Lydia Ayers and Andrew Horner
The Hong Kong University of Science and Technology
Clear Water Bay, Kowloon
Hong Kong
layers@cs.ust.hk
horner@cs.ust.hk

Synthesizing the Chinese Suona

Abstract

This research describes a new Csound additive synthesis design for the ***suona*** *which captures more of its subtle timbral and expressive characteristics. This paper considers the spectral properties of a Chinese double reed instrument: the* ***suona****, and gives a brief description and short musical excerpt with its notation, idiomatic phrasing and typical ornaments. The model produces a tone that is very similar to the original. The source material section describes its spectral characteristics. Following a brief description of the previous wavetable synthesis design, this paper describes an expressive new Csound additive synthesis design used to model the dynamic spectra of the instrument. This design also works for other wind instruments, and allows for expressive modeling of the tones. Finally, a traditional musical excerpt illustrates the expressiveness of the design.*

Figure 1*. Suona*

Introduction

We've modelled the *suona* (see Figure 1) using a wavetable model similar to the one we used for western wind instruments in *Cooking with Csound* (Horner and Ayers 2002). A second design using additive synthesis captures more of the subtle characteristics of each instrument, but at the expense of compilation efficiency and ease of use. We give some background for the *suona* and then describe the acoustic features of the instrument, including its physical properties, spectral response and normal pitch range. The *suona* is designed to be diatonic, though good players can get chromatic pitches by tricky fingering or blowing. The *suona* comes in many sizes because each instrument plays in only one or two keys. D major is one of the most popular keys. We give figures showing the spectra of representative tones. In addition, we focus on making the design expressive. The expression section describes how we tested idiomatic phrasing for the instrument in a short musical excerpt.

Section 2 of this paper gives the background for this research, and Section 3 describes the *suona* and outlines its spectral properties. Section 4 gives an additive synthesis model of the instrument, which is implemented in Csound (Vercoe 1992), and Section 5 briefly describes musical considerations for using the synthesized instrument, including Chinese notation and tuning and a musical example.

Background

Previous work on modeling instruments has primarily focused on simulating western instruments. Much of it modeled the spectra of single wind instrument tones (Risset 1969, Risset and Matthews 1969, Morrill 1977, Horner and Beauchamp 1996, Horner and Ayers 2002). We previously designed wavetable synthesis models (*i.e.*, group additive synthesis models) of many wind instrument tones that are easy to use and sound like the original individual tones (Horner, Ayers and Law 1999, Horner and Ayers 1998a). Researchers have also explored the acoustics of the Japanese *shakuhachi,* the Asian free reed instruments (Cottingham 2004) and large Chinese bells (Braun 2003). Few researchers, apart from Tsai (2003, 2004), have devoted much attention to Chinese wind instruments, and he focused mainly on the *dizi*. We modeled the Chinese wind instruments using a wavetable model similar to the one we used for western wind instruments in *Cooking with Csound* (Horner and Ayers 2002). But musicians don't always play isolated tones in "musical" performances! Rodet and Lefèvre (1997) connected the frequencies of two notes with a line segment and morphed the transition using parameter interpolation to give a smooth slur. We used our previous synthesis design to slur varying numbers of notes in trills for the Chinese *dizi* (Ayers 2003) and to synthesize timbre tremolos and flutter tonguing on wind instruments (Ayers 2004). We have synthesized five Chinese flutes using additive synthesis in Csound (Ayers 2005).

The Suona

The *suona* is a Chinese double reed instrument that sounds like a cross between a trumpet and an oboe. The timbre of the *suona* is brighter and more trumpet-like than that of the oboe because its flared metal bell is large in proportion to its wooden pipe length and diameter. The bell of the *suona* is not attached, allowing it to be removed from the tube for storage. The instrument is played with the entire reed in the mouth rather than between the lips (see Figure 2). The presence of the mouth support on the blowing end of the instrument, along with its trumpet-like sound, has caused foreigners to confuse it with a trumpet.

***Figure 2**. Suona Mouthpiece.*

The low register of the *suona* is solid and rich but a bit breathy. The medium register is strong, firm, bright, easier to control than the low register, and used for showing off. The high register is loud, tense, sharp, and the higher the pitch, the more difficult it is to play softly. The extremely high register sounds very piercing and is extremely difficult to control.

