

**CENTER FOR COMPUTER RESEARCH IN MUSIC AND ACOUSTICS
JULY 1991**

**Department of Music
Report No. STAN-M-72**

**DIGITAL WAVEGUIDE MODELING
AND SIMULATION OF REED WOODWIND
INSTRUMENTS**

Suzanne Eve Hirschman

**CCRMA
DEPARTMENT OF MUSIC
Stanford University
Stanford, California 94305**

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
ENGINEER

© copyright 1991 by Suzanne Eve Hirschman
All Rights Reserved

DIGITAL WAVEGUIDE MODELING AND SIMULATION OF REED WOODWIND INSTRUMENTS

Suzanne Eve Hirschman, Stanford University, 1991

This paper discusses the mathematical modeling and digital simulation of reed woodwind instruments. The acoustical principles governing the general behavior of a reed woodwind are set forth in detail to provide a qualitative basis for a computational version of the acoustics. A more rigorous theoretical treatment of wave propagation modeling through the waveguide digital filter, based primarily on the work of Julius Smith, follows. In addition, models representing reed, tonehole, and bell behavior, again based on Smith's work, are presented. Finally, an interactive computer program is described which was developed on the NeXT computer, based on a SMALLTALK version by Perry Cook and the clarinet modeling approach of Smith. The purpose of this simulation was: 1) as a tool to explore the acoustical behavior of a reed instrument, and 2) to evaluate the acoustical significance of various refinements in the model with respect to musical requirements of a real-time digital instrument. Refinements which were explored included: bell reflection filters with variable order and cutoff frequencies; incorporation of reed dynamics; implementation of a register hole using a 3-port scattering junction; and reed behavioral models of collision elasticity and hydrodynamic effects.

Dedication

This thesis is dedicated to Kit and to Shira, for their inspiration,
and to my parents, for everything.

Acknowledgements

I would like to express my sincere appreciation to the following people: to Professor Julius Smith, my advisor, who got me interested in a new field just as I was ready to finish up and call it a day; to Professor Bernard Widrow, my Electrical Engineering advisor, for encouraging me to persevere along the paths of my interest, and not settle for the status quo; and to all the people at CCRMA for their support and help.

I owe a special debt of gratitude to Dr. Perry Cook, who has been my mentor and sounding board throughout. His musical and acoustical insights have been invaluable to me, as has been all his NeXT acumen and code which he was kind enough to share with me.

Many thanks go to David Hogan Smith and Robert Cronin for their input on chapter 2, and to them, Herbert Myers, Judith Linsenberg, Marilyn Boenau, and Robert Dawson for their insights into practical acoustics from the perspective of the musician. I also thank Mico Hirschberg for his help in understanding the hydrodynamic issues involved in reed behavior. For their loan of instruments, I wish to thank Tom Neuman, Sue McEwen, Lynn Rodoni, Jim Kafka, and Forrest's Music Shop in Berkeley.

I thank Rob Brimmer, Jan Gossett, Hugh Dougherty, Greg Olsen, Gai Saber, Steve Reagan, Bob DeRosier, and all my friends and coworkers for their support and patience while I followed this somewhat unconventional path.

I will always be indebted to my teachers, Kit Higginson and Shira Kammen, who have profoundly influenced me by their example and their encouragement, and have sparked in me an intense passion for the creation of music, a passion which continually surprises and exhilarates me.

And finally, to my parents and family, I owe, simply, everything.

Preface

As the field of digitally produced sound, and its associated computer technology develops, so does the potential for creating realistic real-time computer simulations that model acoustic instruments. Simulation has long been a tool for acoustical analysis; in the field of acoustic modeling, it is rapidly becoming a viable and exciting resource for musical performance.

The primary source of musical sound in an acoustic instrument is the propagation of waves through the instrument medium and the subsequent pressure radiation to the outside air. A good model of the instrument must therefore duplicate the wave behavior, as well as the nonlinear relationships between the medium and its excitation mechanism. Julius Smith has demonstrated that the waveguide digital filter, an offshoot of the normalized ladder/lattice filter, is an ideal structure for this task. The waveguide filter is essentially a two-way delay line interspersed with partially reflective “scattering” junctions which operate on the oncoming wavefronts. Such a filter implements the wave propagation equations exactly. Just as a tonehole or bore diameter change will create an acoustic barrier within an instrument, allowing only part of the wave to penetrate while reflecting the rest, so can the associated changes in impedance be used to calculate the filter reflection coefficients which will result in similar scattering of the propagated, simulated wavefronts. The waveguide section is extremely modular, and can easily be connected to models of excitation and terminating impedance.

This thesis discusses the implementation of Smith’s clarinet model in an interactive environment on the NeXT computer. The clarinet model in its basic form is particularly elegant and efficient, as its cylindrical bore reduces much of the waveguide to a simple delay line. Inelaborate as it is, it provides a powerful demonstration of the principles of wave propagation, and of the essential coupling between reed and air column, as well as being a

potentially rich source of musical sound.

Contents

Dedication	iv
Acknowledgements	v
Preface	vi
1 Introduction	1
1.1 The Sound of the Reed Woodwind Instrument	1
1.2 Modeling of Reed Instrument Behavior	2
1.3 Purpose and Scope of Thesis Research	4
2 Background: Reed Woodwind Acoustics	6
2.1 General Description of Reed Woodwinds	6
2.1.1 Origins of the Reed Woodwind Instruments	6
2.1.2 Definition of the Reed Woodwind	11
2.1.3 Functional Description of a Reed Instrument	13
2.2 Characteristic Reed Tone	14
2.2.1 Single Reed Instruments	16
2.2.2 Double Reed Instruments	20

2.2.3	Other Reeds	30
2.2.4	Reed Tone Conclusions	35
2.3	The Bore / Air Column	37
2.3.1	Sound Wave propagation	37
2.3.2	Acoustical Bore Shapes	46
2.4	The Reed	50
2.4.1	The Reed as a Pressure-Controlled Valve	50
2.4.2	The Reed as a Harmonic Oscillator	53
2.4.3	Additional Nonlinear Effects	55
2.4.4	Double Reeds	59
2.5	The Coupled Reed / Air Column	65
2.5.1	Interaction of reed and air column	65
2.5.2	Energy, Impedance, and Peak Frequency Shifts	66
2.5.3	Nonlinear Effects	70
2.5.4	Harmonics and Regimes of Oscillation	72
2.5.5	Playing in the Upper Registers	75
2.6	Other Acoustical Features	79
2.6.1	Sound Radiation and Reflection: Bell and Toneholes	80
2.6.2	Sources of Loss	89
2.6.3	The player: Vocal Tract	91
3	Digital Modeling of the Reed Woodwind	94
3.1	Frequency versus Time Domain Modeling	95
3.1.1	Frequency-Domain Synthesis	95

3.1.2	Time-Domain Synthesis	96
3.2	Digital Waveguide Bore Model	97
3.2.1	The Wave Equation for Lossless Propagation	98
3.2.2	Impedance and Acoustic Circuits	101
3.2.3	Reflection Functions and the McIntyre Woodhouse Model	104
3.2.4	Waveguide Digital Filters	108
3.2.5	Adaptation of the Waveguide to the Reed Instrument Model	119
3.3	Digital Reed Models	133
3.3.1	Modeling the Reed in a Waveguide Context	134
3.3.2	Nonlinear Refinements to the Reed Model	146
4	Interactive Reed Woodwind Modeling Workbench	157
4.1	NeXT Object-oriented Environment	158
4.2	Modeling Workbench Description	159
4.2.1	Overview of Simulation	162
4.2.2	Main Panel	164
4.2.3	Attack Panel	167
4.2.4	Waveguide Panel	168
4.2.5	Reed Box Panel	172
4.2.6	Bell Panel	179
4.2.7	Spectrum Panel	181
4.3	NeXT SoundEditor & Spectro	183
5	Results	184
5.1	Summary of Results	184

5.2	Basic Acoustic Behavior of the Simulation	185
5.2.1	Observations on A Simple Tone	186
5.2.2	Influence of the Input	189
5.2.3	Influence of the Reed Table Parameters	200
5.3	Scattering Junction Experiments	202
5.3.1	Implementation of a Register Key	202
5.3.2	Experiments with Bore Perturbation	209
5.4	Implementation of the Bell	211
5.4.1	Bell Reflection Filter	212
5.4.2	Transmission Filter	218
5.5	Implementation of a Reed Model	221
5.5.1	Influence of the Time-Domain Properties of the Reed Model	222
5.5.2	Influence of the Reed Resonance on Stable Modes	224
5.5.3	Influence of the Reed Model on Mode Transitions	226
5.5.4	Influence of Elastic Collisions	229
5.6	Implementation of Hydrodynamics (“Bernoulli Effect”)	232
6	Conclusions	237
A	Default Simulation Parameters	241
	Bibliography	244

List of Figures

2.1	Modern Reproduction of an Early Renaissance Soprano Shawm, built by Robert Cronin, 1990 (illustration by author).	8
2.2	Modern Reproduction of an Alto Crumhorn, Design Early Music Shoppe of London, built by the author and Tom Neuman, 1984 (illustration by author).	9
2.3	Functional Diagram of a Reed Instrument	14
2.4	Reed and Mouthpiece of a Clarinet (Instrument courtesy Sue McEwen)(illustration by author)	17
2.5	Waveforms of the Cook experimental clarinet with soft and hard reed, and the McGill Bb clarinet playing G#3	18
2.6	Onset of attack: McGill Bb clarinet playing G#3	18
2.7	Steady-state output spectrum for a clarinet: McGill Bb clarinet playing G#3 (top) and G#4 (bottom)	19
2.8	Clarinet and Saxophone Waveforms for G#3. Sound Source: McGill Collection	19
2.9	Steady-state output spectrum for an alto saxophone. Sound Source: McGill Collection	20
2.10	Reed for a Soprano Shawm, built by Robert Cronin, 1990	21
2.11	Alto Crumhorn C4, Soprano Crumhorn C4, Soprano Crumhorn G4. Alto C4 and Soprano G4 require the same fingering pattern. Sound Source: McGill Collection	22

2.12	Steady-state output spectrum for Alto Crumhorn C4, Soprano Crumhorn C4, and Soprano Crumhorn G4. Sound Source: McGill Collection	23
2.13	Cantus Douçaine with optional windcap, built by Phil and Gayle Neumann, 1989 (illustration by author)	24
2.14	Low note on a Neumann cantus douçaine, capped (top) and open (bottom)	25
2.15	Steady-state output spectra for Neumann cantus douçaine, capped (top) and uncapped (bottom)	26
2.16	Waveforms for a tone on an isolated shawm reed (top) and a low A with the reed attached to a Cronin Soprano Shawm (bottom)	27
2.17	Steady-state output spectrum for a tone on an isolated shawm reed (top) and a low A with the reed attached to a Cronin Soprano Shawm (bottom) .	28
2.18	G# on three orchestral double reeds - the bassoon, English horn, and oboe. Sound Source: McGill Collection	29
2.19	Steady-state output spectrum for G# on three orchestral double reeds - the bassoon, English horn, and oboe. Sound Source: McGill Collection	30
2.20	Party Horn. Courtesy Jim Kafka (illustration by author)	31
2.21	Sound of a Party Horn: Intact, Hole punched in side, Mouthpiece only. Intact party horn courtesy Jim Kafka	32
2.22	Output Spectrum of a Party Horn: Intact, Hole punched in side, Mouthpiece only. Intact party horn courtesy Jim Kafka	33
2.23	Single Reed Mouthpiece for an Oboe. Mouthpiece courtesy Forrest's Music Shop, Berkeley, CA. Illustration by author	34
2.24	Oboe low G, with double reed (top) and single reed (bottom). Chauvet oboe courtesy Lynn Rodoni, Single reed courtesy Forrest's Music Shop, Berkeley	34
2.25	Steady State Spectrum for Oboe low G, with double reed (top) and single reed (bottom). Chauvet oboe courtesy Lynn Rodoni, Single reed courtesy Forrest's Music Shop, Berkeley	35

2.26	Waveform and spectrum of the low G on a Von Heune Workshop Denner Alto recorder, tuned to A415	36
2.27	Displacement wave traveling along a string with fixed end	38
2.28	Reflection and inversion of a pulse on a string at a fixed end	39
2.29	Reflection of a pulse on a string at a free end	40
2.30	Standing wave on a string	41
2.31	First four modes of a string	42
2.32	Compression wave in an acoustic tube	44
2.33	Compression modes in a pipe stopped at one end and open at the other . .	46
2.34	Opposing pressures acting on a single reed	51
2.35	Typical Flow-Pressure curve for a single reed (after Benade)	51
2.36	Mouthpiece pressure waveforms measured by Rocaboy	60
2.37	Waveform of a double reed in a short test cylinder, measured by Van de Laar	60
2.38	Closing aperture of a double reed	64
2.39	Misaligned resonance peaks, after Benade	74
2.40	Nominal location of a register hole on a clarinet	77
2.41	Spectra for an alto shawm playing in two registers (Sound Source: McGill Collection)	79
2.42	Spectra for a Bb clarinet playing in two registers (Sound Source: McGill Collection)	80
2.43	Wave scattering at a tonehole	82
2.44	Wave scattering at a tonehole	83
2.45	Glastonbury pipe, with removable bell. Design Early Music Shoppe of London, built by Tom Neumann, 1981. Instrument courtesy Tom Neuman, illustration by author	86

2.46	Glastonbury pipe, with removable bell, playing Low F, the bottom note on its scale. Left waveform and top spectrum: without bell. Right waveform and bottom spectrum: with bell. Instrument courtesy Tom Neuman	88
2.47	Glastonbury pipe, with removable bell, playing High F. Left waveform and top spectrum: without bell. Right waveform and bottom spectrum: with bell. Instrument courtesy Tom Neuman	89
3.1	Flow associated with a compression wave	102
3.2	Typical Mechanical System	103
3.3	Typical Electrical System	103
3.4	Reflection of an impulse in a clarinet (after McIntyre and Woodhouse) . . .	105
3.5	Waveguide Section (after Smith)	112
3.6	Block Diagram of a waveguide section (after Smith)	113
3.7	Kelly-Lochbaum Scattering for a 2-port Junction (after Smith)	116
3.8	Waveguide Filter Structure (after Smith)	116
3.9	Normalized Ladder/Lattice Structure (after Smith)	117
3.10	Reduction of WGF to a simple delay line for a cylindrical bore (after Smith)	120
3.11	Representation of a conical bore as a series of cylindrical waveguide sections	121
3.12	Delay Line representation of cylindrical and conical bore segments (after Smith)	123
3.13	Two Representations of the bell: 1) As Complementary Highpass Transmission and Lowpass Reflection filters (after Smith) 2) As Lowpass Reflection filter and residual (from Cook)	125
3.14	Three-port junction representation of a tonehole	127
3.15	Wave scattering at a 3-port junction representing a register hole	132
3.16	Standard Reed Lookup Table	139

3.17	The Reed Modelled as a Mechanical Spring/Mass/Damper System (Second Order Harmonic Oscillator)	141
3.18	Bernoulli Pressure Drop in the Glottis (from Hirschberg)	150
3.19	Hypothetical Visualization of Bernoulli Flow and Consequent Pressure Drop in a Reed Mouthpiece	152
4.1	Main Menu: ClariNeXT Workbench.	159
4.2	ClariNeXT Simulation Flow	163
4.3	Main Panel: ClariNeXT Workbench.	165
4.4	Attack Panel: ClariNeXT Workbench.	167
4.5	Waveguide Panel: ClariNeXT Workbench.	168
4.6	ClariNeXT Simulation Flow with Scattering Junction	170
4.7	Reed Box Panel: ClariNeXT Workbench.	172
4.8	Reed Table with Bernoulli Force	179
4.9	Bell Panel: ClariNeXT Workbench.	180
4.10	Spectrum Panel: ClariNeXT Workbench.	182
5.1	Waveform profile for a tone on the ClariNeXT simulation	186
5.2	Attack profile of a tone on the ClariNeXT simulation	187
5.3	Attack portion of true clarinet waveform (Bias removed by recording preamplifier (Sound source: McGill University Master Samples)	187
5.4	Steady-state portion of a tone on the ClariNeXT simulation	188
5.5	Steady-state portion of a tone on a real clarinet	188
5.6	Bore resonances on the ClariNeXT simulation	189
5.7	Output spectrum of the ClariNeXT simulation	190
5.8	Output spectrum of a real clarinet	190

5.9	Input/Output pressure relationship for typical clarinet simulation. Output pressure has been scaled down by 4*1.3	192
5.10	Output pressure profile for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)	194
5.11	Pressure waveforms for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)	195
5.12	Reed Position waveforms for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)	196
5.13	Output spectra for nonbeating regime: 16000 and 17400	197
5.14	Output spectra for beating regime: 20000 and 25000	198
5.15	Pressure profiles for underblown clarinet simulation with attack overshoot of 25% (top) and 100% (bottom)	199
5.16	Transition from attack to steady-state for underblown clarinet simulation	199
5.17	Steady-state pressure waveform and output spectrum for underblown clarinet	200
5.18	Input/Output pressure plots for three stiffnesses of reed in a linear reed table (output pressures scaled down by 4*1.3)	201
5.19	Bore resonances for a register hole G' with diameter ratio 0.1	203
5.20	Output spectrum for a register hole G' with diameter ratio 0.1	204
5.21	Waveform of a register hole G' with diameter ratio 0.1	205
5.22	Bore resonant peak locations vs register hole diameter to bore diameter ratio (α)	206
5.23	Percent change in bore resonant peak locations vs register hole diameter to bore diameter ratio (α)	207
5.24	Bore resonances for a register hole G' with diameter ratio 0.5	207
5.25	Output amplitude as a function of register hole position	208

5.26	Input/Output relationship for a G' achieved through register hole and through delay line truncation (Output pressure scaled down by 4×1.3)	209
5.27	Bore resonances for a high G' achieved through delay line truncation	209
5.28	Output spectrum of a high G' achieved through delay line truncation	210
5.29	Bore Resonances for a pertubed bore	211
5.30	Mode transition from a multiphonic tone (bottom) to a pure high tone (top), achieved with a Kelly-Lochbaum bore perturbation	212
5.31	Multiphonic waveform, achieved with a Kelly-Lochbaum bore perturbation	212
5.32	Bore resonances for four reflection filters: averager, and 2nd, 4th, and 6th order Butterworths with 1500 Hz cutoff	213
5.33	Output spectra for averaging reflection filter and 2nd order Butterworth filter reflection filter with 1500 Hz cutoff	213
5.34	Output spectra for 4th and 6th order Butterworth reflection filters with 1500 Hz cutoff	214
5.35	Output waveforms for 2nd order (top) and 4th order (bottom) Butterworth reflection filters with 1500 Hz cutoff frequencies	215
5.36	Output spectra for 2nd order Butterworth filter reflection filters with 1000 Hz and 1450 Hz cutoffs	217
5.37	Output spectra for 2nd order Butterworth filter reflection filters with 1550 Hz and 2000 Hz cutoffs	217
5.38	Waveforms for a low A: 1) radiated pressure using one-pole highpass filter with radius of 0.9; 2) Internal pressure waveform of the right-going wave; 3) radiated pressure using implicit highpass filter.	219
5.39	Waveforms for a high G, with register hole, using one-pole and implicit transmission filters	220
5.40	Output Spectra for a Low A, using one-pole and implicit transmission filters	221

5.41 Nonbeating waveforms for three reed implementations: reed table, reed model 1, and reed model 2	223
5.42 Nonbeating output spectra for three reed implementations: reed table, reed model 1, and reed model 2	224
5.43 Beating waveforms for three reed implementations: reed table, reed model 1, and reed model 2	225
5.44 Register shifts for the reed model with 2nd, 4th, and 6th order bell reflection filters	227
5.45 Reed resonance mode on ClariNeXT model	229
5.46 Pressure waveform and reed position plot for inelastic and elastic collisions under just-beating conditions	230
5.47 Pressure waveform and reed position plot for inelastic and elastic collisions under strong beating conditions	231
5.48 Output spectra for inelastic and elastic collisions under strong beating conditions	232
5.49 Bernoulli effect: waveform 1 has Bernoulli-type force of level 250; waveform 2 represents the same reed table with no Bernoulli force; waveform 3 represents the reed table with adjusted breakpoint	233
5.50 Bernoulli effect: Top spectrum with, Bottom spectrum without	234
5.51 Bernoulli effect on reed model: waveform 1 has Bernoulli-type force of level 250; waveform 2 represents the same reed table with no Bernoulli force . . .	235
5.52 Bernoulli effect on reed model: spectrum 1 has Bernoulli-type force of level 250; spectrum 2 represents the same reed table with no Bernoulli force . . .	236
5.53 Bernoulli effect: Levels of 500 and 1000 on reed model	236

Chapter 1

Introduction

1.1 The Sound of the Reed Woodwind Instrument

“Oboists, clarinetists and bassoonists are entirely dependent upon a short-lived vegetable matter of merciless capriciousness, with which, however, when it behaves, are wrought perhaps the most tender and expressive sounds in all wind music.” Anthony Baines, 1957 [6]

The reed instrument has long provided man with some of the most delicate, and sometimes some of the most rousing, sources of musical sound. The reed mechanism is at once amazingly elegant in concept while almost unfathomably complicated in operation. The reed itself, at its best a flexible, responsive friend in the lips of the musician, can at its worst be a mutable enemy, prone to harsh sounds and unpredictable behavior. It interacts with an air column bounded by the realities of a physical bore, and a physical player - a lattice of complex acoustical compromises capable of producing breathtaking tones.

The advent of the digital computer in recent years has greatly increased the use of digital computing techniques for electronically emulating the sounds of acoustic instruments. Synthesizers and digital sampling machines have become mainstays in many aspects of modern music, and the term “MIDI” (Musical Instrument Digital Interface) has become, if not a household term, an almost universally recognized acronym among musicians.

Current synthesis and simulation techniques have an important limitation however; although they emulate the sound of the desired instrument, they are not controllable as a real instrument would be. A sound sample is simply a snapshot, a slightly prolonged instant in time. A frequency spectrum, easily replicated by any of various synthesis techniques, is inherently a steady-state representation, free of any of the transients which make a musical sound interesting or even recognizable. Various MIDI-controlled synthesis techniques are available for altering the sound, introducing transients, varying timbre, etc. However, these exist in the electronic world, and have electronic sounding effects. The gestures available to the acoustic musician are, for the most part, out of reach.

A different approach to digital sound production is taken by those in the physical modeling world. Here, rather than take the end product, the output of a black box which exists only in the physical realm, the behavior of the instrument itself is recreated. Wave propagation, tonehole and bell radiation, input modulation - these phenomena which create the sound in an acoustical instrument are set down in mathematical models, implemented by computer code, to create realistic effects in simulation. The musician is provided with the same rich source of gestural control he had with the acoustical instrument. In addition, because the computer model can transcend the physical limitations of its archetype, it provides yet more possibilities, more latitude. Unlike with standard MIDI control techniques, however, these new tones and gestures, though far afield from the traditional acoustic sound, are solidly rooted in physics, and thus, physical reality; they represent a new genre of “electronic” sound.

1.2 Modeling of Reed Instrument Behavior

The study of the physics of the reed instrument is not a recent discipline; it dates at least back to the 19th century, when Weber performed studies of reed organ pipes, and later, when Hermann Helmholtz published his treatise: *On the Sensations of Tone* [25]. Other early contributors to the field included Rayleigh [56], Miller [44], Bouasse [11], and Morse [46]. Worman provides a very good summary of the early history of reed woodwind analysis in [76]. The “modern” history begins perhaps with the landmark paper of Jim Backus, “Small-Vibration Theory of the Clarinet” [5], published in 1963. This study of linearized clarinet behavior, which combined theory with experiment, spawned a new wave of research in the

field, including the comprehensive study by Nederveen in *The Acoustics of Reed Woodwind Instruments* [51]. Other workers included Plitkin, Strong, Stewart, Wilson, and Beavers [55], [70], [75].

The nonlinear relationship between the reed and the bore was explored in depth by Arthur Benade and D. J. Gans, as well as their students, Walter Worman [76], Stephen Thompson [72], and Douglas Keefe [37]. Arthur Benade in particular is considered by many to have been a giant in his field. His textbook, *Fundamentals of Musical Acoustics* [9], is a classic reference, read thoroughly by many of the acoustics buffs with which this author came into contact in the course of the research. Benade and his students described the role of cooperating resonances and regimes of oscillation, which proved useful models for the friendly/hostile responses of musical instruments. Keefe in particular has gone on to publish much work in the area of woodwind simulation, including extensive work in the area of tonehole behavior [34], [35], [33].

From a simulation point of view, the next major step was the 1983 paper by McIntyre, Woodhouse, and Shumacher, "On the Oscillations of Musical Instruments", which presented the concept of describing wave behavior in terms of the dissociated right and left-going waves [43]. McIntyre still perceived this simulation purely as an analytical tool, with some acknowledgement of its "novel" potential for providing the basis of a real instrument. Smith has merged this approach with the well-known normalized ladder/lattice filter structure to yield the Digital Waveguide Filter approach to modeling instruments [64]. In [63], he introduced a simulation algorithm for a clarinet, which included a precomputed reed lookup table for an efficient mouthpiece reflection model and a highpass/lowpass filter combination for the terminating impedance at the bell. On this model, Perry Cook based his Smalltalk reed simulation, which implemented a conical waveguide, as documented in [18]. It is on Smith's work and Cook's code that the much of the analysis and simulation herein is based.

On the experimental side, surprisingly little progress has been made. In particular, little work has been done to characterize the behavior of the reed itself. The author had originally set out to simulate a double reed, but found that there was almost no experimental or theoretical basis for doing so. The few works which do treat the double reed, for example, [57], invariably base their work on modern instruments, comparing the cylindrical single reed clarinet with the conical double reed oboe or bassoon. It is almost impossible to isolate double reed phenomena from the conical bore effects in those results. In addition,

the description of both reed dynamics and hydrodynamics within the reed channel have to date been based primarily on supposition and extremely simple models, rather than experimentation and rigorous theoretical analysis. This trend is currently being reversed by the fluid dynamicist Mico Hirschberg and his fellows in the Netherlands [27], [26].

1.3 Purpose and Scope of Thesis Research

The objectives of the research herein were threefold:

1. To implement Smith's clarinet model in an interactive, evolutionary simulation workbench on the NeXT computer for the exploration of acoustical concepts, models, and parametric sensitivities;
2. To explore the acoustic behavior of a generic reed instrument on the basis of this tool;
3. To isolate those refinements to the model which proved to have musical significance, for eventual incorporation into a real-time instrument.

Refinements to the basic model included:

- Insertion of a scattering junction that divided the bore into two arbitrarily sized sections. This junction could take the form of a three-port junction, representing a dissipative register hole, or a lossless two-port junction, representing a bore perturbation used to study behavior of misaligned bore resonances.
 - Modeling the reed as a simple second order oscillator, and thus introducing reed dynamics into the system. The dynamics of the reed have been modelled similarly in much work to date, from Backus in [5] to Keefe in [34]. In this case, the model was designed to provide a reflection coefficient in a manner compatible with that of the static reed table. The parameters governing the equations were, rather than the more fundamental properties of mass, spring constant, and damping, the resultant, and more musically meaningful, properties of resonant frequency and damping ratio.
 - Modeling hydrodynamics within the reed in terms of a variable "Bernoulli Force", an effect which has been cited frequently in the past, but which is currently under
-

much debate.

- Modeling of variable elasticity in the collision against the lay when the reed beats.

Other minor modifications included: a selectable bell reflection filter, based on Smith's lowpass/highpass representation of the terminating impedance, to allow studies of the effects of cutoff frequency; variable attack envelope, based on Cook's Smalltalk attack objects, to allow studies of the influence of attack on the steady state; and an impulse response, using Cook's SpectrumView object, which calculated the bore resonances.

The discussion proceeds as follows: Chapter 2 provides a somewhat qualitative background of the acoustical principles governing the operation of a reed instrument, with emphasis on the relationship between reed and bore. It is intended to serve as both a literature survey and a tutorial for the modeling and understanding of reed behavior. A good part of chapter 2 is based on Arthur Benade's book, [9]; however, it also folds in the work of many others, both in the research and the musical fields, and attempts to provide a comprehensive summary of the subject.

Chapter 3 presents a more mathematical background. The first half of this chapter is primarily a tutorial on the modeling of wave propagation through an acoustic tube, and summarizes the work of McIntyre, Woodhouse, and Shumacher, and of Julius Smith, with respect to this research. The second half treats Smith's implementation of a clarinet model, and provides detailed mathematical models for the various elements discussed qualitatively in Chapter 2. These include also the refinements discussed above.

Chapter 4 describes the reed workbench simulation itself. It details the function and implementation of each panel element of the tool. The results of various experiments performed with the tool are enumerated in Chapter 5, with conclusions drawn in Chapter 6.

Chapter 2

Background: Reed Woodwind Acoustics

2.1 General Description of Reed Woodwinds

2.1.1 Origins of the Reed Woodwind Instruments

The reed instrument has a long and illustrious history which began when man first found he could coax a blade of grass to sing. Little is left to us from prehistoric times to trace the roots of the reed pipe[2]. Unlike the bone flutes, which tell the early story of a more durable class of instrument, the reed pipes were by necessity constructed of biodegradable material; the reeds were cut directly from the pipe, and so the substance of the pipe had to satisfy the flexibility requirements of the reed . Even if the pipes had through some miracle of ossification survived, Baines speculates that they still might not be plentiful in many parts of the world, noting that "primitive" tribes today, while employing a prodigious quantity and variety of flutes, have developed very little in the way of reed pipes.

Earlier reed instruments were of both the single and double reed variety, meaning that the mouthpiece assembly, or "generator", was constructed of either of a single reed bound to or cut from a rigid surface, or two reeds which could flap against one another. The single reed instrument was not as prevalent, however, and virtually disappeared during the Renaissance period. The most common ancient single reed was used in the drone of the bagpipe, although

a select number of European folk instruments also qualify. Today, many South American folk instruments still use this reed setup. The gourd double pipe, or *tiktiri*, of India, is perhaps the most exotic form of single reed; it is the legendary instrument of the snake charmer [22]. The double reed found its way into a much greater variety of instruments, and out of the realm of folk music into the music of the courts and towns. Many rural cultures had and have their own version of the folk shawm; Robert Dawson, a musician travelling in Nepal, reports that many villages had not only their shawms, but their own unique tune which was considered the only proper sequence of notes on the instrument.

A number of double reeds are recorded illustratively in medieval manuscripts, mostly a variety of the conical *shawm*, notable for its clear tone and exceptionally loud volume. The shawm was, and is, in the hands of most modern players, an instrument whose audiences are happiest outdoors. The shawm, illustrated in figure 2.1, was played with an *open reed*, meaning that the reed assembly was placed directly in the mouth. There is some disagreement about how much embouchure control the player exerted. Some studies of the instruments claimed that the lips were flush against the *pirouette*, the structure from which the reed protruded, too far down to exert much control over the reed, and that the oral cavity played the role of a windcap. Some embouchure control is necessary, as this instrument has no thumbholes, and relies on variations in embouchure for register changes. Use of the pirouette today depends on the design of the reed. D. H. Smith points out that reeds with fan-shaped blades are so sensitive to embouchure changes that use of the pirouette is required for stability [61]. However, modern players frequently play a bit further up the reed, i.e., further toward the end of the reed, but with a looser, more open embouchure than would be used for a modern orchestral double reed instrument such as the oboe. The shawm is, incidentally, the direct ancestor of the oboe, which developed as a quieter conical alternative.

Another type of early double reed instrument employed a wooden windcap which enclosed the reed assembly. The player blew into the windcap, which directed the air in the proper manner against the reed. There are no modern western versions of the direct blown “capped reed”, unless one wants to count the party horns reserved for New Year’s Eve or the practice bagpipe chanter described below, although many reproductions exist for the performance of early music. The most notorious is the crumhorn, illustrated in figure 2.2, a cylindrical instrument shaped like an umbrella handle with tone holes. The crumhorn is actually a

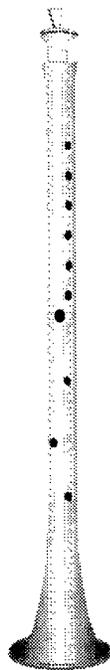


Figure 2.1: Modern Reproduction of an Early Renaissance Soprano Shawm, built by Robert Cronin, 1990 (illustration by author).

member of a large family of capped reeds, often referred to collectively and affectionately as *buzzies*, which contains many straight, cylindrical bore instruments such as *Cornemusens* (a term sometimes applied in general to the family), *Schreierpfeiffen*, and a modern generic reproduction known as the *Glastonbury Pipe*, to be discussed in more detail further on. Earlier versions of the capped reed are referred to somewhat generically as *douçaines*, a term also applied to the *soft shawms*, or open reed instruments which were not shawms. A brief controversy flared in 1986 when the ship *The Marye Rose* was excavated and an original double reed instrument found. The instrument was reported at first to be a shawm, but, according to Herbert Myers, curator of Stanford University's collection of instruments [48], was not a shawm at all, having both a cylindrical bore and a register hole (a vent hole, to be discussed in depth, which allows notes in the upper octave to played).

One category of capped reed which has progressed into modern times is the bagpipe, in

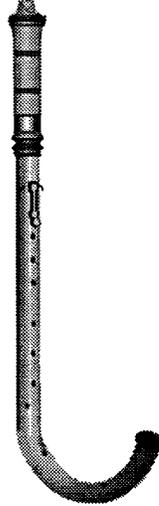


Figure 2.2: Modern Reproduction of an Alto Crumhorn, Design Early Music Shoppe of London, built by the author and Tom Neuman, 1984 (illustration by author).

which the reed pipe, or *chanter*, as well as any number of drones, are driven by a bellows, which in turn is supplied with air by the player or the ambient, and regulated by the elbow or arm. The chanter has tone holes which allow the changing of pitches, while the drones are fixed length pipes capable of playing only one note. Players often work on a practice chanter, which is similar to the actual bagpipe chanter but blown directly like the capped reeds described above. Although this chanter is primarily for less painful practicing, it is an instrument in itself and is sometimes played as such. However, the drones give the bagpipe its characteristic sound, and indeed, drones on variable pitched instruments are frequently used to impart a medieval or pastoral quality to music. Modern Highland bagpipes, like shawms, are very loud outdoor instruments; other pipes, for example the Irish Uilleann pipes, are relatively mellow.

As noted above, the shawm was the direct ancestor of the modern oboe and bassoon. It was modified with a new shape, a softer voice, and eventually, a complicated array of keys on which the more troublesome chromatics could more easily be achieved. The cylindrical bore cornemeuses went by the wayside for the most part, but in eighteenth century, the clarinet came into use. This too was a cylindrical instrument, but with a somewhat novel

mouthpiece. Instead of two reeds beating one against the other, this instrument had but a single reed, vibrating above a *lay* which essentially extended the bore beneath the reed. As was mentioned earlier, primitive single reed instruments were not unknown, but saw little refinement. There is scant evidence of this type of instrument being used in a town setting much before the eighteenth century. Johann Christian Denner is given credit for both inventing the clarinet, and possibly improving on another orchestral early single reed, known as the *chalumeau*, which means "pipe" in the sense of the bagpipe chanter or the rustic reed pipes [6]. It is a common misconception that the ill-fated *chalumeau* was the ancestor of the clarinet; the two in fact coexisted, and little is known about the *chalumeau* before the time of the clarinet [12]. Although Denner may have based the single reed mouthpiece on earlier designs, his great discovery was the vent hole, or register hole, which allowed the instrument to play a twelfth higher. The acoustical implications of the vent hole will be discussed later; it is worth noting that the low register of the modern clarinet is still called the *chalumeau* register, while the upper register is the *clarion*, or *clarinet*. In time, of course, the instrument evolved to include the keys, bell, and form of the modern clarinet.

The next step in single reed development came with the invention of the *saxophone*, by Adolphe Sax in 1841. This new instrument had the single reed of the clarinet but the conical bore walled by the thin metal of the *ophicleide*, an earlier lip reed instrument [12]. As will be discussed, the tonal characteristics of a conical bore are vastly different from those of a cylindrical bore. The saxophone was not the only conical bore single reed; wooden conical bore cousins included the *octavin* (1894) and the Hungarian *Taragato*, which, although not well known outside of Hungary, has been modernized. This latter instrument was a favorite of Benade's [9] for demonstrating acoustic principles. Described by Baines as a sort of wooden soprano saxophone [6], but with a much darker tone, it is still used today by folk bands. The range of the saxophone was originally quite limited, until it was found that the limiting factor was the softness of the reed [71]. When mouthpiece facings designed to allow stiffer reeds were developed, the dynamic capability of the instrument was extended by an octave.

Of course, the single reeds never supplanted the double reeds, just as the conical bore never supplanted the cylindrical. Each development added a new voice to the tonal palette available to the composer, be the work for consort or orchestra. In the 1940's, an attempt

was made to adapt the single reed to the conventionally double reed bassoon and oboe. A single reed is much easier to maintain than a double reed. Although these specialized mouthpieces had some popularity, they did not come into use and are considered oddities today. An example of one of these, for the oboe, will be presented later in this paper.

2.1.2 Definition of the Reed Woodwind

The previous section described some of the historical aspects of reed woodwind development; this section discusses more specifically what constitutes a reed instrument, particularly for the purposes of this paper.

The *reed*, in the acoustical sense, is not necessarily a reed at all, but a valve which responds to fluctuations in pressure by opening and closing, and in so doing, modulating a steadily flowing input stream into a periodic form that intensifies the effects of the fluctuations. It is the coupling between valve and the air column contained within the instrument bore that results in the sustained oscillation necessary for the production of a musical tone. It is the nature of the coupling which characterizes to some extent the timbre of a certain instrument.

When one thinks of reed instruments, one generally conjures up images of the *cane reed* instruments, such as those discussed in the previous section. However, the valving mechanism need not take the form of a cane reed (or even a plastic reed, which substitutes, if imperfectly, for true cane in some instances). The brass instruments, which are played by “buzzing” the lips inside a (usually) metal mouthpiece, are considered *lip reed* instruments. They bear many acoustic similarities to the cane reed instruments; in [34], Keefe treats them mathematically as similarly as possible. The major acoustical difference is that lip reeds are outwardly beating, while cane reeds are inwardly beating. This has important implications with regards to phase relationships. In addition, the lip operates at its resonant frequency, which is a function of the lip tension controlled by the player; we will see that the cane reeds operate, when they are behaving, well below their natural resonance. Of course, the sounds of the brasses are quite distinct from those of the cane reeds, as are their physical forms. In the brass instruments, the metal tubing which contains the bore is long and continuous, with extensions and consequent pitch changes possible through use of valves which redirect the airflow through additional piping. Woodwinds, on the other

hand, generally have tone holes which effectively shorten, rather than lengthen, the bore; these cause additional damping even when closed, and allow sound transmission through other sources than the bell when open. (An interesting hybrid is the *cornetto*, popular in the Renaissance period. This is a relatively small lip reed instrument of wood bound in leather and perforated with tone holes, played through a tiny acorn shaped mouthpiece that is blown like a horn, often on the side of the mouth where the lips are thinner. This instrument is considered one of the hardest to play, but also one of the most beautiful and vocal in sound when played well. Its bass version, the appropriately named *serpent*, is known more for its tortuous shape than its tone.)

The cane reed instruments use a mechanical flap as the coupling device. In the woodwinds, this flap is generally of cane or possibly of plastic. Another instrument which uses a mechanical reed is the organ, which contains many fixed length tongued pipes which operate in much the same way as single reed woodwinds, but without the tonehole lattice. The organ is in fact a much more convenient medium for experimenting with the physics of reeds, in that the geometries are well defined, the structure simple, and the airflow easy to regulate. It is on the organ that Hirschberg [27], [26] is currently studying the hydrodynamics of reed action, and controverting many assumptions common in that field. Even the voice can be considered a reed instrument of this type; the glottal folds act as flaps interacting, albeit weakly, with the bore formed by the vocal tract.

Finally, the *air reed* instruments, such as the recorder and the flute, employ as the coupling device the vortices generated at the sharp edge of the mouthpiece. The action of the pressure fluctuations within the bore vortices periodically direct the vortex shedding to alternate sides of the air stream, modulating the flow into the bore [9]. The most important distinction between this type of generator and both the cane and lip reeds is that the air-reeds are "flow-controlled" rather than "pressure-controlled" valves, and therefore constitute in valving behavior the *dual* of the pressure-controlled cane-reed instruments to be discussed here [9].

In summary, most instruments which are driven by air flow can be considered acoustically to be *reed* instruments. They all generate sound through the use of a nonlinear coupling between resonator and flow source. This paper, however, concerns itself only with the cane reed woodwinds, i.e., those instruments blown manually which contain single or double cane reeds. Throughout, the term *reed instruments* will be synonymous with *cane reed woodwinds*

except where explicitly stated.

2.1.3 Functional Description of a Reed Instrument

The following sections detail the characteristics of the reed instruments. Having defined a reed for the purposes of this paper, a functional description of the instrument which exploits it can be developed. At present, this description will be brief; further sections will give a much more detailed treatment.

Essentially, the reed instrument consists of two parts: the reed, or *generator*, and the air column/bore, or *resonator*. Note that it is the air column *within* the instrument bore which resonates, and not the instrument itself. For this reason, the material from which the instrument is made is not as important acoustically as, say, that in a violin. Although some, including [12], claim that the material is completely immaterial, the softness, porosity, and roughness of the bore surface do influence the boundary layer effects, which are very important in determining internal damping and musical response [49]. However, Benade found that the effects were less than the two percent he cited as being noticeable to a musician. He did note that hardness of the material affects the sharpness of the corners at junctions and toneholes, unless those edges are purposely rounded out; this increases turbulence which can have a very detrimental effect on tone [9]. The question of material importance is a highly controversial and even emotional one. An instrument builder will claim that the material is essential to defining the tone of the instrument, and it is difficult, and in this author's opinion somewhat reckless, to discount the years of empirical knowledge imbedded in the evolution of a craft. This is particularly true in the area of high quality recorder building, where the choice of wood for the flute, for example, boxwood, maple, grenadilla, or ironwood, defines from the start the sort of tone which the craftsman is after. (Early instruments will often be referred to in this paper because they are being actively developed today, and provide many more examples of acoustic phenomena. In addition, the builders today are by necessity experimental acoustic practitioners, because the designs from which they start do not have the time-honored polish of modern-day instruments.) Since the instruments built of more expensive woods are generally much better crafted than those of cheaper materials, conclusions are hard to draw from any sampling of instruments, unless they are built specifically for experimentation under controlled conditions.

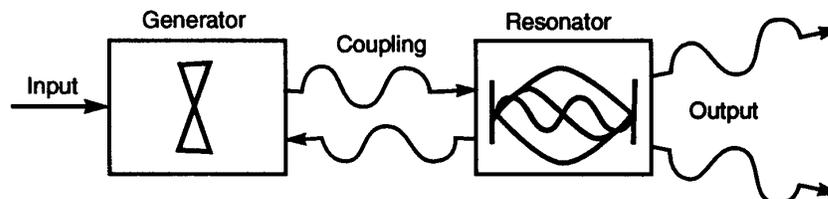


Figure 2.3: Functional Diagram of a Reed Instrument

The reed valve opens and closes to modulate a steady input stream into a pulse train, i.e., a series of compressions and rarefactions which propagate through the air column and, in conjunction with the reflected pulses, form a harmonic progression of standing waves. Once released into the bore, the pulses act according to bore geometry, reflecting and transmitting at any juncture. At the effective end of the bore, part of the pulse energy will reflect back towards the mouthpiece, reinforcing the oscillation, and part will escape into the outside air, transmitting the perceivable tone. Since the period, and therefore the frequency, of the tone is determined primarily by how long it takes to travel from one end of the bore to the other at the speed of sound, the pitch can be changed by altering the effective length of the bore. The player accomplishes this by modifying the configuration of the tonehole lattice, i.e., by opening and closing various holes which have been drilled into the side of the bore. It should be noted that a tube shortened by opening toneholes will not produce the same sound as an unperforated tube of the same acoustic length; the tonehole lattice plays an important role in the quality of timbre. Keefe has explored the tonehole in detail [35] [33]. Toneholes are discussed in more detail in Section 2.6. Except for a special case, the register hole, however, toneholes will not be modelled in the simulation that follows.

2.2 Characteristic Reed Tone

The purpose of this research was to develop a digital simulation of a reed instrument. The chief determinant of the success of this task being aural, it is necessary to determine what constitutes a characteristic reed tone. This is difficult to do precisely, for when we are dealing with sound perception, we enter the very imprecise arena of psychoacoustics. However, although different reeds may sound distinct from one another, they still sound like

reeds, and not like any other type of instrument. Single reeds sound like single reeds, double reeds sound like double reeds, and all sound quite distinct from, for example, the airjet instruments (although the adjective *reedy* is a common descriptor in the almost oenological vocabulary of recorder – an early fipple flute which was extremely important through the baroque period, but which was supplanted in later music by the more dynamically flexible transverse flute – tone description, which also contains such expressive terms as *open* and *nutlike*).

