

## **Chapter 8. A Digital Cousin**

### **Section 8.1. Introduction**

Early in the project, while we were dreaming up solutions to the problem, we often lamented that it was unfortunate that we were restricted to the analog domain. Not only did several methods already exist for hiding extra information in digital audio signals, but we too had several ideas that held promise. Some of these ideas were abandoned immediately once it was obvious that they could not be adapted to the digital realm. However, because of its merits, one idea in particular warranted investigation independent of the sponsored project. In our “free time” Matlab programs were written and simulations run to test the concept’s validity. The results showed that our method worked reasonably well, even for the rough version that was implemented.

The side work expended in developing this method of hiding extra information in a digital audio signal culminated in a conference paper [31]. A reformatted copy of this paper is provided in the next section. It describes the concepts involved, provides simulation results, and indicates areas where future work could perhaps yield improvements. Also, a patent disclosure was filed covering the coding method presented below [32]. In Section 8.3 we discuss other potential enhancements beyond those listed in the paper. Several of these are expected to increase the coding rate without perceptively distorting the audio signal.

## **Section 8.2. Description of Coding Method and Results to Date**

### **Encoding a Hidden Auxiliary Channel onto a Digital Audio Signal Using Psychoacoustic Masking**

John F. Tilki and A. A. (Louis) Beex

DSP Research Laboratory  
The Bradley Department of Electrical Engineering  
VIRGINIA TECH  
Blacksburg, VA 24061-0111, USA  
voice: (540) 231-4877      fax: (540) 231-3362  
e-mail: tilki@vt.edu and beex@vt.edu

*Abstract* - We report on the development of a method of encoding an auxiliary channel onto a digital audio signal such that it is imperceptible to human observers. The encoding is accomplished by imposing slight and controlled changes on the phase spectrum of short-time signal windows. By employing principles of psychoacoustics the auxiliary channel is masked from human perception. Furthermore, the encoding and decoding are achieved via fast algorithms which allow real time processing. The method is applicable for any digitized audio signal containing voice, music, or other acoustic signals to be heard by humans. Possible signal sources include compact discs, digital television, digital radio, digital telephony, and any other source where an audio signal is in a digital format.

#### Introduction

The ability to add extra hidden information to a signal could be useful in many applications. Audio compact discs could be modified to contain artist information, song

titles, song lyrics, karaoke information, still images, and perhaps even video clips. In other contexts an extra channel could be used to transmit control information. The want and need for additional bandwidth is ever-increasing, and the capability to add extra information to digital audio can be quite useful even for archival and storage purposes.

The coding method presented in this paper allows the addition of such an auxiliary channel, and it does so in a manner that does not significantly degrade the perceived quality of the primary audio channel. Furthermore, the method does not require changes in the underlying digital format, and so it completely maintains backward compatibility with existing data formats. One can imagine the possibilities for the audio CD example. Existing CD players could still play CDs encoded with the auxiliary channel. With additional signal processing capability however, new CD players could provide the enhanced functionality described above.

Other researchers have aimed at the same objective using methods such as subband coding coupled with adaptive quantization, and subtractively dithered noise-shaped quantization [3], [4], [5]. However, a direct comparison between the schemes in terms of auxiliary channel data rate versus perceived quality of the primary signal could not be performed. Since the other approaches seem to require more computations, our method could be an excellent alternative especially at high sampling rates and where computational capability is a limiting factor.

## Description of Coding Method

### **Encoding**

It is well known that humans are relatively insensitive to certain types of phase distortions. We use this fact to encode additional information onto the phase of selected bins of an FFT. The signal bins are selected according to human sensitivity in the corresponding frequency bands, and are spaced to facilitate masking by unmodified neighboring bins in the same spectral region. In particular the lower frequencies (below say two kilohertz), which typically contain the most energy in common audio signals, are left untouched. The stronger low frequency components help to mask the changes made in the higher frequency region. Additionally, the unmodified FFT bins in the higher frequency region help to mask their modified neighbors.

The FFT bins immediately preceding the signal bins are used as references. While keeping the digital audio signal's FFT amplitude the same, the phase of the FFT at the signal bins is discarded, and a new phase is assigned based on the neighboring reference bins, a differential phase change, and the digital information to be encoded. For example, a small clockwise rotation of phase relative to the reference can represent a digital one while a counter-clockwise rotation of phase represents a digital zero.