Figure 3 shows the time-varying spectrum of an original A4 tone from a *suona* in Bb. The tone contains a medium amount of instability. Figure 4 shows the average spectrum, which has a strong formant centered on the fourth harmonic. We compared A4 tones from *suonas* in 4 keys, and found that the *suonas* in G and D had the formant centered on the third harmonic, and the *suonas* in Bb and C had the formant centered on the fourth harmonic. In all the tones, the fundamental is much weaker than the first formant, and many of them have progressively weaker formants in an exponential roll-off from the strongest first formant. In the tones with the first formant on the fourth harmonic, the higher formants appear at about every eighth harmonic, in what would be an odd-numbered harmonic series from the formant instead of from the weak fundamental. The tones with the first formant on the third harmonic display a similar effect, with the formants appearing about every six harmonics. The lower formants are more consistently aligned than the upper ones.

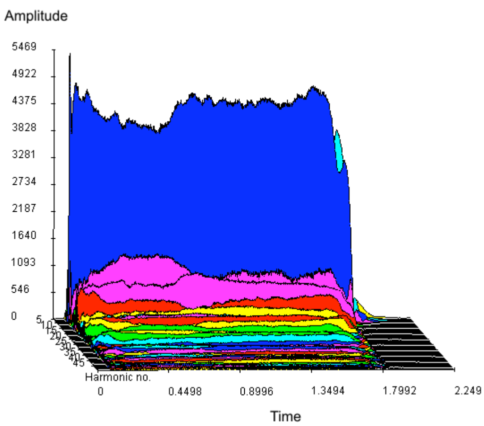

***Figure 3.** Perspective Plot of Suona A4 Tone.*

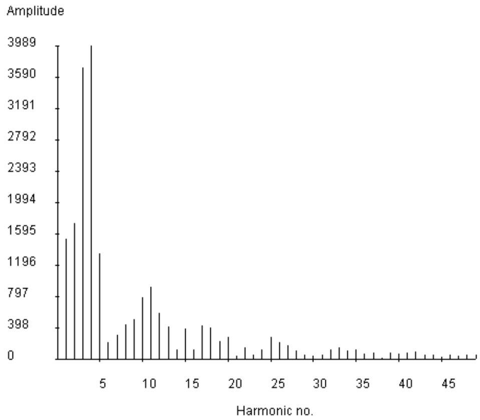

***Figure 4**. Average Spectrum of Suona A4 Tone.*

The Source Material

Guo Ya Zhi recorded two or three *suona* tones without vibrato for each pitch on each of the available instruments in each key. We then listened to and analyzed the tones and chose the best six tones from three instruments to model the range of the *suona* from A3 to D7. Table 1 shows the original tones we modeled for the instruments and the range that the model produced.We used about two notes per octave, with a fairly realistic-sounding range of up to a perfect fifth above and below each of the modeled tones. The highest tone sounds rather pinched and strident, but the other tones have the characteristic *suona* timbre.

Keys of Suonas Used	Tones Modeled	Range
G, Bb, D	D4, A4, D5, A5, D6, A6	A3-D7

***Table 1**. Modeled Tones*

The Wavetable Model

Our first Chinese wind instrument designs used the same wavetable designs we were developing for the western wind instruments in *Cooking with* Csound (Horner and Ayers 2002, see also Horner, Ayers and Law 1999, Horner and Ayers 1998a and 1998b, Ayers 2003, Ayers and Horner 2004a and 2004b). We first modeled the original tones as exactly as possible, cross-fading amplitude envelopes for multiple wavetables on each tone (see Figure 5), but this produced a very complex design which would have been difficult for others to use.

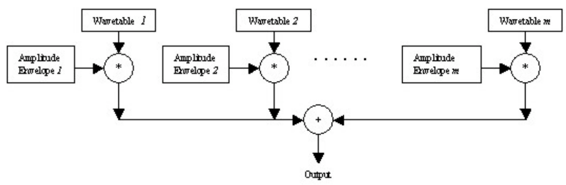

***Figure 5.** The Wavetable Synthesis Model.*

We simplified the design for the *Cooking with Csound* book, so that instead of trying to reproduce the exact time-varying spectra of the tones, we used a generalized amplitude envelope with some random variation. As the tone was too static using a single wavetable, we used a group wavetable synthesis method for each tone. Because the fundamental is the most important harmonic in the tone, we treated it alone as the first group. We used one group for the second and third harmonics, and another group for all the higher harmonics. The three groups allowed some random variation independent of each other, making a more dynamic spectrum, but all the harmonics in each group varied the same way, making the timbres sound slightly synthesized. Figure 6 shows a *suona* A4 tone synthesized using the wavetable design.