Before focusing on reed tone, it is worthwhile to discuss what constitutes any musical tone. All dynamic systems have a set of resonances which can be excited, i.e., frequencies at which the system oscillates easily. The technique of decoupling the response of a system into its normal modes is commonly used in structural analysis and control system design. But while a vibration may cause sound, it may not necessarily sound music. Any sound can be decoupled into its constituent resonant components, or *partials*. For the sound to be musical, these components must be “harmonically related”, that is, arranged in integer multiples of the fundamental, or lowest frequency component. For example, a note perceived as A-440 may also contain in its spectrum some A-880 (the octave, $2 \times f(A)$), E-1320 (the twelfth, or the fifth above the octave, $3 \times f(A)$), A-1660, etc. The steady state timbre of the note, i.e., its tonal quality once the attack transient has died out, is a function of the relative amplitudes of these components. In some cases, the *harmonics*, are much stronger than the actual fundamental to which the pitch is referred. On a bassoon, for example, the fundamental is virtually absent, but psychoacoustic intervention causes the fundamental pitch to be perceived nonetheless. The role of psychoacoustics in tone and pitch perception is an interesting area, but not one to be covered here.

It must be noted that the steady-state timbre of a note is only one part of the distinguishing characteristic of the tone. It has been shown time and again that the transients at the beginnings and ends of notes, the *attack* and *decay*, are at least as important as the steady state in establishing tone color, that is, the perceived tonal character [52].

This section will present some typical reed instrument sound samples, as both a starting point, and to some extent, and endpoint, to the discussions contained within this thesis. The sounds to be discussed came from two sources: the McGill University anthology of sounds [54], and samples taken by this author in admittedly less than ideal laboratory conditions. They are intended as qualitative, rather than quantitative, depictions of the

differences among various instruments.

2.2.1 Single Reed Instruments

Any reed instrument can be classified first by two major characteristics – its reed type and its bore shape. In trying to develop a good model and simulation, it is important to distinguish between the effects of each of these.

The first set of instruments to be discussed are those using a single reed mouthpiece, such as the clarinet mouthpiece, illustrated in Figure 2.4. The salient features of a single reed mouthpiece are a single reed, bound by a ligature over a fairly voluminous mouthpiece channel. The flattened surface which supports the reed is known as the lay, and accounts for little of the surface area beneath the operating portion of the reed. The very edge of the lay near the tip of the reed slopes downward slightly, so that when the reed begins to beat, that is, to touch down on the lay, it must deform at the tip in order to completely close.

Figure 2.5 shows waveforms from an experimental clarinet stub with no toneholes built by Perry Cook, and from the note G#3 on a Bb clarinet from the McGill anthology [54]. The experimental stub had a microphone inserted inside the bore, so that the internal pressure wave at the microphone, rather than the radiated pressure at the bell, was measured; the results for both soft and hard reed are included. All of these waveforms exhibit a symmetric, essentially triangular shape. Figure 2.6 shows the onset of attack on the Bb clarinet. The triangular form starts fairly early on, although the other high frequency effects do not start until the sound approaches steady state. This is consistent with the "blossoming of the spectrum" which occurs at higher amplitudes that will be discussed later in this chapter. Finally, Figure 2.7 shows the corresponding frequency spectrum of the McGill clarinet playing the G#3, and also the G#4 one octave above. Notice that, for the lower note, the first five odd harmonics are very strong, descending about 15 dB from first to fifth. The intervening even harmonics are much weaker, particularly the first even harmonic (the second harmonic), which is down 35 dB. The second and third even harmonics grow somewhat log-linearly from there to a maximum of -6 dB. After the fifth odd harmonic, the distinction in magnitude between even and odd harmonic disappears. It will be shown that a characteristic of a cylindrical bore is the predominance of odd harmonics. Also note a formant-like like peak cresting at about 3500 Hz. From there, the peak sizes decrease

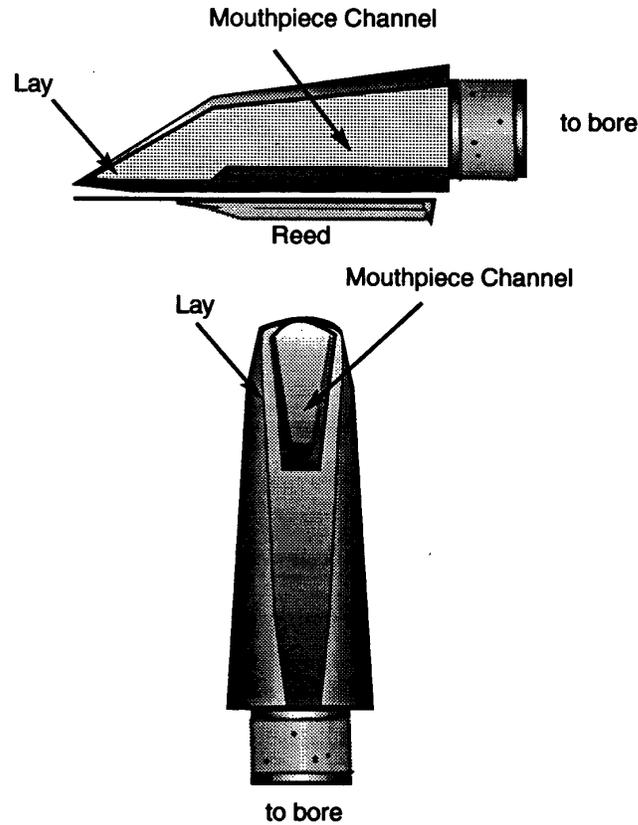


Figure 2.4: Reed and Mouthpiece of a Clarinet (Instrument courtesy Sue McEwen)(illustration by author)

steadily, until they disappear below the -60 dB mark at about 6000 Hz. For the higher note, the peaks are, of course, spread further apart, since the fundamental frequency is twice as high. Now only the first even harmonic is noticeably attenuated with respect to the odds, and the formant at 3000 Hz is less defined. An abrupt cutoff still occurs at about 6000 Hz.

The clarinet is an example of a cylindrically bored instrument. The same sort of mouthpiece is found on the saxophone, which is a conically bored instrument. Figure 2.8 depicts the waveforms of both a B \flat clarinet and an alto saxophone playing G#4, an octave higher than

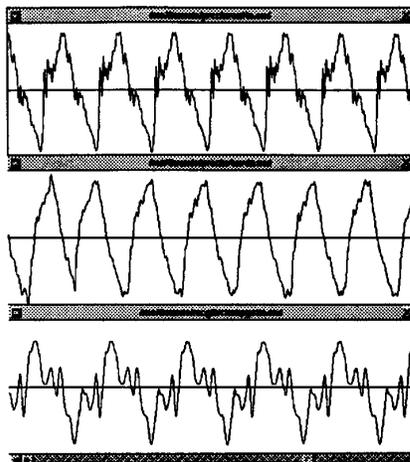


Figure 2.5: Waveforms of the Cook experimental clarinet with soft and hard reed, and the McGill Bb clarinet playing G#3

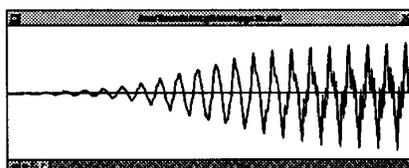


Figure 2.6: Onset of attack: McGill Bb clarinet playing G#3

the tones discussed above. The saxophone has less symmetry in its waveform. It shows a single thick peak corresponding to positive pressure, and a pair of much narrower peaks on the negative side. The spectrum of this tone is shown in Figure 2.9. The most obvious difference between this spectrum and the corresponding clarinet spectrum in Figure 2.7 is the prevalence of the second harmonic, only 5 db down from the first harmonic (this is difficult to discern from the spectrum shown because the peaks are relative to a reference maximum of about 11 dB, and are displayed to only 0 dB. However, the tool which provided the spectrum from which these screen dumps are based, Perry Cook's Spectro program, does also interactively provide the magnitude of the peaks by virtue of a single cursor readout. The cursor readout is the basis for the peak heights cited). The third and fourth peaks are

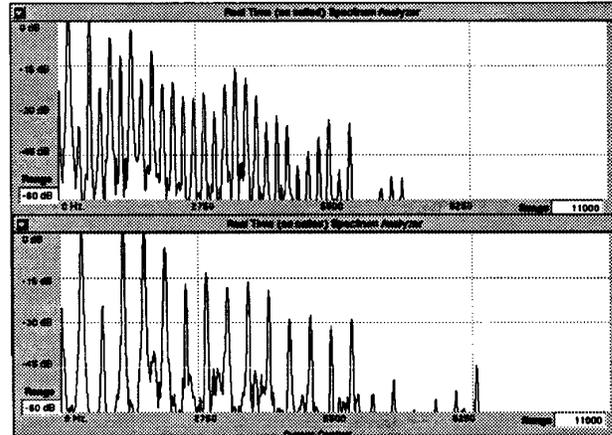


Figure 2.7: Steady-state output spectrum for a clarinet: McGill Bb clarinet playing G#3 (top) and G#4 (bottom)

down about 15 dB, but the fifth peak jumps back 17 dB, from which there is an almost linear decline to the 12th peak at about -35 dB, at which point the spectrum levels off for the most part until about 9000 Hz. The clarinet has some low-level spectral activity out to that point, but not nearly as much as with the saxophone.

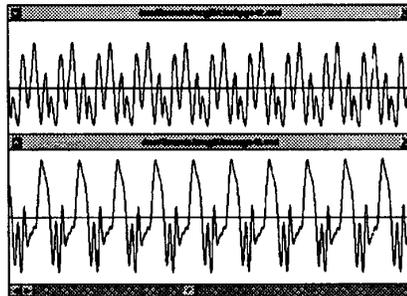


Figure 2.8: Clarinet and Saxophone Waveforms for G#3. Sound Source: McGill Collection

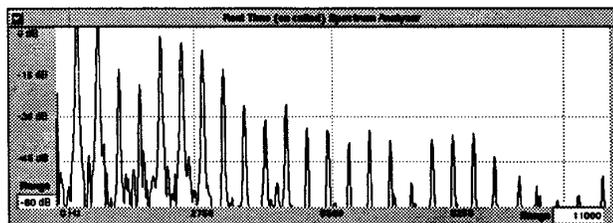


Figure 2.9: Steady-state output spectrum for an alto saxophone. Sound Source: McGill Collection

2.2.2 Double Reed Instruments

This subsection discusses those instruments with a double reed, such as the shawm reed illustrated in Figure 2.10. The double reed is characterized by a pair of reeds, originating from a single strip of reed scored and doubled over, which beat against one another. The volume contained between the reeds is known as the reed channel, and is much narrower than that in the single reed mouthpiece. The reed channel terminates in a tube of small diameter which fits over the staple that leads into the bore. The specific geometry at the tip of the reed during closure is heavily dependent on the *scrape*, or distribution of thickness along the reed. A reed with a stiff *spine* will tend to close at the thinner edges first, and requires additional deformation for complete closure, somewhat like that required for a single reed to close over the lay. A reed with an even *scrape* will tend to close all at once, and is used for a brighter sound [61].

As the previous section noted, there are few cylindrically bored double reed instruments today, save for some folk instruments. Reproductions of historical instruments, however, supply many examples of these. Figure 2.2 illustrated a typical instrument, which had a windcap protecting the reed assembly. This instrument, like its fellows, has a very "buzzy" sound, and crumhorn consorts are frequently likened to kazoo bands. Figure 2.11 illustrates the waveforms for three crumhorn sounds, all from the McGill collection. The first is an alto crumhorn C4, which requires the toneholes in the upper half of the bore to be covered. The second is a soprano crumhorn C4, which requires all toneholes to be covered. The third is a soprano G4, which uses the same fingering pattern as the C4 on the alto instrument. These three are shown together to help isolate the effects of tonehole lattice and instrument size.

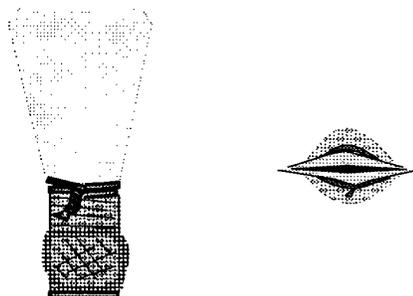


Figure 2.10: Reed for a Soprano Shawm, built by Robert Cronin, 1990

One common feature among all is a pronounced asymmetry, and a great number of narrow spikes. Each waveform does exhibit one narrow long spike on the negative side, probably indicating a brief closing of the reed. There is also spiking on the positive side, but here, the higher spikes are all roughly the same height. Figure 2.12 shows the spectra for these waveforms. As with the clarinet, the first few even harmonics are suppressed with respect to the flanking odd harmonics, although the second harmonic is more prevalent than it had been for the clarinet. In addition, the first few odd harmonics, while larger than the others, are not nearly as dominant as they were before. The general shape of the spectrum is much flatter than it had been for the clarinet, with activity well out to 11000 Hz. There are no well-defined formant-like shapes. Finally, in the spectrum of the higher note on the soprano crumhorn, we see some evidence of some slight multiphonic behavior in the presence of closely spaced pairs of peaks as early as 1800 Hz.

One of the characteristics of the crumhorns described above was the capped reed, which protected the reed from the embouchure damping that normally occurs. The next two examples are of a soprano douçaine, a similar instrument which can be played both capped and uncapped (although in the uncapped state, it plays a semitone flat until thoroughly warmed up). This instrument, which was built by Philip and Gayle Neumann in 1989, is illustrated in Figure 2.13. It has an extremely small diameter bore, and a slight flaring bell at the end, which was an optional feature for increasing the volume. The lower section of the bore, which is unperforated save for two vent holes, is equal in length to the upper section, which contains all of the toneholes. Musically, this instrument is quite similar to the crumhorn, although its design dates back to the Medieval era some centuries earlier.

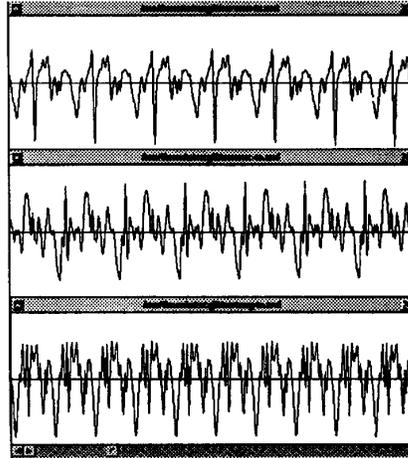


Figure 2.11: Alto Crumhorn C4, Soprano Crumhorn C4, Soprano Crumhorn G4. Alto C4 and Soprano G4 require the same fingering pattern. Sound Source: McGill Collection

It has a slightly more muffled tone than the particular soprano crumhorn it sometimes supplants (both are used by the same ensemble), and blends better with voices, especially when played with the open reed. However, it still has the same buzzy quality characteristic of the cylindrical capped reeds, whether played open or capped.

The waveforms of a low note played on the capped and open instrument are shown in Figure 2.14. The fundamental period in both of these waveforms is somewhat more defined than it had been, with a triangular shape reminiscent of that of the clarinet. The open reed is a bit smoother, as one would expect from the damping provided by the lips, but the essential form remains the same. One interesting feature here is a pair of double peaks on the positive pressure side. For the capped instrument, one of these peaks is noticeably higher than its companion, while for the open instrument, both peaks are the same size. This could possibly be due to the mechanical stop provided by the lip, which in the open case could be limiting the size of the first peak. Both instruments also show a double peak on the negative side. This time, the capped instrument allows both peaks to approach the same magnitude, while in the open instrument, the first peak is much smaller than the second. The tip of the downward triangle is also more defined on the open instrument, as are many of the sharper peaks.

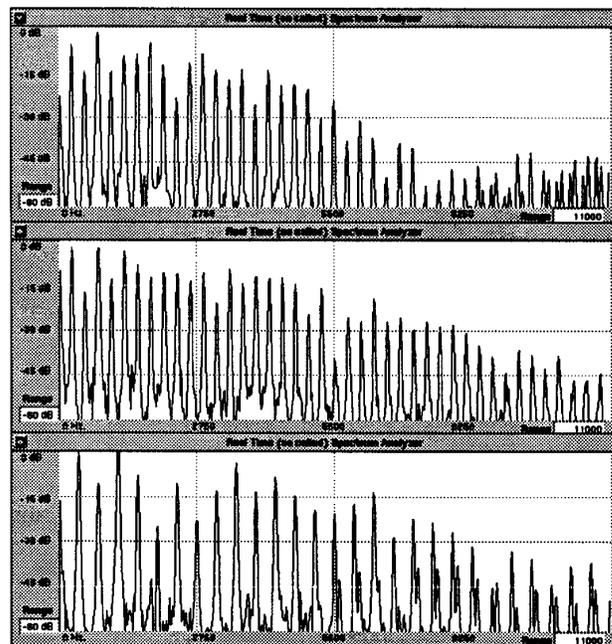


Figure 2.12: Steady-state output spectrum for Alto Crumhorn C4, Soprano Crumhorn C4, and Soprano Crumhorn G4. Sound Source: McGill Collection

The spectra of the two waveforms are shown in Figure 2.15. Here, the main distinction between the two is the shape of the midharmonics, in the second quarter of the graph. Here, the open reed spectrum is extremely flat, while the more resonant capped reed is much less even, with several formant-like peaks. After 8250 Hz, the spectrum of the open reed degrades noticeably, although it still has pronounced high harmonic activity up to that point. The presence of formant peaks in the spectrum, as opposed to a flat, or downward sloping shape, seems to accompany a resonant sound, at least for these simple examples where similar instruments could be compared.

The next example is of the shawm depicted earlier in Figure 2.1, used with the reed depicted in Figure 2.10. Unlike the cylindrical double reeds discussed above, the conical shawm is characterized by a clear, resonant, almost brassy tone. Figure 2.16 shows the associated waveform for a low note on the shawm, and also, for the shawm reed, blown without the

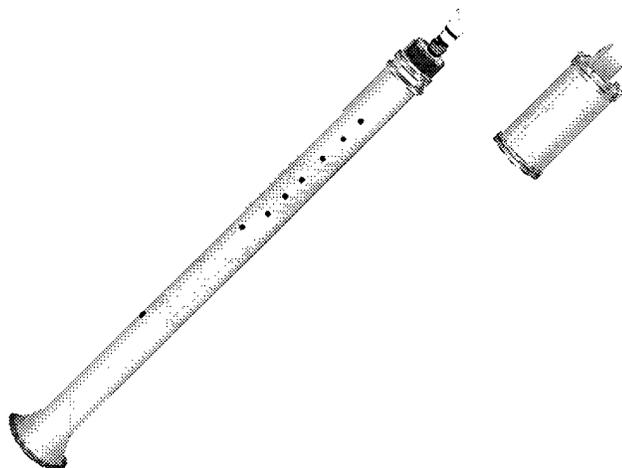


Figure 2.13: Cantus Douçaine with optional windcap, built by Phil and Gayle Neumann, 1989 (illustration by author)

shawm attached. The reed, blown in this fashion, is not playing at its resonance, since it is still attached to the tiny tube at bottom which will serve as a very small bore, but rather, at the "reed-plus-staple" frequency defined by Benade in [9], which is an important factor in the final tone of the whole assemblage. On its own, the reed produces neither a pleasant sounding tone or a pleasant looking waveform, full as it is of sharp spikes. The waveform of the shawm itself is asymmetric; one period looks somewhat like half of a hill, beginning with two small peaks, building up to a maximum three peaks later, and then dropping abruptly back down for the next period. The waveform shown is biased slightly toward the negative side. However, besides that, the waveform is roughly symmetric about its horizontal center. This is somewhat surprising, in that double reeds are generally believed to beat; this would normally cause some vertical asymmetry.

The spectrum of the shawm, illustrated in Figure 2.17, illustrates some classic conical behavior - the increasing dominance of succeeding harmonics for lower frequencies. In this figure, the first harmonic is a good 20 dB lower than the second harmonic. The highest harmonic, the fourth, is yet about 5 dB higher than that. The fifth harmonic drops to the level of the first, but a formant structure brings the next three levels successively higher before a long, slightly uneven dropoff occurs. The shawm reed itself produces a relatively

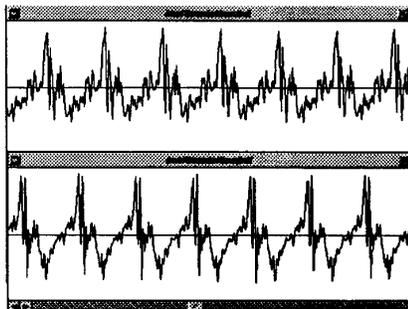


Figure 2.14: Low note on a Neumann cantus douçaine, capped (top) and open (bottom)

flat, wide spectrum from about 1100 to 6000 hz. At that point, the peaks begin losing height, but a second series of multiphonic peaks begins to crop up. The apparent cutoff frequency of the shawm is lower than for the “buzzy” instruments; it is generally accepted that the buzz is due to the strength of the higher harmonics.

Finally, we turn to the modern orchestral double reeds - the oboe, the English Horn, and the bassoon. Figure 2.18 shows the waveforms for three octaves of G#, one from each of these instruments. One similarity among all is the high number of zero crossings, fairly evenly spaced, within each period. All have a dominant positive/negative spike pair, followed by some lower amplitude spike pairs. The relative ratios among these differ for each instrument. For the bassoon, the spikes slope steadily downward from the dominant. For the oboe, they drop off abruptly after the dominant and slope up steadily to the next dominant. On the English horn, the inner spikes form a peak cresting at the middle spike, which is not that much smaller than the dominant spike. Again, there is quite a bit of symmetry about the horizontal. Also, the minor spike widths are roughly uniform.

The corresponding spectra show some similarities as well, as Figure 2.19 illustrates. As with the shawm, the fundamental harmonic is relatively low, with succeeding harmonics increasing steadily up to the fourth or sixth. Soon afterward, there is a sharp drop, followed by a second peak soon after. The bassoon spectrum drops off steadily after the second peak. The English horn displays a third formant structure in the second quadrant. Conversely, even the second formant structure is not well defined in the oboe, whose harmonics attenuate somewhat unevenly after the first peak at the fourth harmonic. All instruments have little

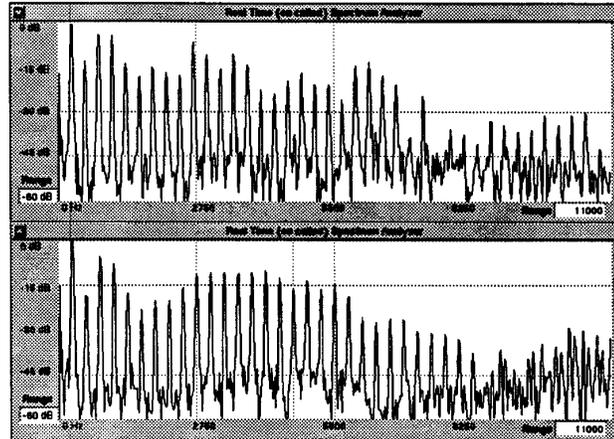


Figure 2.15: Steady-state output spectra for Neumann cantus douçaine, capped (top) and uncapped (bottom)

activity beyond 5500 Hz.

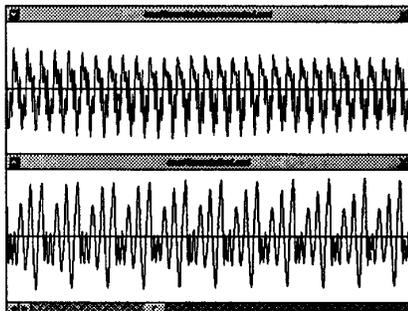


Figure 2.16: Waveforms for a tone on an isolated shawm reed (top) and a low A with the reed attached to a Cronin Soprano Shawm (bottom)

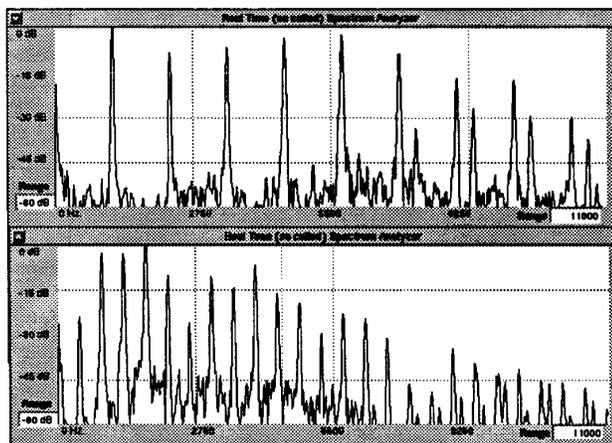


Figure 2.17: Steady-state output spectrum for a tone on an isolated shawm reed (top) and a low A with the reed attached to a Cronin Soprano Shawm (bottom)

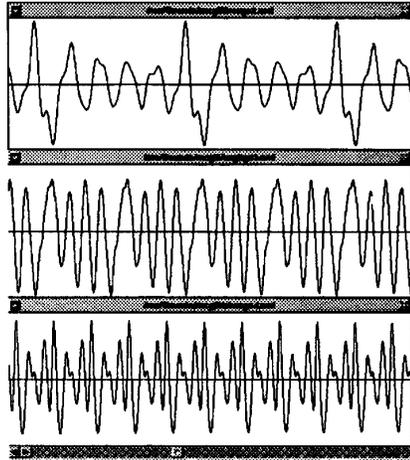


Figure 2.18: G# on three orchestral double reeds - the bassoon, English horn, and oboe.
Sound Source: McGill Collection

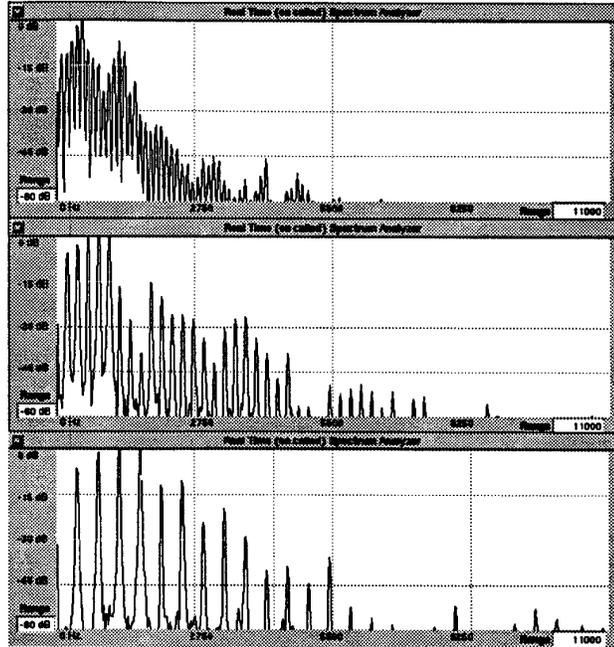


Figure 2.19: Steady-state output spectrum for G# on three orchestral double reeds - the bassoon, English horn, and oboe. Sound Source: McGill Collection

2.2.3 Other Reeds

This section discusses some reed types which do not fall in the above categories. The first is a party horn, of the type used on New Year's eve. The party horn is of interest because of its similarity in sound to the crumhorn, a point often noted in less than benevolent observations on the latter instrument. The party horn tested is illustrated in Figure 2.20. It consists of a plastic capped *single* reed, which flaps against a lay somewhat similar to that of a clarinet, but within the cap, and with no downward slope at the tip. The cap is attached to a paper cone, which, as it turns out, actually has little influence on the sound. Figure 2.21 illustrates the waveforms for an intact horn, a similar horn with a hole punched in the side of the paper cone, and the mouthpiece of the second cone with no cone at all. The three waveforms are strikingly similar. Unlike the shawm reed waveform discussed earlier, which had a much higher frequency than the assembled shawm, all of these waveforms

have roughly the same period and same shape. The shape itself is dominated by a strong negative spike which slopes up to an eventual positive spike of smaller magnitude. The positive spike on the mouthpiece waveform is somewhat truncated. All waveforms are fairly rough looking, indicating a large amount of high frequency influence. This is demonstrated further by the spectra, illustrated in Figure 2.22. The dominant feature on all of these is the strength of the high frequency harmonics, as well as an extremely uneven progression of harmonics.

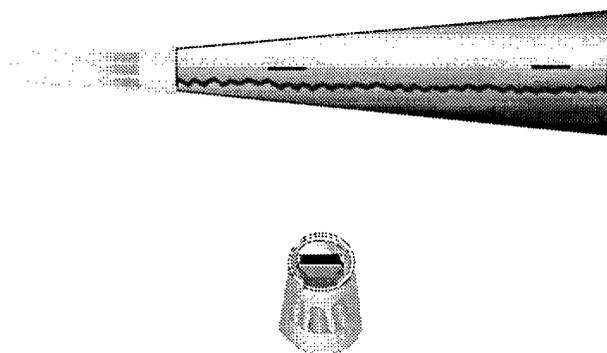


Figure 2.20: Party Horn. Courtesy Jim Kafka (illustration by author)

The next example is of a single reed designed for use with an oboe, illustrated in Figure 2.23. These reeds were originally designed on the assumption that double reed players might prefer the easy maintenance of a single reed. The marketing experiment failed, and the devices have been relegated to the ranks of musical oddities. They do, however, provide the acoustics researcher with an interesting source of comparison. The device shown is really a transitional mouthpiece. It does use a single reed, bound by ligature above a lay. However, the channel beneath the reed, unlike that in a clarinet mouthpiece, is extremely shallow, and connects to a tube that by necessity must be the size of the original staple. The dynamics of the reed may be closer to a single reed, but the hydrodynamics of the airflow through the reed channel is probably much closer to that of a double reed. A final difference between this device and the conventional lay is that the sloping surface at the tip is much less pronounced. This can be expected to produce more high harmonics at loud playing levels, in that the reed no longer has to wrap around the lay to close.

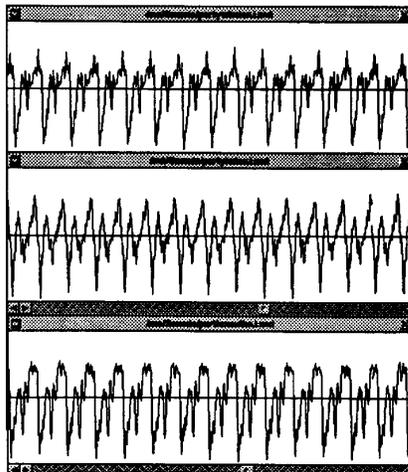


Figure 2.21: Sound of a Party Horn: Intact, Hole punched in side, Mouthpiece only. Intact party horn courtesy Jim Kafka

The single reed sounded somewhat similar to the original double reed in the oboe. It certainly sounded quite oboe-like, although not nearly as refined as the oboe reed, at least at first play. David Hogan Smith, a professional early reed player, kindly assisted in this experiment, and worked the the single reed into a playable form. His major criticism was that the single reed provided almost no latitude in its playing parameters; unlike the double reed, which provides so many variables that the beginner is oftimes confounded, the single reed was difficult to modify through embouchure.

Figure 2.24 illustrates two sample waveforms for a double reed and the single reed on the same oboe body. There are many similarities. The negative peaks on the single reed are fatter than their double reed counterparts, perhaps indicating a different beating characteristic. Response on the positive side is much more similar. Figure 2.25 shows the spectra. On both, the first three harmonics slope upward to a peak held by the third and fourth harmonics. On the double reed, the next few harmonics drop off, whereas on the single reed, the same approximate level is maintained through the seventh harmonic. The next few single reed harmonics drop off considerably from this point, while on the double reed, the level has already built back up. It can probably be concluded that the similarity in the first few

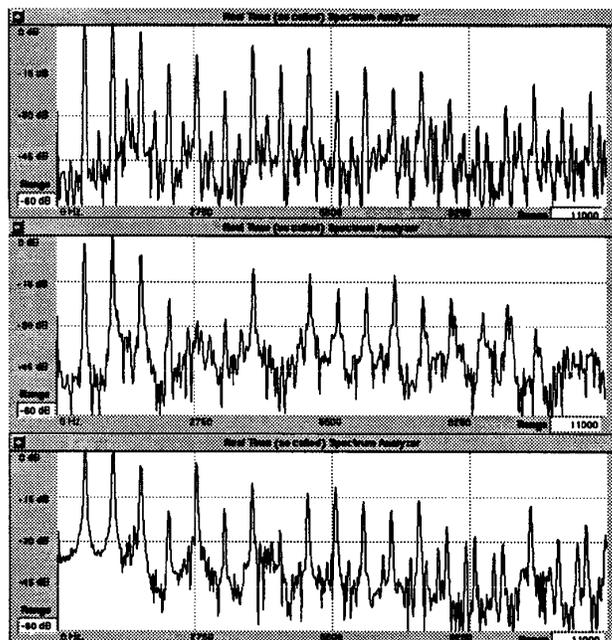


Figure 2.22: Output Spectrum of a Party Horn: Intact, Hole punched in side, Mouthpiece only. Intact party horn courtesy Jim Kafka

harmonics is responsible for the similarity in tone, especially as the pattern matches those in the other conical bore instruments studied here. The subsequent divergence in spectrum in the second quarter can of course be considered the source of the divergence of tone, but a more useful conclusion would require more careful experimentation.

As a final example, the waveform and spectrum of a recorder will be shown. One would expect this waveform to be vastly different from those shown above, but since the topic of discussion is the specific quality of reeds, it is worthwhile to present one other instrument in contrast. The instrument used for this example is a modern replica of an 18th century Denner recorder, built by Von Heune Workshops in Massachusetts, 1990. The waveform shown is of a low note on the instrument. One of the main distinguishing features among instruments is their tone quality in the lower register, where harmonics below cutoff can come into play and make the timbre more interesting. This also results in a much more

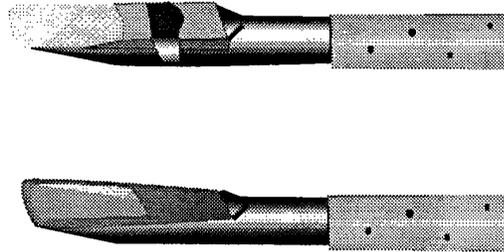


Figure 2.23: Single Reed Mouthpiece for an Oboe. Mouthpiece courtesy Forrest's Music Shop, Berkeley, CA. Illustration by author

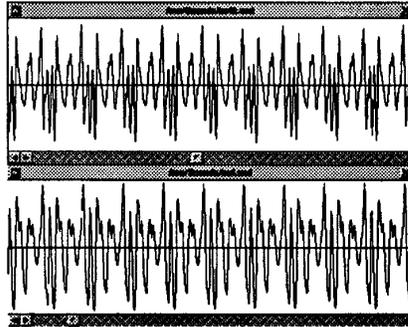


Figure 2.24: Oboe low G, with double reed (top) and single reed (bottom). Chauvet oboe courtesy Lynn Rodoni, Single reed courtesy Forrest's Music Shop, Berkeley

spiky waveform than one would expect from a recorder. Higher notes have the smooth almost sinusoidal shape more common to the instrument. The fundamental is still strong, although the third harmonic is also a major contributor. harmonics drop significantly from there, but there is an additional formant peak in the second quarter. This spectrum is particularly interesting because of its lack of even harmonics. This imparts the instrument with a hollow tone (the term “nutlike” has been applied specifically to this model), listed in [1] as being one of the desirable properties of baroque recorder tone. It will be shown later that the lack of even harmonics, as in the clarinet, is characteristic of a cylindrical bore. In the case of the recorder, it is obtained, according to [1], by the voicing, or sharpness of the labium, and by the focus of the airjet.

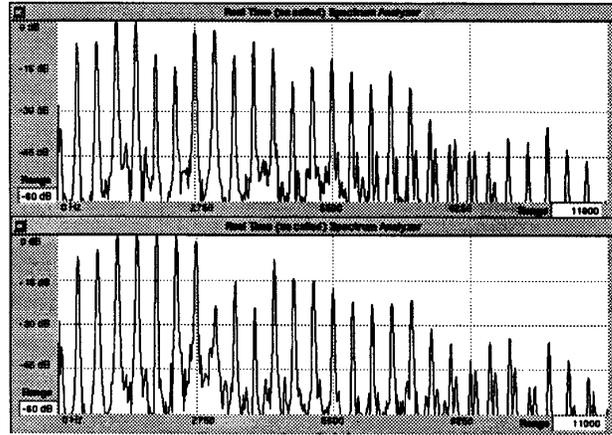


Figure 2.25: Steady State Spectrum for Oboe low G, with double reed (top) and single reed (bottom). Chauvet oboe courtesy Lynn Rodoni, Single reed courtesy Forrest's Music Shop, Berkeley

2.2.4 Reed Tone Conclusions

The discussion above does not pretend to be a definitive or rigorous treatment of reed instrument tone. It is primarily an anecdotal overview of the characteristics of reeds in terms of their pressure waveforms and steady-state spectra. However, some general conclusions can be drawn. First, a "single reed" or "double reed" sound is more related to the bore and reed housing than to the reed itself. The party horn, with its capped single reed, was more similar to the capped double reed than to any other instruments. The oboe with the single reed still sounded oboe-like, although certainly not as pleasant as it did with the double reed. Certainly the reed plays an important role in determining the polished timbre, but is not as important to the fundamental tone as the bore or reed housing.

All instruments had a healthy set of harmonics. The chief distinction between the cylindrical and conical bore instruments was the size of the second harmonic. In cylindrical instruments, the second harmonic was well suppressed. Succeeding even harmonics grew to a level equal to the odd harmonics, however. On the conical bore instruments, the tone was clear and resonant. Their spectra displayed an upward slope on the early harmonics, so that the highest harmonic was often not achieved until the second half of the first decade. Dropoff

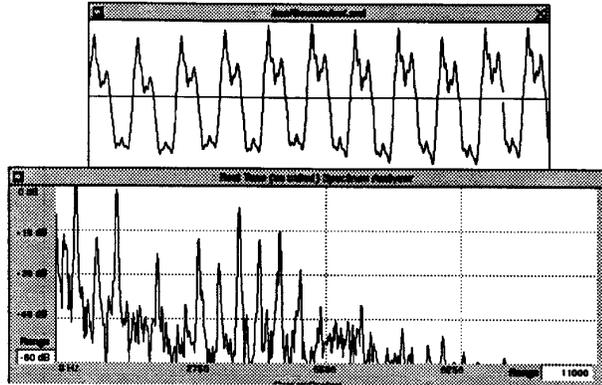


Figure 2.26: Waveform and spectrum of the low G on a Von Heune Workshop Denner Alto recorder, tuned to A415

was somewhat rapid from there.

The cylindrical double reeds were quite different in tone from the clarinet. This can be attributed to the high cutoff frequency of these instruments, which have almost no bell to speak of (an issue to be discussed in a later section). Consequently, the higher harmonics were much more active in the tone, resulting in the buzz for which the crumhorns are famous. The party horn had a similar buzz, and a similar preponderance of high harmonics. It is interesting to note that the sound was not particularly bright, a quality which is often ascribed to timbres with strong higher harmonics.

Finally, the presence of formant-like structures, or hill shaped sets of peaks cropping up after an initial downslope, seemed to correspond with a resonance or refinement of tone. This point will be brought up again in the section on the bell, when the inclusion of an optional bell on an instrument causes an increase in volume and musical resonance, and a corresponding set of formants in the spectrum.

2.3 The Bore / Air Column

As section 2.1 indicated, any reed instrument can be separated into two fundamental components — the generator, provided by the reed, and the resonator, provided by the air column enclosed within the bore. This section examines the air column and how pressure waves propagate through it. Further sections will detail the reed, the coupling between the reed and the air column, and finally, other factors which influence the sound.

2.3.1 Sound Wave propagation

The key to creating musical sound is to produce sustained oscillations in a medium that can somehow transmit that oscillation in the form of pressure pulses to the ambient air. The human (or animal) ear will sense these pressure fluctuations and perceive them as sound. In a stringed instrument, such as a harp or a violin, the vibration of the strings is coupled to the resonating instrument body, which has enough surface area to produce the necessary disturbance to the surrounding air. In a wind instrument, the player provides the raw air pressure that will become sound, with the interactions within the instrument modulating it into the necessary periodic form. This section explores the effects of injecting pressure pulses into an air column.

Waves can be classified by their direction of displacement with respect to their direction of travel. The transverse wave is perpendicular to its direction of travel; for example, waves in the ocean cause a displacement upward that appears to travel horizontally. The torsional wave travels down the axis about which it rotates; a quick twist at the end of a long rod would excite a torsional wave. These are the types of waves propagated in a stringed instrument. On a violin, the action of bowing both twists and displaces the string, and propagates in consequence both torsional and transverse waves. Finally, the longitudinal wave will show displacement along the direction of travel; that ultimate dry wave tank, the Slinky, will propagate longitudinal waves in the form of a spring compression or expansion if it rests lengthwise on a surface and the end is pushed abruptly inward or outward. Sound pressure takes the form of longitudinal waves, with the air acting like a compressible spring. The waves propagated within the bore of a wind instrument are also longitudinal; the action of the reed, to be described later, causes compression and rarefaction pulses to be introduced to the air column, through which they travel at the speed of sound.

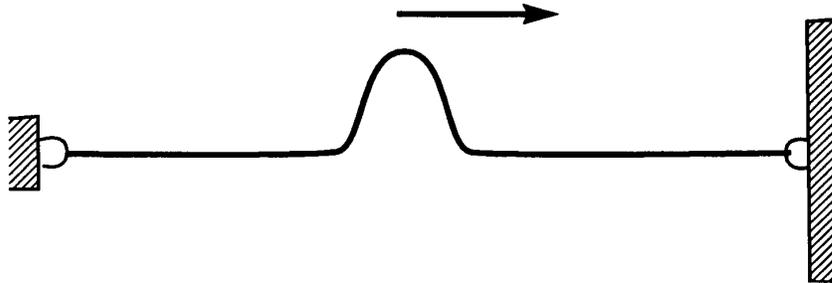
Transverse Waves

Figure 2.27: Displacement wave traveling along a string with fixed end

Although sound waves, and air column waves, are longitudinal, the principles of their propagation apply to all three wave types. It is more convenient to introduce these principles using transverse waves because they are easier to illustrate, and for most people, more intuitive. Figure 2.27 depicts a pulse introduced at one end of a string. This pulse will travel until it reaches the other end. If that end is fixed, the pulse will reflect and begin traveling back toward its start point. However, it does not return in the same form; it inverts upon reflection, and so travels back as an inverted mirror of its former self. This can be explained by treating the original pulse and its reflection as two different waves. Any number of waves can travel along the same string; the displacement of the string at a given point is equal to the sum of the displacements of each of the waves at that point. Figure 2.28 shows what happens when this is applied to a pulse in the process of reflection.

- As pulse A approaches the endpoint, it is the only wave in the summation, and the displacement is positive.
 - As pulse A travels “off the end” of the string, its reflection, pulse B, starts the journey back.
 - The displacement at any point is the sum of the displacements of pulse A and pulse B. However, since the endpoint is fixed, the displacement at the endpoint must be 0.
 - Therefore, assuming instantaneous reflection, the displacements of A and B must sum to 0, which can only occur if B is the negative of A. Thus, the pulse inverts.
-

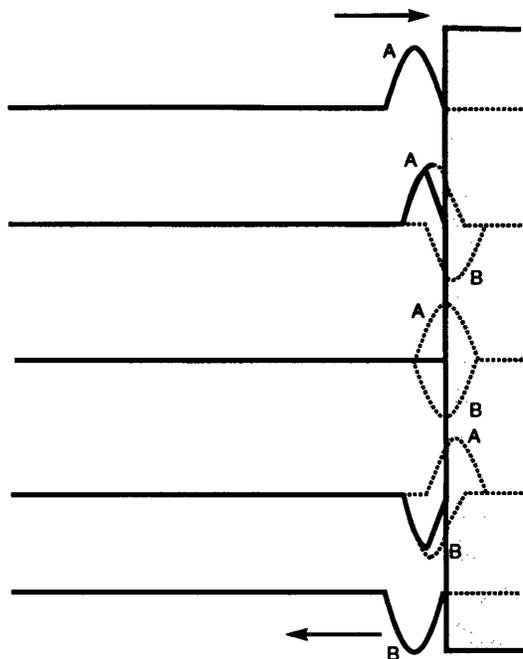


Figure 2.28: Reflection and inversion of a pulse on a string at a fixed end

- When it reaches its original startpoint, assuming that that too is now fixed, the pulse will invert again so that it reverts to its original shape.

Thus, the total period is equal to the time it takes for the pulse to travel twice the length of string. Using c as the speed of sound, this gives a period of $2l/c$ sec [2].

The previous analysis applied for a string with fixed end. For a free end (assuming a gravity-free environment in which the string can be suspended), the boundary conditions are reversed. Figure 2.29 illustrates this case. Now the displacement need no longer be 0 at the endpoint, but rather, will be at its maximum value. Using again the summation of left and right going waves, we get a surprising result; the total deflection at the instant the peak reflects is *twice* the height of the original deflection. This may seem at first to violate some conservation laws, but in fact, does not. Every displacement wave has its complement, a force wave. Since a free end can transmit no force, the force must be zero where the displacement is at maximum. Thus, at the free end, the force wave must invert in the same

way as the displacement wave did at the fixed end. The total instantaneous energy at a point is the product of force and change in displacement, $F\delta d$. Decomposing the energy into its left and right going components, we can see that the total energy is $F\delta d + (-F)\delta d$. As before, the sum is zero, even when the wave is at twice its normal height.

