display (ICAD 2004), Syndey, Australia, July 6-9*, 2004.

[5] E.M. Wenzel, "Localization in virtual acoustic displays," *Presence*, vol. 1, pp. 80–107, 1991.

[6] H. Wallach, "The role of head movements and vestibular and visual cues in sound localization," *Journal of Experimental Psychology*, vol. 27, pp. 339–368, 1940.

[7] J.D. Miller, "Slab: a software-based real-time virtual acoustic environment rendering system," in *Proceedings of the International Conference on Auditory Display (ICAD 2001), Espoo, Finland, July 29-August 1*, 2001.

[8] J. Blauert, *Spatial Hearing*, MIT Press, Cambridge, MA, 1983.

[9] Douglas S. Brungart, Brian D. Simpson, Ronald C. Dallman, Griffin Romigh, Richard Yasky, and John Raquet, "A comparison of head-tracked and vehicle tracked virtual audio cues in an aircraft navigation task," in *Proceedings of the International Conference on Auditory Display (ICAD 2007), Montreal, CA, June 26-29th*, 2007.