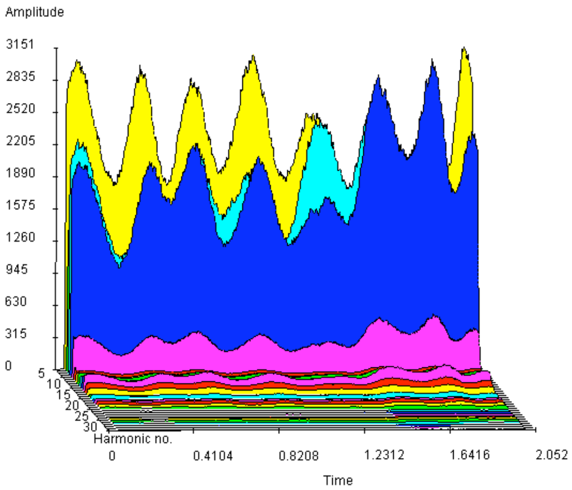

Figure 6. *Wavetable A4 Suona Tone with Vibrato.*

The Additive Synthesis Model

While the wavetable synthesis design produces acceptable matches for the instruments, comparing the instability in the spectra of multiple recordings of the same original tones suggested finding the average spectrum and then randomly varying the amplitude parameters of each harmonic. The large variation in the spectral amplitudes makes the tones sound more lifelike and unique, a feature missing from the preliminary wavetable design, where each tone was so like the original that the effect was similar to sampling. With up to 50 separate harmonic sine wave code blocks to model the very rich spectra of the *suona* tone, the additive synthesis design is cumbersome, and sacrifices the simplicity and efficiency of the second wavetable design. It is tedious to modify it for expressive characteristics. It requires more computation time, but computers have much higher speeds than in the past making it worth using additive synthesis methods to capture the subtle nuances of the *suona* tone.

The Harmonic Code Blocks

Each tone uses the number of harmonic sine wave code blocks required for its spectrum (see Figure 7). Each code block produces a sine wave at the required harmonic frequency, with a slight random inharmonicity. Each harmonic's maximum amplitude is its spectral average amplitude. Each harmonic uses the same overall amplitude envelope, but multiplies the previous amplitude envelope by the original amplitude envelope. The result is that the attack and decay times become longer as the harmonics get higher, which makes the overall tone get brighter as it gets louder and less bright as it decays. Each harmonic signal is added to the signal containing the previous harmonics.

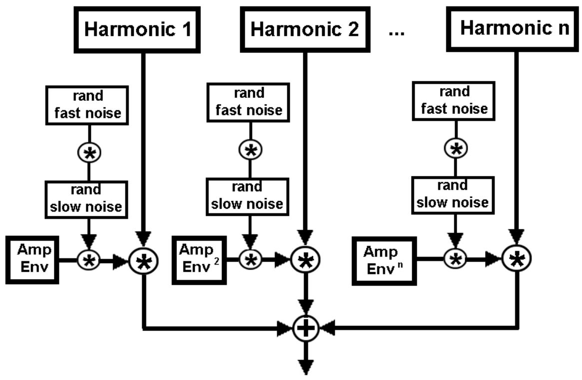

Figure 7. *The Additive Synthesis Model.*

Slow and Fast Noise

Figure 8 shows the time-varying spectrum of an additive synthesis *suona* A4 tone. Each harmonic's amplitude envelope is multiplied by a slow random noise and a fast random noise. The fast noise is essentially jitter.