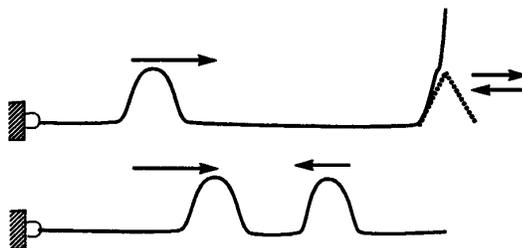


Figure 2.29: Reflection of a pulse on a string at a free end

The preceding paragraph illustrated the propagation of an impulse. The next step is the formation of a standing wave, which is the phenomenon that allows the entire string to serve as an oscillator. Figure 2.30 shows an example of a standing wave. Unlike the pulse traveling back and forth, a standing wave appears to be horizontally stationary; the string appears to move up and down at a fixed rate between points at which it does not move at all. What is actually transpiring is that the traveling wave and its reflection always sum in such a way that the observed motion occurs. At the points of no displacement, the *nodes*, the waves completely cancel one another. At the maximum displacement, the *antinodes* or *loops*, the waves maximally reinforce one another. All points on the string oscillate at the same frequency, even though the amplitude of their swath is a function of their location along the string. The precise mathematical requirements for a standing wave to occur can be found in most physics and acoustics texts, for example, in [46]. The result is a set of vibratory modes, with frequency $nc/2l$ where n is the mode number, which are possible on a string.

Intuitively, we can see that the modes are possible only when the boundary conditions can be satisfied. Figure 2.31 illustrates the first few. The first mode, or the *fundamental*, has only two nodes, at the ends of the string. For this mode, n is equal to 1, so its frequency is equal to $c/2l$ as above. The second mode has a third node in the center, so that the string now shows two opposing peaks. The frequency of this mode is $2c/2l$, or twice the

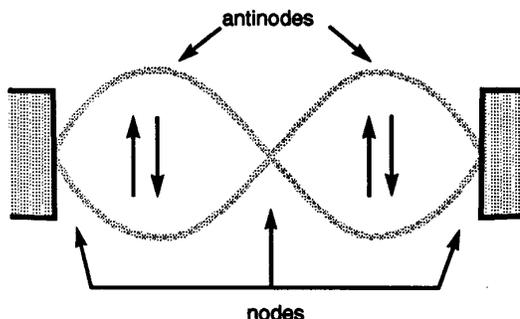


Figure 2.30: Standing wave on a string

frequency of the first mode. Musically, this mode is an octave above the first. The third mode has two nodes along the string, and three loops. It has a frequency of three times the fundamental, and so represents an interval of a twelfth, or an octave and a fifth. The fourth mode has four times the fundamental frequency, or the equivalent of two octaves. The fifth mode has five times, or the equivalent of two octaves and a major third. All these modes can coexist simultaneously. The mathematical significance of all of this is that, not only can a mode always be set up on a string which is some integer multiple of the fundamental frequency, but that the complex motion of a string, or any musical oscillator, can always be decomposed into component modes which have an integer relationship to one another. This is the basis of Fourier series analysis, which treats motion in the frequency domain.

The musical significance of the modal possibilities is that the sound generated by the string will contain components at many different, but related frequencies, even while its pitch is perceived as being at the fundamental frequency. The importance of overtones in determining timbre was touched upon in preceding sections. It is interesting to note that all of the harmonics representing intervals of octaves with the fundamental are even (although not all even harmonics, for instance the sixth harmonic, are at octave intervals with the fundamental). The odd harmonics all represent non-octave intervals. For example, the third harmonic is an octave and a fifth above the fundamental, while the fifth harmonic is a two octaves and a third above. The even harmonics thus have, as a group, a much different effect on timbre than do the odd, since a pitch an octave higher than a second pitch is considered to some extent to be equivalent. It will be shown that under certain

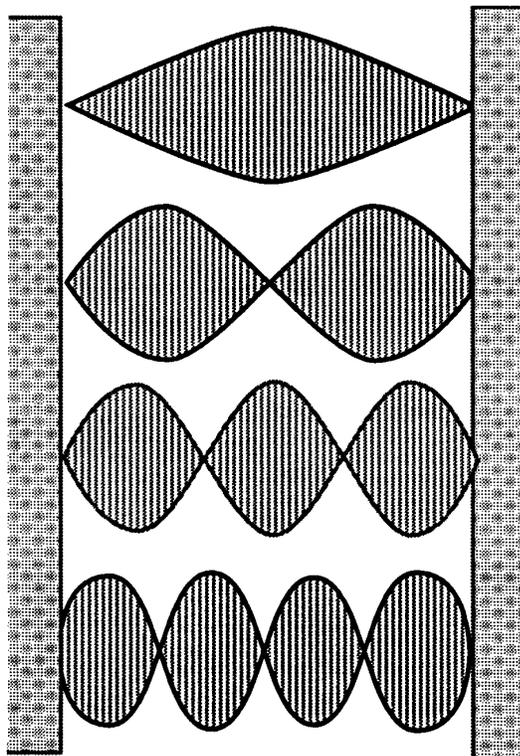


Figure 2.31: First four modes of a string

circumstances, the even harmonics will be suppressed, with important tonal consequences.

Section 2.2 noted that the timbre of an instrument was determined in part by the relative amplitudes of the harmonics. The harmonics of our string can be set arbitrarily by changing the point of attack. Pluck the string at a node, and the corresponding mode will never know. Pluck instead at the antinode, and that mode will be maximally excited. For example, if the pluck occurred directly at the center of the string, no even harmonics would be excited. Plucking exactly $1/4$ of the way from the end would give tremendous energy to the second harmonic, but absolutely none to the fourth. The musical tone of the string would change accordingly. One common special effect for bowed strings is to bow while lightly touching the node of some desired harmonic. All modes are damped out except for those harmonics with nodes at the touch point, and they will ring forth in sometimes haunting ways. Three

caveats must be mentioned here. First, the harmonics of a plucked string are not in reality exact integer multiples of one another; this is an idealized example. Second, a plucked string is a much different case than a bowed string. The oscillation resulting from a pluck is transient, and can bear the slightly stretched harmonics of the string. The oscillation resulting from bowing the string is sustained, a result of the nonlinear coupling between bow and string, and must have exact integer multiples. Coupling limits to some extent the freedom in mode selection that was available in the transient case. Finally, even though the timbre can be changed by changing the pluckpoint, those changes are somewhat subtle. The sound of the pluck and the nature of the decay are much more important in determining the characteristic of the sound than the location of the plectrum, at least in experiments performed by this author on a wire strung psaltery.

Longitudinal Waves

A compression, or longitudinal wave, travels in much the same way as the above transverse wave [2]. Figure 2.32 illustrates a compression wave in a tube. In this case, the medium is the *moving* column of air injected by the player. Whereas the speed of sound in the tube can be considered roughly constant under a given set of playing conditions, it does vary with temperature and moisture. Sound travels more quickly in warm, damp air than in cool, dry air. This gives rise to the necessity for “warming up” a wind instrument; when cool, it should play flat, and must be warmed to its proper pitch.

If we replace the string with an *acoustic tube*, or, in the simplest case, a cylindrical bore, the pressure wave will travel to the end, reflect, and, if the end is open, invert. The inversion is necessary by virtue of the same sort of boundary condition as that described above. At the open end of the tube, which must be a pressure node, the pressure, here the analogue of string displacement, must be equal to the ambient pressure. In other words, there can be neither compression nor rarefaction at that point. Since the pressure is the sum of the pressures of the inbound and outbound waves, the reflected wave will cancel the original wave at the endpoint.

The conversion of a compression into a rarefaction is somewhat harder to intuit than the change in sign of the transverse wave. A partial explanation is that, when the compression reaches the end of the tube, the momentum of the air will carry it, and the corresponding

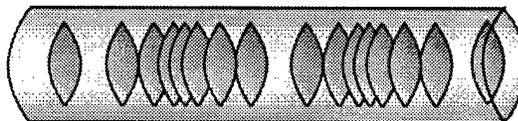


Figure 2.32: Compression wave in an acoustic tube

pressure node, past the physical end of the tube before the reflection occurs. (This is incidentally, also the reason that the acoustic length of a tube is slightly longer than the physical length, and that this acoustic length varies according to wave frequency. It is an important factor in tonehole design.) Once the high pressure air escapes from the confines of the bore, the thermodynamic equation of state, $Pressure \times Volume = Constant$ (at constant temperature) [73], requires that it expand into the environment, which is at a lower pressure than the internal bore pressure (if it were not, there would be no bias, i.e., nonsinusoidal, airflow). The inertia of the air will cause some dynamic overshoot in the expansion process, so that more air escapes than is necessary to equalize the pressure. But this means that upon reflection, the pressure of the pulse itself is reduced to a level lower than the ambient; it has become a rarefaction, or partial vacuum. Conversely, when this rarefaction reaches the opposing open end, the ambient air, now at a higher pressure than the rarefied pulse, will rush in. Once again, the inertia of the air will result in an overshoot, so that the result is a compression. (The air which escapes in this manner does *not* represent the sound energy transmitted. It is the transmitted pressure pulse which propagates through the ambient air, and not this dissipating air.)

Of course, if the end of the tube were a perfect pressure node, it would also emit no sound, since all of the energy would be reflected. As it turns out, physical reality in this case favors the musician. Nothing is perfect, particularly among acoustic instruments; some sound does radiate from the instrument, rather than reflecting internally. This is the function of the bells and toneholes, to be discussed in section 2.6.

A tube with two open ends is analogous to the fixed end string; a pressure wave will reflect back and forth in the same way. It inverts at the far end, and returns as a rarefaction. Upon reaching the start point, it reverts to a compression and the cycle starts anew. Again, the period is equal to $2l/c$. Just as with the transverse wave, the pressure at any point is the

sum of the pressures of the left and right going waves, and a standing wave can be set up in the same manner as that described above.

Now imagine that one end of the tube is stopped, so that air can neither escape nor enter. When the compression reaches the end of the tube and reflects, inversion is not physically possible, and so, it will remain a compression. This can be verified by examining the new boundary condition, which is similar to having a string with a free end. In this case, only a pressure antinode can be supported. This may be better understood by examining the dual of acoustic pressure, the velocity, or flow, wave. Because of the duality, flow nodes occur at the pressure antinodes. If the end of the tube is closed, the flow at that end must be 0; a closed end must be a flow node. Since a flow node implies a pressure antinode, the boundary condition for pressure is explained. Now, cancel, the ingoing and outgoing waves must reinforce each other, and no inversion occurs.

The lack of inversion at a closed end has important implications. First, the period of oscillation is now four times the length of the tube:

- A pulse, beginning as a compression, will reflect off the far, open end as a rarefaction.
- Returning to the closed end, it will reflect once more without inversion. It is still a rarefaction at this point, and the cycle is not yet complete.
- Now it travels back to the open end, inverts into a compression, and finally returns to the closed end, reflects as a compression, and starts a new cycle.

The standing wave set up by this type of oscillation will be an octave lower than one set up in an open pipe of the same length. In addition, the only waves which can be supported will be those with a node at one end and an antinode at the other. Figure 2.33 illustrates this, using transverse waves to represent the longitudinal waves described. Because the tube length encompasses only one fourth the wave, the waves shown in Figure 2.31 are cut in half. We see that the only waves which satisfy the boundary conditions are the odd harmonics, because these waves have an antinode at their midpoints. The even harmonics have nodes at their midpoints, and therefore, nodes at both ends when bisected; our cylindrical bore will not support even harmonics. This is very important, because Helmholtz found that a reed instrument acted like a tube stopped at one end [25], implying that a cylindrical reed instrument, for example, a clarinet, has ideally *no even harmonics* (although in reality,

because of other factors, the instrument does produce weak even harmonics).

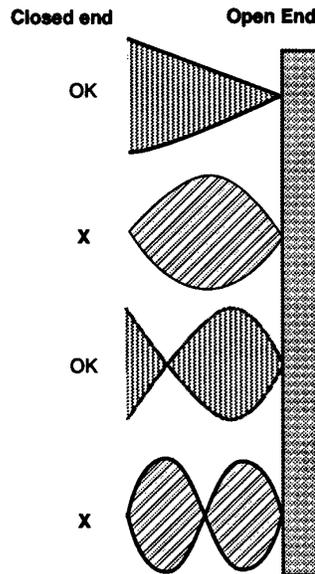


Figure 2.33: Compression modes in a pipe stopped at one end and open at the other

The discussion above essentially ignores the steady component of airflow which results from the difference between mouth and ambient pressures. Recall that the propagating medium is a *moving* column of air, which implies asymmetry in the wave propagation, since the waves traveling away from the player will be augmented by the input flow. It is useful to think of the wave propagation as symmetric, purely oscillatory flow superposed upon the steady flow bias. The wave behavior can then be analyzed independently of the bias, and therefore understood in a less complicated context.

2.3.2 Acoustical Bore Shapes

Cylindrical Bores

The previous section examined the acoustical properties of a cylindrical tube. Presented as the analogue of the stretched string, it is clear that this tube has potential as a musical instrument. In fact, an arbitrarily shaped bore will propagate pressure waves, but only

certain bore shapes will support the harmonically related standing waves. It happens that the cylindrical bore is one of these, which is convenient because it is also the easiest to analyze and to simulate. To summarize, the properties of a cylindrical bore are:

- Fundamental frequency of $c/2l$ if open at both ends, $c/4l$ if stopped at one end
- Supports all harmonics if open at both ends, only odd harmonics if stopped at one end
- Harmonic frequency for a stopped bore is equal to $nc/4l$, where n is always odd and is equal to the mode number. For an open bore, harmonic frequency is equal to $nc/2l$, where n can be any integer.

A later section will address the issue of the effects at the end of the tube, which in fact will act as lowpass filter, and impart a cutoff frequency to the bore beyond which harmonic modes are greatly attenuated.

Although the cylindrical bore is easy to understand, analyze, and build, it is not used for many modern instruments. The clarinet, with its characteristic hollow sound, is the only popular cylindrical bore reed instrument. Many earlier reed instruments, such as those in the capped reed cornemuse family, were cylindrical. However, the suppression of even harmonics puts a major limitation on the tonal coloring. An even more important implication of the even harmonic suppression is how it affects the upper register of the instrument.

The *register* of an instrument indicates which harmonic mode the player has excited. Excitation of the fundamental mode implies a note in the lower register, with the other harmonic modes present as overtones. However, by any number of techniques, to be discussed below, such as embouchure adjustment, overblowing, or opening a register hole, the player can suppress the fundamental and excite instead a higher mode; this higher mode now serves as the fundamental, and provides the basis for the associated harmonic series. In the brass instruments, players ascend through a number of modes by tightening the lips. On valveless brass instruments such as the bugle, the only method for changing pitch is to hop modes, and the only notes available are therefore those in the harmonic series of the lowest note of the instrument. The hunting horn is also valveless; however, by changing the position of his hand in the bell, the player can alter the bore enough to get the intermediate notes in the

scale, at the expense of some uncontrollable tonal changes. In general, only the first two or three modes are used in woodwind instruments. However, an interesting earlier woodwind, the tabor pipe, is played using many of the upper modes. This instrument, as implied by its name, is intended to be played with a tabor, which is a type of drum. The pipe, similar in fipple mouthpiece to a recorder, is played with one hand, the drum beaten with a stick by the other hand. This puts some rather severe constraints on the tonehole configuration of the pipe, which must be both supported and fingered by one hand. In fact, the pipe has only three holes, but with a combination of fingering and blowing pressure changes, the player can accomplish a diatonic scale with a range of an octave and a fifth [45] [6].

The ability to change registers is crucial for an instrument to have any sort of playing range. As Section 2.1 indicated, one of J.C. Denner's great innovations in inventing the clarinet was to include a vent hole that allowed a register change. One of the major limitations of the capped reeds was that, because they had neither vent hole nor a way for the player to exert embouchure pressure, they had a range of just over an octave. Only a small subset of music literature accomodates such instruments. The cylindrical bore has an important drawback in this regard; its second harmonic is not the musically convenient octave, but the twelfth. In other words, the overblown C will become, rather than C', a G', and the fingering patterns for the two registers will consequently be quite different..

Conical Bores

It turns out that the cylindrical bore is not the only bore shape with musical properties; a conical bore will also support standing waves in harmonic proportions with one another. In a conical bore, spherical waves, as opposed to the planar waves in a cylindrical bore, will propagate back and forth as before, with some deformation and acceleration due to the varying cross-section. However, the reflection/inversion characteristics of the conical bore are not quite the same; the conical tube, stopped at the small end, behaves like a *cylindrical tube open at both ends*. In other words, the wave propagation is similar to that discussed in the fixed-end string example. The pulse must cover only twice the length of the tube to complete a period oscillation, rather than four times the length of the tube as in the cylindrical case, so the tube will sound an octave higher. Also, as with the string, the conical tube will support *all* harmonics, both odd and even. The tone will be quite different — this was demonstrated in the previous section — and more importantly, a conical tube

instrument will *overblow the octave*. This is much more convenient. On the shawm, an acoustically elegant instrument, switching registers is simply a matter of tightening the embouchure. Because the new note is exactly an octave above the normally blown note, the fingerings in the upper register are almost (but not quite) the same as those in the lower register. This register change is relatively easy to accomplish, whereas on the open reed cylindrical douçaine, it is virtually impossible. (According to D. H. Smith [62], adjusting the reed resonance properly will allow higher notes to be played even on a cylindrical bored instrument; it is however, not common to do so, and the limited range of the cylindrical instruments is for the most part accepted as a fact of life.)

An intuitive discussion of why a conical bore behaves the way it does is difficult to come by. Indeed, given the explanation of inversion due to airflow in and out of the ends, its behavior seems contradictory. If the period is defined by only two lengths of the tube, then the wave must reflect in the same manner at both ends, i.e., it must invert (or not invert) at both ends. That it behaves in this manner can be proven mathematically, as Morse does in [46]. Backus tried for an intuitive explanation in [2], but admitted failure and referred the reader to the less intuitive mathematics. Suffice it to say that the difference has to do with the fact that the wavefront is spherical and that the bore itself diverges.

In summary, the salient properties of a conical bore instrument are:

- Fundamental frequency of $nc/2l$;
- Supports odd and even harmonics;
- Overblows the octave.

Of course, a conical bore is never completely conical. If it were, there would be no aperture in which to insert a reed. A so-called “conical” instrument is really a truncated cone; the reed and staple to which the bore is coupled must, for the instrument to work properly, complete the cone. The concept of “equivalent volume”, or the acoustic volume of reed and staple, will be discussed below. No matter where the cone is truncated, however, calculations for register hole placement, tonehole configuration, etc., are based on the assumption that the instrument extends to the conical apex; the end of the instrument is considered to be at the apex, even if the physical instrument does not extend that far.

The cylindrical bore and the conical bore are the two basic shapes used for woodwind

instruments. They are modified to some extent with toneholes and bells, but still can be characterized by one of these two shapes. An instrument which departs too much from the ideal shape will not play well, because its resonances will not be harmonically aligned. These two boreshapes are not the only useful ones; brass instruments are hybrids of cylindrical tubing with exponential flare. The vocal tract can take on arbitrary bore shapes, although, since the voice pitch is the result of vocal fold resonance, rather than vocal tract resonance (i.e., pitch is independent of the tract), these shapes need not be acoustically proper [17].

2.4 The Reed

The previous section described the air column, the resonator in the coupled generator/resonator system that comprises a reed instrument. This section examines the reed, which serves as the generator. Both air column and reed will be presented as essentially linear elements. However, these two sections will set the stage for the next section, which will detail the very nonlinear coupling between reed and air column.

2.4.1 The Reed as a Pressure-Controlled Valve

A reed, in the sense used in this paper, is simply a mechanical flap that responds to fluctuations in pressure by changing the size of the orifice through which the player introduces air. Figure 2.4 illustrated the mouthpiece and reed of a clarinet. The player holds the reed in his mouth, and blows air into the aperture between the tip of the reed and the facing of the mouthpiece, known as the *lay*. There is a short section through which the air must pass until it enters the bore, but the bore extends essentially directly beneath the reed up to the opening. Thus, the pressure on the bottom of the reed is the bore pressure, while the pressure on the top is, for the section within the mouth, the mouth pressure, and outside the mouth, the ambient pressure. In addition, the player usually supplies some embouchure pressure, which is really a point force applied at lips.

Figure 2.34 shows a diagram of the opposing pressures acting upon the reed. The reed, being flexible will bend away from resultant force. Thus, if the mouth pressure, which is roughly constant, exceeds the bore pressure, the reed will begin to close, constricting the aperture. As the bore pressure increases, it will force the reed back up, reopening the orifice.

A stiffer reed will require a greater fluctuation in pressure to deflect as much as a soft reed. Figure 2.35, taken from [9], illustrates a typical flow versus pressure curve. In this figure, positive flow denotes flow into the instrument, and the pressure difference is $p_{mouth} - p_{bore}$. Below a threshold pressure difference, here denoted as p_{ft} , the flow increases with pressure difference, as it would through a simple nondeflecting aperture. Beyond p_{ft} , the closure of the reed begins to constrict the flow; as the pressure grows, the flow now decreases. At the point p_c , the reed finally closes completely.

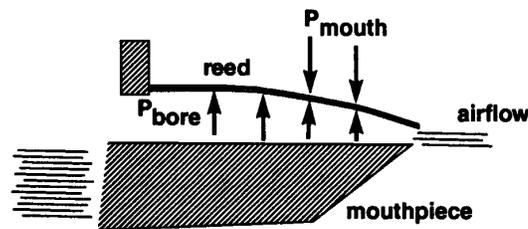


Figure 2.34: Opposing pressures acting on a single reed

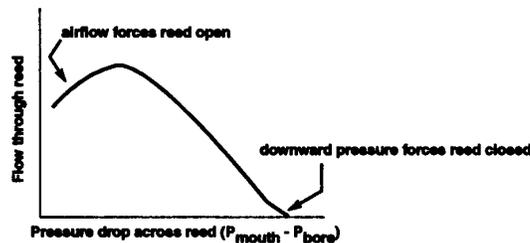


Figure 2.35: Typical Flow-Pressure curve for a single reed (after Benade)

Most literature defines reed deflection in terms of tip displacement, x , and pressure, P , which integrates to force. In fact, these are really small angle approximations. The reed moves rotationally, not translationally, and rotation angle and driving torque provide alternate, more accurate, descriptions of the static and dynamic forces. For example, the deflecting power of a force depends on its moment arm, or distance from the main pivot point. This can be an issue when changing embouchure position, or describing the change in pressure sensitivity as a clarinet reed curls over the sloping lay. For the most part, however, since

reed length stays roughly constant and the angle of its swath is very small, the translational approximation is adequate and convenient.

It is worth noting that the actual flow through the reed is relatively small compared with the flow through an air reed instrument such as the recorder, while required breath pressure is relatively high. Players refer to this as high “back pressure”, or resistance. An early lesson to all reed players is to make sure to breathe out before breathing in, as the lungs are usually full of deoxygenated air that has not yet been expelled. It is usually easy to tell by his obvious discomfort when a reed player has disregarded this rule.

The next section will discuss more specifically the implications of the coupling between reed and air column. However, it is necessary to describe the physical interaction of the two to put the behavior of the reed in context.

- As the player blows into the orifice, he creates the flowing air column that was described in the previous section.
 - At first, if he attacks with a low breath pressure, all he will do is start a dc, i.e., nonoscillatory, flow through the tube.
 - However, as he increases the breath pressure, the mouth pressure acting upon the reed will begin to close it.
 - But this will constrict the air flow, and cause a rarefaction pulse to travel down the air column.
 - As we know from the previous section, the rarefaction will be reflected as a compression.
 - This compression represents an increase in bore pressure; when it reaches the reed end, it will tend to push the reed open.
 - This allows more air to flow in, so that the compression is reflected not only without inversion, but with additional pressure from the added contribution of the input stream.
 - Reflecting again from the far end as a rarefaction, it returns to the reed as a low bore pressure and pulls the reed shut (not necessarily completely). This shutting
-

will constrict the air influx and intensify the rarefaction.

One could start this cycle with the initial compression pulse, depending upon the sharpness of the attack. Either way, the reed intensifies the natural oscillations of the air column, not only by allowing the player to inject energy into the system, but allowing that energy to be injected at the ideal excitation frequency; the reed will oscillate, with a periodic motion, at roughly the frequency of the air column about some null point that is a function of the embouchure set pressure, and modulate the input airstream to that frequency. Initializing the process requires a certain level of breath pressure, the “threshold blowing pressure”, which Backus characterized in his initial clarinet paper [5]. If the input pressure is too high, the reed will be forced shut before any oscillations can begin, and will not reopen. It is worth mentioning that the reed can actually sustain some oscillation without any connection to the bore, i.e., without any acoustic feedback. This is most likely a consequence of the hysteretical hydrodynamic forces (flow dynamics) within the reed channel [27], and is probably not significant in the coupled system.

2.4.2 The Reed as a Harmonic Oscillator

So far, our isolated reed looks like a straightforward valve. None of the dynamic characteristics of the reed, save for the compliance which governs the sensitivity to pressure fluctuation, have played a role in the basic generation of the sound. However, the reed, acted upon by the embouchure, is a dynamic system in its own right. It has mass, the mass of the reed itself. It has compliance, which imparts its springlike character. It has damping. This damping is a characteristic of a cane reed, and is partially responsible for the ability of the reed to vibrate easily at a frequency lower than its own natural frequency. Hall distinguishes between “hard” reeds, generally metal, which have little damping and tend to vibrate at their own resonance frequency, and “soft” reeds, such as the cane reed, whose internal resonant vibrations are quickly dissipated and which can therefore be driven easily at a nonresonant frequency [24]. The damping can be considered to be partially a function of the embouchure, or lips, since the reed is actually pushing against the lip in its vibration. Consequently, a reed has its own resonance, which is generally at a much higher frequency than the fundamental note of the instrument. In Backus’ landmark analysis in [5], the mass of the reed was neglected. Here, studies with varying embouchure pressures revealed the importance of reed damping in tone production. Wilson and Beavers studied

this matter more carefully in [75], including mass in their analytical model, and coming to basically the same conclusion.

One interesting phenomenon related to reed resonance which has received some contradictory explanation is that of the reed “squeaks” that plague inexperienced players. In [75], Wilson claims that insufficient reed damping will cause the reed resonance to be excited, rather than one of the lower bore resonances. In this study, a damping ratio of 0.4 was sufficient to quell any tendency of the reed to jump to its own resonance, even at low blowing pressures. Benade, in [9], modified this to some extent, by defining the “reed plus staple” frequency that would sound in the absence of a bore. Another theory is that the squeaks are torsional waves set up in the reed because of an uneven embouchure, rather than a function of damping (find other reference). From a musical standpoint, squeaking is often said to be a consequence of insufficient breath support [40]. Indeed, it will be shown in the simulation results that under certain conditions, the third register would sound at threshold blowing pressure, with the desired fundamental register coming in solidly as the input pressure was increased. Such a register shift could be perceived as a squeak at an inappropriate time, especially in a cylindrically bored instrument which, recall, would sound an interval of two twelfths. The damping argument is seemingly contradicted by the performance of capped reeds, which do not allow the embouchure to touch, and therefore to damp, the reed. In addition, plastic reeds, which would probably have less damping than a cane reed, are frequently and successfully (some might argue this) used in the less expensive instruments (Myers covers the relative merits and drawbacks of plastic reeds in [49]). This would seem to imply that instruments such as the crumhorn should issue only the resonance frequency, which is not the case. However, Herbert Myers, in [50], has pointed out that standing waves set up in the windcap would couple with the reed in a way not modelled in [75]; the effects of windcap, and its biological counterpart, the vocal tract, is a subject of some debate. More importantly, cornemuse reeds are designed specifically to be less prone to squeaking when undamped than lighter reeds. This design requirement, satisfied by building a reed with a wide throat, is in fact one of the main contributors to the limited dynamic range discussed in the previous section; a reed which responds well in the upper register is also more prone to squeaking [62].

In general, however, the reed of a woodwind oscillates at a much lower frequency than

resonance, and its valving characteristics are much more important than its dynamic characteristics. Studies in [25] with a stroboscope demonstrated a fairly simple flapping motion of the reed tip. This is not the case for all reeds. Metal organ reeds oscillate at their resonant frequencies, and tuning is accomplished by adjusting the effective length of the reed itself with a tuning wire [2]. The oscillation of the party horn discussed in Section 2.2 was dominated by the reed resonance. This was demonstrated by punching a hole in the cardboard cone which formed the “bore”; the hole had no effect on pitch, indicating that the reed resonance, insensitive to bore length, was excited, rather than the bore resonances. Even the complete removal of the cone (accomplished by accidentally stepping on it) had little appreciable effect on the pitch or tone. A more musical example is the human voice; the vocal folds are only weakly coupled to the bore formed by the vocal tract, and oscillate at their own frequency, i.e., their “reed resonance” [17]. If this were not the case, we would sing by contorting our vocal tract and not by changing the tension of our vocal cords, and would not be able to form vowels without affecting our pitch.

Although the dynamics of the reed are of secondary importance in describing basic reed woodwind behavior, they cannot always be totally neglected. From a practical point of view, the reed resonance is important because it represents the highest playing frequency possible. According to Thompson in [71], early saxophones had a limited playing range because the mouthpiece facing designs did not allow the reed resonance to be raised enough. Improved facing designs allowed a higher reed resonance, and the saxophones were given a new octave of playing capability. Another important aspect of reed resonance is the function it performs in the upper register. It turns out that for higher notes, the reed resonance can play a major role in stabilizing tone [72], especially under conditions of otherwise marginal stability. In [60], D. H. Smith frequently emphasizes the importance of the reed resonance in the more practical environment of reed-making. This is an aspect of the modal cooperation to be discussed in Section 2.5.

2.4.3 Additional Nonlinear Effects

The above paragraphs assumed essentially static forces on the reed. However, the air column is formed by air flow, and where there is air flow, there are hydrodynamic effects. This is a subject currently under considerable debate. Earlier researchers had all assumed that the flow through the reed was laminar viscous flow to which Bernoulli’s equation could be

applied. In [76], Worman described and quantified the Bernoulli force, and others, including Keefe in [34], have stated it as fact. The basic idea is that there is pressure drop associated with increased air flow; the faster the particle velocity, the lower is the air pressure. This relationship is defined by the Bernoulli equation

$$P_B + \frac{1}{2}\rho v^2 = \text{const} \quad (2.1)$$

The *Bernoulli Force* is the lift force which allows birds and airplanes to fly, and also, according to some, what causes reed instruments to work the way they do. As the reed tip closes and constricts the opening, the particle velocity through the aperture must increase while the aperture width itself decreases. This will increase the Bernoulli pressure drop, and therefore increase the downward force on the reed. Thus, the nearer the reed tip approaches the lay, the stronger pulls the downward force. This is a nonlinear addition to the simple harmonic oscillator we had before. For a single reed, Worman calculated that the total Bernoulli force was only about 2% of the total force required to close the reed. Benade experimented with the shape of the lay, and found the changes in lay shape had considerable effects on tone; he attributed these effects to changes in the Bernoulli forces which resulted [9]. Schumacher [58] included the Bernoulli effect, as formulated in [51], in his clarinet model, and evaluated its importance by turning the force “on” and “off”. He found substantial transient and spectral effects in the beating regime, including a sharpening of the playing frequency by a half semitone (3%). However, Hirschberg, a true fluid dynamicist, has contested the assumptions on which these findings are based [27]. Hirschberg is currently researching the area of hydrodynamic effects in reed instruments in detail, and can be looked to for important findings on the behavior of airflow through the reed aperture. The Bernoulli force as postulated by Benade and Worman, and Hirschberg’s arguments against this theory, will be presented in more detail in the sections detailing the mathematical models.

Another nonlinear wrinkle in this treatment is what happens if the reed finally touches down on the lay. This event, known as *beating*, causes a considerable change in coupling characteristics, and eliminates the linearity of the oscillator. Beating does not occur in all playing. Indeed, in German style clarinet playing, Benade [9] states that that beating took place in only the very loudest sections, while the French style incurs beating at a much lower dynamic level; this implies a softer reed, since a stiff reed moves less readily, and requires higher pressure fluctuations to deflect all the way to the lay. Beating causes the following

phenomena to occur:

- The wave motion described by the motion of the reed tip is clipped on one side, because once the reed is closed, the tip cannot move any further. This clipping is less abrupt on a clarinet because of the slope of the lay. When the reed first makes contact with the lay surface, it does not cut off that part of the bore aperture sloping away from the reed. As the pressure difference driving the reed increases, the reed will curl around the lay, eventually closing the aperture completely. However, because the moment arm is small, the force required to effect this bending is much higher than that which was needed to drive the original deflection (one could view the reed as a non-linear spring, becoming stiffer as it deflects further). Consequently, the shutoff of air will occur relatively gradually;
 - The flow of air through the mouthpiece stops completely. Any hydrodynamic effects in process will cease abruptly, since they require flow. Just as with the tip motion, the pressure wave representing air influx will be clipped; in addition, the dc component of flow will be temporarily halted.
 - The effective reed resonant frequency will increase as the tips are moved closer together, because they have less distance to travel
 - The reflection coefficient of the reed, which will be described in more detail in Chapter 3, will be equal to one, implying that the returning wave exactly equals the incident wave.
 - The impact of reed against lay could generate a pressure disturbance (i.e., sound) which, in low instruments like the bass saxophone, can be audible. The nature of this sound will depend partially on whether the collision is elastic or inelastic. If elastic, the reed will bounce back off the lay, and by virtue of its own dynamics, rattle against the lay until the oscillation dies out. If inelastic, no bouncing will occur; all kinetic energy is absorbed in the collision. (Stewart and Strong concluded that the collision was inelastic [70], but Hirschberg has found evidence that there is some elasticity [30])
 - As the first point mentioned, the lay itself is a curved surface. Not only does the curving of the reed around the lay round the edges of the clipping process, it
-

also effectively shortens the reed and influences to some extent the reed resonance. However, in light of the stroboscopic evidence cited earlier, which indicated a simple flapping motion, and Hirschberg's studies, this effect is probably negligible. Again, Hirschberg will probably be producing the next important data in this area.

- The pitch will rise with blowing pressure. There are several explanations proposed for this, some based on beating, and some on nonbeating. One theory, postulated by Bak and Domler [7], but controverted experimentally by Hirschberg in [27], is based upon the lay-induced curvature changes described above. Bak's theory suggests that the reed changes not only its own length and resulting resonance, but the length of the overall air column. This would raise the pitch if the change were appreciable. In [27], Hirschberg calculates the beating frequency as a function of tip displacements, and suggests the rise in pitch is a direct result of the shift in tip equilibrium position due to the tightening of the embouchure, the hydrodynamic influences, and the effects of elastic collision against the lay. Another possible explanation, related to Hirschberg's, comes from the field of nonlinear mechanics. In Duffing's equation, the nonlinear spring force is represented by $kx + k'x^3$. Here k is the standard positive linear spring constant and k' is the nonlinear constant, which can be positive or negative. According to Ogata in [53], as the amplitude of oscillation changes, the frequency changes as well, according to the sign of k' . For a reed deforming over the lay, the effective stiffness increases, corresponding to a positive value of k' , which would normally decrease the oscillation frequency. However, in the presence of hydrodynamic forces which "suck" the reed toward the lay or the other reed, the nonlinear spring term might be effectively negative, resulting in an increase in frequency. Finally, Worman cites the rise as being a result of the modal cooperation between misaligned resonances [76]. This view, which if anything assumes the absence of beating, will be discussed more thoroughly in Section 2.5.
 - Rocaboy claims in [57] that when a double reed attached to a conical bore closes completely, the air trapped inside the reed channel, behaving like a compressed spring, begins a free oscillation that forces the reed prematurely open. Although he does give many details, he states that the same result occurs for a single reed with a conical bore. If true, this would add an additional dynamic element only during the interval of the beat. However, this theory is subject to debate.
-

These are just some of the nonlinear effects present in a woodwind instrument. Future chapters will treat them, and their consequences, more rigorously.

2.4.4 Double Reeds

So far, only single reeds have been discussed. Double reeds are a far more complicated case. Figures 2.1 and 2.10 illustrated the shawm and its reed, in form one of the simplest of the double reed instruments. The reed channel is quite a bit different from that in the single reed channel discussed above. Instead of a short channel section, beyond which the reed lies essentially over the bore, the reed channel is now a relatively long passageway bounded on both sides by reed, connected by an even narrower metal tube, the staple, to the bore. The reed valving will still be coupled to the air column oscillation, but the mechanism is much less clear. Whereas with the single reed, it was easy to demonstrate a clear relationship between constant mouth pressure, periodically fluctuating bore pressure, and the aperture variation of the reed, while ignoring hydrodynamic effects for the most part, the role of the reed channel is now much more important; hydrodynamic influences cannot be ignored. The reed channel could be considered an extension of the bore, so that the difference between mouth and bore pressures will produce the force on the reed as before. However, the reed channel is much narrower than the bore, and air will flow through with a different particle velocity. Certainly the pressure fluctuations in the bore will influence the air flow through the reed, but the specific relationships are as yet unknown. Even Hirschberg, in his investigations to date, has considered the double reed too difficult to focus on at first, although some preliminary work has begun [30].

Hydrodynamics aside, it is still possible to discuss the double reed in a more general sense, to lay the foundation for a somewhat intuitive understanding of its operation. The valving is now achieved by the two reeds closing upon one another. Figure 2.36, taken from Rocaboy's study in [57], illustrate the mouthpiece pressure waveforms for a clarinet reed and a bassoon reed. These are reminiscent of the output waveforms discussed in Chapter 2.2, but are closer to the source of reed action. They contrast not only the effects of single and double reeds, but cylindrical and conical bores. Note that the clarinet waveform has a roughly symmetrical shape, closed for about the same amount of time as open. The bassoon reed, on the other hand, is closed for a small fraction of time, a result Rocaboy attributes to the air spring effect described earlier. Figure 2.37 illustrates the waveform of a double reed mounted

within a short cylindrical segment, as measured by Van de Laar with the optical setup described by Backus. The waveform here is quite symmetrical, although the rising slope is steeper than the descending slope. A comparison between this waveform and Rocaboy's bassoon waveform implies the profound effects of the boreshape on reed behavior. One other interesting aspect of Van de Laar's waveform is the obvious beating which occurs. Rocaboy assumed that, on his bassoon pressure waveform, the negative portion of the pulse corresponded to the time that the reed was completely closed. This assumption plays an important role in his air spring theory, but he did not actually measure the deflection of the reed.

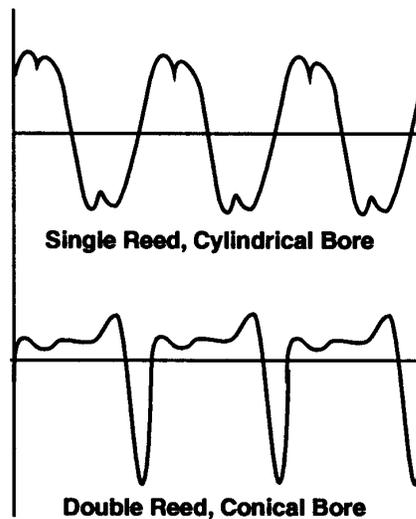


Figure 2.36: Mouthpiece pressure waveforms measured by Rocaboy

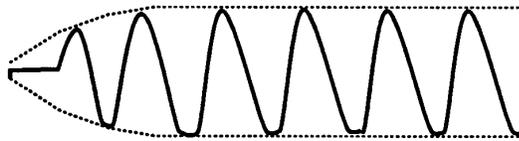


Figure 2.37: Waveform of a double reed in a short test cylinder, measured by Van de Laar

The double reed is a far more sensitive device than the single reed, and reed making occupies

a great deal more time for the oboist or shawmist than for the clarinet player. Part of this is because of the greater number of variables offered by the reed — stiffness, dimensions, scrape; the more options one has, the more opportunities one has for problems. Another consideration is the conical air column to which most double reeds are attached. In order to fulfill the musical requirements for the instrument, the reed assembly must “complete the cone”. Again, because the assembly is so unconstrained, preparing the reed is an art in itself. In [61], David Hogan Smith lists the three basic requirements of a double reed:

- The equivalent volume of the reed and the resonant frequency of the reed assembly must be matched to the bore of the instrument;
- The reed should be free blowing and responsive in all registers;
- The reed should produce a desirable tone color.

The first point deals with the completion of the cone. This is not, of course, simply a reed issue, but a reed plus bore issue. However, it is discussed here because adjustments are limited to variations of the reed. As Section 2.3 discussed, the reed assembly is attached to the truncated cone of the bore. For the instrument to work well, the reed and its connecting tube, the staple, must provide the missing volume. A quick glance at any double reed instrument would indicate that no instrument could possibly meet this requirement. However, Benade, in [9], claimed that the *acoustic* volume of the compliant reed was much greater than its *geometric* volume, and defined the *equivalent* volume as the volume of a “hard-walled cavity whose shape is chosen to make the acoustical effect on the natural frequencies of an air column identical with the composite effect of the actual reed-plus-cavity on the playing frequencies of the air column”. This is a somewhat empirical definition. It basically says that, since we assume from our acoustical knowledge that the reed must complete the cone, it must represent more than its measurable geometric volume.

In [4], Backus cites Nederveen for calculating the equivalent volume due to compliance as being equal to $0.9 V_0$, with V_0 the missing cone volume. Backus tested experimentally the effects of compliance on the equivalent volume and resulting frequency shift. He found that compliance had only one fifth of Nederveen’s predicted effect. In this same paper, Backus challenged several other theoretical assumptions. Some of these findings are discussed in Section 2.7, which covers the discrepancies between theory and reality.

A reed connected to a cylindrical tube will also have an equivalent volume which will affect the playing frequency. Here, the reed volume is not quite as critical because it is not responsible for fulfilling the need of an acoustical bore. However, since the toneholes have been designed for a given reed volume, the instrument would not work as well, even ignoring the pitch deviation, because the resonances would not be as well aligned. The following section discusses the importance of cooperating resonances.

Benade worked out the equivalent volumes for one reed by attaching it to various lengths of cylindrical tubing, and, from the resulting playing frequency, calculating the equivalent length of tubing for required for that note. He found that the equivalent volume remained roughly constant over a wide playing range. In addition, he found that softening or thinning the walls of the reed cavity enlarged the equivalent volume. In the case of a cone truncated only to a small degree, changes in the equivalent volume had a much more pronounced effect on succeeding modes; a slight change in the fundamental would have four times that change in the second mode, and nine times that change in the third. Thus, a slight sharpening of the fundamental, a result of too small a cavity, would result in an even sharper octave; this is referred to as *widening*. Conversely, a cavity with too much volume would have *narrowed* octaves; the octave would be flattened much more than the fundamental. This could be interpreted as a reedmaking tip. If the cone has been properly completed by the reed, the second mode of the instrument, will be precisely one octave higher than the first mode. Otherwise, the second mode will be either sharp (too small a volume) or flat (too great a volume). A wide octave can therefore be fixed by thinning the blades of the reed, and thus increasing the volume. A narrow octave requires a shortening of the reed, which effectively stiffens it and decreases the equivalent volume. However, D. H. Smith, in [60] believes that the reed resonance plays a much more important role than the equivalent volume in determining the pitch of the octave. Of course, the embouchure plays an important role in the overall sound of the instrument, as well as in the final adjustments.

The equivalent volume requirement represents a low-frequency constraint on reed design. An analogous high-frequency requirement for completing the cone is that the playing frequency of the reed assembly, Benade's *reed-plus-staple* frequency, F_{rs} , is correct. This frequency, which is not the same as reed resonance frequency, but which is affected by the same reed parameters, should be equal to that of the missing part of the cone. Thus, if x_0 is the length of the missing cone section, then, using the period defined in Section 2.1 of $2x/c$, we get the

relationship: $F_{rs} = c/2x_0$, where c is, again, the speed of sound. Like the equivalent volume, F_{rs} is heavily dependent on the scrape of the reed; A thin blade will have a lower frequency than a stiff one. In fact, we will see in Chapter 3 that the undamped reed resonance is equal to $\sqrt{\text{stiffness}/\text{mass}}$. A thinning of the blades will cause a reduction in both stiffness and mass. If these effects occurred in parallel, the resonance would not be affected. However, practical experience has proven that a thinner blade does have a lower resonant frequency, and that the scrape therefore influences the reed stiffness are much more than the reed mass. Benade found that, as with equivalent volume, F_{rs} stayed constant on a given instrument over a wide range. D. H. Smith provides practical advice for adjusting the reed. For a lower F_{rs} , one can cut a fan-shaped reed with wide tip, or, if the reed is already cut, thin the blades. Maintaining a loose embouchure, and playing nearer the tip of the reed, will also lower F_{rs} . Conversely, F_{rs} can be raised by tightening the embouchure or playing further down the reed. Unfortunately, one cannot add material to the blades, so a low F_{rs} can be fixed only by playing control.

The second of Smith's reed requirements is that it both play freely and be responsive in all registers. This is a trade in stiffness and mass. According to Smith, a reed with less mass will be freer blowing, and will "speak" more readily in the lower octave. In the upper register, stiffness is required, especially on an instrument such as shawm with no register hole, where register shifts are achieved entirely through embouchure control. Generally mass and stiffness go hand in hand in the reedmaking process. However, by using the stiffest part of the cane, directly beneath the bark, a lighter reed will result. Myers suggests in addition [49] that the tips can be thinned while the bulk of the blades are left relatively thick. This provides the necessary stiffness, while allowing a crisper attack and freer speaking.