Suppose that the FFT has a complex value of  $X_r$  at a reference bin. In polar form this value is represented as

$$X_r = R_r e^{jq_r} \tag{8.1}$$

where  $R_r$  is the magnitude and  $q_r$  is the phase. Similarly the value of the FFT at the neighboring signal bin is

$$X_s = R_s e^{jq_s} \quad (8.2)$$

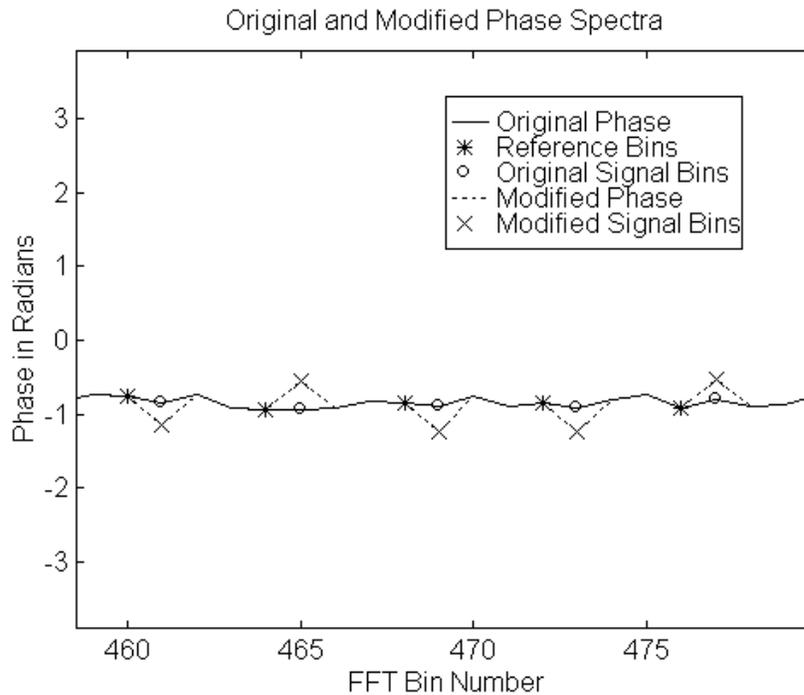
Then if  $j$  is the differential phase change, the new FFT value for the signal bin is assigned as

$$\widehat{X}_s = R_s e^{j(q_r \pm f)} \quad (8.3)$$

Note that it is the phase of the *reference* bin that is used along with the differential phase to replace the phase of the signal bin. Addition of the differential phase represents a digital one and subtraction a digital zero.

Once the phases of the signal bins have been modified, the conjugate phase is assigned to the matching negative frequency bins. An inverse FFT then yields the modified signal block containing both the primary audio and the secondary channel. The process is then repeated for subsequent signal blocks.

Figure 8.1 below shows a small region of an example phase spectrum and its modified version after encoding. The phase at the reference bins is denoted by asterisks, the original phase at the signal bins is denoted by 'o', and the modified phase at the signal bins is denoted by 'x'. This example shows a digital pattern of [0 1 0 0 1].



**Figure 8.1 Example Phase Spectra**

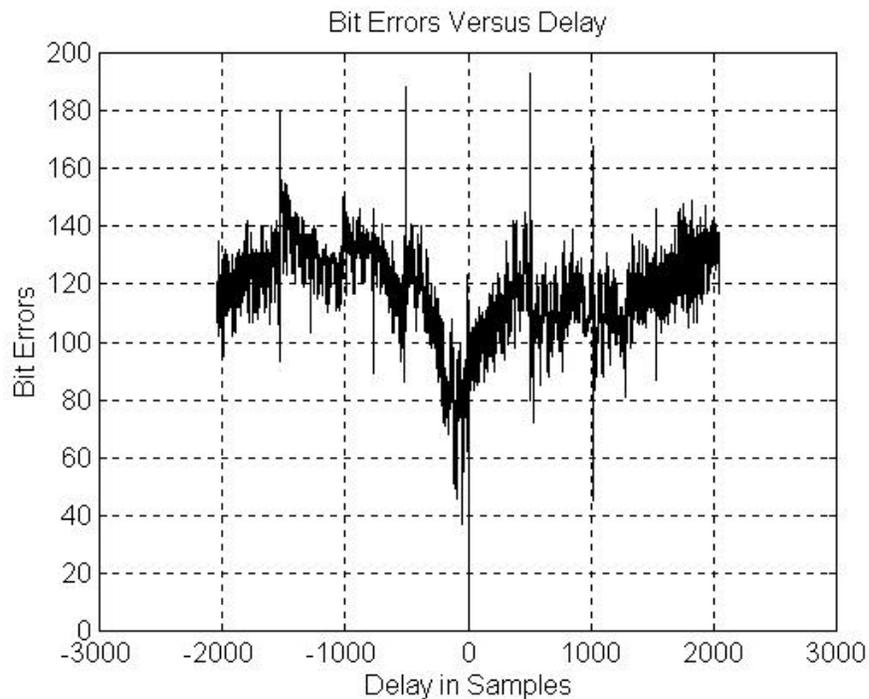
For a given source sample rate the FFT size is chosen to be short enough in time to maintain the imperceptibility of the introduced phase distortions. However, the FFT size should conversely be chosen as large as possible under the above constraint such that the FFT bin frequencies are as dense as possible. In our limited simulations, FFT sizes corresponding to block lengths of ten to fifty milliseconds worked well. Smaller blocks often yielded “granular” noise, while larger blocks usually created noticeable distortion of the primary signal.