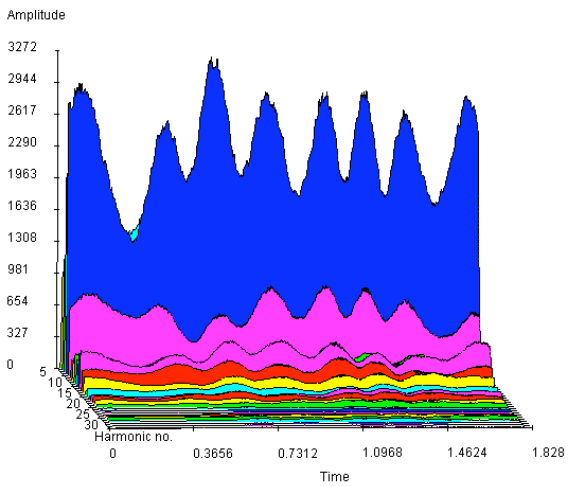

Figure 8. *Additive Synthesis Suona A4 Tone with Vibrato.*

Vibrato

We used a time-varying FM vibrato on the frequency of the tone. The design adds a strong and variable amplitude/frequency modulation vibrato over the whole tone which is not present in the original tone. We used a 10 percent random variation on the original frequency, the vibrato rate and the vibrato depth (see Figure 9).

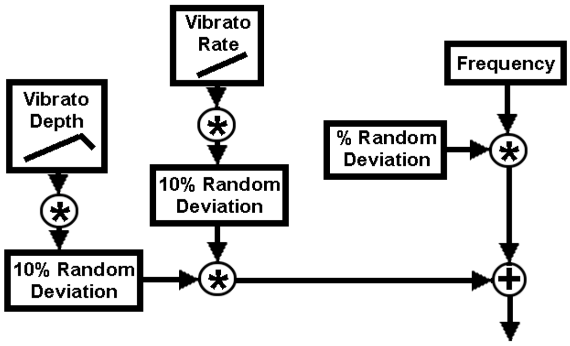

Figure 9. *Vibrato Design.*

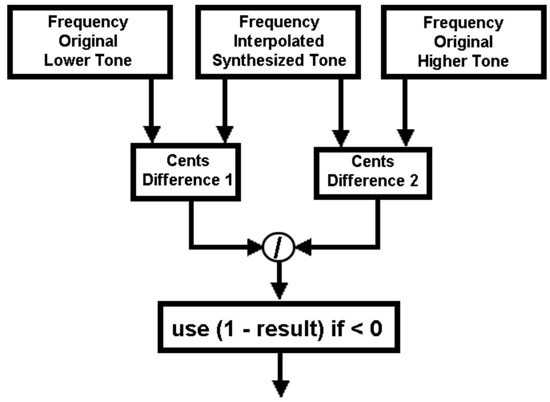

Figure 10. *Spectral Interpolation to Eliminate Range Breaks.*

Avoiding Aliasing

The additive synthesis design improves on the wavetable design because it avoids aliasing by simple code which multiplies the number of harmonics by the frequency of the note and compares the result to the sampling rate; the design decreases the number of harmonics if necessary.

Eliminating Range Breaks

Additive synthesis also has the ability to interpolate parameters to eliminate the range breaks between the spectra of modeled tones. The interpolation method weights the cents differences between the desired tone and each of its neighboring tones to produce the average amplitude for each of the harmonics.

The design first chooses the original tone closest to the tone to be synthesized, then it finds the frequency of the original tone. The design determines if the frequency of the tone to be synthesized is higher or lower than the original tone, and finds the frequency of the original tone on the other side of the tone to be synthesized. It calculates the proportional weight of the synthesized tone spectrum from the weighted average between the modeled original tone and the tone on the other side of the frequency of the synthesized tone (see Figure 10). It uses the cents differences between the frequency of the tone to be synthesized and the two original tones because frequency is logarithmic. The proportional weight of the synthesized spectrum comes from the proportional weight of the spectra of each of the original tones. Normalization of the tone is also proportional to the spectral weight. The weight is then applied to the amplitudes for each of the harmonics, producing a smooth spectral interpolation not only for single tones, but during the transitions in the slurs and trills as well. The musical examples are both expressive and realistic.

Notation and Tuning

This section describes Chinese cipher notation and Ling Lun's tuning which we used to model the *suona*.

Range and Notation

The normal range of the *suona* is a major scale of about two octaves, but experienced players can get chromatic pitches and higher notes. Chinese instruments normally use a cipher notation with numbers from 1 to 7. Pitch 1 is the tonic of a major key. For example, Figure 11 shows the range of a *suona* in G in both Chinese and western notation. In Chinese notation, the numbers represent pitch with no need for staves. A dot below a number means that the pitch is in the lower octave, no dot means that the pitch is in the middle octave, and a dot above a number means that the pitch is in the higher octave. In instruments with extended ranges, extra dots indicate pitches that are in even higher octaves.