Finally, the tone color of the instrument can be adjusted through the reed. According to reed makers [60],[50], the critical factor in tone color is the way in which the tip closes. Figure 2.38 illustrates the tip of the reed. If the reed is scraped evenly, the tip will close, ideally, at one time. This will result in a more square waveform, and consequently, higher harmonics and a brighter tone. If the reed is given a "spine", i.e., scraped more on the sides than at the center, the tip will not close at once, but rather, close first at the tips and then work in toward the center. This rounds the waveform, and results in a darker tone. That this type of tone control is considered an essential part of reed making indicates an important observation about double reed mechanics; the reed must, in normal operation,

close, or "beat". If it did not, than the degree of spine in the reed would not have developed into such an important factor in reed design. (Hirschberg and Van de Laar have in their preliminary experiments verified that the double reed does beat under normal operation [30]).

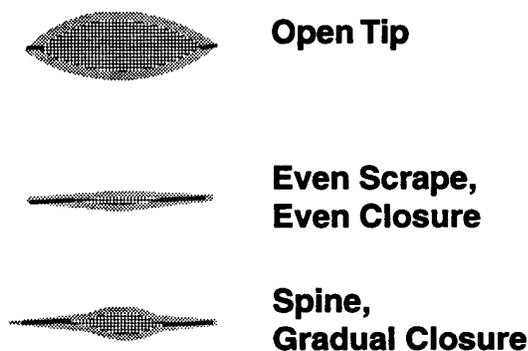


Figure 2.38: Closing aperture of a double reed

One other variable which effects both tone color and playing comfort is tip displacement [49]. A larger displacement will require a larger breath pressure to drive the system, and will result in greater dynamic volume. As was described in the previous section, blowing harder will raise the pitch. However, this sharpening is offset by the lowering in pitch due to the increase in tip displacement, and consequently, the distance the reed has to cover to complete a cycle of oscillation if it is beating. (This empirically derived explanation fits nicely with Hirschberg's theory of the dependency of pitch on breath pressure mentioned above). Another interesting consequence of tip displacement, together with stiffness, is the resulting stability of the tone and sensitivity to crossfingering. Myers states, again in [49], that larger tip displacements result in greater pitch fluctuations and less stable tone. The result can actually be desirable, because a less stable tone will be more sensitive to crossfingerings, which are often required for producing semitone intervals, or *chromatics*. David Hogan Smith has noted the same effect for reed stiffness; an instrument with a stiff reed will be less sensitive to crossfingering [62]. His explanation is that an incorrectly adjusted reed will have a mismatched equivalent volume and reed resonance, resulting in a less stable tone. However, if this were the case, it seems there would be an optimum operating point somewhere in the middle of the range, and that crossfingering sensitivity increase on both

sides of this point, rather than towards the open, flexible side alone. Another potential cause is the lowered slit resistance of a more wide open reed; the importance of resistance and impedance in determining note stability are covered in the following section.

We now have the following qualitative model for a double reed. A valve, similar to the single reed valve, controls the aperture through which the musician provides a steady flow of air. The flow of air through the aperture is modulated by the pressure fluctuations in the bore to which the reed channel leads. The resulting pressure fluctuations within the reed channel, countering the mouth pressure of the musician, result in a fluctuating force on the reed, deforms the reed, modulating the aperture size. The reeds can again be modelled as harmonic oscillators, although this model is complicated by considerations of equivalent volume and frequency, especially when connected with a conical bore. Because the embouchure now entirely surrounds the compliant reed structure, embouchure effects are more pronounced. Finally, we can assume that beating does take place in normal level playing. This can be tested using some additional facts to be covered in the following section.

2.5 The Coupled Reed / Air Column

The previous sections examined the air column and the bore as essentially separate entities, with the relationship between the two mentioned primarily to provide context. This section examines in detail the nonlinear coupling between two, and what that coupling implies. First, the physical interaction between the two are recapitulated. The notion of impedance and phase is introduced in preparation for a more mathematical discussion in following sections. This is followed by a discussion of the consequences of nonlinear coupling. Finally, the section describes the “Regimes of Oscillation” present in a musical instrument.

2.5.1 Interaction of reed and air column

Previous sections have described the essentials of the interaction between the reed, which modulates the input flow, and the air column, which provides the necessary oscillation. To summarize:

- The air column contained within an arbitrarily shaped bore can transmit compression pulses as waves. The frequencies at which standing waves can exist are the normal modes, i.e., the resonances, of the air column;
- Cylindrical and conical bores will support normal modes that are integer multiples of one another. This is the relationship required for the production of musical tones;
- The reed reacts to fluctuations in pressure by deflecting away from the resultant force, thus changing the size of the aperture. The periodic action on the aperture size, due to the periodic pressure fluctuation in the bore, modulates the d.c. input pressure into a similar periodic function;
- The periodic input pulses excite the normal modes of the air column, causing a sustained oscillation which is perceived as musical sound. In turn, the standing wave thus excited provides the fluctuation necessary to drive the reed-induced input modulation.

There exists therefore a symbiotic relationship between reed and air column which results in a sustained oscillation, or “sound regeneration”. Unlike the plucked vibrating string, whose modes could be excited almost arbitrarily by plucking at different points, the coupled reed/air column has more rigid phase/mode relationships.

An important point not yet mentioned is that, even in a linear representation of the system, there exists a *threshold blowing pressure*, below which the instrument will not sound. In [5], Backus calculated this pressure analytically, and verified his results experimentally. This has important implications both in dynamic flexibility and tone coloration.

2.5.2 Energy, Impedance, and Peak Frequency Shifts

The discussions up to this point have been primarily qualitative. The purpose of this section is to set the stage for a more quantitative analysis of the physical system.

Previous sections have described the standing waves which represent the harmonic oscillations of the air column. A convenient representation of this oscillation is the *input impedance* of the bore. This impedance is, technically, the frequency-dependent ratio of pressure to flow, p/v , at the entrance to the bore for a given tonehole configuration. It consists of

both: 1) a reactive portion, the imaginary part, which represents the periodic exchange of potential and kinetic energy due primarily to the inertia of the air; and 2) a resistive portion, the real part, due mostly to boundary layer effects, which is purely energy dissipative. High impedances, or more specifically, impedances with high imaginary parts, imply high pressures and little flow, and thus, a buildup of energy. Anyone who has pumped up a bicycle tire can appreciate the energy buildup at the expense of the pumper as the pressure in the tire becomes higher, provided there are no leaks in the inner tube to provide an unwanted source of flow. Low impedances imply the converse, low pressures and high flow, and consequently, an energy drain [76]. It is very easy, at least from a muscular rather than an aerobic standpoint, to pump a tire with a sizable leak, but all the energy transferred from the pumper to the tire goes right out the hole. Because impedance can be measured, and because it is commonly used in “acoustic circuits”, i.e., electrical analogues of acoustical systems, and can thus be manipulated in the same way as any electrical impedance, it is frequently used to characterize acoustical relationships.

The impedance function mimics the natural modes of the system; the resonances are defined in this context by the frequency of the impedance peaks. In a reed instrument, the impedance relationship leads to the coupling that results in sound. The input flow is at a maximum when the aperture is at its widest. Since the aperture opening is a function of bore pressure, this corresponds to the point in time when the bore pressure is highest. At the frequencies corresponding to the impedance peaks, we have the best opportunity for energy buildup, the required condition for sustaining the oscillation. Of course, if energy is constantly flowing in, it must be also flowing out. The losses in the system are discussed in a later section; suffice it to say at this point that a stable oscillation results from an equilibrium between energy input by the player and energy loss from the instrument. (In a flute, whose generator is a flow-controlled valve, the opposite is true; harmonic generation takes place at the impedance *troughs* [9], or conversely, the *admittance* peaks.)

In the bore, the impedance is due to the combined effects of the bore itself, the toneholes and bell, and any roughness or discontinuities along the path of the air. The bore itself can be treated as a duct, which has a real *characteristic impedance* corresponding to its cross-sectional area of

$$z_c = \frac{\rho c}{A} \quad (2.2)$$

where ρ is the density of air and A is the cross-sectional area. A simple cylindrical section of

tube therefore has a real characteristic impedance associated with it, which implies that flow in and out are in phase. However, an aggregate of such tubes, with additional contributions of reflecting terminators such as the toneholes and the bell, provides the complex acoustic impedance discussed above.

The bore is not the only source of impedance; the reed provides its own resistance to airflow. Backus cites in [5] the Bernoulli relationship:

$$U = A\sqrt{2p/\rho} \quad (2.3)$$

where A is the area of the slit. This results in the following relationship for resistance R :

$$R = \frac{p}{U} = f(\sqrt{p}, 1/A) \quad (2.4)$$

In addition to its nonlinearity, being dependent on the square root of P , the resistance is also proportional to the reciprocal of the area, which in turn, is proportional to the tip displacement; the more open the reed, the less it resists the flow of air. We will see that this is true for toneholes as well; the larger the opening, the lower the resistance. Backus also provides an expression for reactance in a slit, but notes that its effect is negligible.

In [76], Worman describes the total system impedance as being the parallel combination of the reed impedance and the bore impedance. The justification for this is that the air that does not flow through the aperture will flow instead into the gap left outside the aperture by the deflected reed. This is a somewhat suspect explanation, in this author's opinion, although it has been quoted almost verbatim in several other papers as fact. A better justification can be had by looking at the reed/bore junction during the process of wave propagation. When a compression pulse is initially introduced by the player, it must pass first through the reed, then through the bore. This would suggest a series combination. However, at the junction between reed and bore, part of the pulse will reflect back out the mouthpiece. In addition, any pulses returning from the bell will, upon reaching this junction, partially transmit out of the mouthpiece, and partially reflect back toward the bell. A division of flow is therefore taking place, which implies a parallel arrangement.

The relationship between reed and bore impedance and their effect on coupling is elucidated by Keefe in [34] by examining the reciprocal of impedance, the *admittance*. Like the impedance, the admittance has its real and imaginary parts representing energy dissipation and storage. The bore, which cannot create energy, must have a positive resistance, and a

correspondingly positive real part. In order to inject energy, the reed must therefore have a negative resistance so that the real parts counter one another; this is consistent with its role as a generator. Since a sustained oscillation requires that more energy enters than exits at the reed/bore junction, the reed admittance must exceed the bore admittance for sound regeneration at that frequency to occur. The admittance troughs are aligned with the bore resonances; this implies that the condition will be satisfied most readily at or near the bore resonances. An additional condition for oscillation at the fundamental has to do with the phase relationships. It turns out that, because of a change of sign in phase, oscillation is possible only on the lower frequency side of the impedance peaks.

Another way of looking at the reed in its role as a generator is through the energy relationship implied by the flow control curve in Fig 2.35 of the previous section, as McIntyre does in [43]. The basic requirement for the maintenance of oscillation is that the energy supplied by the generator compensate for the energy losses throughout the system. The energy supplied is equal to the product pu , where p is the pressure drop and u is the flow. The average energy must be positive. Taking the time average over one period, the energy average, \bar{E} , is:

$$\bar{E} = \bar{p}(t)\bar{u}(t) \quad (2.5)$$

According to McIntyre, this is satisfied when the operating point is such that the pressure and flow are *positively* correlated, i.e., an increase in pressure causes an increase in flow. The pressure drop is $p_{\Delta} = p_{mouth} - p_{bore}$; the mouth pressure is constant while the bore pressure fluctuates. An increase in p_{bore} will correspond to an increase in flow to the *right* of the crest. The operating point must therefore be somewhere along the *negatively* sloped portion of the flow-control curve. Stability is guaranteed because large bore pressures will cause an excursion into either the positively sloping side, implying negative correlation and energy dissipation, or reed closure, which will temporarily prevent further introduction of energy.

One result of the coupling between reed and air column is that the resonances of the system are lower than the resonances of the air column alone. This flattening was observed by several early acousticians, and properly attributed to several causes. According to Benade[9], Weber showed that a yielding air-column termination lowers the natural frequency and Helmholtz added the effects of the flow control action in the explanation. This result is not surprising if one considers that the dynamic nature of the reed implies an additional

phase lag at that end. A phase lag implies a time delay, and this would lower the natural frequency. In addition, any damping in the reed would damp the resonance as well. The phase relationship described in the preceding paragraph also defines a trend toward flattening; since oscillation can only occur below the resonance peak, and not above it, the pitch will tend downward. This is an issue for reed simulation if tuning is to be an issue; if the delay line length is based solely on the desired pitch frequency, without taking into consideration reed-induced delays, the “instrument” will play flat.

2.5.3 Nonlinear Effects

The nonlinear coupling described above does more than simply sustain an oscillation; it implies several important effects on tone quality.

One important effect can be heard by playing a crescendo on a clarinet, starting at a level when the reed is not beating. As the volume grows, so does the influence of the higher harmonics. This “blossoming of the spectrum”, as Keefe phrases it in [34], is a result of the dependencies among harmonics. Benade [9] cites the non-beating harmonic relationship:

$$p_n \propto p_1^n \tag{2.6}$$

where p_n is the amplitude of the n 'th harmonic, and the constant of proportionality is the height of the n 'th impedance peak.

Keefe extends this to any wind instrument oscillation whose flow control function is weakly nonlinear (e.g., a non-beating reed). The constant of proportionality is the magnitude of the impedance at the harmonic. Equation 2.6 implies that each harmonic helps to drive the succeeding harmonic. Worman proves this in [76]. The “blossoming” is not merely a function of the exponential terms. The threshold pressure described above is actually unique to each harmonic; a harmonic pressure which reaches this threshold becomes capable of regenerating itself, and stabilizing the overall oscillation. The cooperation among modes, discussed at length by Benade in [9], is the topic of the next section.

When the reed begins to beat, the flow control function becomes strongly nonlinear, and the harmonic relationships change. Now all the harmonics grow at an equal rate [9]; a change in dynamic level will not result in a change of tone, at least up to a point. This has

important bearing on reed selection. A stiff reed, which will not beat readily, will provide the dynamically influenced tonal changes. A softer reed which beats easily will have a much more constant tone and feel.

Heterodyning is another nonlinear effect, the result of the dynamic system being driven by a periodic source. In a musical sense, heterodyning is the process of combining tones into new tones whose frequencies are the sums or differences of those tones already in existence. For example, following Benade in [9], the combinations of tones P and Q will produce the heterodyne components $2P$, $(P+Q)$, $(P-Q)$, and $(2Q)$. These components will then combine with each other, and with the original components, to form new terms such as $3P$ and $(2Q+P)$. Difference tones can be very audible, particularly in high register instruments such as recorder where the waveforms are not particularly complex, and can in fact be used for tuning. Although heterodyning is similar to partial formation, it bears a significant distinction: heterodyned components do not have to be harmonically related, whereas harmonics do. Thus, the production of heterodynes can result in a *multiphonic* sound, the which can be either jarring or effective, depending on the context. In general, the results of heterodyning are most pronounced at high volume.

For a reed-driven system, heterodyning has the effect of producing oscillations that normally would not be supported by the air column. If the reed oscillation frequency (not its resonance frequency) is P , then it would drive the system with components of $2P$ and $3P$ as well. On a conical bore instrument, as we have seen, these components are part of the normal spectrum anyway. However, on a cylindrical bore instrument such as the clarinet, the $2P$ component would not be supported by the air column; recall that a cylindrical bore supports only odd harmonics. Because the $2P$ component is part of the driving mechanism, it will be present in the resultant oscillation, at some dissipative expense to the overall energy balance. The result is that the clarinet spectrum will have weak even harmonics, although at too low a level to support overblowing at the octave. There was, incidentally, a great hue and cry in 1877 when these even harmonics were discovered, according to Worman in [76]. Helmholtz had stated that clarinets would favor odd harmonics, for the reasons described in Section 2.3, and that he had been unable to find even harmonics experimentally. This was, apparently, interpreted by some as implying that even harmonics could not exist in the clarinet. When even harmonics were found by Blaikley, Helmholtz' reputation came under severe fire. Worman states other examples of the unfair treatment of Helmholtz'

work, results of misinterpretations of his findings, and concluded: “The shallowness of many of Helmholtz’ critics is appalling”. For better or for worse, the reexamination of Helmholtz continues: Hirschberg, a current leader in reed hydrodynamic research, started a recent presentation with the chart :“What does Helmholtz Neglect?” [28].

In addition to even harmonics, heterodyning can also produce components beyond the cutoff frequency of the instrument. The cutoff frequency will be discussed in more detail, but basically, it is the frequency at which the air column ceases to be capable of sound regeneration. The cutoff frequency has a significant influence on tone. Just as heterodyning can drive even harmonics in a column that cannot regenerate them, it can drive similarly unsupported high harmonics. This, according to Benade, is the source of the “buzzing” that one frequently hears when instruments are played at a loud dynamic level. Recall the buzzing of the capped reed instruments discussed in Chapter 2, which had a set of fairly ragged looking high harmonics far into the frequency range. For both cases, the unregenerated frequencies represent a drag on the system.

Of course, these are not the only nonlinear effects present in the reed instruments. The effects of hydrodynamics, embouchure influences, and other variation are critical in the resulting sound. A key area in research is trying to characterize the many nonlinear relationships, and add more reality into the existing models. Many, including Lindeman in [39], are investigating chaotic analysis as a potential tool. Results so far do not match well with musical experience; however, chaos is seen as having great potential, particularly for explaining nonperiodic instrument behavior such as multiphonics [34].

2.5.4 Harmonics and Regimes of Oscillation

So far, the coupling between air column and bore that drives an oscillation have been described, as well as the existence and influence of harmonics. In addition, the relationship in power between higher harmonics and the fundamental has been treated, and a symbiotic relationship among all harmonics suggested. This section examines this relationship of “modal cooperation” in more detail.

The dynamic response of a linear system is often described in terms of a “normal mode” analysis, in which the various vibrational modes, the number of which is determined by the order of the system, are decoupled. This technique is used on systems from circuits to

spacecraft. Unlike the normal modes in such systems, which operate more or less independently of one another, the modes in a musical system are codependent; being harmonically related, they can excite and reinforce one another, and cannot be fully treated as individuals. To deal with this state of affairs, Benade defined “regimes of oscillation”, combinations of coexisting vibrations which exist in a sustained oscillation. The basic features of this system have already been suggested in earlier chapters:

- From Bouasse, quoted in [9]: “If the reed-valve is nonlinear, then oscillation is favored if the air column has one or more natural frequencies that correspond to one or more of the higher partials of the tone being produced.”
- From Benade, in [9]: “A *regime of oscillation* is that state of collective motion of an air column in which a nonlinear excitation mechanism (the reed) collaborates with a *set* of air-column modes to maintain a steady oscillation containing several harmonically related frequency components, each with its own definite amplitude”. Thompson added to the set of air column modes the modes of the excitation mechanism in this definition, in accordance with his proposition that the reed resonance also plays an important role in upper register note stabilization.

One explanation of this cooperation follows from the requirements for a self-sustained oscillation. To reiterate, when conditions are such (threshold pressure reached, reed admittance higher than bore admittance, etc.) that energy can be injected into the system, sound regeneration will occur. The harmonics are present whether or not they can regenerate on their own. If they cannot regenerate, then they represent a dissipative load on the system. Once regeneration occurs for a given mode, additional energy can be brought in. Thus, as each harmonic comes into its own, it adds to the communal pool of energy that supports all the harmonics. The more harmonics which are suitably excited, the more energy in the system, and the more stable the note [76]. Note that unsupported heterodyne components, such as the even clarinet harmonics discussed above, always represent energy dissipation.

This coupled set of oscillations explains not only the stability of notes, but also some of the curious pitch behavior that occurs during changes in loudness (one is tempted to use the more musically appropriate term, *dynamics*, or at least *volume*, instead of *loudness*. Unfortunately, both of these terms have other specific engineering connotations which are used elsewhere in this thesis), particularly on nonoptimal instruments. Following Worman’s

example in [76], the two-peak impedance curve, illustrated in Figure 2.39, can be used to demonstrate the interactions. For low pressures, the dominant oscillation is the fundamental. As the previous section described, as the pressure increases, the energy input into the harmonics will increase. Therefore, a crescendo will soon bring the second harmonic, in conjunction with the second impedance peak into play. If the impedance peaks are perfectly aligned, the result will be an even more stable, richer tone. However, if the impedance peaks, and therefore the resonances, are not perfectly aligned, a pitch change will result as well. This is because the instrument, being a physical system, is essentially self-optimizing, and will try to maximize the energy input into the system.

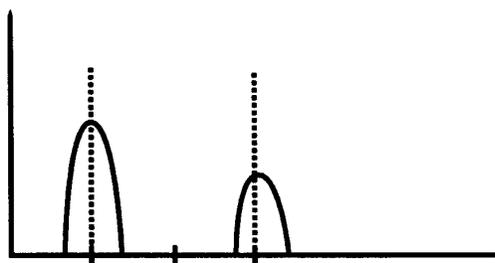


Figure 2.39: Misaligned resonance peaks, after Benade

Imagine now that the impedance peak is a bit wide, i.e., the second resonance is sharp. The original second harmonic will fall a bit before the peak, and its sound regeneration will be less efficient. However, if the fundamental rises just a bit in pitch, the second harmonic will increase by twice that much and fall closer to the second impedance peak. The most stable note will therefore be, approximately, the maximum energy solution for the two harmonics. As the crescendo advances, and the strength of the second harmonic increases, the pitch will climb higher. The player can, and generally will, combat this trend with changes in embouchure, but at the expense of additional energy dissipation and a more strangled tone.

Note that the “blossoming of the spectrum” which produces the above effect is present only in the absence of beating; during beating, the harmonics grow in parallel. In the Section 2.4 on reeds, Hirschberg provided a different explanation for the dependence of pitch on the playing volume dependent tip offset, which required that the reed do beat.

In the preceding example, not only will the pitch climb, but the note, even at its stable

point, will not be bringing in as much energy as with the well-aligned instrument, and will be more difficult to play. Instruments with badly aligned resonances will therefore require much more energy. It might be noted that in flute-like instruments, which use an airstream generator, Coltman [16] found that the optimal modal cooperation, evidenced by greatest sound power radiation and tonal quality, occurred when the octaves were stretched by about 25 cents, rather than when they were perfect integer multiples. Benade explained this in [9] for the case of louder playing by noting that, whereas the fundamental mode is sensitive to breath pressure changes, the second mode is not. In order for the modes to be properly aligned for mezzoforte playing, they must be stretched in the steady state. This stretch will not be obvious at the pianissimo levels for which it exists because at those levels, the fundamental mode is dominant anyway.

The preceding paragraphs described the importance of aligning resonances in the production of a stable tone. It must be noted at this point that perfect alignment is not necessarily the best musical arrangement. An example of too strong a cooperation is given by Myers in [49]. He describes the “honking” of early cylindrical recorders because the first and second modes were aligned *too* accurately. Design improvements by c1400 involved *degrading* the modal cooperation, by constricting the bore and consequently spreading the modes (This may seem to be a contradiction in interpretation to the paragraph above on Coltman’s work. However, the case is really opposite to that cited above; the modes on the untapered early recorder were actually found by a present-day builder to be too well-aligned under playing conditions. [50].) Most recorders used today, whether of modern design or modelled after a prototype, are conically tapered toward the bell end.

2.5.5 Playing in the Upper Registers

The previous section illustrated how the different modes of oscillation cooperate with one another, in accordance with their relative impedances, to stabilize a tone. This section describes how that relationship can be exploited to produce upper register tones.

Recall that an upper register tone is achieved by exciting the second harmonic of the fundamental, without exciting the fundamental itself. From the discussion above, this would seem like an impossible task; the cooperation among modes is so strong that excitation of

the second harmonic implies excitation of the fundamental. However, the musician has several tricks in his bag to accomplish the job. On some instruments, increased volume alone will cause a register change. Certain notes on a recorder, for example, can be played an octave higher simply by blowing harder. According to Backus, this is because the period of oscillation is indirectly related to the velocity of the air, which increases with blowing pressure. When the blowing pressure becomes high enough, the corresponding period becomes too small to be supported by the fundamental mode, and the second harmonic, which has a shorter period, becomes the new fundamental of the tone [2].

With reed instruments, the musician must somehow change the resonance relationships. On an open reed instrument such as a shawm, the player changes registers by tightening his embouchure and blowing harder. As we have seen, increasing pressure increases the strength of the higher harmonics. However, assuming that the shawm, a double reed instrument, exhibits beating over most of its dynamic range, the second harmonic does not grow with respect to the first. The embouchure-induced change in register is not well understood, partially because the technique is not used much in modern-day instruments, and therefore, has not yet been thoroughly investigated. One possible explanation is that when the reed tips are beating so closely together, they simply cannot support the slower mode [49]. This is similar to what happens in the flow-controlled flute. Another theory has to do with the roles of equivalent volume and reed resonance. By tightening the embouchure, the player alters the equivalent volume of the reed, possibly to the point where the reed can no longer support the fundamental tone. In the upper register, however, the reed resonance is more important than the equivalent volume, and can help to stabilize the higher octave note. In this manner, the player selectively prevents one mode from being excited, while allowing the desired mode to ring [62].

On most modern reed instruments, register shifts are accomplished by special toneholes placed to alter the resonances of the instrument. The purpose of a *register hole*, or *vent hole*, is to allow the second resonance to sound while somehow repressing the first resonance. This is the same concept as the embouchure tightening described above for the shawm. Section 2.3 laid the groundwork for describing how the selective suppression can be accomplished. Recall that in a plucked string, the relationships among the harmonics depend upon the pluck point. Modes with nodes at the pluck point will not be excited at all, while modes with loops will be excited most strongly. An extension of this is the bowed string technique

of producing harmonics. The player can gently put his finger to the location of the node of a desired harmonic. Rather than suppressing the mode, as in the case of the pluck, he now damps out all other modes, and drives only those oscillations unaffected by the damping (the process of bowing a violin involves a coupling similar to the reed instrument, incidentally, and was studied by Julius Smith in detail in [66] and [63]). A similar action is performed by opening a register hole. The added tonehole, which will not in itself completely damp out affected vibrational modes, is placed ideally at the node of the second harmonic. Thus, the second harmonic will be unaffected by the change. Figure 2.40 illustrates the two harmonics of a clarinet, which are, the fundamental and its twelfth, since the octave is not supported. The second harmonic has a node one third of the way down the bore; this is the optimal place for a register hole. Note that the optimal register hole location is a function of borelength, and therefore, of fundamental pitch. A clarinet cannot have a register hole for each note; in fact, it has only one. Benade discusses the consequences of this insufficiency [9].

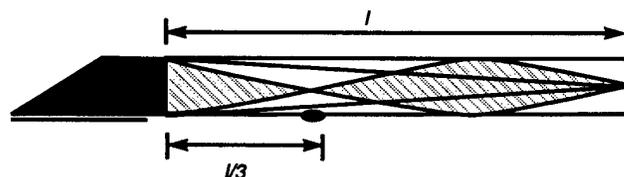


Figure 2.40: Nominal location of a register hole on a clarinet

The effects of the register hole are more subtle than those of the violinist's mode damping. Its efficacy is very much a function of the required cooperation among modes, as described by Benade [9]. Recall that, for a note to ring, the resonances must be aligned. This sets up the regime of oscillation which feeds energy into the system. The purpose of the register hole is to change the resonance pattern so that the first mode cannot be supported by any of the other modes, and therefore not drive the frequency of the oscillation. If this can be achieved, the next available mode will become the fundamental, and the sounding pitch will rise accordingly. The hole accomplishes this both through its resistance, which reduces the level of the peak, and through its reactance, which relocates the peak. The optimal energy solution will then be based on the regeneration of the second mode, and the now inharmonic first mode will be silent.

Future chapters will describe a model of the register hole which was incorporated in the simulation. There were many interesting results which illustrate the interaction of the modes, even for tests where the register key was unsuccessful.

Figures 2.41 and 2.42 show the spectra for a shawm and a clarinet blown in the lower and upper registers. As described, the clarinet register shift is achieved via the register key, while the shawm shift is effected via embouchure changes. Note that the high note will be necessarily be much less rich in harmonics than the low note, although the higher frequency harmonics it has will be more pronounced. This is the consequence of the cutoff frequency, which puts an absolute limit on the modes supported. The fundamental frequency of the clarion register is 2.5 times that of the chalumeau register, and the new first harmonic 2.5 times that of the new fundamental. Clearly, fewer harmonics can be supported; the spectrum of the clarinet shows the number of harmonics above -30 dB decreasing from 17 to 5. The shawm fares similarly, in that it has an apparent cutoff frequency close to that of the clarinet. Because the shawm is conical, it does have a few more harmonics on the low-frequency end, in that the even harmonics are more pronounced. It is notable that on the shawm, the fundamental peak is relatively small, while in the upper register, it is quite significant; there is none of the upward sloping that the lower register displayed. The upshot is that upper register tones are somewhat less interesting than the lower register tones in terms of harmonic content. The decreasing sensitivity of the ear in the upper range contributes to this effect. In recorders, the easiest distinctions among models can be made in the low registers where the harmonic structure is most flexible. In the upper ranges, it is much more difficult to tell the instruments apart spectrally. However, better instruments have much nicer transient properties in the upper range, even if their steady-state properties are undistinguished.

The reduction in harmonics implies not only a change in timbre, but a reduction in the amount of modal cooperation, and the consequent stability of the note. There are simply very few modes present to stabilize one another. This is where the role of reed resonance is believed to become prominent. Stephen Thompson studied the influence of reed resonance on modal cooperation in detail in [72]. The resonance of the reed, even while exceeding the cutoff frequency, can still interact with supported modes and stabilize them. (D. H. Smith found that, for double reeds, a reed resonance just above the instrument cutoff worked best [60]). For this reason, a musician will continually adjust his embouchure as he ascends the

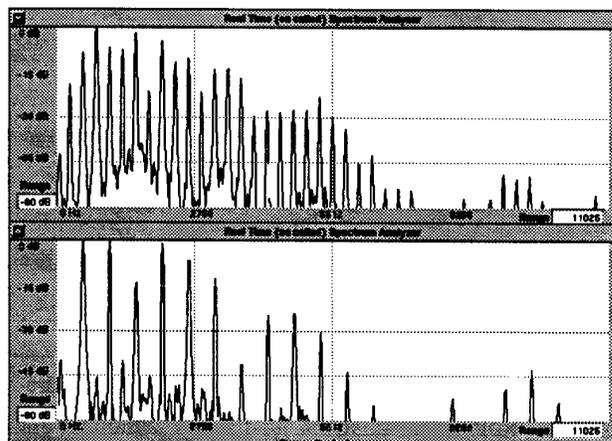


Figure 2.41: Spectra for an alto shawm playing in two registers (Sound Source: McGill Collection)

scale in the second register; he actually is adjusting the reed resonance to match a harmonic of the note played, and thus enjoys an extra measure of stability. This makes sense for a single reed instrument, and gives justification for the inclusion of a dynamic reed model for any simulation expected to operate in the upper register. The double reed case is a bit more complex, because as the musician adjusts his embouchure, he adjusts not only the damping/resonance of the reed, but the equivalent volume as well. However, if the equivalent volume is truly important only in the fundamental register, than the resonance matching argument still holds in high register [62].

2.6 Other Acoustical Features

The preceding sections have developed the fundamental relationships that allow a reed instrument to work. There are many other aspects to the instrument that affect its sound and stability, such as the toneholes, the bell, and even the vocal tract of the player. This section describes the effects of these features.

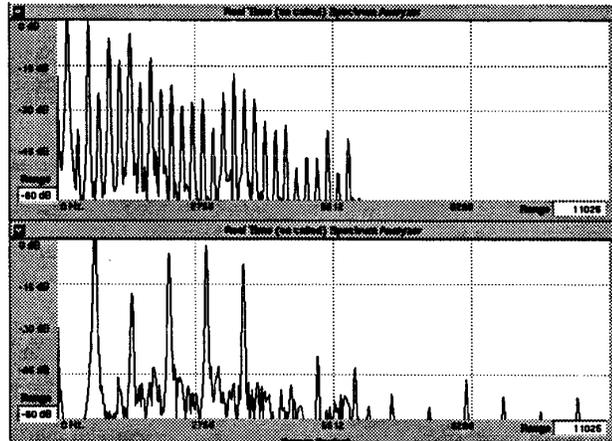


Figure 2.42: Spectra for a Bb clarinet playing in two registers (Sound Source: McGill Collection)

2.6.1 Sound Radiation and Reflection: Bell and Toneholes

Until now, the reflection of pressure waves at the open end of the instrument has been idealized. In Section 2.3, the wave reflection discussion implied a total reflection of wave energy back toward the reed. If this were the case, the instrument would be extremely stable (assuming boundary layer losses would offset the energy input) but produce absolutely no sound at all; sound production requires wave transmission into the surrounding air. The dimensions of the bell and toneholes play a major role in determining how much energy is reflected — allowing coupling and oscillation regeneration — and how much is transmitted as sound.

The open end of the tube was defined in Section 2.3 as essentially an *acoustic barrier*, with perfect nodes and antinodes. It is possible to define a less rigid barrier by extending the concept of impedance discussed earlier. Recall that impedances were used to characterize the acoustic tube response. They can also characterize the partial reflection/transmission of waves that occur at various locations in instruments. The presence of toneholes, the bell, or changes in the bore diameter all result in changes in column impedance. An impedance change acts like a flexible acoustic barrier, reflecting some energy and transmitting the rest.

The relationship between the two can be represented by *reflection coefficient*, a measure of the fraction reflected. The concept of impedances and reflection coefficients is a basic component of the waveguide models in the following chapter, and will be described more mathematically there.

Another way to view the toneholes and bell are as radiating pistons of air. This is how Stewart and Strong modelled their bell in [70]. Keefe has done extensive work in this type of modeling with toneholes [35],[37]. The piston action in the toneholes depends upon the thickness of the walls of the instrument.

The purpose of a tonehole is to shorten the acoustical length of the tube by inserting an acoustic barrier inside the bore against which waves can reflect. The hole allows the air exchange required for wave inversion, as well as providing an outlet for the transmission of sound power. Once again, an “end correction” is in effect, and the location of the tone hole will not absolutely indicate the new, frequency dependent acoustic length of the tube. This length, manifested in the pitch of the fundamental, is a function of hole diameters and wall thickness. Benade discusses the subtleties of tonehole placement and spacing in [9]. Myers expands on this in [49]. Tuning a woodwind instrument is accomplished only coarsely by sizing the walls and drilling the holes. Fine-tuning is a process of iteratively widening the holes (or filling them in with paraffin when the tuner overshoots the mark) until all the notes are in proper pitch.

The action of the tonehole can be explained in terms of the impedance change it induces. The amount of reflection at an impedance junction depends upon the ratio of impedances. We can define a term r which represents the impedance ratio Z_2/Z_1 , where Z_1 and Z_2 are the impedances before and after the junction. As illustrated in Figure 2.43, a wave approaching this junction is *scattered* according to r ; Part of it is reflected and the rest is transmitted. The following chapter describes the exact relationship between r and the reflection which occurs. A value of r near unity implies little reflection; A value far from unity implies strong reflection.

The dimensions of the tonehole determine its impedance. A very large tonehole has very low impedance; air will flow in and out of it easily. The consequent impedance *drop* is very large, and strong reflection will occur (Herbert Myers noted in [49] the practical relationship between hole diameter and note stability). A large tonehole approaches the ideal end of

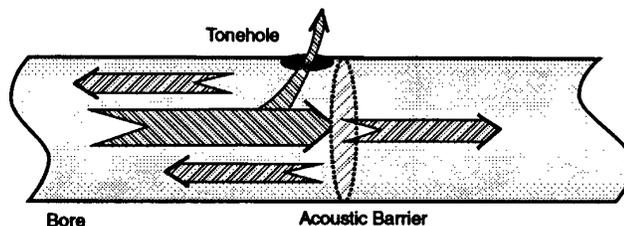


Figure 2.43: Wave scattering at a tonehole

the tube, and the effective end of the tube will be very near the location of the hole. A smaller tonehole will have a higher impedance, resulting in a smaller impedance drop and weaker reflection. The end correction for this hole will be larger, that is, the effective bore length is much longer than the length of the bore terminated by the hole. Thus, a tonehole that is too small will cause a flattening of the fundamental it was designed to elicit. The register hole described earlier is an example of a high impedance tonehole. The passage of air in and out of it dissipates energy and causes the reduction in fundamental peak size. In addition, the reduction in effective length which it causes is just enough to move the fundamental peak to an inharmonic location.

Figure 2.44 illustrates the circuit equivalent of the tonehole junction. Looking again at the scattering that takes place at the tonehole, the wave actually splits into three parts. One part reflects back toward the reed, and is the part that sustains the oscillation. The transmitted portion will travel either through the tonehole or further on down the bore. The amount reflected depends actually on the *parallel* combination of the bore impedance (assuming a cylindrical instrument where the bore impedance is the same throughout most of the instrument) and the hole impedance. The expression for impedances in parallel is:

$$\frac{1}{Z_{bt}} = \frac{1}{Z_b} + \frac{1}{Z_t} \quad (2.7)$$

When the tonehole impedance, Z_t is large, its contribution to the parallel impedance, Z_{bt} , is small; since the bore impedance, Z_b , is constant through the junction, impedance change will be small and the reflection correspondingly weak. When Z_t is low, it dominates Z_{bt} ; the impedance drop across the junction is now large, and the reflection correspondingly strong. The transmitted wave will split between the parallel paths, again according to Z_b and Z_t . Part will transmit as sound from the hole; the rest will propagate further down the bore.

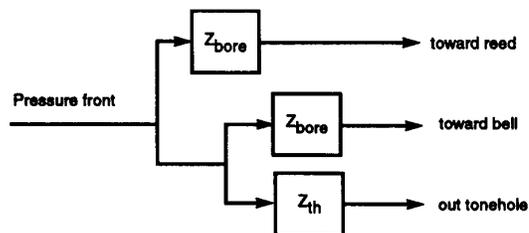


Figure 2.44: Wave scattering at a tonehole

Thus, not only does the tonehole shorten the instrument, it provides a highpass transmission port for the sound. A general rule of thumb is that a tonehole can transmit any wave with half wavelength less than the hole diameter. A very large tonehole can radiate a significant portion of the sound, while a smaller tonehole will be much less efficient. Whatever sound energy does not transmit will propagate down the bore to next tonehole, which will transmit another fraction, and so on to the bell, where any energy still remaining will again either radiate or reflect[24]. The higher the frequency of the harmonic, the further down the bore will its standing wave penetrate [9]. Notes sounded near the top of the tube, where the unreflected portion must run a considerable gauntlet of toneholes before reaching the bell, will be transmitted almost entirely through these holes. Lower notes, which use fewer open holes (although some instruments contain vent holes which are never covered in order to provide a lattice for the low notes), will rely more on the properties of the bell. The bell is, essentially, the final tonehole, and, as implied above, and, according to some acoustics texts [24], affects primarily the lower notes (in brass instruments, on the other hand, the bell is the only outlet for radiation, and is therefore always active).

There is somewhat more flexibility in designing the bell, which has flare as a parameter in addition to thickness and diameter. The bell is often described as being “impedance matched”, meaning that the impedance at the end of the bell is quite close to that of the ambient. A truly impedance matched bell would be a perfect transmitter. However, an instrument with such a bell would lose its lowest note, because there would be absolutely no internal reflection. Given that the bell is present partially to strengthen the lower notes, this would be unfortunate. Instead, the bell is designed to be a compromise between the ideal end of the tube, where all energy would be reflected, and the ideal impedance matched

transmitter, where all energy would be transmitted. A good bell mimics the radiation properties of the tonehole lattice, so that there is continuity in tone from lower to higher notes.

The influence of the bell can actually be quite profound. According to D. H. Smith, an important evolutionary step in the development of the shawm was the lengthening of the bell section, which stabilized many notes considerably. Apparently, even on the shawm, which has relatively large toneholes, enough of the wave energy makes it to the bell that, even for higher notes, the reflection characteristics at the bell are critical [62].

Not all of the wave energy becomes transmitted sound; losses take up a considerable amount. The tonehole lattice in particular provides many opportunities for energy loss, because the toneholes have relatively sharp edges, which can result in loss-causing turbulence. Even when closed, they form ducts which impede the air flow in the tube [9]. The additional damping helps create the characteristic woodwind sound. Keefe has done quite a bit of work in characterizing the tonehole and tonehole lattice [35], [33], [32]. The effects of the tonehole lattice will not be discussed here, as it was not incorporated into the clarinet model. Again, Myers has given a thorough treatment of the practical issues of tonehole alignment in [49].

A major consequence of the combination of tonehole lattice and bell which can be easily modelled is the *cutoff frequency* of the instrument [9]. The cutoff frequency, implied by the partial reflection/transmission described above, is the limit at which sustained oscillations can no longer be supported. Each tonehole, including the bell, has its unique cutoff frequency; one objective in tonehole design would ostensibly be to equalize the cutoff [9]. However, Myers points out that, at least in earlier woodwinds, which had none of the mechanical keywork advantages of modern woodwinds, cutoff frequency equalization of toneholes was the least of the builder's worries [49]. The flare of the bell is designed specifically for this type of balancing; "mimicking the tonehole lattice" really implies duplicating the tonehole cutoff frequency.

As an example, the cutoff frequency of a clarinet is about 1500 Hz. All of the strong harmonics will have frequencies below that point. Notice in the clarinet spectrum presented in Figure 2.7 that this cutoff frequency corresponds roughly to -15dB of attenuation in the steady state sound. Thus, the bell/tonehole lattice can be represented as a pair of lumped, complementary lowpass reflection and highpass transmission filters. Because of

the importance of the internal coupling process, a rather nonintuitive situation results. The clarinet has no problem transmitting the higher harmonics; without the coupling and modal cooperation, however, they will be so weak that even fully transmitted, they will be almost imperceptible. This is one of the paradoxes of the acoustic sound transmission process; since efficient transmission implies inefficient reflection, and reflection is required for coupling and reinforcement, a very efficient transmitter will produce a very weak sound [9]. Benade studied the importance of the cutoff frequency by altering that value by $\pm 2\%$ on two identical clarinets. The instrument with the lower value had a dark sound well suited to his classical friends but useless to his jazz friends. That with the higher value had a bright sound that his classical friends spurned but his jazz friends borrowed frequently. Note that the instrument with the bright sound had a high cutoff frequency, meaning that it was an inefficient transmitter for relatively high frequency harmonics. The other instrument could transmit a larger bandwidth, but without internal support, these harmonics were relatively weak.

The term *radiation efficiency* is often used to describe the sounding power of an instrument [24]. The radiation efficiency describes the *end-to-end* sound transmission capability instrument, i.e., the power radiated out versus the player's power input. This term encompasses both the internal mode sustaining mechanisms and the individual transmission and loss terms within the instrument. A high radiation efficiency indicates an optimal compromise between the reflection and transmission needs described above.

A final example of the importance of the bell in a woodwind instrument and its effects on the output spectrum is the instrument illustrated in Figure 2.45. This is a modern reconstruction of a cylindrical capped reed of the cornemeuse family, named the *Glastonbury Pipe* because it was designed after a medieval tapestry found in Glastonbury. The pipe has the convenience of a removable bell; removing the bell effectively muffles the instrument without altering the pitch. With the bell attached, the instrument has a fairly resonant, loud tone suitable for playing with other instruments. With the bell off, the tone is quieter and less forward, more suitable for blending with the voice or softer instruments. The toneholes on this instrument are quite small, so that one can expect much energy to pass through to the bell. In addition, the throat of the bell itself is perforated with tiny holes, which maintain the effective borelength at its bellless dimension, and prevents the normal lowering of pitch that would occur when an extra section is added on.

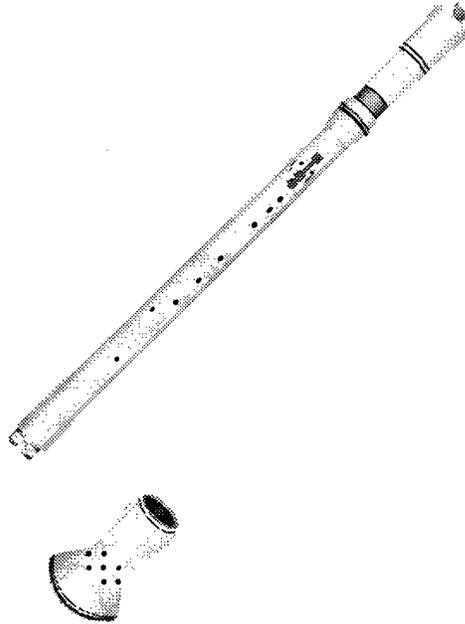


Figure 2.45: Glastonbury pipe, with removable bell. Design Early Music Shoppe of London, built by Tom Neumann, 1981. Instrument courtesy Tom Neuman, illustration by author

Figures 2.46 and 2.47 show the waveforms and output spectra for the instrument, without and with bell, playing a low F, the lowest note on the instrument, and a high F, which requires the closing of only two toneholes. As could be predicted from the associated tonehole configurations, the waveforms differ more for the low than the high note. The waveform without the bell has only one well-defined peak, whereas the bell waveform has about 7 peaks per period. The high F waveforms are a bit closer, in that there are more defined peaks in the bellless version. However, this version still has a single pronounced peak, while its counterpart has many peaks of nearly the same height. These waveforms hint at a reduction in the size of the fundamental, which is verified from the output spectra. In both cases, the bell boosts the higher harmonics in the first quarter with respect to the fundamental, so that there is a sloping upward in the spectral shape, as we saw in Chapter 2.2 with the conical reeds. Also, the bell spectra both show a slight formant structure in

the second quarter, where the bellless spectra are relatively flat. Finally, the very high harmonics in the final quarter are smaller for the bell spectra, which follows from the reduction in cutoff frequency which a bell would induce. The slight reduction in buzzing due to this greater attenuation may also be responsible for a smoother, more resonant tone.