### **Decoding**

The signal is decoded by comparing the phase at the (known) signal bins to the phase at the neighboring reference bins, and assigning bits based on the direction of relative phase rotation.

## **Synchronization**

The decoding process requires perfect synchronization with the encoded signal blocks. For a storage medium such as a compact disc synchronization is simple and only requires knowledge of the offset of the first signal block. For real-time data streams however, synchronization is more difficult. Several options are available, including decoding candidate blocks until an error detection criterion is satisfied. Error detection can be accomplished with cyclic redundancy checks (CRC) or other efficient error detection schemes [19]. The candidate block alignments are changed one sample at a time until a valid data transmission is detected. The offset at the point of valid data detection provides synchronization for all subsequent signal blocks. Figure 8.2 below shows how zero bit errors result when perfect synchronization is achieved. When synchronization is incorrect by even a single sample delay or advance several bit errors result ( $\geq 38$  in the case shown).



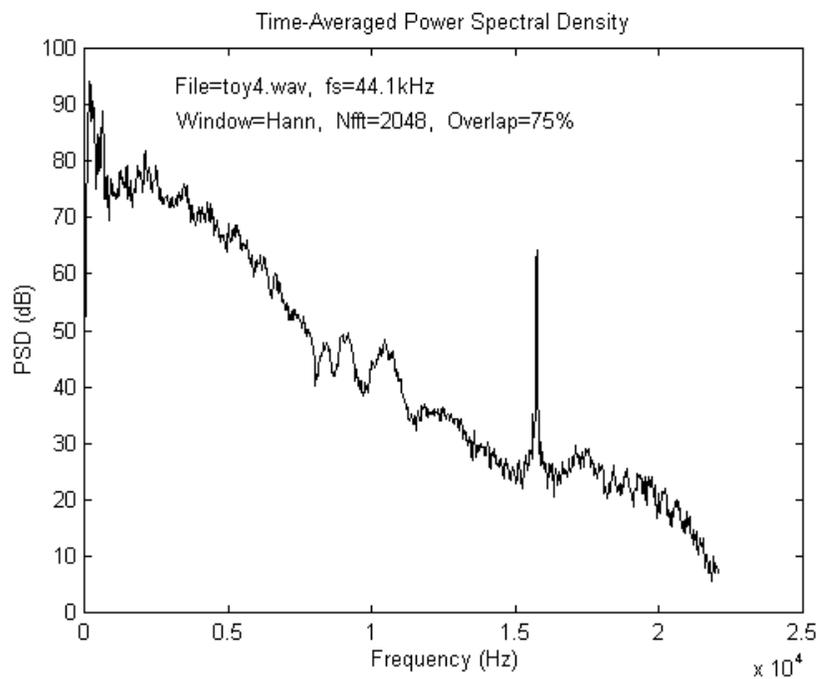
**Figure 8.2 Block Synchronization by Error Detection**

Another possibility for synchronization is the inclusion of known bit patterns in the data stream. This situation also requires shifts by a single sample at a time until the embedded bit pattern is detected. Such a “control” function for synchronization has been used in a similar context in the analog channel case [33].