Figure 11. *Range of Suona in G.*

Figure 12 shows an excerpt from Guo Ya Zhi's recording of the *suona* solo, *Entanglement of the Fenyang Tune and Baban* (Guo 2002). An alternate version by Liu Wing (1989) is somewhat different, as each performer plays a piece with a personal style. The Chinese notation of rhythm and pitch shown below the western notation corresponds to a *suona* in A.

Figure 12. Excerpt from the Suona Solo, Entanglement of the Fenyang and Baban as performed by Guo Ya Zhi.

A number usually represents one beat. The time signatures work as they do in western music. The line under the numbers is similar to the western beam but with no stems connected to the notes. Two eighth notes grouped together on a beat have a single line under the two numbers. Four sixteenth notes have a double line under the numbers. Dots after numbers and ties work the same as in western music. A 0 means a rest.

Tuning

Some Chinese wind instruments had a more or less equal spacing between the tone holes, which allowed for the convenient placement of the musician's fingers on the instrument. This practice resulted in a tuning where the pitches become more widely spaced as the tube gets shorter. This tuning was replaced by Ling Lun's (Pythagorean) tuning, or a string of perfect 5ths, which we used to model the *suona*. Nowadays electronic tuners are available and some instruments may be tuned in western tuning. However, we prefer Ling Lun's traditional tuning.

Chinese Musical Expression – Charateristic Ornaments

This section describes the excerpt from the solo*, Entanglement of the Fenyang Tune and Baban* that we used to test idiomatic phrasing for the *suona*. This excerpt features several expressive techniques that demonstrate the musical potential of the design. The design is successful with long and short tones, slurs, grace notes, trills and flutter tonguing.

Slurs and Grace Notes

In our previous wavetable synthesis instrument (Ayers 2003, 2004), we used a line segment to connect the note parameters for slurs and morph the transition using parameter interpolation (Rodet and Lefèvre 1997). The wavetable synthesis instrument cross-fades two unison signals as with overlapped notes, but they share one frequency line segment and phase, and use their correct wavetables.

In the additive synthesis instrument, each harmonic is one component signal for the spectrum, with its own frequency and amplitude line segment for the group of slurred notes. Then all the harmonics are added together. Figure 13 shows a Spectrogram analysis of one harmonic in a synthesized four-note slur, with a trill on the last note.

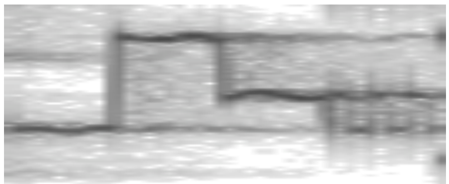

Figure 13. Four-Note Slur with Trill.

Chinese performers often add grace notes according to their taste and the mood of the music, and can add different grace notes on different performances of the same piece. As grace notes are a special case of a slur, we use the same line segment method to model them.

Pitch Bend

An alternative ornament to grace notes is bending the pitch, usually down, though it may also slide up. The pitch bend effect requires only deciding how much to slow down the sliding transitions already between the pitches of the four-note slur design. Figure 14 shows a returning pitch bend from the end of the *Entanglement of the Fenyang Tune and Baban* example.

Figure 14. Returning Pitch Bend.

Figure 15 shows a Spectrogram analysis of the last phrase of the original performance of the piece. This 15-second phrase uses pitch bend when changing among many of the notes. In the original recording, the *suona* is accompanied by other instruments. Figure 16 shows a Spectrogram analysis of the same phrase in the synthesized performance of the piece. The graph of the synthesized version is much cleaner than for the original because the synthetic version does not have accompanying instruments or reverb. Although the synthesized pitch bend is slightly simpler than the original, the contours in the graph are reasonably close to the original contours.