The conclusion one can draw from this example is that the bell affects more than the cutoff frequency; it influences the lower harmonics as well, in ways that are quite significant musically. This is not really news. Although Benade emphasized the effects of the cutoff frequency, he also treated the influence on independent, low order modes of perturbations in the bore, of which the bell is one [9]. Another conclusion is that the musical quality of "resonance" appears to be a function of the mid-range harmonics. Two phenomena distinguished the resonant tones from the muffled tones — the boosting of the harmonics in the latter first quarter, and the slight formant formed in the second quarter. These effects were consistent for all notes recorded from the instrument. A good model of the bell would reproduce these affects, as well as the more easily obtained cutoff frequency.

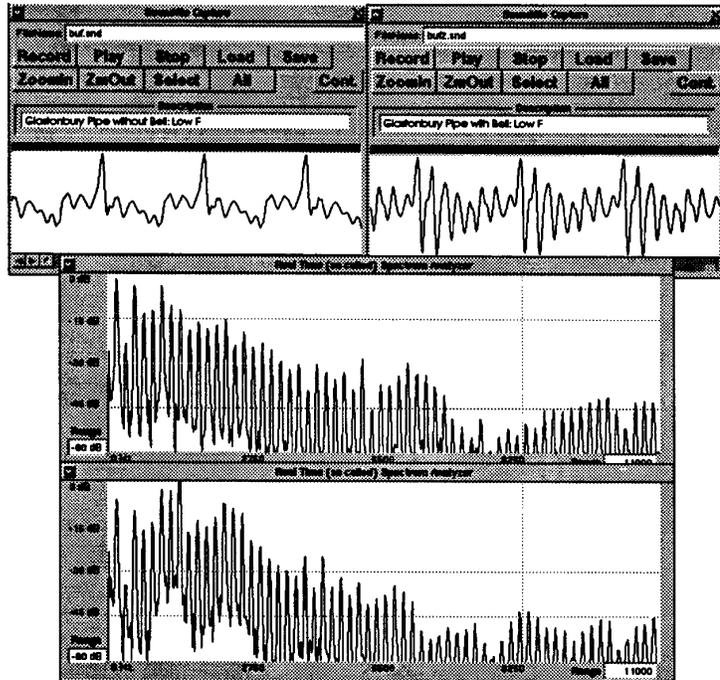


Figure 2.46: Glastonbury pipe, with removable bell, playing Low F, the bottom note on its scale. Left waveform and top spectrum: without bell. Right waveform and bottom spectrum: with bell. Instrument courtesy Tom Neuman

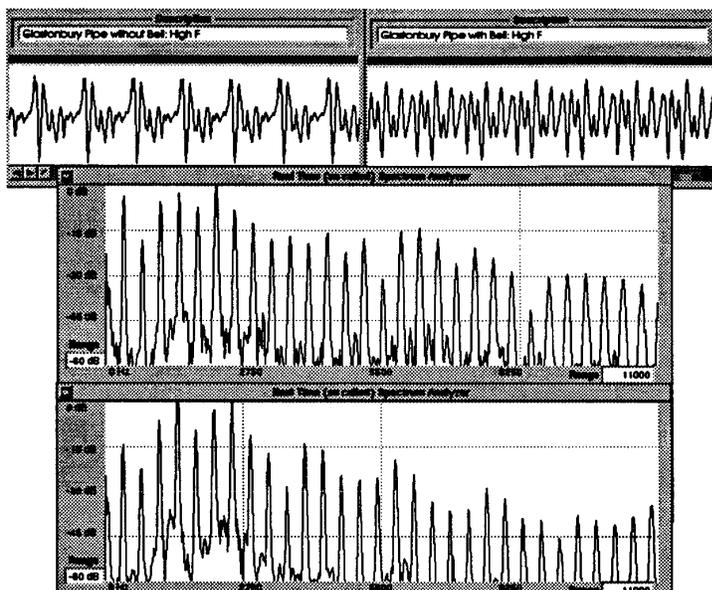


Figure 2.47: Glastonbury pipe, with removable bell, playing High F. Left waveform and top spectrum: without bell. Right waveform and bottom spectrum: with bell. Instrument courtesy Tom Neuman

2.6.2 Sources of Loss

If the radiation efficiency of an instrument were perfect, all the blowing power input by the player would be converted into sound power. In fact, the sound energy radiated is minuscule, according to Nederveen in [51], compared to the losses internal to the instrument. The efforts of the player, and the efforts of the instrument builder to align air column resonances, go mainly into overcoming the energy dissipation. Any attempt at understanding the energy equilibrium implicit in a stable note must take into account the deficit side of the equation. In addition, any loss or internal impedance will lower the fundamental frequency of the note played — as well as change the harmonic relationships among resonances — and therefore has important tuning implications. This subsection discusses some of the sources of loss

within the reed woodwind.

The major sources of loss are the boundary layer effects of friction and heat exchange at the wall of the instrument. The friction is due to the viscosity of the air, and will bring particle velocity near zero at the walls. Thermal losses occur because the pressure fluctuations in the air column induce temperature fluctuations. The walls of the instrument are of roughly constant temperature, and continuously exchange heat with the enclosed air. The thickness of the viscous boundary layer is a function of fluid viscosity, fluid density, and the frequency of the oscillation, in the following relationship from [51]:

$$t = \sqrt{\frac{\eta}{\omega \rho}} \quad (2.8)$$

where η =viscosity = $18 \times 10^{-6} \text{ N} \cdot \text{s}/\text{m}^2$, ρ =density= $1.2 \text{ kg}/\text{m}^3$, and ω = frequency in Hz. This leads to a boundary layer of only 0.05 mm for a frequency of 1000 Hz. Most of the losses in the system are assumed to be in this very thin layer, while the losses in the remainder of the air column are minimal. In a cylindrical tube this implies, again citing [51], that particle velocity is constant over the air column cross section. A one-dimensional plane wave is the result.

The presence of the boundary layer has two beneficent consequences for the modeller. The first is that, because boundary layer loss dominates the real part of the impedance term, that is, the nonoscillatory resistive part, losses can be lumped together in a simpler mathematical formulation. Second, because boundary layer damping increases with vibration amplitude, this loss helps to stabilize the oscillation [76]. Just as a parachutist in free fall reaches a terminal velocity at which aerodynamic drag matches gravitational acceleration, the instrument has a loudness limit at which boundary layer drag prevents further crescendo.

The presence of toneholes, closed or open, wall roughness, and any other departures from the smooth-walled tube will increase loss. Based on various experimental findings, Nederveen increased the boundary layer effect by 60 percent. The effect of wall material on losses was discussed in an earlier section. Reiterating, experiments have shown that the material surrounding a bore has been shown experimentally to have less than a 2 percent effect on the internal losses, and are not considered a significant factor in this area [9].

2.6.3 The player: Vocal Tract

A most important acoustical element in any instrument is its player. An exquisitely designed instrument can sound no better than a toy in the hands of an inexperienced player. A badly designed instrument can still sound exquisite in the hands of a musician. The player controls not only the input to the instrument, through his breath, and the tonehole lattice configuration, through his fingering. He may play a physiological role in the actual sound generation within. The effects of the embouchure on the reed — damping and sweep — have already been discussed. An additional, much contested control, is that which the player effects by configuring his facial and throat muscles, that of the vocal tract resonance.

The significance of the vocal tract resonance in tone control has undergone significant debate. The use of *resonances*, whether scientifically justified or no, is used by many instrumentalists, at least as an imaging technique, for getting across a certain tone. The performer may strive to “find the resonance” in various parts of the face, from cheekbone to eyeball. In addition, the degree to which the cheeks are puffed out and the openness of the throat are, from the performer’s point of view, important factors in tone production. Whether the changes in vocal tract resonance induced by these methods directly influence the tone is difficult to assess. The thinking is that the player can tune “himself” to a harmonic of the note played, and stabilize that note. Going further, he can change the harmonic content of the tone, the brightness or darkness, by appropriate adjustments in the vocal tract, that is, the shape of the mouth and throat.

That the player may have some measure of tone control through his tract shape can be demonstrated by listening to the tone of an inexperienced and experienced player. Even on a recorder, whose enclosed mouthpiece allows little additional control, the differences are profound. The real question is whether the vocal tract air column is physically interacting with the bore air column. An equally plausible explanation is that the shaping of the tract helps the player to modulate and focus his airflow into the instrument, and that the resulting airstream velocity and shape are what bring about the desired effect. Between these two hypotheses is the possibility that the tract air column is not coupled directly to the bore air column, but still bears the resonances of the tract and influences the input excitation. The *psychoacoustic* coupling is important in all cases, particularly for the technique of “finding resonances” in the facial cavities. Just as the player adjusts his

embouchure somewhat unconsciously to bring about good tone, he can adjust his vocal tract to strengthen noticeable effects, such as the ringing of a harmonic in a desired place. Whether it is air velocity or air column resonance that accomplishes the task is immaterial to him. His is essentially a “position feedback system” rather than a “force feedback system”; how he reduces the system error is transparent.

Backus, a bassoon player himself who probably had encountered more than his share of resonance pedagogy in his life, set out to test whether vocal tract resonances could have appreciable effect in a reed woodwind [3]. He approached the problem experimentally, by examining the possible influence of upstream, vocal tract impedance, on downstream, instrumental, tone. His conclusion was that the upstream pressure fluctuations due to resonance could not possibly be strong enough to influence the tone. (He also stated that in years of bassoon playing, he had never noticed a tonal change induced by a change in mouth volume). However, Keefe [34] and others have refuted Backus’ results, claiming that vocal tract resonances have a significant effect. Keefe includes the vocal tract impedance as an important term in his model formulation [34].

A vocal tract model is beyond the scope of this paper, save for its effect on the reed, and will not be included in the simulation. However, a check on its potential importance can be obtained by examining the windcapped instruments. Recall that in the windcap instruments, the reed is surrounded by a cap with a fixed aperture for the player to blow through. Some believe that a standing wave can be set up in the cap, and that this can contribute to the tone [50]. In support of this, the author has had a related experience with a tenor *Schrierpfeife*, a windcapped instrument whose lowest note is a C'. When the instrument is cold, it plays well throughout its lower register; the lowest notes are stable and strong. When it becomes warm and moist, it begins to demonstrate an unfortunate resonance somewhere between C and D when C is fingered, and between D and E when D is fingered. This resonance is not only out of tune, but it has a choked, muffled quality, consistent with it being inharmonic with the natural resonances of the instrument. Nonetheless, it tends to dominate over the much clearer proper tones, and the instrument has at times been rendered useless. The insidious resonance cannot be obtained with the windcap off; playing on the open reed always produces the clean proper tone. This indicates a clear contribution on the part of the windcap (although one might argue that the additional damping caused by playing the open reed, and not the removal of the windcap air cavity,

eliminates the resonance). An additional piece of information is the way the resonance is combatted by different musicians. One succeeds by puffing his cheeks out “Dizzy Gillespie” style. Another (this author), takes pains to start the note softly and focus the air a bit “from the top”. Another almost never encounters the problem and has difficulty understanding the complaints of her colleagues. Still another conquers the problem by refusing to play the instrument. Clearly each of these performers has a way of dealing with an instrument that offers relatively little opportunity for player control. The connection appears almost indisputably to be the combined effect of embouchure and windcap on the internal standing wave pattern.

This conclusion should be taken with a grain of salt, however. D. H. Smith has pointed out that the bottom vent hole on the instrument, which comes into play only for the lowest notes, could well be misplaced, so that the resonances for those notes were not aligned in the first place. The varying compensation used by the different players may simply point to the normal variation in the way different people combat instrument deficiencies. Some experimentation with the effects of the windcap would be useful here, and could provide important data on the influence of the upstream pressures.

Chapter 3

Digital Modeling of the Reed Woodwind

Chapter 2 presented a primarily qualitative discussion of the acoustic workings of the reed woodwind, and sought to lay a somewhat intuitive foundation for the theory. This chapter takes a more mathematical approach, providing both the specific theoretical background and the modeling equations for the woodwind simulation.

The ultimate purpose of the model developed here is to provide a realistic computer-generated simulation of a reed woodwind instrument. Not only should it be capable of sounding genuine, but it should offer the musician a genuine interface that allows him the kind of control available on an acoustic instrument. This realism is a starting point. Unfettered by the physical constraints of a conventional instrument, the musician can vary parameters formerly unavailable to him, or at best fixed for a given performance. But even these “unphysical” variations should be associated with a physical reality; their effect should be the same as it would be on a physical instrument could the musician control, say, reed stiffness, as easily as he controls reed damping or input pressure.

On the flip side of the need for acoustic realism is the liberty to discard from the model acoustically insignificant details. A twelfth-order transmission filter may be more accurate than a second-order filter, but if the ear cannot discern the difference between the two, then the difference does not matter. The limited perception of the human ear can greatly simplify the computational requirements of an acoustic model.

3.1 Frequency versus Time Domain Modeling

The term *simulation* usually implies a time domain model, while the term *synthesis* generally refers to the frequency domain (although it is often applied as well to sampling, which is technically in the time domain). The first choice one must make in developing a model is whether to design in the frequency domain or the time domain. Each has its advantages and disadvantages. This section examines the essential elements of each approach.

3.1.1 Frequency-Domain Synthesis

In this day and age, most people are implicitly familiar with frequency domain music synthesis, the source of the majority of “electronic” sound. Most musicians are in general also familiar with the concept of overtones and resonances, if in an intuitive rather than a mathematical sense. In addition, much audio research is in the frequency domain, and physical explanations for acoustic behavior — cutoff frequency, harmonically aligned resonances, etc. — exist in that domain.

Until recently, all synthesizers used some sort of frequency-domain model to represent instruments. In the earliest versions, the characteristic set of overtones were constructed using *additive* synthesis, which requires a separate generator for each harmonic. The converse approach is *subtractive* synthesis, which requires the generation of only one waveform, such as a square wave. Subsequent filtering produces the desired spectrum. Some state-of-the-art synthesizers, such as the Yamaha X-series, use John Chowning’s Frequency Modulation (FM) synthesis technique, which can synthesize complex waveforms using only two components, the carrier and modulator signals. [15] [14].

And yet, the frequency domain approach has a considerable drawback; it assumes the steady state. Time and again it has been demonstrated that, while the steady state behavior forms the backbone of a perceived tone, the transients — the attack and decay, the fluctuations — truly define it to the human listener. Psychoacoustic experiments have shown that people often cannot identify the contributor of many steady-state tones even if the spectral implementation is perfect [52]. For this reason, modern synthesizers must tack on samples of attack and decay to a looping steady-state section. In addition, synthesizers can add vibrato (although technically a frequency-based phenomenon, vibrato is perceived as a time-varying)

to help define the tone to the human listener. These time-domain enhancements are not merely embellishments; they are necessities.

Another problem with frequency domain modeling is that it is neither physical nor modular. The model begins with the end result, the emitted tone. One cannot easily model different components and connect them in a meaningful way, and an evolutionary construct is difficult to come by.

Finally, nonlinear effects cannot be easily described in the inherently linear frequency domain. The results of these effects are frequency-based; this paper often makes reference to regimes of oscillation, mode locking, reed and instrument resonances, coupling, and impedance, all at least partially in terms of their spectral descriptions. But these descriptions are after the fact; the model, the coupling, the interactions, are all executed in the time domain. Just as the musician plays his instrument and compensates for errors in the time domain, but assesses his results partially in the frequency domain, so must a good simulator program be allow the user to move freely from one domain to the other, to evaluate and improve on his output.

3.1.2 Time-Domain Synthesis

An alternative to frequency domain synthesis that is currently enjoying much popularity is the sampling approach. Sampled sound is simply digitally recorded sound, stored at various pitches in a wave table. Ideally, there is a sound sample for every pitch of the instrument being “synthesized”. In reality, sample rate conversion can be used to interpolate between more widely spaced pitch samples. Crossfading serves to splice attacks onto steady state, merge between pitch samples, and accomplish other temporal tasks. Memory requirements are naturally sizable. Even more so than the frequency synthesis, sampling is a “black box” approach. The musician is constrained to the particular instrument that was sampled. He does not have the freedom even to modify the spectrum as he did with true frequency synthesis. The results are very realistic, but very inflexible.

On the other end of the time-domain spectrum is physical modeling and simulation. Although time domain simulation has been used as a tool for analyzing acoustic behavior since the 1960's, it has not yet entered the realm of real-time synthesis. Unlike the relatively efficient frequency domain synthesis, time simulation is very computationally expensive. The

earliest simulations required days for a few seconds of data — in 1979 Stewart and Strong [70] reported a real-time ratio of 250,000 to 1, or 67 hours per second of simulation. Even today, only the simplest models can be used in a truly interactive manner. As computer technology develops, and machine speed accelerates, the possibility of using time-domain synthesis is becoming viable.

One of the major strengths of the time domain is also one of its greatest potential drawbacks — its ability and its need to incorporate physical models. No longer is the sound behavior of the instrument alone enough to characterize it. The time domain simulation developer must understand the inner workings of the “black box”. Every degree of freedom has its price. Omit an important component, and the synthesized sound will lack. Just as the aspiring musician learns to adjust nonintuitive physical parameters to enhance his sound, the modeller, and the musician who will eventually play the model, must learn how the different components of his simulation relate to his output. This is a much less straightforward path than adjusting a harmonic or tacking on an attack sample. The creation of truly realistic sounding time-synthesized sound relies upon accurate knowledge of how the system works mathematically; unfortunately, much of the experimental research necessary for the development of a comprehensive model has simply not yet been performed. Many assumptions common throughout the literature are being refuted in new experiments. Hirschberg has stated explicitly that no one really understands what makes a double reed function, and his is the most recent work to date.

The time domain approach is the course taken for this research. Specifically, this paper investigates the use of *waveguide digital filters* an efficient, modular modeling approach, in incorporating various understood features of the sound generation process.

3.2 Digital Waveguide Bore Model

The previous chapter described the physical characteristics of the air column contained within the bore of a woodwind. To summarize, this air column has the following salient features:

- Essentially linear behavior;
 - Supports compression waves traveling in both directions at the speed of sound;
-

- Frequency of oscillation dependent on length of air column;
- Partially reflects waves at impedance junctions. Reflections will invert at open ends, but not at closed end;
- Highpass transmission, lowpass reflection at junctions where the impedance drops, such as at toneholes and the bore terminus;
- Ability to couple with a nonlinear element such as a reed.

This section develops a mathematical model which will reproduce these features.

3.2.1 The Wave Equation for Lossless Propagation

The first step to developing the model is to state the fundamental equations of planar, that is, one-dimensional, lossless wave propagation. This discussion is taken from [41], [46] and [47].

The wave variables for an air column in an acoustic tube can be defined as p , the pressure, and u , the volume velocity. If A represents the cross-sectional area at point x , and ρ represents the density of air, then the familiar *force = mass \times acceleration* relationship, which implies conservation of momentum, can be written as:

$$A(x) \frac{\partial p(x, t)}{\partial x} = -\rho \frac{\partial u(x, t)}{\partial t} \quad (3.1)$$

Conservation of mass is implied by the equivalence of the mass outflow to the change in density, δ [47]:

$$A(x) \frac{\partial \delta}{\partial t} = -\frac{\partial}{\partial x}[(\rho + \delta)u] = -(\rho + \delta) \frac{\partial u}{\partial x} - u \frac{\partial \delta}{\partial x} \quad (3.2)$$

For adiabatic compression, the change in density can be related to the change in pressure by the approximation [46]:

$$\delta = \frac{1}{\gamma c P_o} p \quad (3.3)$$

where γ_c is the specific heat ratio, generally equal to about 1.4 for air, and P_0 is the ambient pressure.

It turns out that the speed of sound through the medium, c , can also be expressed as a function of γ_c and P_0 :

$$c = \sqrt{\frac{\gamma P_0}{\rho}} \quad (3.4)$$

or

$$\gamma P_0 = \rho c^2 \quad (3.5)$$

Then the pressure/density relationship becomes:

$$\delta = \frac{1}{\rho c^2} p \quad (3.6)$$

Substituting into 3.2:

$$\frac{A(x)}{\rho c^2} \frac{\partial p}{\partial t} = - \left(\rho + \frac{p}{\rho c^2} \right) \frac{\partial u}{\partial x} - \frac{u}{\rho c^2} \frac{\partial p}{\partial x} \quad (3.7)$$

Eliminating second order terms, canceling ρ , and exchanging sides yields:

$$\rho \frac{\partial u(x,t)}{\partial x} = - \frac{A(x)}{c^2} \frac{\partial p(x,t)}{\partial t} \quad (3.8)$$

This is the continuity of mass equation in [41].

Equations 3.1 and 3.8 can be combined to yield the *Webster Horn Equation* :

$$\frac{\partial}{\partial x} \left[\frac{1}{A(x)} \frac{\partial u(x,t)}{\partial x} \right] = \frac{1}{c^2 A(x)} \frac{\partial^2 u(x,t)}{\partial t^2} \quad (3.9)$$

In a cylindrical tube, the cross-sectional area is constant, and $A(x)$ can be simplified to A , which then cancels itself out. Equation 3.9 reduces in this case to:

$$\frac{\partial^2 u(x,t)}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 u(x,t)}{\partial t^2} \quad (3.10)$$

Similarly for pressure:

$$\frac{\partial^2 p(x, t)}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 p(x, t)}{\partial t^2} \quad (3.11)$$

These equations can be satisfied by decomposing the waves into their right and left-going components. Recall that this decomposition was described in the previous chapter for the purpose of justifying the boundary conditions affecting wave reflection. The pressure and flow functions now become:

$$p(x, t) = p^+(t - x/c) + p^-(t + x/c) \quad (3.12)$$

$$u(x, t) = u^+(t - x/c) + u^-(t + x/c) \quad (3.13)$$

where the two components of each are the right and left-going waves respectively.

The solution to the wave equations above then takes the oscillatory form:

$$p = p^+ + p^- = P_0^+ e^{ik(x-ct)} + P_0^- e^{ik(x+ct)} \quad (3.14)$$

With the left and right-going formulation, the momentum conservation equation 3.1 can be rewritten as:

$$A(x) \left[\frac{\partial p^+(x, t)}{\partial x} + \frac{\partial p^-(x, t)}{\partial x} \right] = -\rho \left[\frac{\partial u^+(x, t)}{\partial t} + \frac{\partial u^-(x, t)}{\partial t} \right] \quad (3.15)$$

Since the time taken to travel distance x is linked to the speed of sound:

$$\Delta t = \frac{\Delta x}{c} \quad (3.16)$$

the time and spatial derivatives of the components are similarly related as:

$$\frac{\partial f(t + x/c)}{\partial t} = (-c) \frac{\partial f(t \pm x/c)}{\partial x} \quad (3.17)$$

$$\frac{\partial f(t - x/c)}{\partial t} = (c) \frac{\partial f(t \pm x/c)}{\partial x} \quad (3.18)$$

Using this in equation 3.15, and dividing through by A to get all constants on the right hand side:

$$\left[\frac{\partial p^+(x,t)}{\partial x} + \frac{\partial p^-(x,t)}{\partial x} \right] = \frac{\rho c}{A} \left[\frac{\partial u^+(x,t)}{\partial x} - \frac{\partial u^-(x,t)}{\partial x} \right] \quad (3.19)$$

This same result can be obtained through similar manipulation of the mass conservation relationship in equation 3.8. Now the constants on the right hand side of the equation can be grouped into one constant, R :

$$R = \frac{\rho c}{A} = \frac{\sqrt{\rho \gamma P_0}}{A} \quad (3.20)$$

Examining again the relationship in equation 3.19, we see that R relates pressure to flow. Simplifying for a unit section, we have:

$$p^\pm = \pm R u \quad (3.21)$$

This expression is almost identical to the impedance relationship first presented in Chapter 2. This term R is in fact defined as the *characteristic impedance*, or, more appropriately, the *characteristic resistance*, of the air column. Whereas the *input*, or *acoustic*, impedance discussed in Chapter 2 was a complex quantity, the *characteristic impedance* is always real and positive (the complex acoustic impedance at a point is often normalized by the characteristic impedance to form the *specific acoustic impedance*).

One subtlety in equation 3.21 is sign reversal for the left-going wave. This is because of the flow direction implied by a rarefaction or compression, as illustrated in 3.1. Flow will always head into a compression, and away from a rarefaction. Because the right going and left going flow will accordingly travel toward the same pressure antinode, where the flow goes to zero, the sign of the left-going flow wave must be opposite that of the right, or, equivalently, opposite that of the pressure.

3.2.2 Impedance and Acoustic Circuits

Chapter 2 discussed the parallels between wave propagation along a string and in a tube, and the complementary relationship between pressure and flow wave. In fact, there are many

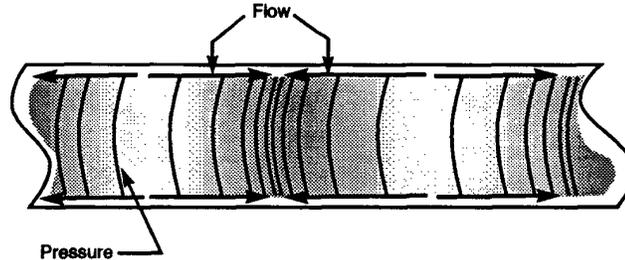


Figure 3.1: Flow associated with a compression wave

other wave propagation phenomena in nature. All exhibit essentially the same mathematical behavior. All have the dual relationship between a *force* variable, such as pressure, and a *flow* variable, such as volume velocity. Shearer et.al. categorizes these variables even more generically in [59]. The *through* variable is represented in an absolute sense by, for example, flow or current. It has the same value at the two ends of an element, satisfying mass conservation. The *across* variable is represented in a relative sense by, for instance, pressure or voltage drop. The across variable represents a potential which can change across an element. Once the through and across variables of any dynamic system have been defined, that system can be distilled to a generic network containing elements and variables with well-known relationships. The physical particulars of the system are no longer important. For this reason, the well known Kirchoff's Law and Ohm's Law, developed for electrical systems, apply equally well to acoustic, thermal, or mechanical systems.

The conditions satisfied by Kirchoff's laws are those of *compatibility* and *continuity*. The *compatibility* constraint requires that the across variable be the same for all ports of a junction, as represented by Kirchoff's Voltage law. Figure 3.2 illustrates a mechanical system, whose across variable is displacement, or actually, velocity. It is easy to see that the displacement x_{damper} must be equal to the x_{spring} , regardless of the particulars of damper and spring. The same is true for the equivalent electrical system in Figure 3.3. The voltage drop across section A must be the same as that across section B. This implies that *the total voltage drop around a loop must be zero*.

Kirchoff's current law satisfies the *continuity* constraint, which requires that the through variable be conserved. In Figure 3.2, this implies that, while the force F_0 can be distributed

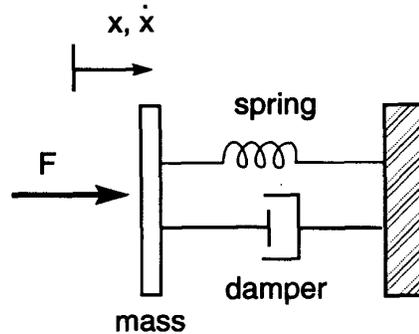


Figure 3.2: Typical Mechanical System

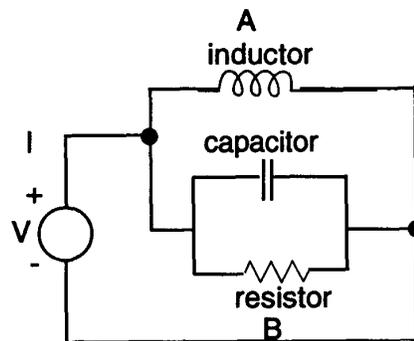


Figure 3.3: Typical Electrical System

unequally between damper and spring, the sum of the two components F_{damper} and F_{spring} must equal F_0 . In the electrical example, the current must be conserved, so that the currents flowing through loop A and loop B sum to the current supplied at point I. Restating, *the through variables must sum to 0 at a junction.*

Finally, Ohm's Law for electrical circuits states the relationship between through and across variables in terms of *resistance* (there is also an entirely different acoustic Ohm's Law which states dependence of sound perception on relative harmonic amplitude rather than phase [24]). Extending the purely dissipative resistance to the complex, frequency-dependent impedance term Z , we have again the relationship:

$$Z = \textit{Through}/\textit{Across} \quad (3.22)$$

These generalized relationships can be restated for an acoustic system, where the through variable is *flow* and the across variable *pressure*:

- The acoustic pressure drops around a loop must sum to zero;
- The airflows entering a junction must sum to zero;
- The pressure is related to flow by: $P = ZU$.

Most acoustical studies have involved the reduction of the system into this sort of acoustic circuit. Some, such as [57], have attempted to create specific electrical analogs for each acoustical element, by transforming air inertance, compressibility, and dissipation into equivalent inductance, capacitance, and resistance. While this places the problem in a more familiar realm for the electronically oriented, it removes some of the physical sense of the system. A more general approach is to model the system purely in terms of its generalized impedance, rather than by a point-for-point analog. This is the avenue taken by waveguide / transmission line theory. Not only is it simpler conceptually, but it allows elegant mathematical modularity and efficiency.

3.2.3 Reflection Functions and the McIntyre Woodhouse Model

One of the cornerstones of the waveguide technique is the use of *reflection coefficients*. The reflection terms directly model the physical consequences of impedance changes, that being the partial reflection / partial transmission of a wave at a junction.

The groundbreaking work of McIntyre, Schumacher, and Woodhouse, documented in [43] and [58] introduced the concept of formulating the acoustic model by separating the right and leftgoing waves. In their simulation, which was intended as an aid in studying and quantifying acoustic behaviour, the use of reflection functions was vital for reducing the computational load.

Traditionally, linear time-invariant (LTI) dynamic systems are characterized by their impulse response, or *Green's* function. If the clarinet response is represented by $g(t)$, the

mouthpiece pressure by $p(t)$, and the air volume velocity by $u(t)$, then the equation for mouthpiece pressure, following [58], is the convolution integral:

$$p(t) = \int_0^{\infty} g(t')u(t-t')dt' = g * u \quad (3.23)$$

Because of all the internal reflection which bounces the wave back and forth, the impulse response $g(t)$ takes some time to decay, and therefore has many terms. This causes the above equation to be quite computationally intensive.

A more fundamental, more compact, dynamic term is the reflection that occurs at the clarinet bell (this simple model does not include toneholes). As we have seen in the discussion of cutoff frequency in the previous chapter, the bell acts as a lowpass filter. An impulse, which contains in theory infinite bandwidth, will be reflected as a finite bandwidth hump; in the time domain, this implies that the impulse reflection is smoothed, or “temporally smeared” in the words of [43], as shown in Figure 3.4. The reflection function, $r(t)$, is simply the impulse response of the equivalent lowpass filter. Note that $r(t)$ decays much more rapidly than did the total system impulse response, $g(t)$.

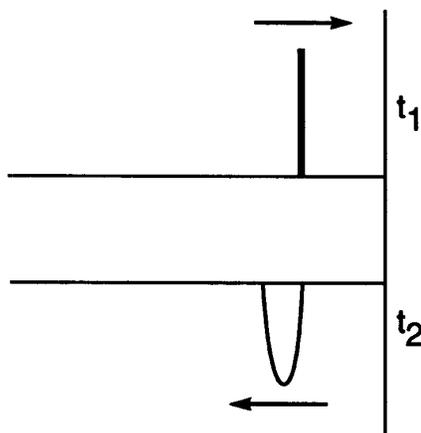


Figure 3.4: Reflection of an impulse in a clarinet (after McIntyre and Woodhouse)

The left-going wave, i.e., the reflected wave, is the result of the convolution of the right-going wave with the reflection function:

$$p_i(t) = \int_0^{\infty} r(t')p_o(t-t')dt' = p_o(t) * r(t) \quad (3.24)$$

Because the reflection function decays quickly, this integral is computationally much easier to evaluate than that in equation 3.23.

For his reflection function, McIntyre [43] chose to use a simple time-domain Gauss function which approximated a cylindrical tube without toneholes:

$$r(t) = \begin{cases} ae^{-b(t-T)^2}, & (t \geq 0); \\ 0, & (t < 0) \end{cases} \quad (3.25)$$

In his related additional work, Schumacher [58] chose to define the reflection function in the frequency domain, $R(\omega)$. The relationship between $R(\omega)$ and the associated tube impedances can be derived by examining the pressure/flow relationships at the reed/bore junction. Following the discussion of McIntyre, the flow and the pressures are related at the reed by the real characteristic impedance, Z_c :

$$Z_c u(t) = p^+(t) - p^-(t) \quad (3.26)$$

The total pressure is the sum of right and left-going waves, as described above:

$$p(x, t) = p^+(t - x/c) + p^-(t + x/c) \quad (3.27)$$

At the reed, $x = 0$, so this simplifies to:

$$p(t) = p^+(t) + p^-(t) \quad (3.28)$$

Eliminating p^+ in equations 3.26 and 3.28, and substituting the resulting term for p^- in 3.24:

$$2p^+(t) = r(t) * \{Z_c u(t) + p(t)\} \quad (3.29)$$

Again from equations 3.26 and 3.28, the total pressure p can be expressed as:

$$p = 2p^+(t) + Z_c u \quad (3.30)$$

Combining:

$$p = r(t) * \{Z_c u(t) + p(t)\} + Z_c u(t) \quad (3.31)$$

This is the equation implemented by McIntyre. In the frequency domain, using the convolution theorem, equation 3.31 transforms to:

$$P(\omega) = R(\omega)\{P(\omega) + Z_c U(\omega)\} \quad (3.32)$$

Solving for the reflection function, $R(\omega)$:

$$R(\omega) = \frac{P(\omega) - Z_c U(\omega)}{P(\omega) + Z_c U(\omega)} \quad (3.33)$$

At this point, we can define the complex acoustic impedance, Z_a , with the standard impedance ratio definition:

$$Z_a(\omega) = \frac{P(\omega)}{U(\omega)} \quad (3.34)$$

Equation 3.33 now becomes:

$$R(\omega) = \frac{Z_a(\omega) - Z_c}{Z_a(\omega) + Z_c} \quad (3.35)$$

A similar result can be found at the terminating junction. In his derivation, Schumacher replaced impedance with admittance, $Z = 1/Y$, so that at the bell:

$$R(\omega) = \frac{Y_0 - Y(\omega)}{Y_0 + Y(\omega)} \quad (3.36)$$

where Y_0 is the characteristic admittance (i.e., the reciprocal of the characteristic impedance discussed above), and $Y(\omega)$ is the complex input admittance, which Schumacher calculated using the one-dimensional Schrödinger wave equation.

Taking the Fourier transform of $R(\omega)$ yields $r(t)$, the reflection function in the time domain:

$$r(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{r}(\omega) e^{i\omega t} d\omega \quad (3.37)$$

The preceding formulation was a physicist's approach to analyzing acoustical behavior. The purpose of the simulation was to better understand the workings of the clarinet, but not to replicate it. McIntyre's comment on the procedure indicates much about his perspective (*italics mine*):

"Indeed, if the convolution integral were done by hardware using integrated circuits available for the purpose, and the remaining programming done as efficiently as possible in assembly language, a fast minicomputer could produce results at a cycle rate in the audible range. *The result would perhaps have some novelty: an electronic musical instrument based on a mathematical model of an acoustic instrument.*"

3.2.4 Waveguide Digital Filters

The concepts employing right and left-going waves along with reflection functions to represent wave propagation through impedance junctions, introduced by McIntyre and Woodhouse to model musical instrument behavior, extend easily to a filter architecture known as the *Waveguide Filter* (WGF). This approach has been developed extensively by Julius O. Smith, and much of this section follows directly from a compendium of Smith's WGF work in [64].

History

The term *filter* is generally used to denote any medium through which a substance may be passed and operated upon. For example, liquid or gas filters can be used to purify incoming material by separating components. A color filter can separate spectral components of light by selective absorption. In electronic or digital systems, filters are used for shaping an input signal, either through their frequency-dependent amplitude response, such as for lowpass or bandpass filtering, or phase response, such as for communication channel equalization. These filters are usually time-invariant and linear, with fixed phase and amplitude response.

Although originally conceived for the purpose of signal shaping, the filter definition in digital systems has been extended to any network which produces some desired system response to a defined input. In particular, it can be applied to networks constructed to implement a mathematical model of a physical system. In this case, the filter represents a *simulation* of a real system or *plant*, and is designed with the intention of duplicating the input/output characteristics of the plant. Feedback control systems frequently rely on such models to shape their response to errors. Because plant characteristics are seldom time-invariant, neither are the filters designed to model them. Self-adapting filters, such as Kalman filters and all forms of adaptive filters have the capability of adjusting their own parameters based upon some feedback error measure. However, these operate under the assumption that the coefficient changes will be very slow. After an initial *convergence* period of rapid change, the filters are expected to be “quasi-static”, i.e., seemingly time-invariant with respect to the filter impulse response. Should rapid changes occur, performance degradation can result.

One problem with conventional filters is that, while they can be used to model the input/output behavior of physical systems, their forms often bear little resemblance to the reality which they represent. These filters are usually designed in the frequency domain, to reproduce the resonances and damping, or poles and zeros, of the modelled plant. In their time domain forms, they have little meaning. Even if an “acoustic circuit” of capacitor, inductor, and resistor equivalents are set up, the resulting parameters are not intuitive, and cannot be varied in a physically meaningful way. Also, because the implementation is a digital realization of a continuous system, there is an implicit continuous-to-digital mapping which can be exact at only a limited set of frequencies.

The *Waveguide Digital Filter* structure was enlisted to combat the above problems by modeling physical wave propagation directly. The WGF was not the first filter structure used for this task. Fettweis developed the *Wave Digital Filter*, which combined the notion of *wave variables* and scattering layers with the use of equivalent circuit elements [23]. Markel and Gray, in their work on speech synthesis in [41], introduced the use of the “normalized lattice filter for simulating the properties of the “acoustic tube”. These employed cascades of equal length cylindrical tubes, each with its own unique diameter and corresponding characteristic impedance. As we have seen in earlier discussions of wave propagations, scattering would occur at each junction according to the ratio of impedances. While computationally stable, this formulation had two disadvantages: 1) it required a perfectly

reflecting (and therefore nonphysical) termination, which limited its modularity; and 2) it was computationally intensive, requiring four multiplications per pole, as opposed to the one multiplication per pole necessary for a one-multiplier lattice filter. Smith has extended this work to the WGF, by creating a highly modular, computationally efficient and stable, framework for acoustic modeling [64].

Description / Advantages

The Waveguide Digital Filter is a network of component waveguides, or *transmission lines*, which represent the medium through which waves can propagate. Again, following Smith in [64], it has the following features:

- Designed in the time domain to realistically simulate acoustic wave propagation;
- Component waveguides are cylindrical sections with fixed characteristic impedance. At this fundamental level, the flow and pressure is in phase and the impedance is *real* and *positive*, which is very convenient;
- Component waveguides support both right and leftgoing waves, using the wave propagation rules discussed above;
- Lossless scattering is implemented at junctions between sections, based on scattering coefficients which have been computed from cross-sectional area;
- Waveguide networks follow the laws of continuity and compatibility discussed in the previous section;
- Any waveguide network may be connected to another network, or terminated by some other sort of filter;
- WGF's are linear, finite-order, and recursive. In special cases they reduce to NLF form.

The following section will discuss these features in more detail. The advantages of the WGF, based upon the basic description above, include:

- Precise, physically meaningful representation of the plant. Time-domain design
-

obviates digital-to-continuous mapping. Parameters can be varied in a realistic manner. This is a necessity for a musical instrument with an interface to a live performer, who will want to use intuitive gestures for controlling the musical output;

- **Computationally stable.** The inherent stability is a strength of normalized lattice/ladder filters in general. Because they have unity power gain at each node, there can be no overflow within a stable filter, and consequently, no zero-input overflow oscillations. With magnitude-truncation arithmetic, which guarantees that the signal will never be increased by rounding, there will be no zero-input limit cycles. In addition, the roundoff noise properties are good;
- **Algorithmically stable.** This is also a feature of all normalized lattice filters. As long as the input signal is bounded, and the reflection coefficients are less than unity, the filter will be stable;
- **Extremely modular.** Components can be combined in series or in parallel, junctions can have any number of ports, with no sacrifice in stability and power decoupling. This allows instruments models to “evolve” easily;
- Because the WGF reduces in certain cases to the NLF, it can be analyzed and manipulated with time-honored NLF techniques.

These advantages are particularly important for the application of WGF's to the modeling of musical instruments. Smith points out other more general advantages in his discussion in [64]. In particular, Smith covers the n -dimensional waveguide, with arbitrary connections and m -port junctions. The treatment here will be limited primarily to the 2-dimensional waveguide, which adequately models the propagation of waves through a linear medium such as an acoustic tube.

Basic Waveguide Section

The components of the WGF are, as suggested above, extremely simple. The basic waveguide section, taken again from Smith [64], is illustrated in Figure 3.5.

- The characteristic impedance of the section, Z_2 , is based on its cross section.
-

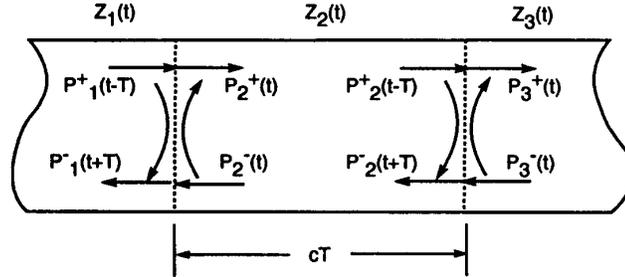


Figure 3.5: Waveguide Section (after Smith)

- Right and left-going waves propagate from, respectively, adjoining sections with impedances Z_1 and Z_3 .
- As the right-going wave, $P_1^+(t - T)$, reaches the 1-2 junction, it partially reflects back through section 1, according to the impedance ratio.
- The portion which is transmitted, $P_2^+(t)$, is scattered again at the 2-3 junction. The reflected portion here is superimposed on the left-going wave, $P_2^-(t + T)$, which is taking a similar route from section 3.
- The waveguide section represents one time delay, T . Its “length” is correspondingly cT , where c , as always, is the speed of sound.

Note that these steps are essentially the same as those discussed earlier for wave propagation. This arrangement can be represented by Smith’s simulation flow diagram in Figure 3.6. Here, the transmission gains are represented by g_i and the reflection gains by ρ_i . The traversal time delay, T , takes the form z^{-T} .

We now have a system characterized by a network of simple sections and junctions, which can be described completely by the set of transmission and reflection gains. In fact, the next section will show that this set of gains — four per section — can be distilled to one reflection coefficient, ρ , which characterizes the junction.