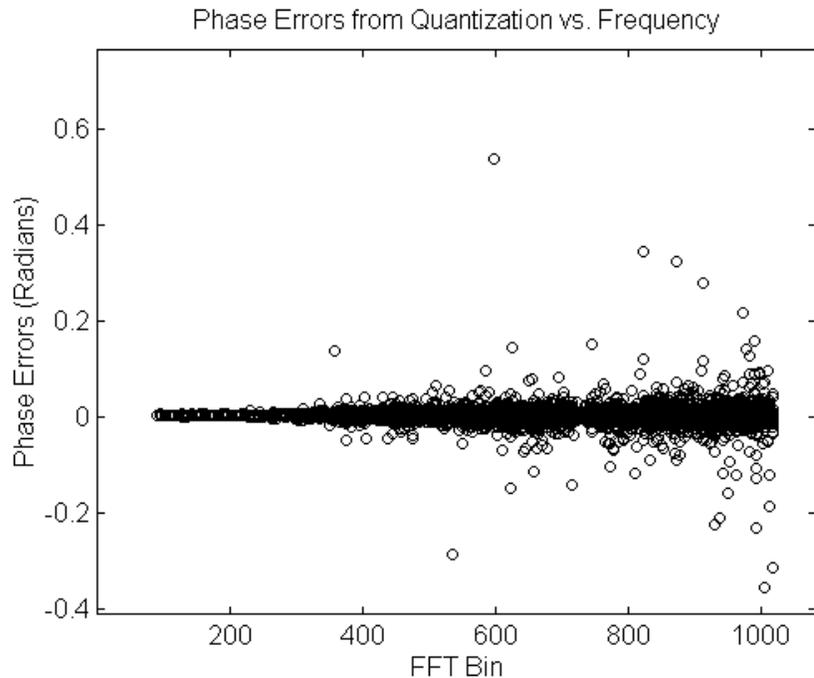
### Simulation Results

The coding method has been tested and works well. The example provided in the paper is based on four seconds of the audio component of a TV commercial, sampled monophonically at a rate of 44.1 kHz, consistent with CD quality audio. The lowpass nature of this audio signal is shown in Figure 8.3 below. (Note that the sinusoid at 15.734 kHz is interference from the horizontal synchronization pulse for the video electron beam.) The decoded phase errors would be zero were it not for the quantization

back to sixteen bits that occurs after the inverse FFT. The phase errors due to quantization are shown in Figure 8.4. A differential phase of  $\pi/8$  was used to modify the phase of every fourth FFT bin above 2.0 kHz in 2048 point FFT blocks. A secondary channel resulted with a capacity of over 5000 bits per second with little or no degradation in the perceived quality of the primary audio signal.



**Figure 8.3 Time-Averaged PSD of Audio Signal**



**Figure 8.4 Phase Errors from Quantization vs. Frequency**

Higher data rates can be achieved by expanding the frequency region used for coding and using smaller spacing between signal bins in the FFT. However, as would be expected, the perceived quality of the primary signal suffers as the data rate of the auxiliary channel increases.

### Future Improvements

Future improvements include using several increments of the differential phase for the forward and reverse rotations of phase, allowing for the coding of multiple bits per FFT bin. Also, pre-quantization of encoded FFT blocks should allow pre-compensation for quantization effects during the encoding process. In addition, the audio signal can be decomposed into subbands, with each subband independently used for coding. This

would allow more bits to be encoded in selected subbands based on the relative influence the individual subbands have on overall perception.

Finally, adaptive encoding based on local power levels and frequency content in the primary audio signal can perhaps increase the effective capacity of the auxiliary channel. A related variation would be adaptation of the differential phase change based on the local frequency content and on some measure of phase deviation in the region. This could be used to compensate for the fact that frequency regions with lower energy are more sensitive to quantization effects than regions with more energy. Note that the phase errors in Figure 8.4 above are largest in the regions of lowest energy, as seen from comparison with Figure 8.3.

### Conclusion

In conclusion, we have developed a method of successfully encoding an auxiliary channel onto a digital audio signal. This channel is encoded using psychoacoustic principles such that it is imperceptible to human observers. Both the encoding and decoding are accomplished with the FFT, allowing real time implementations. The coding process is appropriate for any application containing a digitized audio signal.

### **Section 8.3. Other Future Work**

In addition to the potential improvements mentioned in the paper in Section 8.2 above, other possibilities exist for increasing the data rate of the hidden channel and/or better maintaining the quality of the original audio. If we modify the amplitude of the signal bins, as well as the phase, many more bits could be encoded per FFT bin. Like the

phase, the assigned amplitude would be relative to a neighboring reference bin. As would be expected, if the reference bins are much stronger than the signal bins, very little perceived distortion would occur since the reference bins (and others in the vicinity) would mask the distortion introduced at the signal bins. Furthermore, the new amplitudes could be set so that the signal bins reside under the local masking threshold (the reader can refer back to Section 3.2.1 for a review of psychoacoustic masking). However, if the converse is true, and the signal bins contain too much energy relative to their neighbors, then the introduced distortion is likely to be perceived, and perhaps even considered objectionable.

This dependency on local masking suggests adaptive selection of signal and reference bins, which is another possibility for improvement. Within each frequency region, the strongest bin (the bin with maximum amplitude) could be designated as the reference for the subband. Then some percentage of the weakest in the subband could be changed both in magnitude and phase to encode the auxiliary channel information. This would ensure that the strongest signal components are preserved and the weakest are modified. Furthermore, the magnitudes would be assigned expressly to fall under the local masking threshold. Thus it is expected that the data rates of the hidden channel could be significantly increased while simultaneously maintaining the perceived quality of the original audio signal.

One potential difficulty with this approach, however, could be encountered with quantization. As described in Section 8.2 above, phase errors due to quantization are most problematic for frequencies with little energy content. Here we are explicitly

choosing the weakest frequency components to use in the encoding. Care would have to be taken to certify that the block encoding was still valid after the inverse FFT and quantization. Alternatively, it might be possible to introduce pre-compensation in anticipation of the quantization effects.