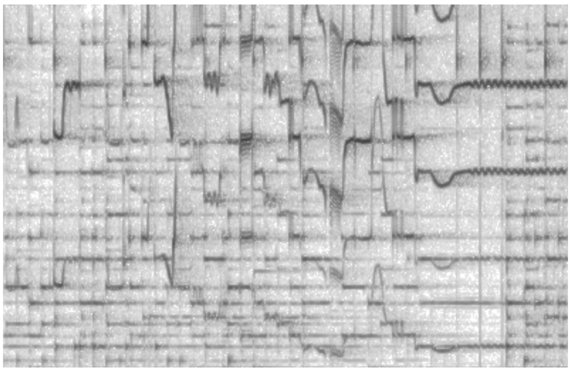

Figure 15. *Spectrogram Analysis of Original Recording.*

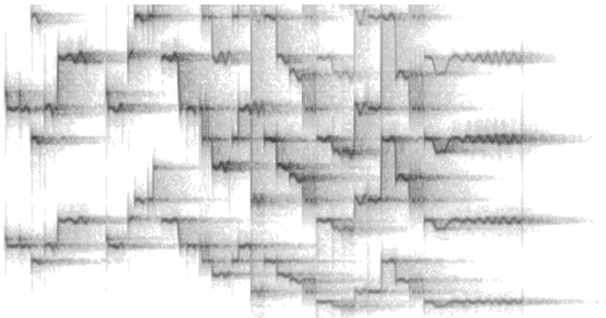

Figure 16. *Spectrogram Analysis of Synthesized Version.*

Trills

The excerpt from *Entanglement of the Fenyang Tune and Baban* contains some very fast trill ornaments. Trills are a special case of a slur. Our previous work used amplitude and frequency modulation functions to cross-fade the trills in the wavetable instrument designs (Ayers 2003, 2004, 2005, Ayers and Horner 2004(a) and (b). Of course, the additive synthesis design improves the transitions between the trilled notes as it does for the slurred notes.

The repeating trill function models the frequency change for the trill. The function does not need to model pitch variation of the average tone, change of speed or jitter, so it can represent one average cycle of the trill, and adjusting the parameters randomly within their typical ranges can vary each cycle. The frequency modulator must use a function that can approximate the frequency analysis graph of the trill. Averaging one cycle provides a good shape for a trill frequency function. The function oscillates between the frequency of the lower note and the frequency of the higher note (see Figure 17). The cycles begin with a slight overshooting of the required frequency, perhaps 20%, and we randomly vary the trill rate. Since trills usually begin with the lower note, we use a phase shift of 180˚ on the frequency function modulating the difference between the frequencies of the lower and higher notes.

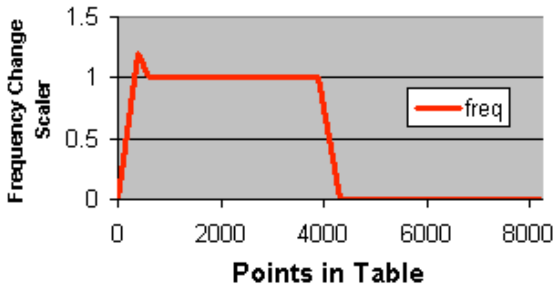

Figure 17. *Trill Frequency Function.*

The wavetable design used an amplitude function containing a spike in the transitions similar to the drop-off in amplitude in the transitions between the trilled notes in the waveform of two cycles of the trill (Ayers 2003, 2004). However, the interpolation in the transition improved the new additive synthesis design so that it does not need the amplitude spike. Using the same function for amplitude and frequency, but with a 0˚ phase for the amplitude, gave a realistic trill. To give each note in the trill the correct spectrum, the amplitude function modulates each of the *differences* between the harmonic amplitudes of the lower and higher notes. One of the notes may be a little bit louder or softer than the other.

As with the previous wavetable design, the amplitude and frequency modulating functions permit several score parameters to control the trill rate using a line segment. For example, our design uses four score parameters to control an initial, middle and final tremolo rate, and the time required to change from the first tremolo rate to the second. A separate score parameter controls the changing amplitude of the trill as it would for a single sustained tone.

Combining Notes with Trills

As *Entanglement of the Fenyang Tune and Baban* includes various combinations of slurred notes with trills, we designed a special instrument to combine a trill with a group of slurred notes. We first set the line segments for the frequency and all the harmonic amplitudes, then use a delay on the frequency and amplitude modulators to cause the modulation to begin at the same time as the note requiring the trill occurs in the line segment. Using the delay instead of just an amplitude envelope ensures that the trill begins on the desired note. Finally, we add all the modulators to all their corresponding signals whose parameters have been set using line segments.

Flutter Tonguing

The * symbol in *Entanglement of the Fenyang Tune and Baban* indicates flutter tonguing, which the player accomplishes by gargling in the back of the throat as rolling the r interferes with the vibration of the *suona* reed inside the mouth. Opening and closing the performer's throat produces amplitude modulation.