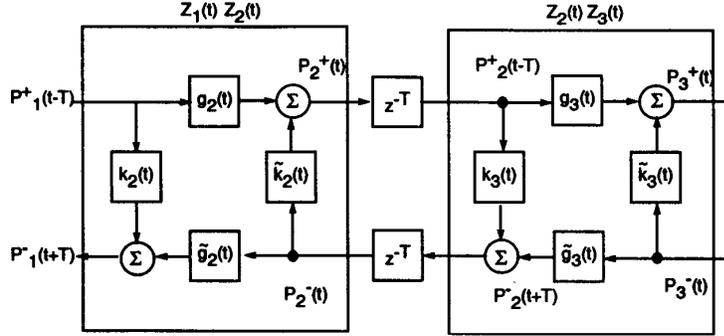


Figure 3.6: Block Diagram of a waveguide section (after Smith)

Scattering Equations

The manipulation rules from which the basic waveguide equations derive are the same as those discussed earlier for wave propagation - the pressure/flow/impedance relationship, continuity, or conservation of mass, and compatibility, or conservation of energy. Restating and simplifying from earlier sections:

- Pressure/Flow/Characteristic Impedance

$$P_i^+ = Z_i U_i^+ \quad (3.38)$$

$$P_i^- = -Z_i U_i^- \quad (3.39)$$

where P_i^\pm is left/right-going pressure in the i 'th section, U_i^\pm is left/right-going flow in the i 'th section, and Z_i is the characteristic impedance of the i 'th section;

- Decomposition into Left and Right-going Waves

$$P_i = P_i^+ + P_i^- \quad (3.40)$$

$$U_i = U_i^+ + U_i^- \quad (3.41)$$

where P_i and U_i are the instantaneous pressure and velocity in section i , with respect to both space and time;

- Conservation of Mass and Energy (Continuity across a Junction)

$$P_{i-1}(cT, t) = P_i(0, t) \quad (3.42)$$

$$U_{i-1}(cT, t) + U_i(0, t) = 0 \quad (3.43)$$

The continuity equations can be modified by noting the relationships in Figure 3.5. By definition in this figure, the pressure at the entrance to a section, that is, the extreme right where $x_i = 0$, is:

$$P_i(0, t) = P_i^+(t) + P_i^-(t) \quad (3.44)$$

$$U_i(0, t) = U_i^+(t) + U_i^-(t) \quad (3.45)$$

Similarly, at the extreme left of a section, where $x = cT$ and T is the time required to traverse a section:

$$P_{i-1}(cT, t) = P_{i-1}^+(t - T) + P_{i-1}^-(t + T) \quad (3.46)$$

$$U_{i-1}(cT, t) = U_{i-1}^+(t - T) + U_{i-1}^-(t + T) \quad (3.47)$$

Then equation 3.43 can be rewritten as:

$$0 = U_{i-1}^+(t - T) + U_{i-1}^-(t + T) + U_i^+(t) + U_i^-(t) \quad (3.48)$$

For the sake of brevity, the t arguments will be dropped for now, remembering that:

$$i \Rightarrow t; \quad i_{-1}^+ \Rightarrow t - T; \quad i_{-1}^- \Rightarrow t + T; \quad (3.49)$$

Simplifying equation 3.48 in this manner:

$$0 = U_{i-1}^+ + U_{i-1}^- + U_i^+ + U_i^- \quad (3.50)$$

In order to put this in terms of pressure, we can use equation 3.39, and define the admittance $Y = 1/Z$. Then:

$$0 = Y_{i-1}P_{i-1}^+ - Y_{i-1}P_{i-1}^- + Y_iP_i^+ - Y_iP_i^- \quad (3.51)$$

Again from equation 3.43, the pressures P_{i-1} and P_i must be equal to each other. Defining P_J as this junction pressure, using equation 3.41 to eliminate the right-going wave components for now, and solving for P_J :

$$P_J = 2 \frac{Y_{i-1}P_{i-1}^- + Y_iP_i^-}{Y_{i-1} + Y_i} \quad (3.52)$$

Again using equation 3.41 to reintroduce P^+ and eliminate P_J , and solving for P_i^+ :

$$P_i^+ = 2 \frac{Y_{i-1}P_{i-1}^- + Y_iP_i^-}{Y_{i-1} + Y_i} - P_i^- \quad (3.53)$$

$$P_i^+ = 2 \frac{Y_{i-1}P_{i-1}^-}{Y_{i-1} + Y_i} + \frac{P_i^-(-Y_{i-1} + Y_i)}{Y_{i-1} + Y_i} \quad (3.54)$$

We can now define the *reflection coefficient*, ρ_i , as:

$$\rho_i = \frac{Y_{i-1} - Y_i}{Y_i + Y_{i-1}} = \frac{Z_i - Z_{i-1}}{Z_i + Z_{i-1}} \quad (3.55)$$

Using a similar procedure to derive P_{i-1}^- , and bringing the t arguments back in, we get the *scattering equations*:

$$P_i^+(t) = [1 + \rho_i(t)] P_{i-1}^+(t - T) - \rho_i(t) P_i^-(t) \quad (3.56)$$

$$P_{i-1}^-(t + T) = \rho_i(t) P_{i-1}^+(t - T) + [1 - \rho_i(t)] P_i^-(t) \quad (3.57)$$

Figure 3.7 illustrates the implementation of these equations. This configuration is known as the *Kelly-Lochbaum* junction. Note that, as promised, the set of four transmission reflection gains for each section can now be characterized by a single reflection coefficient. Again, these apply to the simplest 2-port junction. Smith provides the more general equations in matrix form in [64].

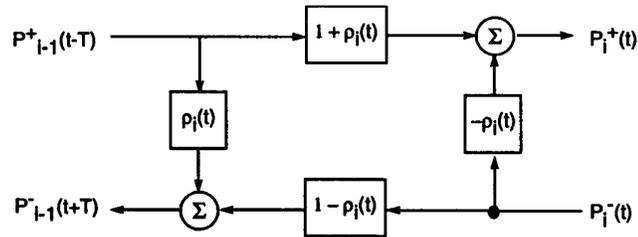


Figure 3.7: Kelly-Lochbaum Scattering for a 2-port Junction (after Smith)

Cascaded 2-Port Waveguides

The final element necessary for connecting the sections defined above to form a complete filter is the sample delay, denoted in conventional signal processing nomenclature as z^{-T} . This represents the transit time across one section. These sections can now be linked as illustrated in Figure 3.8, taken from [64]. Figure 3.9, taken from the same source, depicts the same system in classical ladder/lattice structure. It can be shown that the two are equivalent; the delays can be “pushed” around the loop so that they are limited to the “lower rail”. From these two structures, we can see the advantages of Smith’s approach. The ladder structure has only half the number of delays, which may be more computationally efficient. However, it no longer internally represents the physical wave propagation of left and right going waves (although the input/output is identical). Also, it must have a purely reflective termination, which, as we shall see, precludes the possibility of a realistic radiation model.

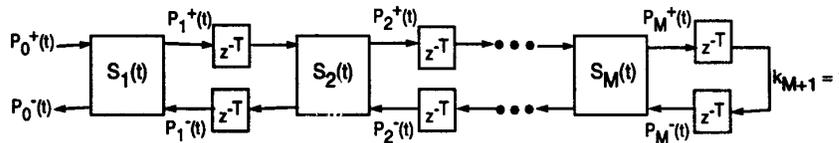


Figure 3.8: Waveguide Filter Structure (after Smith)

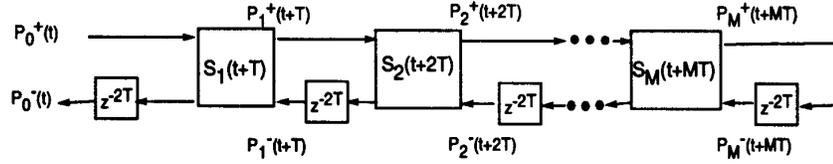


Figure 3.9: Normalized Ladder/Lattice Structure (after Smith)

Signal Power and Junction Passivity

Waveguide sections constructed in the above manner have an instantaneous signal power (\mathcal{P}) of:

$$\mathcal{P}_i^+(t) = P_i^+(t)U_i^+(t) = \frac{[P_i^+(t)]^2}{Z_i(t)} \quad (3.58)$$

$$\mathcal{P}_i^-(t) = P_i^-(t)U_i^-(t) = \frac{[P_i^-(t)]^2}{Z_i(t)} \quad (3.59)$$

$$\mathcal{P}_i(0, t) = \mathcal{P}^+(t) + \mathcal{P}_i^-(t) \quad (3.60)$$

For the 2-port WGF network defined above, Smith shows that power is conserved by summing the ingoing and outgoing power components:

$$\mathcal{P}_j \triangleq \sum_{i=1}^2 P_i U_i \quad (3.61)$$

Since the pressures, P_i , seen by all contributors must be equal:

$$\mathcal{P}_j(0, t) = P_j [U_{i-1}(cT, t) + U_i(0, t)] \quad (3.62)$$

But from equation 3.43, the flows within the brackets must sum to 0. Then:

$$\mathcal{P}_j(0, t) = 0 \quad (3.63)$$

Thus, power is conserved, and the junction is, ideally, *lossless*.

An important power-related issue is computational stability in the presence of numerical errors due to rounding. In this case, equation 3.43 may not be exactly satisfied. However, junction passivity can be assured if the incoming power always bounds the outgoing power, i.e.:

$$|\mathcal{P}_{i-1}(cT, t)| \geq |\mathcal{P}_i(0, t)| \quad (3.64)$$

For the waveguide, we can denote \hat{P} as the finite-precision form of the ideal P . Using the definition of signal power in equation 3.60 with the passivity requirement leads to the constraints:

$$|\hat{P}_i^+(t)| \leq |P_i^+(t)| \quad (3.65)$$

$$|\hat{P}_{i-1}^+(t+T)| \leq |P_{i-1}^+(t+T)| \quad (3.66)$$

This is equivalent to requiring that rounding always be implemented as a truncation toward zero.

Implementation Issues

The above equations present the basic waveguide representation of the acoustic tube. These can be reformulated for more efficiency, and for other computational benefits. Salient modifications suggested by Smith include:

- Reconstruction of the Kelly-Lochbaum junction, which requires 4 multiplies and 2 additions, into the *one-multiply* scattering junction, which entails 3 additions. This can be more efficient computationally. In addition, it reduces the number of multiplicative error sources to one.
- Formulation of the “half-rate waveguide filter”, which lies between the WGF and ladder structures described in Figure 3.8 and 3.9 above. This structure combines the reduction in delays with the extendibility of the full-rate WGF. It is still,

however, conceptually a step away from the wave propagation being modelled.

- Power-normalizing the filter, which implies that the signal power in a section will be fixed, regardless of characteristic impedance. Smith suggests several methods for accomplishing this:
 1. Normalizing the waveguide by scaling the wave down (or up) at impedance junctions (specifically, at the delay line outputs) according to the square root of the impedance ratio. This conserves signal power, but introduces two additional multiplies per second. Also, a change in characteristic impedance entails a change in the reflection coefficients of both the affected section and its neighbors;
 2. Propagating rms-normalized waves by modifying the reflection coefficients. Although this approach is not more costly in the conventional Kelly-Lochbaum structure, the resulting scattering cannot be reduced to the one-multiply form. However, it makes transforming between duals (e.g., pressure and flow) much easier, because the reflection coefficients will be the same for both. Changes in characteristic impedance in a section will be isolated to the reflection coefficients of that section, rather than propagating to its neighbors;
 3. Using transformer-coupled waveguides. This technique implements the normalization gains just at the entrance (left-hand side) of the section. This allows the one-multiply scattering junction to be used, but retains the advantage of isolated characteristic impedance / reflection coefficient changes.

3.2.5 Adaptation of the Waveguide to the Reed Instrument Model

So far we have defined a two-dimensional cascaded waveguide filter with undefined terminators. This filter can now be applied to an acoustic model of the woodwind bore similar to that discussed in previous sections. The terminators will represent, at one end, the power input provided by player through the reed, and at the other end, the radiation of sound from the bell and toneholes. Just as described in Chapter 2, the combination of bore and bell forms a resonator which will couple nonlinearly with a generator, the reed model. The remainder of this section discusses the waveguide construction of the woodwind resonator,

which could be extended to any wind instrument regardless of generator type. The following section will treat the reed generator model, which we will see can be formulated as a time-varying reflection coefficient.

The Bore

The acoustic tube model represented by the WDF is in itself an accurate depiction of the bore of an instrument. In the simplest case, the cylindrical bore with no toneholes, there is a constant cross-sectional area, and therefore impedance, throughout the bore. The waveguide then reduces to a simple pair of delay lines, representing right and left-going “rails”, as shown in Figure 3.10. The length of the delay line depends upon the fundamental frequency of the desired note. Recalling that, for a cylindrical bore, the period of the oscillation is equal to 4 times the borelength, the number of delays required, n_d can be computed as:

$$n_d = \frac{\omega_s}{4\omega_f} \quad (3.67)$$

where ω_f is the desired fundamental frequency ω_s is the sampling frequency.

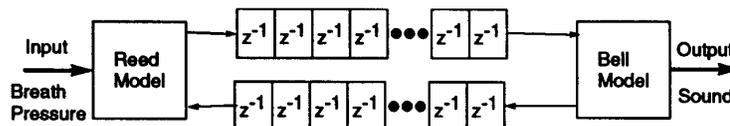


Figure 3.10: Reduction of WGF to a simple delay line for a cylindrical bore (after Smith)

For example, with a 22050 Hz sampling frequency, typical in sound applications, the precise “length” required for the standard A440 is 12.53, or somewhere between 12 and 13 delay lines. Note that this number is very small, and that errors due to integer truncation will be appreciable, particularly in the upper harmonics. This problem can be countered in a number of ways. One would be to increase the sampling rate. The sampling rate, however, is often fixed. The NeXT computer does offer a higher sampling rate of 44100 Hz, which would double the length of the delay line. However, this is at the expense of doubled storage and computational requirements. The delay line can also be doubled in length by dropping

the note an octave, i.e., halving its frequency. This is the option selected for the studies here, and in fact, well matches the chalumeau register of the Bb clarinet. The lower tone turned out to be the better choice from the standpoint of researcher comfort as well; the higher tone was more annoying to listen to. This was particularly important for the register hole experiments, which raised the fundamental by a 12th, or 2.5 times the original frequency.

It turns out that this computation will yield only an approximation of the desired frequency. The presence of the reed should, as Chapter 2 stated, lower the sounding frequency. In addition, the bell model, which will be discussed shortly, can introduce a significant amount of phase lag, and correspondingly, a flattening of the pitch.

A conical bore is much more complicated, because it entails a continuous change of cross-section. This can be modelled in waveguide form by using a “finite-element” approach, cascading a series of cylindrical sections, each with increasingly large cross-section. Figure 3.11 illustrates the conical bore in sampled form. For an accurate representation, the length of a single section should be shorter than the wavelength of the highest frequency oscillation expected.

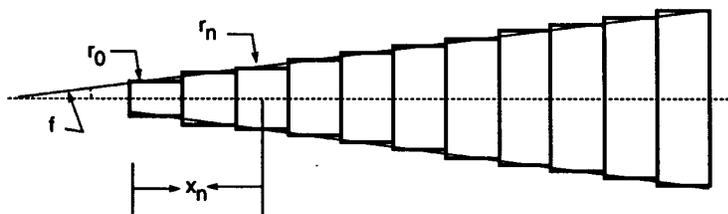


Figure 3.11: Representation of a conical bore as a series of cylindrical waveguide sections

In [18], Cook modelled the conical bore of a saxophone in this manner. He calculated the cross-sectional area of a bore section, a_n , based on the bore radius, r_0 , the flare coefficient, f , and the distance along the bore, x_n , as:

$$a_n = \pi(r_0 + fx_n)^2 \quad (3.68)$$

Because cross-sectional area is inversely related to impedance, the corresponding reflection coefficient is:

$$\rho_n = \frac{a_n - a_{n-1}}{a_n + a_{n-1}} \quad (3.69)$$

Note that, because the bore is conical, the relationship between fundamental and sampling frequency is now:

$$n_d = \frac{\omega_s}{2\omega_f} \quad (3.70)$$

By including junctions at every spatial sample, i.e., every unit delay element, an exact simulation at the sample points is obtained. Cook also experimented with reducing the number of junctions from his nominal value of 40, running cases with 20, 10, and 5 evenly spaced sections. As could be expected, the coarser models exhibited degraded performance, although even the 5-section bore had roughly the same form of FFT as the 40-section bore. Cook also tried dividing the bore into unequal sections; the result was a waveform reminiscent of the evenly divided examples, but with a unique harmonic series for each section that manifested itself in a stable, multiphonic sound.

Although Cook's model reportedly worked well, it required quite a bit of computational time. Julius Smith [68] has recently been exploring a more efficient model of the conical bore, using spherical coordinates to describe the spherical wave which propagates. Recall that for a cylindrical section, the real characteristic impedance, unscaled by area, is expressed as:

$$Z_{cyl} = \rho c \quad (3.71)$$

where ρ is the density of air and c is the speed of sound.

Smith describes the characteristic impedance of a conical section as:

$$Z_{con}^+ = \frac{\rho c}{1 - \frac{1}{jkr}} = \frac{Z_{cyl}}{1 - \frac{1}{jkr}} \quad (3.72)$$

$$Z_{con}^- = \frac{-\rho c}{1 + \frac{1}{jkr}} = \frac{-Z_{cyl}}{1 + \frac{1}{jkr}} \quad (3.73)$$

where r is the distance from the apex of the cone.

With this approach, a "finite-element" sampling is unnecessary. A single conical section can extend from output point to output point, with no approximation errors. In the simplest case of the conical bore without toneholes, this reduces to a simple delay line similar to that of the cylindrical bore above. More generally, the cylindrical and the conical waveguide sections can be depicted as shown in Figure 3.12, taken from [68]. These can extend all the way from reed to bell, or cover the distance between two toneholes. The two are quite similar in implementation, save that the pressure output in the conical section must be divided through by r , the distance of the output point from the apex.

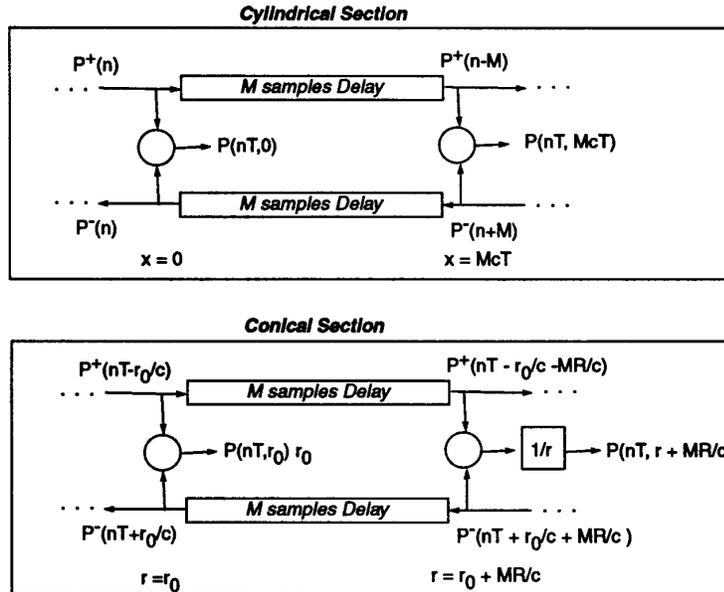


Figure 3.12: Delay Line representation of cylindrical and conical bore segments (after Smith)

The Bell: Models of Radiation and Reflection

The termination of the waveguide bore must model the wave reflection and sound radiation at the bell. To summarize these functions:

- The bell acts as a highpass filter, transmitting high frequencies and reflecting low frequencies;
- The *cutoff frequency* of the instrument, influenced by both bell and toneholes, determines the reflection/transmission proportions. On a well-designed instrument, the cutoff frequencies of toneholes and bell are all roughly equal;
- Very large toneholes approximate a bell, and imply strong reflection and sound radiation primarily at the first open tonehole. Crossfingerings below this hole will have little effect. Thus, a simple bore with only a bell terminator and a time-varying borelength based on desired pitch can be considered to accurately represent an instrument with large toneholes. Smaller toneholes imply weaker reflection at the first opening, and consequently a more distributed radiation pattern as well as more sensitivity to tonehole lattice configuration. A simple bore/bell model alone cannot be used to model the effects of the toneholes in this case;
- The presence of the bell can also affect the amplitudes of the lower and midrange harmonics, as was demonstrated in Section 2.6. This may be seen as its contribution to the overall air column and air column resonances, or as a formant-like filtering operation on the incident pressure wave. More experimental and theoretical work would have to underlay any meaningful model of this effect;

To fulfill the reflection/transmission function, a pair of complementary lowpass/highpass filters serve to terminate the waveguide bore, as shown in Figure 3.13. The right-going wave, P_N^+ , enters both filters. The reflected wave, P_N^- , is given by:

$$P_N^- = -f_{LP}(P_N^+) \quad (3.74)$$

The radiated sound, P_R , which is written to a sound file, is:

$$P_R = f_{HP}(P_N^+) \quad (3.75)$$

Alternatively, the output can be taken as the total pressure at the end of the bore, which is the equivalent of the residual left from the lowpass wave reflection [20], or:

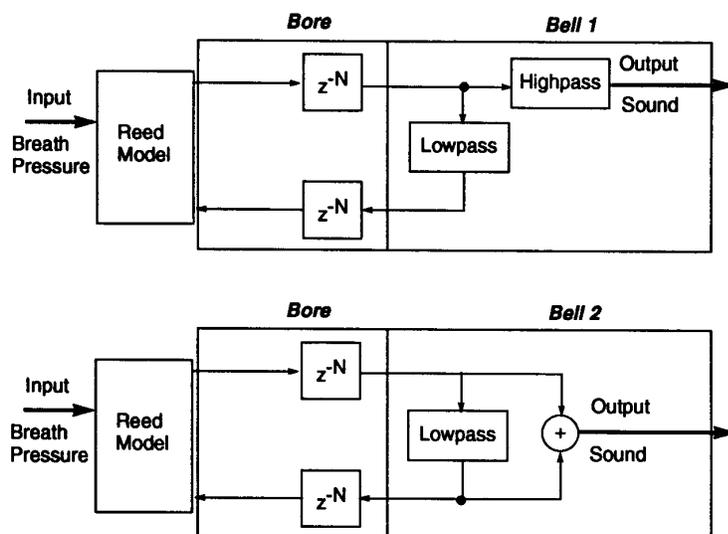


Figure 3.13: Two Representations of the bell: 1) As Complementary Highpass Transmission and Lowpass Reflection filters (after Smith) 2) As Lowpass Reflection filter and residual (from Cook)

$$P_R = P_N^+ + P_N^- = P_N^+ + f_{LP}(P_N^+) \quad (3.76)$$

This formulation is also shown in Figure 3.13. In eliminating the highpass filter, it provides less independent control over the sound radiation; if the reflection/transmission is truly complementary, however, no information is lost. This implementation also fits in nicely with the definition of an "imperfect" pressure node, as described in Chapter 2. At an ideal node, the entire wave will reflect, the pressure will be zero at the end of the tube, and no sound will radiate. The sound radiation is a consequence of incomplete reflection and a nonideal node; the deviation of the pressure at the "node" from 0 is precisely the residual pressure which must be radiated as sound.

Toneholes and Register Holes

The effect of sideholes drilled into the bore is to induce an impedance drop, which, if large enough, will induce a new, shorter standing wave at higher pitch. A waveguide scattering junction serves this function well. Figure 2.43 illustrated the wave behavior at a tonehole junction. Figure 3.7 showed a lossless 2-port scattering junction, which will serve as our starting point for the implementation of a tonehole.

The reflection coefficient can be determined directly from the ratio of tonehole diameter (d_t) to bore diameter (d_b). Let:

$$\gamma \triangleq d_{th}/d_b \quad (3.77)$$

denote the ratio of diameters.

Since the characteristic admittance, Y , is proportional to the hole diameter:

$$Y_t = \gamma Y_b \quad (3.78)$$

Recall that for a 2-port scattering junction, we can define the reflection coefficient in terms of the impedance change:

$$\rho = \frac{r - 1}{r + 1} \quad (3.79)$$

where

$$r = \frac{Z_2}{Z_1} = \frac{Y_1}{Y_2} \quad (3.80)$$

The impedance and admittance of the first section are simply those of the bore, Z_b and Y_b . As Chapter 2 noted, the impedance of the second section is the *parallel* combination of tonehole and bore impedance. But this means that the corresponding admittance is the simple *sum* of Y_b and Y_t . Thus:

$$Y_1 = Y_b$$

$$Y_2 = Y_b + Y_t = Y_b(1 + \gamma)$$

Then the reflection ratio, r , becomes:

$$r = \frac{Y_1}{Y_2} = \frac{1}{1 + \gamma} \quad (3.81)$$

and the reflection coefficient, from Equation 3.79, becomes:

$$\rho = \frac{\frac{1}{1+\gamma} - 1}{\frac{1}{1+\gamma} + 1} = -\frac{\gamma}{2 + \gamma} \quad (3.82)$$

This is the reflection coefficient for implementation in Equation 3.57.

Note that, thus far, loss has not been included in this scattering model, as it must be to implement the tonehole radiation that occurs. To effect this, an attenuating gain could be inserted at the output of the scattering junction to represent that portion of the wave that would have been lost through the hole. Equivalently, a 3-port, rather than a 2-port junction, can be modelled, as shown in Figure 3.14. Smith developed the general equations for an n -port junction in [64]. They will be duplicated here without proof, as the derivation follows the lines of the 2-port junction already discussed and is well covered in [64].

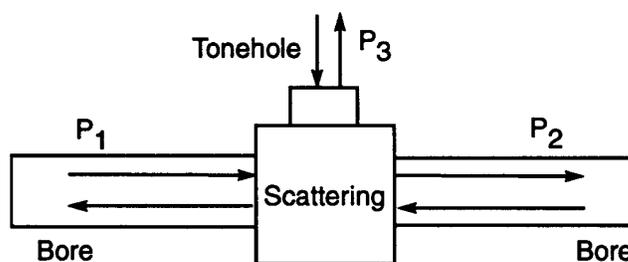


Figure 3.14: Three-port junction representation of a tonehole

The following variables can be defined for the n -port junction:

- *junction admittance*

$$Y_J \triangleq \sum_{i=1}^N Y_i \quad (3.83)$$

- *junction impedance*

$$Z_J \triangleq Y_J^{-1} \quad (3.84)$$

- *junction flow*

$$U_J \triangleq 2 \sum_{i=1}^N Y_i P_i^+ \quad (3.85)$$

In addition, the parameter vector $\bar{\alpha}$ can be defined that linearly combines the N incoming pressure waves into the single junction pressure:

$$\alpha_i = 2Z_J Y_i \quad (3.86)$$

$$P_J = \sum_{i=1}^N \alpha_i P_i^+ \quad (3.87)$$

As we have seen before, the outgoing pressure wave, P_i^- , is equal to:

$$P_i^- = P_J - P_i^+ \quad (3.88)$$

Note that the nomenclature has changed somewhat from that used for the cascaded waveguides. Now the superscript $+$ denotes waves entering the junction, and $-$ denotes waves exiting. In the earlier notation, $+$ entered the junction for the forward delay line, but exited the junction for the reverse delay line. This is so that the notation is compatible with the matrix formulation to follow. At the end of this development, these differences will be reconciled.

Using Equation 3.87, this can be expanded to matrix form:

$$\begin{bmatrix} P_1^- \\ P_2^- \\ \vdots \\ P_N^- \end{bmatrix} = \begin{bmatrix} \alpha_1 - 1 & \alpha_2 & \dots & \alpha_N \\ \alpha_1 & \alpha_2 - 1 & \dots & \alpha_N \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1 & \alpha_2 & \dots & \alpha_N - 1 \end{bmatrix} \begin{bmatrix} P_1^+ \\ P_2^+ \\ \vdots \\ P_N^+ \end{bmatrix} \quad (3.89)$$

This is the generalized *scattering matrix* of the junction. It is worth repeating Smith's derivation of the general 2-port junction, as a familiar example, before going on to the tonehole 3-port junction.

The values of α for the 2-port junction are, from 3.86 and 3.84:

$$\alpha_1 = 2 \frac{Y_1}{Y_1 + Y_2}$$

$$\alpha_2 = 2 \frac{Y_2}{Y_1 + Y_2}$$

Using the scattering matrix, we get the following equations for pressure:

$$P_1^- = (\alpha_1 - 1)P_1^+ + \alpha_2 P_2^+ \quad (3.90)$$

$$P_2^- = \alpha_1 P_1^+ + (\alpha_2 - 1)P_2^+ \quad (3.91)$$

The reflection coefficients, ρ_i , can be defined as the reflected value of P_i^+ if all other $P_j^+ = 0$. Then the values of ρ are related to the α -parameters by:

$$\rho_i = \alpha_i - 1 \quad (3.92)$$

In this case, we get:

$$\rho_1 = \alpha_1 - 1 = \frac{Y_1 - Y_2}{Y_1 + Y_2} = \frac{Z_2 - Z_1}{Z_1 + Z_2} \quad (3.93)$$

This is the same value of ρ we found earlier. The value for ρ_2 can be found by noting that, since

$$\sum_{i=1}^N Y_i \triangleq Y_J$$

the α -parameters must sum as:

$$\sum_{i=1}^N \alpha_i = \sum_{i=1}^N 2Z_J Y_i = 2Z_J \sum_{i=1}^N Y_i = 2$$

Then, in the 2-port case:

$$\alpha_1 + \alpha_2 = 2 \quad (3.94)$$

Using this to solve for ρ_2

$$\rho_2 = \alpha_2 - 1 = 1 - \alpha_1 = -\rho_1 \quad (3.95)$$

Thus, the scattering equations of Equation 3.57 are duplicated.

The procedure is similar for the 3-port junction. Here:

$$\begin{aligned} \alpha_1 &= 2 \frac{Y_1}{Y_1 + Y_2 + Y_3} \\ \alpha_2 &= 2 \frac{Y_2}{Y_1 + Y_2 + Y_3} \\ \alpha_3 &= 2 \frac{Y_3}{Y_1 + Y_2 + Y_3} \end{aligned}$$

These yield the reflection coefficients:

$$\begin{aligned} \rho_1 &= \frac{Y_1 - Y_2 - Y_3}{Y_1 + Y_2 + Y_3} \\ \rho_2 &= \frac{Y_2 - Y_1 - Y_3}{Y_1 + Y_2 + Y_3} \\ \rho_3 &= \frac{Y_3 - Y_2 - Y_1}{Y_1 + Y_2 + Y_3} \end{aligned}$$

Note that ρ_1 , which corresponds to our original single reflection coefficient, can be grouped as:

$$\rho_1 = \frac{Y_1 - (Y_2 + Y_3)}{Y_1 + (Y_2 + Y_3)}$$

If $Y_1 = Y_2 = Y_b$ and $Y_3 = Y_t$, as before, we get the same parallel combination of bore and tonehole. This validates the assumption that the tonehole and bore impedances could be treated as such. Reducing the admittances to Y_b and γY_b , the α -parameters simplify to:

$$\alpha_1 = \frac{2}{2 + \gamma}$$

$$\alpha_2 = \alpha_1 = \frac{2}{2 + \gamma}$$

$$\alpha_3 = \frac{\gamma}{2 + \gamma}$$

These are equivalent to ρ -values of:

$$\rho_1 = \frac{-\alpha}{2 + \gamma} = \rho$$

$$\rho_2 = \rho_1 = \frac{-\gamma}{2 + \gamma}$$

$$\rho_3 = \frac{-2}{2 + \gamma}$$

The pressure equations follow from Equation 3.89. Note that, for a realistic representation, the tonehole port must be connected to some kind of terminator, for example, a lowpass reflection filter. This will provide the function for P_3^+ , which models the piston action of the air in the tonehole, and looks much like a miniature bell. Like the bell, the tonehole radiation function looks like a highpass filter, intensified by the fact that most of the low-frequency energy of the incident wave was reflected back into the bore.

A special case of the tonehole is the register hole, i.e., the venthole which suppresses the fundamental mode and allows the second harmonic to sound. Recall from Chapter 2 that the optimal location of a register hole in a clarinet is 1/3 of the way down the effective bore. For a purely resistive register hole, i.e., one that attenuates the fundamental mode without displacing it, the scattering junction equations become almost as simple as in the 2-port case. Now the register hole is assumed to reflect no energy at all, and P_3^+ is set to 0. The pressure P_3^- represents the energy which is lost through the hole, but does not need to be calculated unless the register hole radiation pressure is desired for output. The scattering equations now become:

$$P_1^- = \rho P_1^+ + (\rho + 1)P_2^+$$

$$P_2^- = (\rho + 1)P_1^+ + \rho P_1^-$$

The radiated pressure is:

$$P_3^- = (\rho + 1)P_1^+ + (\rho + 1)P_2^+$$

Finally, returning to the notation of the cascaded waveguide, with the $i - 1$ 'th and the i 'th sections flanking the register hole, and combining multiplication operations for efficiency:

$$P_{i-1}^- = \rho P_{i-1}^+ + (\rho + 1)P_i^- \tag{3.96}$$

$$P_i^+ = (\rho + 1)P_{i-1}^+ + \rho P_i^- \tag{3.97}$$

$$P_{hole}^+ = (\rho + 1)(P_{i-1}^+ + P_i^-) \tag{3.98}$$

These equations, illustrated in block diagram form in Figure 3.15, are similar in form to the lossless Kelly-Lochbaum junction, but profoundly different because of the sign differences. It would probably be possible to find some sort of efficient computational form, such as the 1-multiply junction, for their implementation.

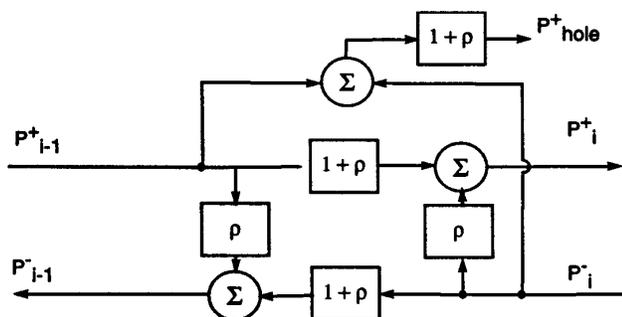


Figure 3.15: Wave scattering at a 3-port junction representing a register hole

Note that theoretically, the upper register should also be attained without any mode suppression at all when the hole, placed 1/3 of the way down the bore, becomes large enough to effectively terminate the air column. Such a hole would result in a fundamental 3 times higher in pitch than the original tone, exactly where the first harmonic lay. However, this cannot be achieved with a purely resistive register hole model, such as that described above.

A full tonehole model, with both resistance and reactance, would be necessary for any kind of realistic fundamental pitch control with a fixed-length delay line.

A register hole was included in the ClariNeXT testbed, and experimental results will follow.

3.3 Digital Reed Models

While the output end of the WDF bore is connected to a simple linear reflection/transmission filter pair, the input end must be connected to a nonlinear element representing the reed. Summarizing from the previous chapter, the reed has the following properties:

- Sensitivity to pressure fluctuations that results in aperture size changes that regulate the input airflow;
- Beating and nonbeating modes of operation;
- Characteristic impedance based upon the instantaneous aperture opening;
- Linear resonant frequency, based upon stiffness, mass, and a combination of internal and embouchure-induced damping, of roughly 10 times the cutoff frequency of the instrument;
- Modified beating frequency based on tip displacement;
- Influenced to some extent by the hydrodynamic forces associated with the airflow through the reed.

The fundamental reed model must provide at least the first property, which allows the critical coupling between bore and reed. The other properties are enhancements which have, in many cases, still not been measured and properly assessed, and are therefore more difficult to model accurately.

3.3.1 Modeling the Reed in a Waveguide Context

Survey of Reed Models

The reed has generally been treated as a simple harmonic oscillator coupling in the manner described above with the bore. In early models, the oscillator was considered to be purely linear; no beating was admitted into the analysis. Backus [5] measured the acoustic mass experimentally and included it in his initial model, but concluded that its influence on the resonant frequency was negligible; Nederveen [51] agreed with this conclusion, although he retained the term in his linear model. Stewart and Strong [70] modelled the reed as a damped, nonuniform cantilevered bar. In his independent work in [58], Schumacher used the traditional harmonic oscillator equation and included as well a differential equation for the flow through the slit. However, his cooperative effort with McIntyre and Woodhouse [43] neglected dynamics completely, and relied on the empirical pressure/flow curve of Figure 2.35, which implicitly includes a beating model, to provide values for the flow into the instrument. Smith, whose work will be developed in the following sections, also used a lookup table, but in a formulation that required no calculation of flow at all.

The WGF Reed as a Time-Varying Reflection Coefficient

The utility of reflection functions in modeling acoustic systems has already been demonstrated. This section develops the adaptation of the reed action to the WGF model, and is based entirely on the work of Smith in [64] et. al..

As we know from earlier sections, the flexible reed will change the size of the aperture in response to pressure fluctuations. From our knowledge of impedance changes and reflection functions:

- When the reed is fully closed, its associated impedance, Z_r , is effectively infinite; The wave will reflect, uninverted, in its entirety. However, the player can inject no additional energy at this point;
 - When the reed is partially open, its impedance will go to a finite, non-zero value. An oncoming wave will partially reflect back into the bore, and partially transmit through the reed, in proportion to the degree of reed closure. For this model, we
-

will assume that the transmitted portion represents an energy loss. In addition, the aperture will admit flow from the mouth, thus introducing added energy. Although one might expect some of this flow to be reflected at the reed/bore junction, consistent with earlier statements that the reed and bore act in parallel, Smith's model considers the two elements to be entirely in *series*;

- When the reed is fully open, none of the oncoming pressure wave is reflected. However, the maximum amount of flow energy is injected at this point by the player.

Intuitively, we can predict a reflection coefficient function based, through its dependency on the instantaneous aperture size, on the instantaneous pressure difference between mouth and bore. In addition, a good formulation would elegantly handle the essentially complementary relationship between wave reflection and energy input.

- Flow into the mouthpiece

As always, the derivation starts with the continuity requirements for pressure and flow. Defining the pressure acting on the reed as P_Δ :

$$P_\Delta = P_b - P_m \quad (3.99)$$

the flow admitted into the the mouthpiece, $U_m(P_\Delta)$, is:

$$U_m(P_\Delta) = \frac{\bar{P}_\Delta}{Z_r(P_\Delta)} \quad (3.100)$$

Note that the reed impedance, Z_r , is indicated as a function of P_Δ .

- Flow into the bore

The flow into the bore follows from the pressure/flow relationships for the left and right-going components of the wave:

$$P_b = P_b^+ + P_b^- \quad (3.101)$$

$$U_b^+ = Z_b P_b^+ \quad (3.102)$$

$$U_b^- = -Z_b P_b^- \quad (3.103)$$

Summing:

$$U_b = \frac{P_b^+ - P_b^-}{Z_b} \quad (3.104)$$

- Continuity of Flow

Assuming a series connection between mouthpiece and bore, the flow into the mouthpiece must equal the flow into the bore. Then, substituting in for P_Δ :

$$\frac{P_b^+ + P_b^- - P_m}{Z_r(P_\Delta)} = \frac{P_b^+ - P_b^-}{Z_b} \quad (3.105)$$

Solving for P_b^- , the returning portion of the wave:

$$P_b^- = \left(\frac{Z_r + Z_b}{Z_r - Z_b} \right) P_b^+ - Z_b P_m \quad (3.106)$$

- Definition of Reed Coefficient

We can now find a term expressing the reflection and energy admission of the system based upon the reed and bore impedances. First, define the impedance ratio, τ , as:

$$\tau(P_\Delta) \triangleq \frac{Z_b}{Z_r(P_\Delta)} \quad (3.107)$$

Then, using equation 3.106 as a model, the reed coefficient, ρ_r , can be similarly defined by:

$$\rho_r(P_\Delta) = \frac{1 - \tau(P_\Delta)}{1 + \tau(P_\Delta)} \quad (3.108)$$

Now equation 3.106 is in the simpler form of:

$$P_b^- = \rho_r(P_\Delta) P_b^+ + \frac{1 - \rho_r(P_\Delta)}{2} P_m \quad (3.109)$$

Notice how both energy input and wave reflection are governed by the same coefficient, $\rho_r(P_\Delta)$. When $\rho_r(P_\Delta)$ is 0, corresponding to a fully open reed, the left-going wave is lost to the system while the maximum portion of the incoming wave is propagated. At the other extreme, a value of 1 for $\rho_r(P_\Delta)$ will result in complete reflection of the left-going wave, but no incoming augmentation. At levels in between, where both the reflected wave and the incoming pressure can contribute to the reflected pulse, augmentation can be considerable.

- Reduction of reed reflection equation to one-multiply form

The above expression is complete, but not efficient, as it contains a subtract and two multiplies. Ideally, we would like a more elegant expression of the form:

$$P_b^- = \hat{\rho}_r()P_b^+ + \hat{P}_m \quad (3.110)$$

where $\hat{\rho}_r()$ is a modified reed coefficient and \hat{P}_m is a precalculated function of the input pressure.

First, define a new term, P_Δ^+ , which is independent of the returning pressure:

$$P_\Delta^+ = 2P_b^+ - P_m \quad (3.111)$$

Solving for P_b^+ and substituting into equation 3.109:

$$P_b^- = \rho_r(P_\Delta) \frac{P_\Delta^+}{2} + \frac{P_m}{2} \quad (3.112)$$

The factors of two can be imbedded in, respectively, the reed coefficient, which as we shall see can be represented by a precomputed lookup table, and the driving pressure. Also, if the reflection coefficient is a function of Δ , it can also be expressed as a function of Δ^+ , so that, for a given P_m , P_b^+ , and P_b^- :

$$\rho_r(P_\Delta) = 2\hat{\rho}_r(P_\Delta^+) \quad (3.113)$$

Rewriting 3.112 accordingly, using P_h to denote half mouth pressure:

$$P_b^- = \hat{\rho}_r(P_\Delta^+) \frac{P_\Delta^+}{2} + P_h \quad (3.114)$$

Thus we have an efficient expression relating the reflected wave to the incident and input waves, based indirectly on the size of the reed aperture. In Smith's original formulation, and Cook's implementation work in [18], a simple lookup table representing reed stiffness and valving, but not reed dynamics, was incorporated. As will be shown in the next section, reed dynamics and other phenomena can be modelled through the calculation of coefficient ρ_r , at the expense, of course, of added computation.

Reed Tables

The previous section demonstrated the representation of the reed action in terms of a time-varying reflection coefficient. The reed model is represented in the way that the reflection coefficient is calculated.

The most direct way to implement the static reed model is by a simple normalized lookup table. The table employed by Smith [64] and Cook [18], illustrated in Figure 3.16, was of the form:

$$\rho_r = \begin{cases} 1.0, & P_\Delta^N \leq P_{closed}; \\ 1.0 - m(P_\Delta^N), & P_\Delta^N \leq P_{open}; \\ 0.0, & P_\Delta^N > P_{open} \end{cases} \quad (3.115)$$

Here, the reed table index, i_Δ^N , is computed as:

$$i_\Delta^N = \text{int} \left(\frac{P_\Delta^+}{f_s} + P_{offset} \right) \quad (3.116)$$

where F_s is the normalizing pressure scale factor and P_{offset} is the pressure bias exerted by the combination of embouchure and reed tip setting.

and the slope of the table is defined by the difference between the breakpoints P_{open} and P_{closed} .

With these parameters, the following reed properties can be represented:

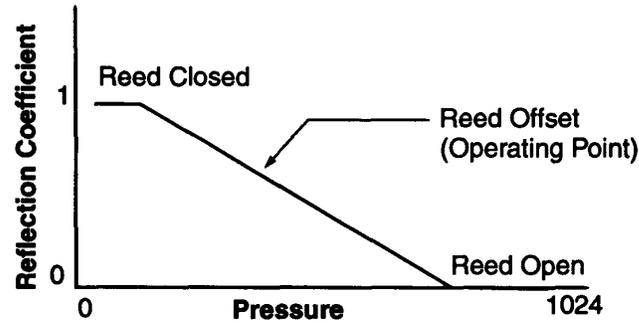


Figure 3.16: Standard Reed Lookup Table

- **Reed Closure:** Reed closure corresponds to a value of ρ_r of 1.0, or everything to the left of the sloped line (since P_{Δ}^+ is equal to $2P_b^+ - P_m$, the larger the value of P_{Δ}^+ , the more the reed is pushed open). The breakpoint, P_{closed} , indicates the pressure at which the reed just begins to open and admit airflow;
- **Reed Opening:** Conversely, the fully open reed is represented by a reflection coefficient of 0, to the right of the sloped line. Its breakpoint, P_{open} , is the point at which the reed just begins to constrict airflow;
- **Reed Stiffness:** reed stiffness, or the sensitivity to pressure fluctuation, corresponds to the slope of the operating line, or alternatively, the pressure difference between P_{closed} and P_{open} . Reed stiffness can also be varied by changing scale factor f_s , which is intended to map from the Soundfile-normalized pressure difference to the 4-byte reed table scale;
- **Reed Position / Offset:** The initial reed position, a function of reed setting and embouchure pressure, results in the value of P_{offset} . The lower P_{offset} , the closer is the reed to being closed in its null position.

This representation cannot incorporate any hysteretical or dynamic properties. However, it can represent the most important nonlinear phenomenon of beating (in this case, necessarily inelastic). The representation is elegant, efficient, and easily parameterized.

Several additional features could be incorporated to the reed table without severely impacting the computational requirements. For example, the nonlinear stiffness of the spring which characterizes deformation over the lay, or uneven closing of the tips in a double reed, can be modelled by reducing the slope as the reed nears closure, and the reflection coefficient accordingly nears 1. A later section will discuss the computation of a Bernoulli-like force which tends to suck the reeds closed: this corresponds to a steepening of the slope near closure. To be even more elaborate and realistic, a table synthesized from experimental data could be incorporated. There would be no additional computational expense, since the table is precomputed. It would be more difficult to parameterize such a table, but with sufficient experimentation to understand how different reed properties affect the curve, a reasonable approach might be developed. Alternatively, the reed program could choose from several different reed tables, each representing a single reed sample. However, this approach is more in the spirit of sampling than simulation.

A feature which would incur additional computation would be to allow hysteresis, i.e., to have a unique table for ascending and for descending P_{Δ}^{\dagger} . The program would have to determine the direction of pressure change. However, effects such as elastic beating and hydrodynamic hysteresis could be incorporated.

Reed Dynamics

A much more computationally intensive approach is to dynamically model the reed. Such a model would allow the reed resonance stabilization of the upper notes discussed by Thompson [72]. The simplest form of this model is the damped harmonic oscillator, generically represented by:

$$\ddot{x} + 2\zeta_r\omega_r\dot{x} + \omega_r^2x = \omega_r^2F \quad (3.117)$$

where ζ_r is the damping ratio, ω_r is the reed resonance, and F is the driving force, $A(P_{\Delta} + P_e)$ where A is area.

In this conventional form of the equation, the force term is scaled by ω^2 in order to provide unity gain at steady state, i.e.:

$$\frac{x_{ss}}{F} = 1 \text{ for } \dot{x} = \ddot{x} = 0$$

For the reed, the oscillator takes the form of a mechanical mass/spring/damper system, as shown in Figure 3.17. For this case, the dynamic equation can be written as:

$$m\ddot{x} + b\dot{x} + kx = F \quad (3.118)$$

where m , b , and k are, respectively, the mass, damping, and spring constant of the reed.