On examination of recorded flutter-tongued *suona* tones, we found that the frequency of the amplitude modulation is approximately 30-35 Hertz (see Figure 18). The next step was finding a function to represent the average amplitude envelope of the individual modulations. We found that a simple amplitude envelope would suffice, with a minimum amplitude of about 0% of the total amplitude (see Figure 19), rather than the 20% used for the *dizi* fluttertongue in our previous work (Ayers 2003, 2004). Then we apply the amplitude modulation to the currect note within the slur using a delay, or apply it to all the notes in the slur.

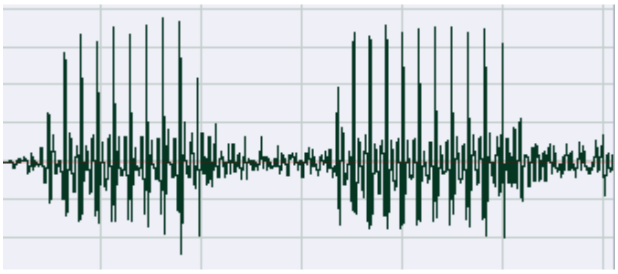

Figure 18. *Flutter Tongued D#5 Suona Tone.*

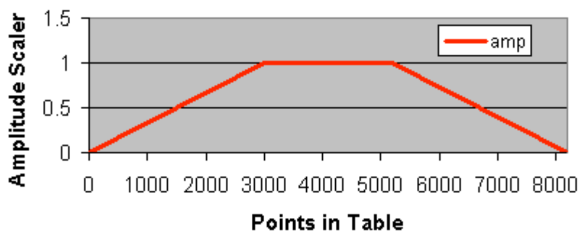

Figure 19. *Flutter Tongue Amplitude Envelope Function.*

Figures 20 and 21 compare the phase vocoder analysis (Dolson 1986) of an original D#5 flutter tone with the same synthesized tone. Although the synthesized tone is not quite as rich as the original, the perceived effect is still quite realistic.

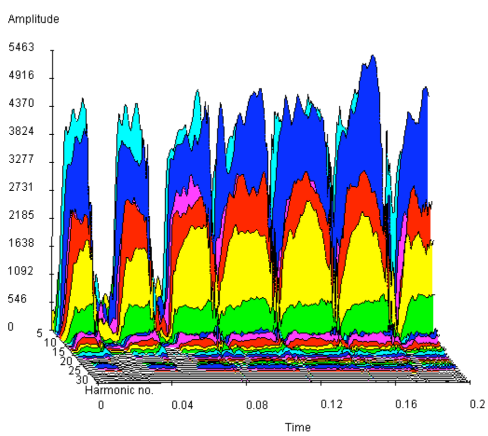

Figure 20. *PV Analysis of Original Suona D#5 Flutter.*

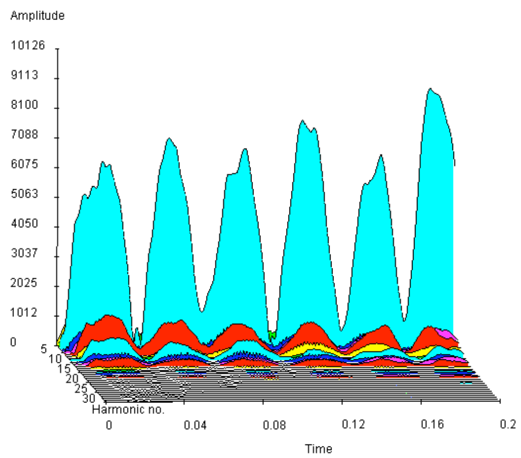

Figure 21. *PV Analysis of Synthesized Suona D#5 Flutter.*

Finishing Touches

We wrote a short utility program which writes the tempo statement with subtle beat scaling to give the beat a more human feeling. We also added reverb for a large occupied hall, which works well for the Chinese wind instruments.

Conclusion and Future Work

Composers are now experimenting with extended techniques for Chinese instruments, though their construction limits what the acoustic instruments can produce. Guo Ya Zhi has added a sliding mechanism near the reed of a *suona* which makes production of semitones much easier. Future work includes vocal sounds combined with instrumental sounds, such as humming while playing and the *suona* technique called *kaxiao*, which uses vocal sounds with a special metal 'reed' within the mouth.