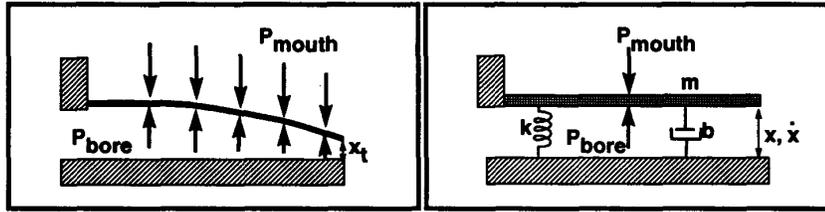


Figure 3.17: The Reed Modelled as a Mechanical Spring/Mass/Damper System (Second Order Harmonic Oscillator)

There is therefore the following relationship between reed static and dynamic parameters:

$$\omega_r = \sqrt{\frac{k}{m}} \quad (3.119)$$

$$\zeta_r = \frac{b}{2\sqrt{km}} \quad (3.120)$$

Rewriting in the form of 3.117:

$$\ddot{x} + \frac{b}{m}\dot{x} + \frac{k}{m}x = f_l(F) \quad (3.121)$$

where $f_l(F)$ is a scaling of the input force which will now be determined.

From equation[3.118], the steady-state position is driven solely by the spring constant and the driving force:

$$F = kx_{ss} \text{ for } \dot{x} = \ddot{x} = 0 \quad (3.122)$$

In order to satisfy this condition, the driving force must be scaled accordingly. From equations[3.121] and [3.117]:

$$\frac{k}{m}x_{ss} = \omega^2 x_{ss} = f_l(F) \quad (3.123)$$

This gives:

$$x_{ss} = \frac{1}{\omega^2} f_l(f) \quad (3.124)$$

From equation[3.122]:

$$x_{ss} = \frac{1}{k} F \quad (3.125)$$

Then:

$$f_l(F) = \frac{\omega^2}{k} F \quad (3.126)$$

Finally, the equation of motion for this model of the reed, in terms of resonant frequency, damping, and stiffness, is:

$$\ddot{x} + 2\zeta_r \omega_r \dot{x} + \omega_r^2 x = \frac{\omega_r^2}{k} F \quad (3.127)$$

An additional relationship is the *damped*, or playing frequency of the reed, ω_d , which represents the player control over the effective reed resonance. This can be expressed by:

$$\omega_d = \omega_n \sqrt{1 - \zeta_r^2} \quad (3.128)$$

Note that the player can only lower the playing resonance (although in the case of beating, the playing resonance can be raised by decreasing tip displacement, as was discussed in Chapter 2). The critical damping ratio, ζ_c , has a value of 1. A higher value indicates

overdamping, and a tendency toward slow response. A lower value indicates underdamping, and a tendency toward overshoot, and resonant response to discontinuities in the forcing function such as collisions against the lay. In fact, the usual range for a control system is between 0.4 and 0.8 [53]; this turned out in the simulation to follow to be a good working range for the reed as well.

As an example, consider a reed with a free resonance of 2620, which is 10 times the frequency of C-262. Thompson, in [72] specified the adjustable resonance range as 2000-3000 Hz, so this is a reasonable number. Note that the “free” resonance, which can be measured by plucking the reed blades, is already internally damped by the inherent reed fiber friction, so that ω_r is already below the true natural frequency. In order to play A-440, the player will want to adjust the reed frequency to a harmonic of that note. The closest applicable value to 2620 is 2200, the fifth harmonic. The required damping ratio can be found by solving 3.128 for ζ_r :

$$\zeta_r = \sqrt{1 - \left(\frac{\omega_d}{\omega_r}\right)^2} \quad (3.129)$$

Here, the value of ζ_r is therefore 0.54, a relatively low number. For the next lower note on the even tempered, diatonic scale, G-392, the sixth harmonic can be used, at 2352 Hz. The player can loosen his embouchure somewhat, to bring ζ_r to 0.44, a change of 10 percent. Alternatively, if he is transitioning from the the more heavily damped note, he may choose to tighten up the embouchure and tune to the fifth harmonic at 1960 Hz. The corresponding ζ_r for this choice is 0.66, now 19 percent higher than before. In this case, the player had to compensate roughly the same amount in each direction.

A final consideration in implementing the reed equation is the boundary conditions. If x is the displacement from the reed equilibrium position, and a_c is the reed closure displacement, then for inelastic collision:

$$\text{For } x = a_c : \dot{x} = 0 \quad (3.130)$$

For elastic collision (which Stewart in [70] claims does not occur, but which Hirschberg claims does [27]), instead of nulling velocity upon collision, the velocity is negated. This also doubles the change in momentum at that point. Then:

$$\text{For } x = a_c : \dot{x} = -k_e \dot{x} \quad (3.131)$$

where k_e is a scalar from 0 to 1 representing the degree of elasticity.

Because of the flexibility of the lips, it will be assumed that the reed is unconstrained at that end of its swath.

The reed displacement, x , maps directly to the reflection coefficient, ρ_r . The dynamic model can be equated to the static model by analyzing the steady state.

- Reed closure displacement

$$\ddot{x} = \dot{x} = 0 \quad (3.132)$$

$$x = a_c \quad (3.133)$$

$$\omega_n^2 x = \omega_n^2 a_c = \omega_n^2 P_{closed} A_r \quad (3.134)$$

$$a_c = -P_{closed} A_r \quad (3.135)$$

This gives the value for the maximum positive offset, a_c . Note that the sign of displacements a and x is opposite that of pressure P . This is because, by the definitions above, x is *positive* toward closure. By convention, however, a positive value of P implies compression, which would push the reed outward in its *negative* direction. Conversely, a negative value of P is a rarefaction pulling the reed closed.

- Free flow displacement

As mentioned above, the reed does not have a maximum outward limit; it is assumed that the lips maintain a constant, nonterminating force. However, the offset at which unrestricted flow begins, although not a limit, can still be expressed in terms of the early parameter P_{open} , as:

$$a_o = -P_{open}A_r \quad (3.136)$$

- Mapping to reflection coefficient, ρ_r

Between the extremes, the piecewise linear mapping is similar to that for the reed table. The total swath, a_s , is the absolute sum:

$$a_s = |a_o| + |a_c| \quad (3.137)$$

Then the general relationship between x and ρ_r is:

$$\rho_r = \begin{cases} 0.0, & x \geq a_c; \\ \frac{x-a_c}{a_s}, & a_c < x \leq a_o; \\ 1.0, & x < a_o \end{cases} \quad (3.138)$$

A lookup table for this relationship could be generated to save a multiply and a subtract. However, given the number of computations involved in propagating the dynamic equations (5 multiplies and 4 additions for a rectangular integration scheme), the savings here is less significant than it was above.

The reflection coefficient computed using this approach is really a function of P_Δ and not the desired P_Δ^+ . In other words, the true pressure difference driving the reed depends partially on the value of the left-going wave, as discussed earlier. The only time $P_\Delta = P_\Delta^+$ is when the reed is completely closed so that $P_b^+ = P_b^-$. Assuming small incremental changes in pressure, ρ_r could be computed in any of the following ways:

- In exactly the same way as the reed table, so P_Δ is considered approximately equal to P_Δ^+ ;
 - Using the current value of P_b^+ and the previous value of P_b^- . This implies a delay error;
 - Using the previous value of P_b^+ , implying a full unit delay in response. Since the reflection at the yielding reed is not truly instantaneous, this may represent reality better than the instantaneous reed table;
-

- Using the current value of P_b^+ and extrapolating from previous values of P_b^- to estimate a “current” value of P_b^- . This approach is most consistent, but requires additional computation.

3.3.2 Nonlinear Refinements to the Reed Model

Both reed models discussed above were in themselves linear, although they represent the nonlinear element in the coupled system. In this section, nonlinear aspects of the reed will be discussed. Note that so far, no distinction has been made between the single and double reed. It is assumed that the basic models are essentially the same, even though the parameters of the models and the subsequent behaviour can be quite different. Even in the nonlinear discussion, the basic phenomena discussed will be assumed to be present in both reed types, while the degree to which they manifest themselves will vary.

Beating

The first phenomenon to be discussed is beating, either of the reed against the lay or two double reeds against each other. This phenomenon was already covered to a certain extent in the previous section.

For the reed table implementation, inelastic beating is implicitly modelled by the constant unity portion of the reflection coefficient, which will result in the complete cutoff of incoming flow. Elastic beating cannot easily be represented in this form, since reed transients are not modelled.

The dynamic reed model does allow both types of beating, by virtue of the choice of boundary conditions at a_c discussed above. Another way to model a varyingly elastic impact would be to apply an instantaneous force of $F_b = 2m_e\ddot{x}$ at the time x reaches a_c , as well as limiting x to a_c . The kinetic energy conserved could be adjusted by the choice of m_e . This is a more physical modeling of what occurs than simply reversing the reed velocity.

Hydrodynamics

The modeling of hydrodynamics is a far more complicated problem, as it requires a knowledge of the actual airflow, an issue that our model has thus far been able to avoid. Although the derivations have all been based on impedance arguments, and impedance is by definition the complex pressure/flow relationship, the actual impedance term, and the related flow, all dropped neatly out of the equations by the time the derivation was complete.

Trying to incorporate flow into the model means stepping into a very gray area of woodwind knowledge. Conventional theory has it that a Bernoulli “lift” force, dependent upon the instantaneous tip separation, acts upon the reed to pull it down increasingly more strongly as the reed approaches the lay. Worman [76] attributed 2 to 3 percent of the total force acting on the reed to the Bernoulli force. Schumacher [58] noted a pronounced difference in results when he switched his Bernoulli model in and out of his simulation. The contention is that the Bernoulli effect would be much more pronounced for the double reed because of the extended length of the reed channel, along which the two reeds lie quite close to one another.

This section will present the traditional theory, offer some modifications, and discuss aspects of its implementation. Incorporation into the two reed models discussed above will be covered in Chapter 4. Uncertainties and imprecisions in the assumptions on which this model is based will be treated as well. For much of the analysis presented here, it is assumed that most of the pressure drop for the instrument occurs at the mouthpiece, so that the the steady-state pressure within the bore is equal effectively to zero throughout. Only the steady-state flow originating at the mouth is treated; the pressure fluctuations in the mouthpiece due to the internal reflection of compression and rarefaction pulses do not factor in.

The so-called “Bernoulli Force” results from the pressure-flow relationship specified in the “Bernoulli Flow Equation” [74]:

$$p + \frac{\rho\nu^2}{2} + \rho gz = P_a \quad (3.139)$$

where p is the applied pressure, P_a is the “stagnation” or ambient pressure, ρ is fluid density, ν is the particle velocity, and gz is the gravity vector which we can neglect in the short

woodwind. This assumes an instantaneous pressure drop to ambient. In the more general case, we can replace P_a with P_c , the pressure of the gas stream in the channel, which for Bernoulli flow is constant everywhere.

The corresponding volume flux, Φ_B , is simply the product of particle velocity and cross-sectional area A :

$$\Phi_B = A\nu = aw\nu \quad (3.140)$$

where a is the tip displacement and w is the width.

The pressure drop, $\Delta p = P_m - P_c$, is then:

$$\Delta p = \frac{\rho\nu^2}{2} \quad (3.141)$$

Application of this equation assumes the following conditions for Bernoulli flow [27][26]:

- moderately high Reynolds number: The Reynolds number, a dimensionless parameter cited in [74] as the most important number in fluid flow, is essentially the ratio of inertia to viscosity. A low Reynolds number implies very viscous flow with significant boundary layer effects. Hirschberg states that for Reynolds numbers below 10, *Poiseuille* flow, rather than Bernoulli flow, occurs. As the Reynolds number becomes higher, the boundary layer thins out. The boundary layer is normally assumed to be thin enough to justify the use of planar wave propagation [51]. At very high Reynolds numbers, above about 1000, the flow starts turning turbulent. [74].
- frictionless flow
- flow separation with no reattachment: this is a function of the viscosity, as represented by the Reynolds number, discussed above.
- steady, incompressible, irrotational fluid flow. This implies that mechanical energy is conserved.

Backus considered the Bernoulli equation to be a valid representation of flow through the slit aperture formed by the reed [5]. Hirschberg agrees that, in the case of the clarinet, the

Bernoulli equation provides an acceptable model [26],[27] as the rapid increase in cross section at the end of the reed channel, and the short length of the channel, assures the necessary flow separation without subsequent reattachment. However, he does include a *contraction coefficient* which accounts for the difference in size between the separated airstream and the reed channel.

It is worth noting that the empirically derived relationships between flow and pressure all take the form of the power law [34]:

$$\Phi = Bp^{\mu_1} a^{\mu_2} \quad (3.142)$$

For the clarinet, Backus found values of $B = 37$, $\mu_1 = 2/3$, and $\mu_2 = 4/3$. These deviate quite a bit from the ideal Bernoulli relationship. The bassoon reed, measured by Nederveen, came closer, with values of $\mu_1 = 1/2$ and $\mu_2 = 1$.

Note that the nonlinear relationship between u and p seems to be at odds with the linear impedance relationship of $Z = P/U$. This is actually not a contradiction, as Backus pointed out in [5], because the value for Z is instantaneously equal to P/U , even if it grows nonlinearly with U .

The “Bernoulli Force” occurs whenever there is a localized increase in particle velocity as a result of a constriction in the flow duct. It plays an important role in the functioning of the vocal cords. Figure 3.18, taken from [29], illustrates the glottis. Here, the volume flux is determined by the separation point. At the constriction, in order to maintain the volume flux, the particle flow must increase. From equation 3.141 above, the increase in particle flow implies a simultaneous decrease in the pressure at the constriction. The additional pressure drop is manifested as a suction force which pulls the vocal folds closer together, intensifying the effect.

A similar force is believed by many to play a role in the functioning of a reed, and especially so in the double reed, where the reed channel is very long and narrow. The generally accepted model for the clarinet reed, introduced by Worman in [76], defines the additional downward pressure on the reed as identically the pressure in the Bernoulli equation (here inverted so the downward pressure is expressed as positive):

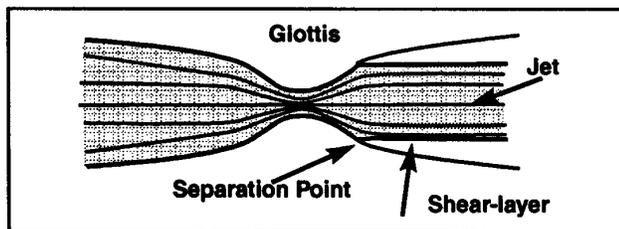


Figure 3.18: Bernoulli Pressure Drop in the Glottis (from Hirschberg)

$$P_B = P_m - P_s = \frac{\rho v^2}{2} \quad (3.143)$$

The reed channel, which was illustrated for the clarinet in Figure 2.4, has a cross-sectional area which varies with the distance along the reed, y . The volume velocity is then defined as

$$\Phi(y) = v(y)A(y) \quad (3.144)$$

where

$$A(y) = wy \tan \theta_r \quad (3.145)$$

The Bernoulli pressure thus acts upon the entire reed, and can be lumped into a single force term by integrating over the length of the reed:

$$F_b = \frac{\rho u^2}{2} \int [A(y)]^{-2} dy \quad (3.146)$$

For the clarinet reed, in which the reed bounds a wedge-shaped area defined by the reed angle, θ_R , Worman determined the area function to be:

$$A(y) = wy \tan \theta_r \quad (3.147)$$

The Bernoulli force in this case is, ostensibly:

$$F_b = \frac{\rho \Phi^2}{2w \tan \theta x(t)} \quad (3.148)$$

This force is dependent on both flow and tip displacement. Worman then substitutes in the experimental relationship for flow described in 3.142 to calculate the Bernoulli force. This derivation has recently been questioned by Hirschberg in [27], because, in assuming a pressure gradient along the length of the reed channel, it implies no flow separation until at least the end of the reed channel. This condition, however, invalidates the use of the Bernoulli's equation at all.

One seemingly inconsistent aspect of this theory is our knowledge that the airflow occurs in the first place as a result of the pressure drop between mouth and ambient, $P_m - P_a$. It would seem that, without any constrictions at all, the Bernoulli pressure is necessarily equivalent to the pressure drop driving the flow, rather than 0. In fact, this is true. The P_m portion of the $P_m - P_b$ force on the reed defined in the preceding sections is actually the Bernoulli pressure in the ideal case, where the volume flux is determined at the tip and the pressure drops instantaneously to the ambient of zero. What is generally referred to as the "Bernoulli Pressure", or "Bernoulli Force", is the difference between this nominal pressure and the pressure derived from the Bernoulli equation. In other words, it is the difference between the true channel pressure and the ambient. For our purposes, then, the Bernoulli pressure is defined as the channel pressure, rather than the pressure drop:

$$P_B = P_c = \frac{\rho v^2}{2} - P_m \quad (3.149)$$

For the nominal case discussed above, since:

$$P_m = \frac{\rho v^2}{2}$$

we have $P_B = 0$ where there are no constrictions.

A different situation occurs if the volume flux is not determined at the tip alone. This is in fact what Worman was assuming when he substituted in the empirical relationship for flux rather than the Bernoulli flux. If the true volume flux is greater than the Bernoulli flux associated with the region in question, then the region represents a constriction in which

the relative pressure is reduced, i.e., the channel pressure drops below the ambient. In the glottis example above, the volume flux, $\Phi_s = A_s \sqrt{\frac{2P_0}{\rho}}$ was determined by the driving pressure and the cross-sectional area at separation. In order to maintain that flux in the constriction, the particle flow had to increase and the pressure drop. Were the flow instead to separate at the narrowest part of the constriction, the volume flux itself would have been lower — $\Phi_c = A_c \sqrt{\frac{2P_0}{\rho}}$, or $\frac{A_c}{A_s} \Phi_s$. The only pressure would be the source pressure, P_0 , with no additional “Bernoulli” contribution.

Similarly for the reed, if the volume flux exceeds the Bernoulli flux associated with the reed tip, there will be an increased particle velocity and downward force in the reed channel. The easiest representation of this situation is one which parallels the glottis, where the separation point occurs somewhere beyond the tip (Hirschberg’s flow visualization on a double reed model indicated clear separation at the tip [30]; this was, however, just a model, and may not have included the downstream conditions). Figure 3.19 illustrates a reed channel where the reed channel is a wedge beginning at the tip and diverging linearly to the mouthpiece entrance. Here, the separation point occurs at separation a_s . The volume flux is therefore, assuming Bernoulli flow throughout the channel and instrument:

$$\Phi = a_s w_s \sqrt{\frac{2P_m}{\rho}} = a_s w_s v_s^2 \tag{3.150}$$

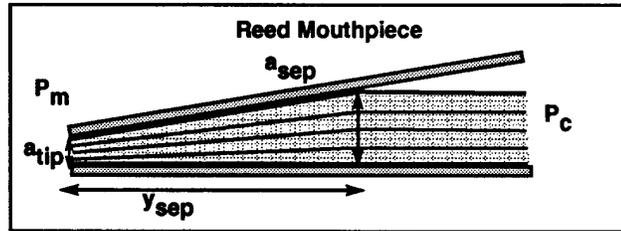


Figure 3.19: Hypothetical Visualization of Bernoulli Flow and Consequent Pressure Drop in a Reed Mouthpiece

At the tip, the particle velocity is:

$$v_t = \frac{\Phi}{a_t w_t} = \frac{a_s w_s}{a_t w_t} v_s \tag{3.151}$$

Assuming that the reed channel is rectangular so that $w_t = w_s$, the associated pressure drop is then:

$$P_t = P_m - P_c = \frac{1}{2}\rho v_t^2 = \frac{1}{2}\rho \left(\frac{a_s}{a_t}\right)^2 v_s^2 \quad (3.152)$$

The Bernoulli pressure acting at the tip of the reed is the difference in pressures between tip and separation point, or the channel pressure if an ambient pressure at separation is assumed:

$$P_B = P_t - P_o = P_c = -P_m + \frac{1}{2}\rho v_s^2 \left(\frac{a_s}{a_t}\right)^2 \quad (3.153)$$

or, since $P_m = \frac{1}{2}\rho v_s^2$:

$$P_B = P_m \left[1 - \left(\frac{a_s}{a_t}\right)^2\right] \quad (3.154)$$

which, for $a_t \ll a_s$, approximates to:

$$P_B = -P_m \left(\frac{a_s}{a_t}\right)^2 \quad (3.155)$$

The equivalent force would be this integral evaluated from tip to separation point, or:

$$f_B = -a_s^2 P_m \int_0^{x_s} \frac{1}{a} dx \quad (3.156)$$

where:

$$a = a_t + \frac{a_s - a_t}{x_s} x \quad (3.157)$$

A similar approach would be to base the volume flux on the experimentally determined flux, as Worman did. This is not particularly easy to justify, as the pressure-flow relationship over the entire instrument is nothing like the Bernoulli relationship, indicating that there is much more complicated flow behaviour occurring. Bearing this in mind, but proceeding anyway, we define the particle velocity in the mouthpiece as:

$$\nu_t = \frac{\Phi}{w a_t} \quad (3.158)$$

where Φ was defined in equation 3.142.

The pressure drop is then:

$$P_m - P_s = \frac{1}{2} \frac{\rho \Phi^2}{w^2 a_t^2} \quad (3.159)$$

and the Bernoulli pressure is:

$$P_B = P_s + P_m = P_m + \frac{1}{2} \frac{\rho \Phi^2}{w^2 a_t^2} \quad (3.160)$$

One example of the incorporation of this pressure into the reed models discussed above will be given in Chapter 4.

Double Reeds

One of the initial goals of thesis research had been to distinguish the double reed from the single reed. This has proven to be a difficult task because of the dearth of experimental data. Although several have studied the reed action in oboes and bassoons, the cross has always been from the cylindrical bore single reed to the conical bore double reed. It is almost impossible to draw conclusions about the contrast between double reed and single reed because of the profound contribution of the boreshape to the experimental results.

Some of the unique properties of double reeds can be isolated:

- **Two dynamic components:** The most obvious distinction between the two reed types is that the double reed has two, rather than one dynamic component. Although the two reeds are related, being the result of creasing and bending a single strip of cane, the two sides can still have different mass and stiffness. One can lump the dynamic properties into a single component that then implies symmetry, but then one loses any effects of asymmetry which might be present in a real instrument.
- **Normal operation in a beating state:** In general, double reeds seem to operate best

in a beating state, although, depending on the spine of the reed, complete closure may not occur. The nonlinear stiffness described for the reed as it curls over the lay likely applies to the double reed as tip closure begins as well. It is quite possible that the beating helps to cancel out the effects of asymmetry described above, by resynchronizing the two reeds with every period of oscillation. In addition, the extra impulse provided by the collision, if elastic, could aid in maintaining the oscillation. It is interesting to note that when D. H. Smith experimented with the oboe single reed described earlier, he found that it too required beating for proper operation. It is difficult to say, however, that this was a peculiarity of the geometry or simple the fact that the threshold blowing pressure was close to the beating pressure.

- **Extended, narrow reed channel:** Unlike the clarinet reed channel, which is very short and terminates in an abrupt change in cross-section inside the mouthpiece, the double reed is relatively long and narrow. In addition, it terminates, not in the bore, but in a staple of much smaller radius which leads into the bore. Because of the long channel and the uncertainty of the downstream conditions, the hydrodynamic behavior is difficult to characterize [27]. It is generally believed that the Bernoulli effect is much more important in the double reed instrument because of its channel geometry, but nothing has yet been proven. The tendency of double reeds to “snap shut” was noted by Benade, who accepted the Bernoulli effect as a given [9]. Even if the Bernoulli force is not the proper model for whatever hydrodynamic forces act on the reed, as Hirschberg claims, it seems clear from experimental evidence that something unique to double reeds is at work. However, since all the double reeds with which Benade experimented were attached to conical bores, which can be a critical factor in reed behavior, such conclusions must be drawn cautiously.
 - **More sensitive to the embouchure.** Although the player has some control of the behavior of the single reed, he has much more influence over that of a double reed. By changing embouchure, he can profoundly change the tone quality and the pitch of the instrument. He has access to both reed resonance and equivalent volume, and can vary the various parameters as he wishes. Again citing D. H. Smith’s experiences with the oboe single reed, he said the most frustrating aspect of the device was the reduced control one had over tone production. Even were the tone quality acceptable, it was still somewhat rigid.
-

Rocaboy found profound differences between the reed behavior of a clarinet and a bassoon [57]. His pressure profiles were presented earlier. For the clarinet, a very symmetric response was found, with an equal amount of reed opening and closure time. The bassoon reed, on the other hand, was closed for a very short amount of time, resulting in a highly asymmetric waveform and its characteristic tone. Rocaboy attributed this behavior to the bore, rather than the reed, however, and claimed the same results for a single reed fitted to a conical bore (quite likely of the same type used with the oboe for this research). Rocaboy's theory is that the air trapped in the reed cavity when closure occurs acts like a spring, forcing the reed back open, and that this is a function of the accelerated air in a conical bore. His discussion is beyond the scope of this paper, which does not deal in general with boreshape. However, it is worth mentioning, because it indicates that some of the characteristics considered unique to the double reed are really unique to the conical bore. It just so happens that in modern instruments, the two often seem synonymous (except for the saxophone, which Rocaboy did not examine).

One interesting feature of Rocaboy's bassoon waveform was the ratio between opening time and closure time for the reed. He observed that the opening time was equal to the double transit time for the wave to propagate to the end of the bore and back, while the brief closure time corresponded well to the time it would take for the wave to travel to and from the fictitious apex of the cone (the acoustic length of a truncated cone is the distance from projected apex, not from the point of truncation, to end). Rocaboy interpreted the short closure time as being a function of the air spring activation. In fact, an explanation more consistent with his observations on the timing of the pulses would be the nature of wave reflection in a conical tube. Recall that a conical bore acts like a tube open on both ends, and supports both odd and even harmonics. This implies that a pulse inverts both at the bell and at the reed. A rarefaction pulse, traveling to the reed, will tend to force the reed closed. But soon after, this pulse will reflect off the apex, invert, and return as a compression, canceling out the effects of its predecessor. The closure time in this case will be equal to the transit time to and from the apex, exactly as Rocaboy noted. If the reed has a large enough offset to bias it toward being open, this one spike will be the only closure time involved. This explanation is a gross oversimplification of the behavior of a spherical wave at the conical apex, which is not at all well understood. However, it is an alternate explanation for some classical conical behavior.

Chapter 4

Interactive Reed Woodwind Modeling Workbench

Chapters 2 and 3 described the basic acoustics of a woodwind instrument. These models were implemented on the NeXT computer in an environment suitable for integrated and flexible testing. The purposes of the digital Workbench were:

- To provide a means of testing the various models that were proposed;
- To test how well the model obeyed various theoretical predictions and empirical observations about how the instrument should work;
- To survey the relative importance of various model refinements with respect to their worth in terms of computational cost;
- To perform interactive parametric studies linking simulation parameters with resulting sound;
- To form a basis for a future instrument implementation on the NeXT DSP chip, for access through the NeXT MusicKit.
- To provide an evolutionary programming structure for future programming enhancements and modeling studies.

The NeXT computer provided an ideal platform for these goals. It was designed to handle

sound files and audio output. Its graphical Interface Builder program allowed an intuitive, flexible user interface. Its object-oriented development tools allowed rapid prototyping, and rapid incorporation of other related tools. Finally, its ability to run many interactive jobs simultaneously allowed the cooperative use of several tools at once without having to link them in a single program.

4.1 NeXT Object-oriented Environment

Before getting into the specific models and interfaces developed for this research, it is worth mentioning the context for which they were designed. This is not intended to be an introduction to object-oriented programming, but simply a discussion of its utility in developmental acoustic research.

The basic element of an object-oriented environment is the *object*. An object can communicate with another object by passing *messages*, which invoke an action on the part of the receiving object called a *method* in NeXT Objective C. Methods can be compared roughly to subroutines, although a true object-oriented programmer would cringe at the thought. Just as one subroutine can call another which might reside in some library that the programmer has never opened, one object can message another object which is totally independent.

The NeXT Interface Builder is a set of interface primitives which the developer can link together and assign objective properties to. For example, a button can be defined to trigger a method in an application object. When the button is "pressed" by the user, the method fires. Code within the method may message a graphical object in the interface to display a result of the calculation. In performing its computations, it may message other interface objects to determine their values. In addition, it may message nongraphical objects to execute some of their defined tasks. For example, the *Play Note* button on the ClariNeXT Workbench triggers the method *runprog* in object *clarinext*, which is defined as the main application. In determining the clarinet parameters, *runprog* interrogates the appropriate graphical objects for user input and uses them to send an initialization message to the reed object, *Reed*, so that it knows to read the graphical objects appropriate to it and set up its constants and reed table. *runprog* then loops through the number of desired samples, iteratively messaging the reed and filter routines. Finally, it writes the simulation results

to the soundfile indicated in another graphical objects, and some summary information in the appropriate form objects.

The workbench was not designed for computational efficiency. Because of the many options to be tested, there are many if-statements and tests necessary that would be eliminated from a streamlined version. Thus, computations take much longer than they would in a dedicated instrument program that used a more rigidly defined set of methods. However, the models themselves were developed with efficiency in mind so that meaningful comparisons among them could be made.

4.2 Modeling Workbench Description

The purpose of this section is to describe each panel in the ClariNeXT workbench, both in terms of its interface and the specific implementation of the element represented. The workbench consists of the following panels, selected from the menu displayed in Figure 4.1:

clarinet2	
Info...	
Edit	e
Play note	p
Go Panel	g
Attack	a
Waveguide	w
Reed Box	r
Bell	b
Spectrum	s
Hide	h
Quit	q

Figure 4.1: Main Menu: ClariNeXT Workbench.

Play Panel

- Main Go button which initiates simulation
- Noise and Vibrato level selection
- Output normalization control
- Soundfile specification

Attack Panel

- Attack envelope specification
- Attack envelope graphical display
- Decay envelope specification

Waveguide Configuration Panel

- Note (Delay Line length) Selection
- Octave drop option
- Register Key Select
- Scattering junction parameters

Reed Box

- Reed Table/Model select
 - Reed Table parameters
 - Reed Table graphical display
 - Bernoulli Pressure select / scaling
-

- Beat elasticity specification
- Reed Model reed resonance calculator
- Reed stiffness specification
- Reed Model embouchure parameters

Bell

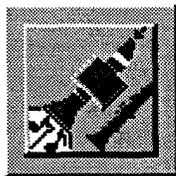
- Bell Reflection Filter type select
- Bell Reflection Filter Cutoff Frequency
- Bell Reflection Filter Gain
- Bell Transmission Filter Select
- Bell Transmission Filter Parameters

Spectrum

- Impulse Response trigger
- Graphical spectrum display (courtesy P.R. Cook)
- Resonance peak frequency and relative gains
- Cooperatios (interpeak ratios)

Certain elements which were not incorporated into this workbench were sound playback and analysis tools. However, the NeXT Sound Editor and Spectro, an interactive spectrum analyzer developed by P. R. Cook were used in conjunction with the Workbench to provide post-simulation capabilities.

4.2.1 Overview of Simulation



The simulation was based on the Smalltalk clarinet developed by Perry Cook, on which he based his waveguide report in [18]. Cook's parameters, which derived from experimental work in his normalized waveguide environment, served as the starting point for the parameters here, rather than any physical parameters. A later section will discuss the equivalence between the normalized simulation parameters and their physical counterparts.

One of the objectives of the normalized clarinet was to allow as much integer operation as possible. Thus, pressures were computed and propagated in integer form wherever feasible. The only exception was in the delay lines of the terminating Butterworth filters. There, the truncation errors degraded performance to such an extent that intermediate filter values were calculated and retained in floating point form.

Figure 4.2 illustrates the flow of the simulation. The basic steps taken, following models described in previous chapters, were:

1. Advance circular buffer pointers and buffer delay line outputs:

$$i_{del} = \text{mod}(i_{del} + 1, N_{del}) \quad (4.1)$$

where i_{del} is the buffer pointer and N_{del} is the number of delays in the line, specified by the user. The implementation of the register key necessitated the addition of an extra delay line, so this procedure was extended.

$$dl_{out} = dl(i_{del}) \quad (4.2)$$

where dl_{out} is any of the delay lines used.

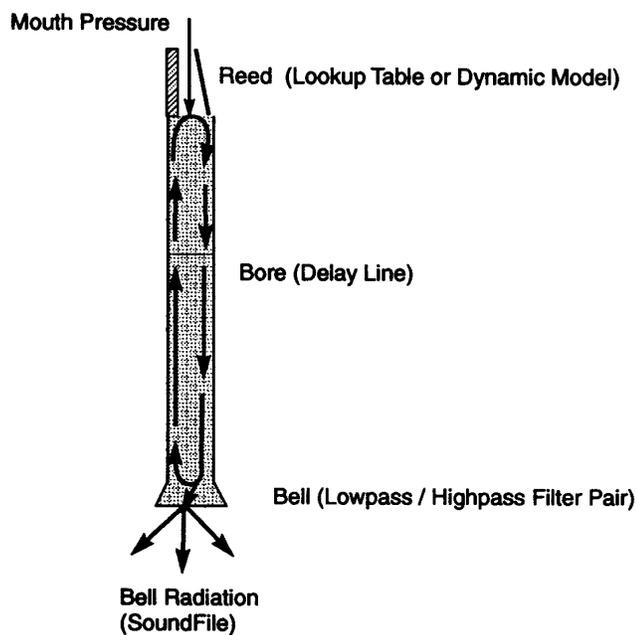


Figure 4.2: ClariNeXT Simulation Flow

2. Compute input pressure based on input envelope, noise and vibrato specifications:

$$P_{input} = P_{env} + P_{noise} + P_{vibrato} \quad (4.3)$$

3. Compute the reflected, left-going wave at the bell (reflection filter) and insert into reverse delay line:

$$P_{reverse}[i_{del}] = \mathbf{reflect}(P_{forwardout}) \quad (4.4)$$

4. Compute the right-going wave reflected at the reed (reflection method takes care of wave inversion) and insert into forward delay line:

$$P_{forward}[i_{del}] = \mathbf{reedReflect}(P_{reverseout}) \quad (4.5)$$

5. Compute the radiated sound at the bell (transmission filter)

$$P_{sound} = \text{transmit}(P_{forwardout}) * K_{sound} \quad (4.6)$$

or alternatively

$$P_{sound} = (P_{forwardout} + P_{forward[idel]}) * K_{sound} \quad (4.7)$$

where K_{sound} is the user-defined output scaling gain.

6. Play the sound and write the data to a Soundfile.

Because the delay line outputs are buffered at the start of each iteration, these steps are commutative. The parameters governing these operations are controlled by the user through graphical interface panels. The remainder of this section will discuss the various aspects of the integrated models and user control over their operation and selection.

4.2.2 Main Panel

The purpose of the main panel is to specify the basic input/output parameters for the run. The input options are:

Noise

Noise is implemented as additive white noise in the input pressure. The user specifies the level of the noise, f_n , as a fraction in the range [0,1] of the nominal input level, P_{in} , so that:

$$P_{in+nz} = P_{in}[1 + (1 - RU()) * f_n] \quad (4.8)$$

where RU is a uniformly distributed random number between -1 and 1.

Chris Chafe and Perry Cook have explored better ways of modeling noise in a musical signal [13],[19].

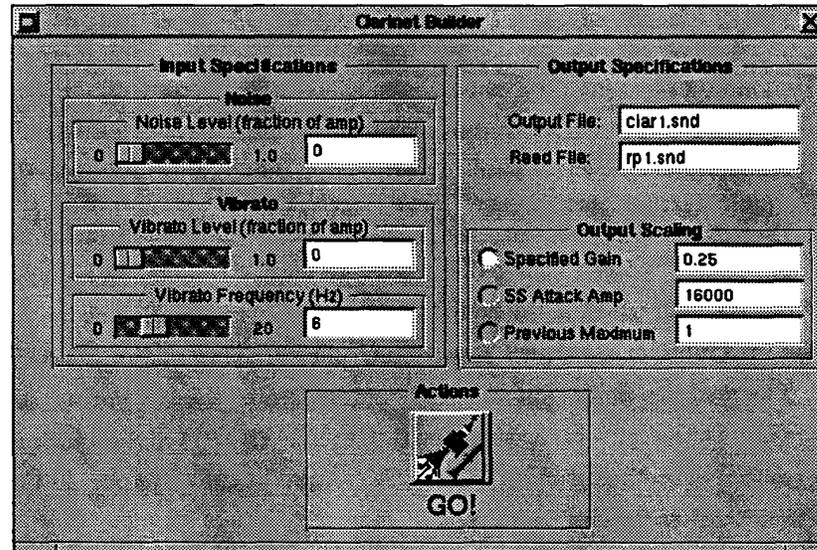


Figure 4.3: Main Panel: ClariNeXT Workbench.

Vibrato

Vibrato is implemented as an additive amplitude modulation of P_{in+nz} . The user can specify both the fractional level, f_v , and the modulation frequency, ω_v . The modulated pressure signal, P_{i+n+v} , or simply, P_{input} is then:

$$P_{input} = P_{in+nz} + f_v P_{in+nz} \sin(i\omega) \quad (4.9)$$

Note that, unlike a frequency domain-based system, the vibrato is implemented via amplitude modulation only, and not through any frequency modulation. This is, of course, the way a player of a real wind instrument produces vibrato - by periodically varying his breath pressure (in recorders there is also a direct frequency modulation through the technique of *flattement*, or finger vibrato; this does not extend to reeds however, except possibly for some types of folk music [10]). Because, for various reasons that have been discussed, the played pitch is amplitude dependent, the resulting tone on a true instrument will be frequency modulated as well as amplitude modulated.

Output File

The NeXT computer provides a convenient Sound class which permits the writing of data to a Soundfile. This Soundfile can be read as such by other programs. Because the ClariNeXT program is intended to be run in tandem with extant sound analysis tools, there is little built-in capability for replaying and editing the simulated sound.

The program actually writes to three different Soundfiles. To the main Soundfile goes the radiation from the bell, which is the Soundfile played upon completion of the run. In addition, the reed position, based on the reflection coefficient computed in the Reed methods, is written. The reed position is automatically normalized to the scale of the soundfile and inverted so that a value of 0 is closed (reflection coefficient = 1) and a maximum value is fully open (reflection coefficient = 0). Finally, the output at the third port in the scattering junction (“tonehole” radiation), is written.

Output Scaling

The ClariNeXT Workbench uses a Soundfile with 16-bit data format and 22050 Hz sampling implying a range of ± 32767 . This can be compared with a normalized input pressure within the range [13000 - 20000]. If an output point exceeds this range, it will saturate the value; the resultant clipping will corrupt both the sound and any subsequent analysis. Therefore, some output scaling options are included to allow the user to normalize the amplitude:

- Constant gain: User-defined gain on the output. This is really equivalent to changing the gain on the transmission filter;
 - Attack amplitude: This option is appropriate for parametric studies which involve the variation of the input amplitude, so that results can be correlated independently of the input pressure;
 - Previous maximum normalization: Because the Soundfile is so large (22050 samples per second of sound), the output data is not buffered before being converted to Soundfile type. However, the program does track the maximum amplitude, **max-amp** as the data are computed, and saves the value. If the same set of parameters is run with the Previous maximum option set, the program will this time scale the
-

output by $\text{maxamp} / \text{SOUNDMAX}$, where **SOUNDMAX** is the desired maximum amplitude. This is the option of choice when coping with numerous simulations where clipping is common.

4.2.3 Attack Panel

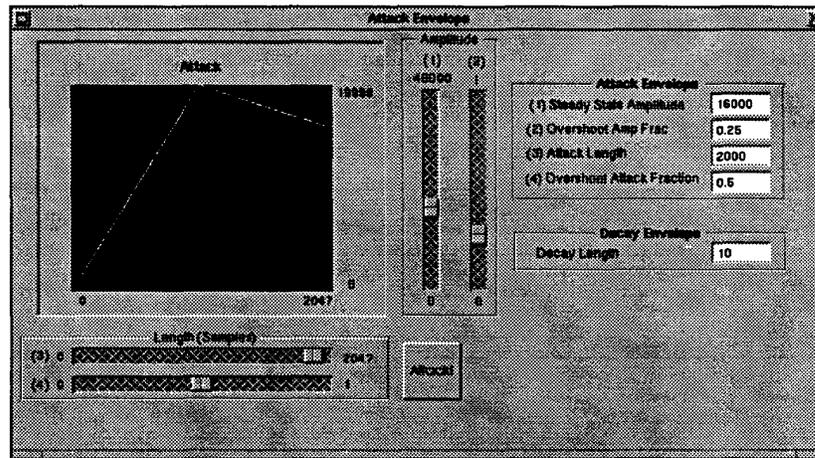


Figure 4.4: Attack Panel: ClariNeXT Workbench.

The transients, and even the distribution of the steady-state modes of oscillation, depend upon the nature of the initial attack. This panel allows the user to set the breakpoints for a two-segment attack envelope, which is then displayed in the slider-bounded window. The attack parameters, which were based on the Perry Cook's results [18] are:

- Steady-state attack amplitude: the pressure level at the conclusion of the attack. Attack amplitudes generally ranged from about 13,000 to 20,000.
- Overshoot amplitude: the fractional midattack overshoot, which produces a *sforzando* effect. This value can be set to 0 to produce a smooth, ramped attack;
- Attack length: length of the attack, in samples. The attack length affects the slope of the attack ramp, and therefore the perceived articulation. The nominal value was set to 2000, or about 0.1 second;

- Overshoot segment length: length of the first attack segment. Also affects perceived articulation.

The resulting envelope is written to an attack lookup table, which is used by the simulation to set the input pressure for the length of the attack. From that point until the point of decay, which is based on the user-defined decay length, the steady-state pressure is used. Finally, the decay table, which is simply a ramp from the steady-state pressure to zero in the time specified by the user, terminates the tone.

4.2.4 Waveguide Panel

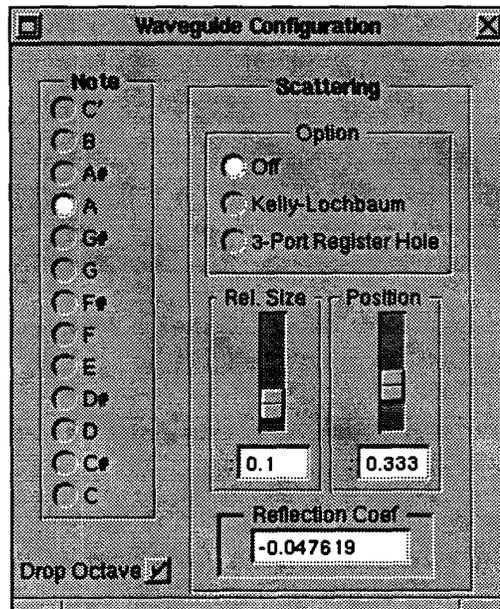


Figure 4.5: Waveguide Panel: ClariNeXT Workbench.

The waveguide panel defines the configuration of the waveguide, which represents the resonating air column. The user can set both the nominal fundamental tone, which defines the overall delay line length, and the register key. Options are:

- Note and Octave Drop: This option sets the length of the delay line, based on an

equal-tempered scale starting at a C with database frequency ω_{base} . The frequency of the note, indexed chromatically from C, is:

$$\omega_{note} = (\sqrt[12]{2})^{i_{note}} * \omega_{base} \quad (4.10)$$

For reasons mentioned earlier, it is often convenient to drop the octave so that there are more elements in the delay line and so that the resulting sound is less grating on the user. This could be done by dividing ω_{base} by 2, but that value is transparent to the interactive user, as it is not expected to be an experimental variable. Instead, the octave drop flag, i_{oct} , is used in the calculation of the delay length.

The number of delay elements is then:

$$n_{del} = (\text{int}) \frac{(\omega_s)}{\frac{4}{1+i_{oct}} \omega_{base}} \quad (4.11)$$

- Register Key/Bore Perturbation: Inclusion of an intermediate scattering junction representing either a single bore perturbation or a register hole at an arbitrary point along the junction. The user can set the position of the hole/perturbation (on or off), the type of junction used (lossless 2-port or 3-port), the relative size of the register hole or perturbed bore with respect to the bore diameter, and the fractional position along the bore, f_{rh} .

Incorporating tonehole capability entailed a major change in the structure of the waveguide, as illustrated in Figure 4.6. Now the forward and reverse delay lines are separated into two parts, joined by the scattering junction. The lengths of these new segments, n_{del0} and n_{del1} , are computed as:

$$n_{del0} = n_{del} * f_{rh} \quad (4.12)$$

$$n_{del1} = n_{del} - n_{del0} \quad (4.13)$$

To implement the scattering junction, an object, Scatter, was defined which had methods for both the 2-port and 3-port scattering junctions defined in Chapter 3. The waveguide propagation steps discussed above now include the intermediate step:

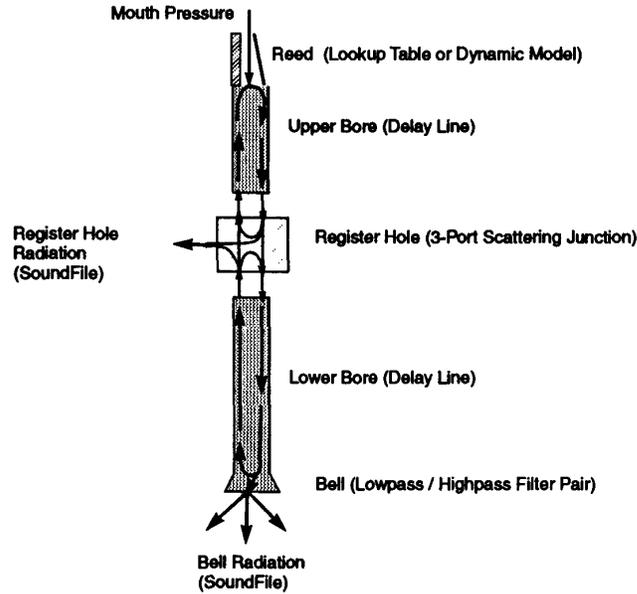


Figure 4.6: ClariNeXT Simulation Flow with Scattering Junction

3a. Calculate the reflections at the scattering junction if the register key or perturbation is set:

if regkey set

$$P_{r0in}, P_{f1in} = \text{scatter}(P_{r1out}, P_{f0out}) \quad (4.14)$$

else

$$P_{r0in}, P_{f1in} = P_{r1out}, P_{f0out} \quad (4.15)$$

The scattering junction was computed following the method described in Chapter 3. If the 2-port junction was selected, the Kelly-Lochbaum reflection coefficient, ρ , was computed as:

$$\rho = -\frac{\gamma}{2 + \gamma} \quad (4.16)$$

where γ is set by the user in the waveguide panel. The panel displays the value for ρ which is calculated. In addition, the scattering coefficients $(\rho + 1)$ and $(\rho - 1)$ are computed *a priori*. The one multiply junction could be used for a more efficient scattering, but has not yet been incorporated.