Acknowledgements

We would like to express our appreciation to Guo Ya-Zhi and Chan Chun-Nin for recording the suona samples and consultating on performance, to Jenny Lim for notating the sound example for this research and to RGC Grant Projects HKUST6167/03E, 613505 and HKUST 6020/02H for funding this research.

References

Ayers, L. 2003. "Synthesizing Trills for the Chinese Dizi" Proceedings of the *International Computer Music Conference*, Singapore.

Ayers, L. 2004. "Synthesizing Timbre Tremolos and Flutter Tonguing on Wind Instruments" Proceedings of the *International Computer Music Conference,* Miami, Florida, USA.

Ayers, L. 2005. "Synthesizing Chinese Flutes Using Csound" *Organized Sound.* 10, 1.

Ayers, L. and Horner, A. 2004a. "Synthesis of Chinese Dizi Ornaments" *Ghost in the Machine: Performance Practice in Electronic Music*, Australasian

Computer Music Association Conference. Wellington, New Zealand.

Ayers, L. and Horner, A. 2004b. "Expressive modeling of Chinese folk wind instruments" *Journal of the Acoustical Society of America,* San Diego, California, USA, 116, 4, part 2.

Braun, M. 2003. *Bell tuning in ancient China: a six-tone scale in a 12-tone system based on fifths and thirds.* http://w1.570.telia.com/~u57011259/Zengbells.htm.

Cottingham, J. 2004. "The acoustics of the khaen, bawu and gourd pipe," *Journal of the Acoustical Society of America,* San Diego, California, 116, 4, part 2.

Dolson, M. 1986. "The Phase Vocoder: A Tutorial" *Computer Music Journal.* 10, 4.

Guo, Y-Z. 2002. *Entanglement of the Fenyang Tune and Baban* on *A Hundred Birds Worshipping the Phoenix.* Starlight Music International Ltd.

Horner, A. Ayers, L. and Law, D. 1999. "Synthesis Modeling of the Chinese Dizi, Bawu, and Sheng" *Journal of the Audio Engineering Society.* 47, 12.

Horner, A. and Ayers, L. 1998(a). "Modeling Acoustic Wind Instruments with Contiguous Group Synthesis" *Journal of the Audio Engineering Society.* 46, 10.

Horner, A. and Ayers, L. 1998(b). "Modeling Chinese Musical Instruments" Proceedings of the *135th Meeting of the Acoustical Society of America*. Seattle, WA, 4.

Horner, A. and Ayers, L. 2002. *Cooking with CSound, Part 1: Woodwind and Brass Recipes*. AR Editions.

Horner, A. and Beauchamp, B. 1996. "Piecewise Linear Approximation of Additive Synthesis Envelopes: A Comparison of Various Methods" *Computer Music Journal.* 20, 2.

Liu, W. 1989. *Entanglement of the Fenyang Tune and Baban* on *Chinese Wind Instrumental Music Vol. 2*. Hugo Productions (HK) Ltd.

Morrill, D. 1977. "Trumpet Algorithms for Computer Composition" *Computer Music Journal*, 1, 1.

Risset, J.-C. 1969. *Introductory Catalogue of Computer-Synthesized Sounds*. Murray Hill, N.J: Bell Telephone Laboratories.

Risset, J.-C. and Matthews, M. 1969. "Analysis of Instrument Tones" *Physics Today*, 22, 2. 23-30.

Rodet, X. and Lefèvre, A. 1997. "The Diphone program: New features, new synthesis methods and experience of musical use" Proceedings of the *International Computer Music Conference.* Thessaloniki, Greece.

Tsai, C-G. 2003. *The Chinese Membrane Flute (dizi): Physics and Perception of its Tones.* Ph.D. dissertation under the supervision of Prof. Dr. Wolfgang Auhagen and Dr. René Caussé, Humboldt-University Berlin, Germany.

Tsai, C-G. 2004. "The timbre space of the Chinese membrane flute (dizi): Physical basis and psychoacoustical effects" *Journal of the Acoustical Society of America,* San Diego, California, 116, 4, part 2.

Vercoe, B. 1992. *Csound: A Manual for the Audio Processing System and Supporting Programs with Tutorials.* Cambridge, MA: MIT Media Lab.