Rewriting the scattering equations presented earlier to use f and r to denote forward and reverse lines, and *in* and *out* to denote incoming and outgoing paths:

$$P_{out}^r = \rho P_{in}^f + [1 - \rho] P_{in}^r \quad (4.17)$$

$$P_{out}^f = [1 + \rho] P_{in}^f - \rho P_{in}^r \quad (4.18)$$

The three port junction used a purely resistive port to represent the register hole, as described in Chapter 3. The reflection coefficient was the same as for the Kelly-Lochbaum junction, but only the additional value $(\rho + 1)$ was needed for computing the reflections.

In the 3-port case, the scattering equations were:

$$P_{out}^r = \rho P_{in}^f + [1 + \rho] P_{in}^r \quad (4.19)$$

$$P_{out}^f = [1 + \rho] P_{in}^f + \rho P_{in}^r \quad (4.20)$$

In addition, the pressure escaping through the tonehole was computed for program output purposes:

$$P_{out}^{th} = (1 + \rho) P_{in}^f + (1 + \rho) P_{in}^r \quad (4.21)$$

Nominal values for scattering parameters were derived experimentally, and will be discussed in the next chapter.

4.2.5 Reed Box Panel

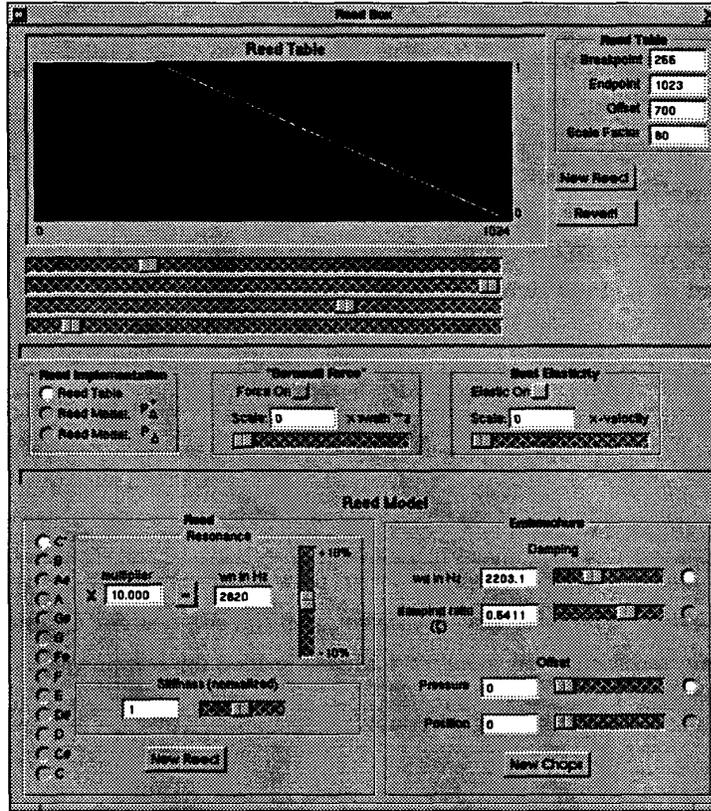


Figure 4.7: Reed Box Panel: ClariNeXT Workbench.

The reed box panel allows the user to define the parameters of the reed. Two implementations are offered — reed table and reed model — as were discussed in the previous section.

Reed Table

The reed table is defined in much the same way as the attack envelope. The user sets, in normalized pressure samples, the closure breakpoint (P_{closed}), the endpoint, or opening breakpoint (P_{open}), the zero-input operating point (P_{offset}), and the scale factor, f_s , which maps from the normalized pressure sample space to "true" pressure. The nominal values,

which were taken from Perry Cook's Smalltalk version [18], are set to:

$$P_{closed} = 255$$

$$P_{open} = 1023$$

$$P_{offset} = 700$$

$$f_s = 80$$

Restating equation 3.116 for convenience:

$$i_{\Delta}^N = \text{int} \left(\frac{P_{\Delta}^+}{f_s} + P_{offset} \right) \quad (4.22)$$

where $P_{\Delta}^+ = P_b - P_m$. Note that this definition of P_{Δ}^+ differs from that described in Chapter 3 by a factor of 2. This is because that factor is already imbedded in the value of the mouth pressure (which is really the half pressure) and in the reed table scale factor.

As Chapter 3 discussed, the user can change the effective stiffness by varying either f_s or the operating slope, which is defined by P_{open} and P_{closed} . The reflection coefficient used in the reflection equation:

$$P_b^- = \hat{\rho}_r(P_{\Delta}^+) \frac{P_{\Delta}^+}{2} + P_h \quad (4.23)$$

is set to 1 for pressures less than P_{closed} and 0 for pressures less than P_{open} .

Reed Model

The dynamic reed model implements the equation:

$$\ddot{x} + 2\zeta_r\omega_r\dot{x} + \omega_r^2x = -\frac{\omega_r^2}{k}\Delta P \quad (4.24)$$

This is a particularly convenient form because of the implied steady-state equivalence of displacement, x , and negated pressure, $-P$ (normalized area is assumed, so that $\int PAdA = P = F$) for a unity value of stiffness k_r .

The dynamic equations were implemented in state space form with simple rectangular integration as follows:

1. Set up first order differential equations

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= \frac{\omega_n^2}{k_r}(-\Delta P + P_{offset}) \end{aligned}$$

where ΔP is the pressure drop across the reed. Note that ΔP must be negated to account for the definition of positive tip displacement for closure, and positive pressure for compression.

2. Propagate states

$$\begin{aligned} x_1 &= \int_t^{t+\Delta t} x_1 dt \approx x_1 + x_1 \Delta t \\ x_2 &= \int_t^{t+\Delta t} x_2 dt \approx x_2 + x_2 \Delta t \end{aligned}$$

3. Limit tip excursion on positive side for beating

$$x \geq a_{closed} \text{ and } x_2 > 0.0 \text{ then}$$

$$x_1 = a_{closed}$$

$$x_2 = -k_{elast}x_2$$

The model allows for elastic or inelastic collisions. The user can set the elasticity constant, k_{elast} , to define the fraction of momentum conserved in the collision. The normal range of this constant is [0 1], where 0 represents total inelasticity, i.e., the reed comes to a complete stop, and 1 represents total elasticity, i.e., the reed bounces off with the same absolute velocity as it had when it collided.

4. Map x_1 to the equivalent reflection coefficient, ρ_r , as described in Chapter 3.

$$\rho_r = \begin{cases} 0.0, & x \geq a_{closed}; \\ \frac{x - a_{open}}{a_{swath}}, & a_{closed} < x \leq a_{open}; \\ 1.0, & x < a_{open} \end{cases} \quad (4.25)$$

5. Calculate the wave reflection

$$P_b^- = \rho_r(\Delta P) \frac{P_\Delta^+}{2} + P_h \quad (4.26)$$

- Driving Pressure

As Chapter 3 mentioned, while the dynamic reed model can be implemented in a manner fully compatible with that of the reed table, so that the pressure determining the reed displacement, ΔP is P_Δ^+ . However, since the dynamic model is more physical than the reed table, this nonphysical term poses a potential problem. An alternative would be to base the driving pressure on the true junction pressure drop, P_Δ , to determine the tip opening and resultant reflection coefficient, while maintaining the use of P_Δ^+ to calculate the actual wave reflection:

$$P_\Delta = P_b^+ + P_b^- - P_m$$

and

$$P_\Delta^+ = 2P_b^+ - P_m$$

Unfortunately, the value of P_Δ is not accurately known, since it relies upon the value of P_b^- , which is the quantity being calculated in the first place. Instead, the program uses the *previous* value of P_b^- , so that there is a sample delay in P_b^- . This is still a more accurate estimate than was P_Δ^+ . A better estimate would result from calculating the rate of change in P_b^- and using the most recent rate to extrapolate to the current expected value of P_b^- . Note that this delay affects only the value of the reflection coefficient, and not of the wave itself, which is still based on the current P_Δ^+ .

- Maximum Resonance (Reed stiffness)

The dynamic reed model allows the user to define both stiffness and damping in terms of the reed resonance. The interface is designed with the assumption that the user is most interested in the harmonic collaboration of the reed with the bore resonances. Therefore, the default values allow the user define parameters in terms of the waveguide options selected.

- Maximum Resonance (Reed stiffness)
-

The nominal reed resonance, in the absence of embouchure damping, is the natural frequency of the reed. The user can set:

1. The playing pitch on which the reed resonance will be based. The default value is C', the highest available chalumeau register pitch on the instrument (actual value depends on the user-defined ω_{base}). The resulting factor is computed, as the playing pitch was, as:

$$\omega_{reed} = \frac{(\sqrt[12]{2})^{i_{res}} \omega_{base}}{1 + i_{oct}} \quad (4.27)$$

2. The harmonic multiplier, f_h which scales the base reed resonance. The default value is 10, which reflects the 10x factor in a clarinet reed.

3. The natural frequency in Hz. This can be calculated automatically as the product of the above two factors or set directly. In addition, the user can use the adjoining slider to vary the frequency within a ± 20 percent range.

- Reed Damping

When the "New Reed" option is selected, the program will automatically compute the damping that will set the damped resonant frequency, ω_d , to an integer multiple of the playing pitch selected in the Waveguide Panel:

$$\omega_d = \omega_{note} f_h \quad (4.28)$$

To handle a register key, which will result in a tone 2.5 times the playing frequency, the equation is modified to:

$$\omega_d = (\omega_{note} + \frac{1.5i_{reg}}{1 + i_{oct}})(f_h - (1 + 3i_{reg})) \quad (4.29)$$

where i_{reg} is set to 1 if the register key is set.

The damping ratio is then computed by:

$$\zeta_r = \sqrt{1 - \left(\frac{\omega_d}{\omega_r}\right)^2} \quad (4.30)$$

The user can also set the either damped frequency or the damping ratio directly. The “New Chops” button will cause the calculation of one from the other, based on which value is selected as the driver in the panel. If the “New Chops” option is not selected, the program will use whatever values are present for ζ_r and ω_d without correlating the two.

- Embouchure Pressure

The initial pressure offset, or alternatively, initial tip displacement, affects the playing of the instrument by changing the mean reed impedance. In the beating regime, the tip offset also determines the swath of the reed, and consequently, affects the beating frequency. Unlike a physical reed instrument, the user can set this parameter independently of damping. Either the offset pressure or the offset displacement can be specified. Again, the “New Chops” button will calculate one from the other in accordance with the user selection. From Equation 4.24:

$$\omega_r^2 x = \frac{\omega_r^2}{k} P \quad (4.31)$$

if

$$\dot{x} = \ddot{x} = 0$$

For k normalized, the nominal condition, offset pressure and displacement are equivalent.

Bernoulli Model

Chapter 3 developed the equations describing the “Bernoulli effect”. The equations for the moveable separation point were incorporated into the simulation. Recall that this models a condition where the reed tip represents a constriction where flow is accelerated and the related pressure drops below ambient. Rather than try to determine the location of the separation point, it was assumed that the pressure would drop off quickly enough that it would be appreciable only along a small, finite segment at the tip. The scaling of the force implies the distribution of the resulting pressure across the entire reed.

For the reed table model of the reed, the effect can be most efficiently designed directly into

the reed table itself. This is consistent with the streamlined approach on which the reed table model is based.

The Bernoulli force for the reed table is based on the reed tip displacement represented by the nominal table. The nominal table, ρ_r can be computed as discussed earlier. The Bernoulli force is of the form:

$$P_B = \frac{K_B}{a^2} \quad (4.32)$$

The scaling constant, K_B , represents both the separation displacement, a_s , and the distribution factor which spreads the point pressure at the tip over the reed.

This can be mapped to a reed table index by:

$$i_B = \frac{(\text{int})P_b}{f_s} \quad (4.33)$$

The idea is to replace the reflection coefficient at pressure point i_p with the reflection coefficient at $i_{p'}$, where p' is the pressure augmented by the Bernoulli pressure. The reed table index can then be adjusted by

$$i_{\Delta}^n = i_{\Delta}^n + i_B \quad (4.34)$$

bounding i_{Δ}^n by 0 and N .

Figure 4.8 illustrates the implementation of such a table. Note that the slope is steeper at reed closure. In addition, the closing point moves to right, reflecting the lower pressure drop required for complete closure. This will bring about the desired effect of an increased closing rate near reed closure. No additional steps need be taken for this implementation.

For the dynamic reed model, the force was calculated based on the current position as:

$$p_B = K_B \frac{a_{swath}}{a_{closed} - x_1} (1 - \gamma) P_m \quad (4.35)$$

The $(1 - \gamma)P_m$ term accounts for the flow restriction due to the reed closure.

and subtracted from the forcing pressure in the dynamic equations.

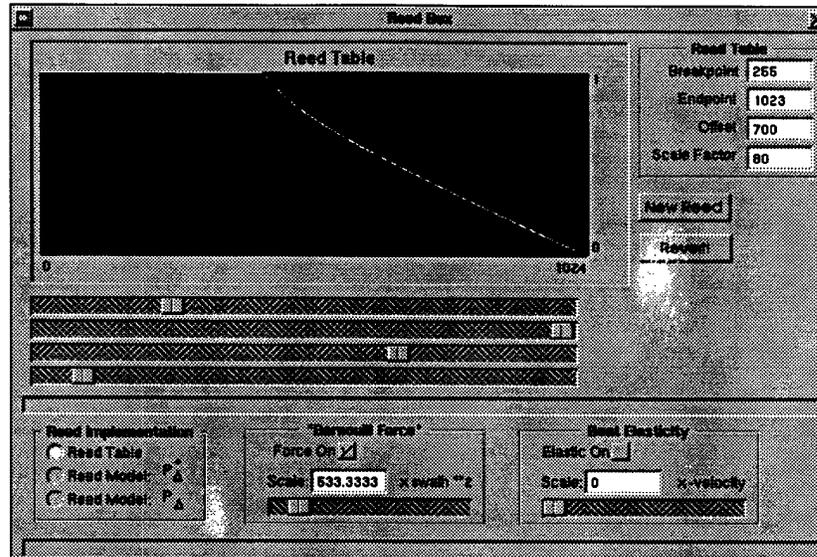


Figure 4.8: Reed Table with Bernoulli Force

4.2.6 Bell Panel

The bell panel determines the implementation of the reflection and transmission filters at the end of the delay line. Because the reflection filter is critical in determining the internal wave behaviour, while the transmission filter affects only the output, the reflection filter was allowed much more complexity. Because of this, the filters do not form the ideal complementary pair. When the radiation was modelled as the identically the pressure at the bore, i.e., the sum of left and rightgoing waves, or equivalently, the unreflected residual of the rightgoing wave, the ideal complementary relationship was restored.

All filters are of the form:

$$\frac{Y}{X} = k \frac{B}{A} \quad (4.36)$$

where Y , X , A , and B are all vectors with length determined by the order of the filter. The delay lines Y and X represent, respectively, the previous outputs and states. The filter

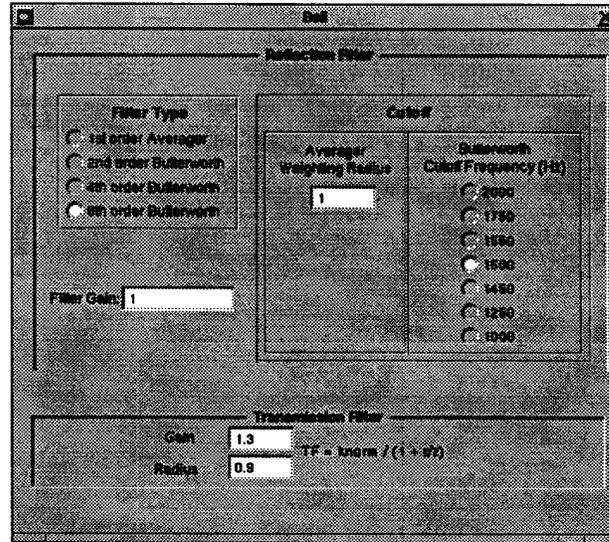


Figure 4.9: Bell Panel: ClariNeXT Workbench.

coefficients, B and A , are based upon the filter selections of the user, and were precalculated using the interactive program MATLAB [42].

The bell reflection is then computed and propagated using:

$$P_{bell} = y_{new} = k(B_0 x_{new} + B_{1,norder} \bar{x}_{past}) - A \bar{y}_{past}$$

- Filter Type

Two basic filter types are offered for reflection: a first order averager, which can be implemented in the most computationally efficient manner; and a Butterworth filter, for which 2nd, 4th, and 6th order versions can be selected. The Butterworth filters are the more physical, and allow much more flexibility in determining the cutoff frequency, at the expense of added computational burden and phase delay.

- Filter Cutoff

For the Butterworth filters, a range of cutoff frequencies from 1000-5000 Hz was offered. In particular, a pair of cutoff frequencies at 1550 Hz and 1450 Hz were defined specifically to

represent the two percent change which Benade found to be so important [9]. The radius of the simple averaging filter can also be adjusted to weight the past value when calculating the average.

- Filter Gain

Nominally, the filter gains are set to 1. However, the user can set the gain arbitrarily to study the effects of losses within the system.

- Transmission Filter parameters

The transmission filter is of the form of a single pole filter:

$$k_{norm} = \frac{k}{(1 - |r|)} \quad (4.37)$$

$$y = k_n x - r y^{-1} \quad (4.38)$$

The user can set both filter gain and weighting radius. The transmission filter gain is not nearly as important as the reflection filter gain, as it has no effect on internal reflection. Rather, it simply scales the output, which may be scaled again by the normalization procedure described earlier.

It is worth repeating that the transmission filter is not considered to be an accurate model for sound radiation. Attention for this research was focused on the internal reflections, and not so much on the actual sound quality. A more comprehensive model would contain the complementary filter pair. Also, a filter pair is necessary only for a cylindrical bore with relatively small bell, where the impedance ratio at the bell is close to infinite. For a conical bore with impedance matched flaring bell, the associated scattering junctions can be extended all the way to the end of the instrument, and no additional filtering is necessary. In the current mode, the sound radiation out of the register hole provides a somewhat more realistic example.

4.2.7 Spectrum Panel

The purpose of the Spectrum panel is to generate the impulse response of the bore and

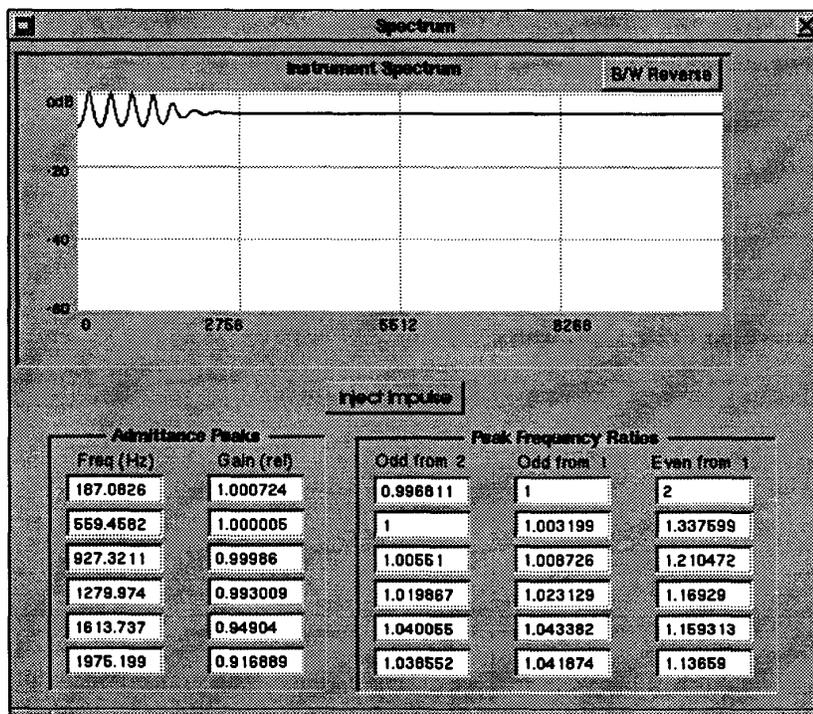


Figure 4.10: Spectrum Panel: ClariNeXT Workbench.

display its spectrum, using the SpectrumView object developed by Perry Cook. The impulse response is calculated using code that parallels the main simulation. This allows the user to compare the bore resonances with the final sound.

The impulse response is generated by:

1. Initializing the first element of the forward delay line to the impulse pressure. The normal input pressure is set to 0.
2. Freezing the reed at its steady state offset, as defined by the reed table or reed parameters. The reed reflection coefficient is thus fixed at one value, so that the nonlinear coupling of the reed does not interfere with the impulse response.
3. Propagating the delay lines as was done earlier, but for only 1024 samples.

4. Shipping the output data to the program spectrumView, which calculates the Hartley transform and plots the resulting spectrum in the panel. The panel also displays the locations and relative sizes of the peaks, which then correspond to the impedance peaks discussed in the literature. A final display in the output is the Peak Frequency ratios, or "cooperatios", which indicate how closely the resonances are aligned with one another. This particular capability will become much more valuable when enhancements to the bore, such as conicity or a tonehole lattice, are incorporated.

4.3 NeXT SoundEditor & Spectro

As was mentioned earlier, the program was designed to be used in conjunction with existing signal processing software on the next, and no effort was made to duplicate extant tools. The most useful were the NeXT SoundEditor, which allows the user to play SoundFiles, as well as cut and paste parts of the waveforms within, and Perry Cook's Spectro program, which computes a cascade of frequency spectra for a time-varying signal by windowing portions of the input array and taking the Hartley transform. These tools were invaluable to the development of the reed workbench, and to the subsequent analysis and simulation work performed with it.

Chapter 5

Results

5.1 Summary of Results

This section reports some of the results obtained with the simulation workbench described in this thesis. To reiterate from Chapter 1, the functions of the simulation were:

- To implement J. Smith's clarinet model in an interactive environment on the NeXT computer;
- To explore the acoustic behavior of a generic reed instrument with the help of this model;
- To assess those refinements in the model which might be useful in a real-time musical instrument.

In general, the simulation performed very well, considering its simplicity. The essential behavior of a reed instrument was duplicated somewhat realistically, although the developers of many earlier clarinet models have been able to claim the same. The register key, which, to the knowledge of the author, was new to this simulation, worked quite nicely; the second register was achieved cleanly and easily. The dynamic reed model, which is not a new concept but is new to this particular WGF formulation, exhibited some interesting, if musically undesirable phenomena, including "squeaking" at both the reed resonance and in the third register of the instrument. This model had some advantages under normal playing

conditions, and offered some aural improvement in sound quality in the non-beating regime; however, when special conditions were imposed, such as elastic collisions and hydrodynamic effects, which resulted in frequent discontinuities such as hard beating against the lay, the model suffered from too slow a sampling rate, and possibly too simplistic an integration scheme. Because of this, it was difficult to aurally assess the success of the modeling of these conditions; the chattering due to the limit cycling of the reed model dominated the sound. However, some results could be garnered by focusing on appearance of the overall waveform, from which the dynamic problems could be filtered out. Although it would benefit from a number of enhancements, the simulation proved to be a useful tool, not only for developing sounds, but for providing a better understanding of the acoustic interactions which occur in a reed woodwind.

5.2 Basic Acoustic Behavior of the Simulation

This section discusses the various aspects of the behavior of the simulation. The specific parameters which served as the default for all the runs to follow are provided in Appendix A. The nominal model provided:

- Sampling rate of 22050 Hz;
 - Reed Table implementation of the reed model, with parameters based on Perry Cook's SMALLTALK simulation;
 - Fundamental tone set approximately to A220;
 - Second order Butterworth filter with a cutoff frequency of 1500 Hz to model reflection at the bell. The radius of the zero in the one-zero transmission filter was set to 0.9.
 - Attack amplitude of 16700 soundfile-normalized pressure units (see Chapter 4), which, for the default reed parameters, was situated in the thick of the non-beating regime, but well above threshold pressure. A 25 percent attack overshoot was provided for the default attack. Total attack time was 2000 samples, or about 0.09 sec;
-

- No noise or vibrato added to the input;
- Most simulations were rerun with output gain normalization to equalize the results for comparison in terms of spectral content and general shape (see Chapter 4).

5.2.1 Observations on A Simple Tone

Figure 5.1 illustrates the overall waveform of a typical tone. In it, one can see the regeneration process develop; the steady state is reached at about 7000 samples (0.32 sec), about 5000 samples (0.23 sec) after the conclusion of the actual attack portion of the input envelope. This means that the transient portion of the tone far exceeds the length of the attack. Zooming in on the attack envelope in Figure 5.2, we see a smooth, nonlinear envelope, starting from the initial bias output from the attack through the repeated augmentation of successive compression pulses. This can be compared to the true clarinet attack that was originally shown in Section 2.2., and which is repeated here for reference in Figure 5.3. The development of the attack was a bit more interesting for the true clarinet because its steady-state waveform, to be discussed presently, was much less smooth than the simulated clarinet, and the blossoming of the higher harmonics is much more apparent.

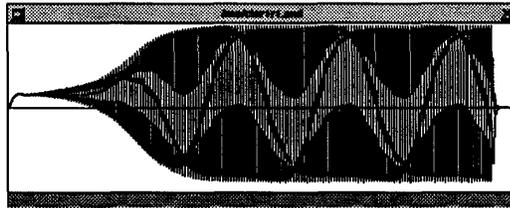


Figure 5.1: Waveform profile for a tone on the ClariNeXT simulation

The steady state portion of the waveform is magnified in Figure 5.4. Although there is evidence of some higher harmonic activity, there is not nearly as much as there was for the true clarinet, whose wave form is also duplicated here, from Section 2.2, in Figure 5.5. Whereas the simulated wave is close to a square wave, the true clarinet exhibits an output closer to the triangular wave. Recall that this triangular shape was evident as well with Cook's clarinet stub, for which the *internal* waveform was measured by fitting a microphone

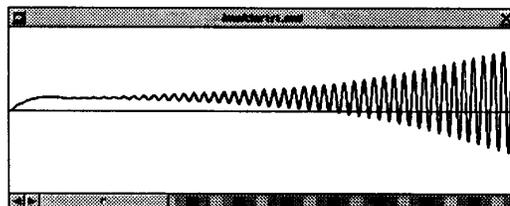


Figure 5.2: Attack profile of a tone on the ClariNeXT simulation

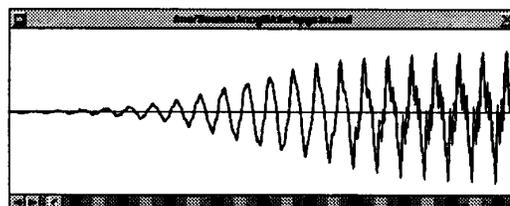


Figure 5.3: Attack portion of true clarinet waveform (Bias removed by recording preamplifier (Sound source: McGill University Master Samples))

in a hole drilled in the side. The transmission filter with which the bell is modelled has some influence on the interpretation of the output, but the correlation between Cook's internal waveform and the McGill clarinet external waveform indicates that both share the major characteristics. This was true with the simulated model as well; at an earlier stage in the development, the internal pressure fluctuations were tapped. It was found that, with the transmission filter described here, the internal and external waveforms had fairly similar shapes. The "ideally complementary" transmission filter, that is, the total pressure at the bell node, actually results in a very different looking waveform, as will be discussed. Even so, it is still convenient to use a filter that mimics the internal wave behavior, since it is the internal behavior that really governs the interactions that characterize the sound. More careful quantitative studies of this behavior would require the tapping of the actual internal waveform, rather than the bell output currently provided.

The spectrum of the bore resonance is illustrated in Figure 5.6. The effects of the cutoff frequency are apparent. Notice, too, that the even harmonics are completely absent, as

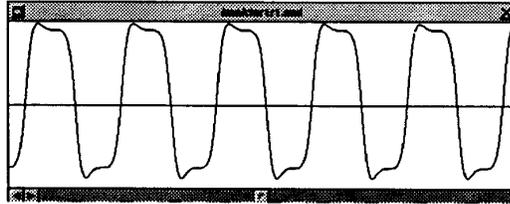


Figure 5.4: Steady-state portion of a tone on the ClariNeXT simulation

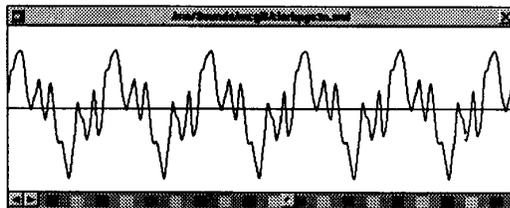


Figure 5.5: Steady-state portion of a tone on a real clarinet

theory would suggest. Figure 5.7 shows the corresponding output spectrum, including an approximation of the peak locations for comparison with the bore resonances. This spectrum has some important features. First, the even harmonics are now present, and the difference between intervening even and odd harmonics decreases as the frequency increases. This is somewhat similar to the true clarinet spectrum, shown again in Figure 5.8. For the true clarinet, the second harmonic was much lower than for the simulation, and succeeding harmonics grew, so that the fourth harmonic is at a level close to that of the simulation, while the sixth harmonic exceeds the level of the simulation.

A second feature, which does not obey theoretical prediction, is the close agreement between the first resonance and the first output harmonic. Chapter 2 noted that the presence of the reed should lower the playing frequency, but here, it appears to have negligible effect. The simulation diverged from experience for other pitch-related phenomena as well. The presence of the bell does affect the frequency of both resonances and output harmonics. The note displayed, whose nominal value was A220, is actually closer to a G#. In order to match more closely, the G# sample of the true clarinet was chosen for comparison.

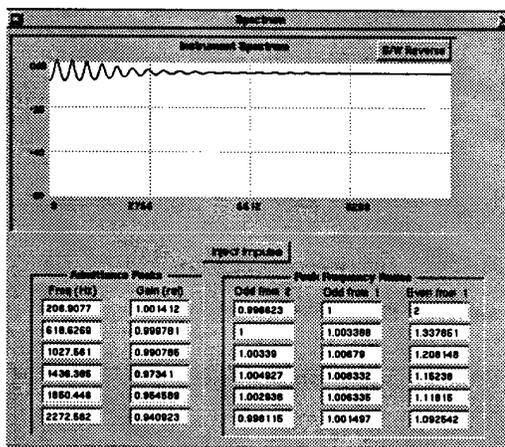


Figure 5.6: Bore resonances on the ClariNeXT simulation

Finally, we can see an attenuation of about 40 dB at the bore resonance cutoff frequency of 1500 Hz (the first peak is about 9dB, so the -30dB line really indicates an attenuation of -39 dB). The spectrum does not exhibit the high frequency behavior of its true clarinet counterpart. However, this is due at least partly to the modeling of the highpass transmission filter discussed earlier.

5.2.2 Influence of the Input

In this section, the sensitivity of the simulation to the input profile is discussed. The influence of the input pressure on the output tone is extremely important for a potential real-time instrument, because the input pressure is one of the musician's primary sources of control during performance. The ability to control tone with breath pressure and articulation is a gesture that, for the most part, is not available with conventional synthesis techniques; a sampler would have to provide a waveform, not only for varying notes, but for varying levels of pressure on each one. Even then, the implementation of mode transition would require nonphysical workarounds.

Both the transient and steady-state phases of the attack envelope were investigated. First,

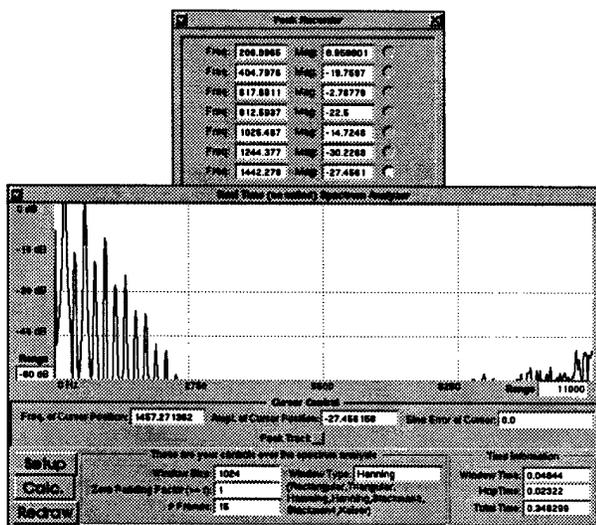


Figure 5.7: Output spectrum of the ClariNeXT simulation

the steady-state amplitude of the input pressure was varied, to determine if the amplitude-dependent generation of harmonics would occur as described by Worman [76] and observed on real clarinets. This also allowed definition of the nonbeating and beating regimes, as well as the total pressure range between thresholds, that is, between where the reed just begins to support the oscillation and where the reed is blown completely shut before oscillation can begin. The next subsection examines the effects of the attack envelope itself. Articulation — the way a musician shapes his attack by tonguing and breath pressure — is critical

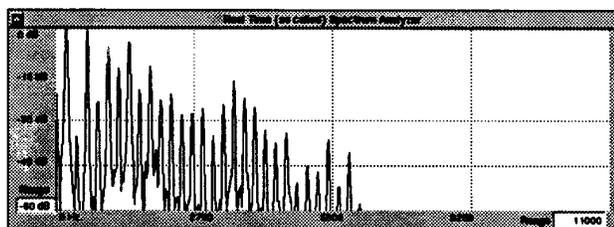


Figure 5.8: Output spectrum of a real clarinet

from a musical standpoint. It is important therefore to understand how the simulation will respond to attack envelope variations.

Sensitivity to Input Pressure Amplitude

To summarize from Chapter 2, a true clarinet exhibits the following input-amplitude dependent behavior:

- Threshold blowing pressure beneath which a sustained oscillation cannot be produced;
- Non-beating regime where increased input pressure generates increasingly higher harmonics;
- Beating regime where increased input pressure generates parallel growth in harmonics;
- Maximum blowing pressure where the reed blows shut;
- Pitch fluctuations with input amplitude fluctuations.

On the first four points, the simulation performed quite well. Threshold blowing pressure, the “blossoming of the spectrum,” the onset of beating, and the blowing shut of the reed, were all attained. For the nominal model, the pressure ranges, in soundfile-normalized ($\times 32768$) pressure units, were as follows (normalized pressure input):

- Threshold Blowing Pressure: 15630
- Onset of Beating: 17400
- Maximum Blowing Pressure: 36200

Figure 5.9 shows the input/output pressure relationship for the default parameters, but with the delay line set for a G rather than an A. Here, the break at the threshold blowing pressure is quite apparent. The initial slope, once a successful oscillation has been achieved, is quite steep until it catches up to the fairly linear portion of the graph which begins roughly at the non-beating region. Often, under the threshold pressure, there was some oscillation, but at a very low level that was far exceeded by the bias portion of the input. Also, the threshold

marked a more precipitous climb in output pressure for some cases than others. Where the threshold pressure was high, the initial slope was extremely steep. Where the pressure was lower, the transition from nonplaying to playing was less defined, with an intermediate oscillation level that was too weak to be considered well-supported tone, but strong enough to produce some small amount of sound. In these cases, it was difficult to define a true threshold.

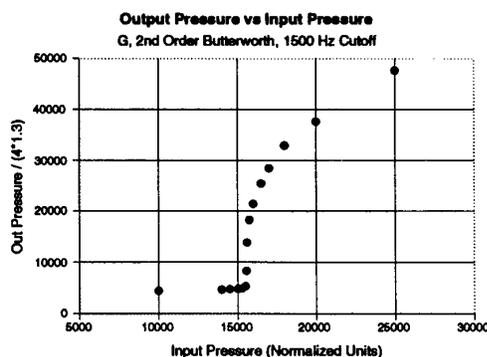


Figure 5.9: Input/Output pressure relationship for typical clarinet simulation. Output pressure has been scaled down by $4 \cdot 1.3$

These pressures are in the integerized, soundfile normalized units of the simulation. However, in terms of relative range from threshold to beating, they agree quite well with Worman's results [76]. He cited a threshold blowing pressure of around 2000 Nt/m^2 , and an additional pressure of around 300 Nt/m^2 to bring the reed to the beating state; this equates to a 15 percent change in pressure between threshold and beating (these were actually analytical results, but he claimed that they reflected experimental results). For the simulation pressures above, the equivalent pressure change is 11.3 percent, which is extremely close to Worman's number; the simulation behaves quite realistically in this regard. No experimental data were available for the upper threshold pressure; many early studies, from which much of the experimental data derived, stopped at the edge of the beating regime.

Figure 5.10 and Figure 5.11 illustrates the normalized wave profile and individual waveforms for three levels of input: near threshold (16000), onset of beating (17400), and moderately strong beating (20000). The associated reed positions, computed from the reed reflection

coefficients, are shown in Figure 5.12. Here, the transition of the reed from nonbeating through hard beating can easily be seen. Note that the steady-state is achieved much more quickly as pressure increases; the third wave profile, for the relatively high input pressure of 20000, approaches the size of the actual attack envelope with its transient of just over 2000 samples. The spectra of the two non-beating cases are shown in Figure 5.13. The generation of harmonics is evident; the lower frequency harmonics will move as little as 2dB, while the higher harmonics grow quite a bit more than that. As could be expected, the tone became brighter as the higher harmonics became more prevalent. The spectrum for the beating case is shown in Figure 5.14, along with the spectrum for a relatively high input pressure of 25000. Note that these two spectra are much more similar than the two nonbeating spectra were, even though the difference in the driving pressure is over three times as large. This follows from the parallel growth of harmonics that occurs once beating begins. The relative sizes of the harmonics should not change, and the normalized spectra should match closely, as they do. Note too that, as the reed beats harder, the even harmonics shrink with respect to the odd harmonics. It appears that the even harmonics are growing weaker, but these are normalized spectra; the odd harmonics are growing stronger.

One aspect of predicted behavior which is absent here is the dependency of pitch on input amplitude. Recall that this effect has been attributed to 1) misaligned harmonics in the nonbeating regime [9], 2) changing tip offset in the presence of elastic collisions [27], and 3) amplitude-frequency dependency for a nonlinear spring [53]. None of these features are present in this nominal case, and so, the pitch dependency cannot be expected. This aberration will, however, be a recurring theme in this chapter.

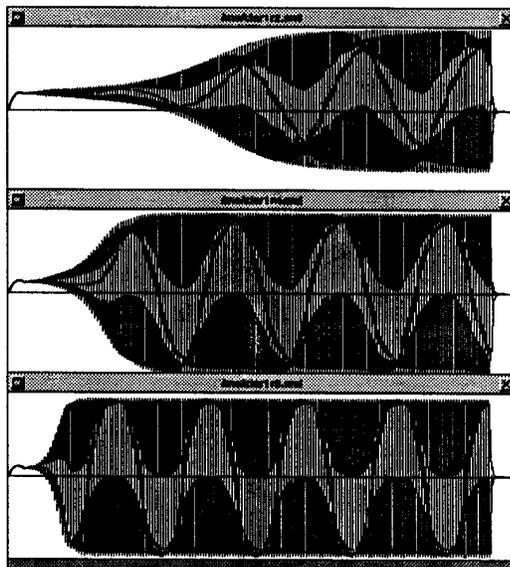


Figure 5.10: Output pressure profile for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)

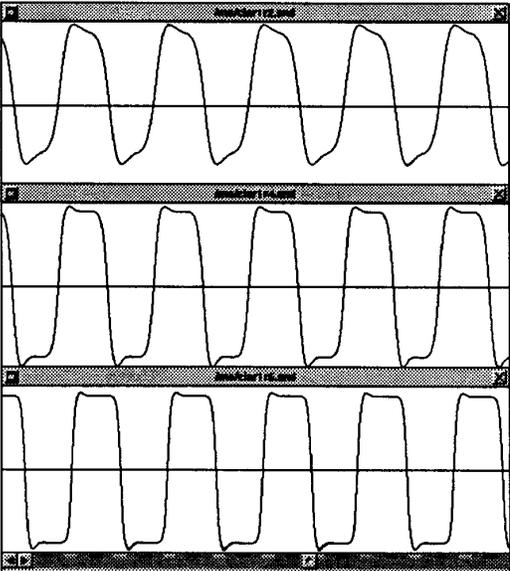


Figure 5.11: Pressure waveforms for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)

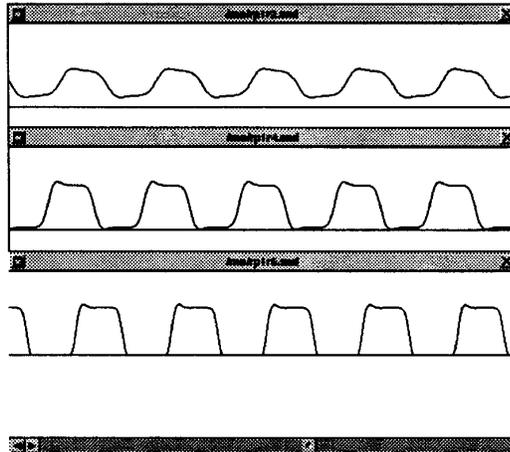


Figure 5.12: Reed Position waveforms for clarinet simulation blown at 16000 (non-beating), 17400 (beating threshold), and 20000 (moderately hard beating)

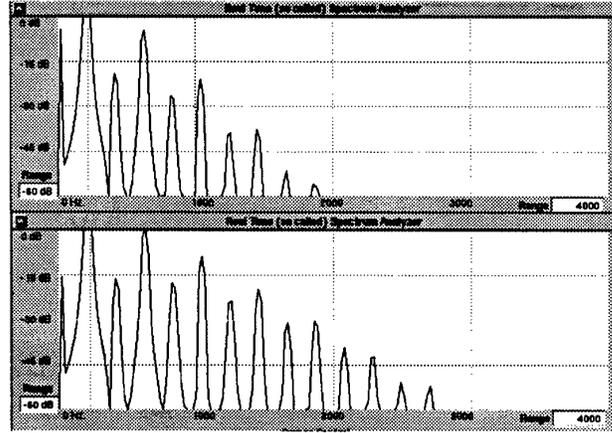


Figure 5.13: Output spectra for nonbeating regime: 16000 and 17400

Sensitivity to Attack

The previous section discussed the influence of the steady state pressure on the waveform and output spectrum. Another issue of great importance in playing a reed instrument is the form of the attack, as defined by the player's articulation. The simulation attack envelope provides for an initial overshoot, or *sforzando* stage, before settling down to the steady state. The default attack envelope, illustrated in the panel description in Chapter 4, specifies a 25 percent overshoot which peaks 30 percent of the way through the attack. For this experiment, both the overshoot slope and the overshoot peak size were varied to ascertain their influence on the steady state waveform.

First, the slope of the overshoot was varied by moving the breakpoint from 0, for an almost infinite slope, to the end of the attack at 2000. The overshoot peak was maintained at 25 percent. The final steady-state showed no sensitivity at all any of the variations. Although the transient, of course, changed, both aurally and visually, the final output pressure converged on the nominal steady-state waveform. Next, the degree of overshoot was changed, from 0 to 90 percent. Again, there was no influence on the steady state.

Although the attack did not appear to have much affect beyond the transient portion of the note in the preceding example, the variations were performed under an extremely stable

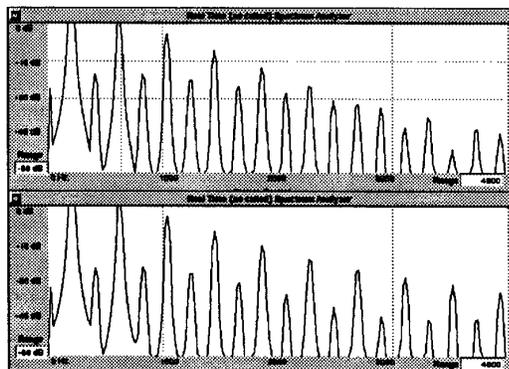


Figure 5.14: Output spectra for beating regime: 20000 and 25000

set of conditions, where there really was only one possible outcome. Where attacks play particularly important roles are in situations where the system is at least bimodal, that is, there is more than one different stable steady-state which can be achieved with the same instrument configuration. A good example of this is provided by certain notes in the upper register on the recorder, which cannot be played without a solid, accurate attack. There are notes on some shawms that cannot be played at all if articulated; the player must slur up or down from a nearby note to excite the desired mode. Misalignment of resonances can also lead to a bimodal condition.

A similar situation exists at the threshold blowing pressure; a strong attack to initiate the note can be followed by a reduced blowing pressure sufficient to sustain it. This case was tested on the ClariNext simulation. The steady-state input pressure was set at 15500, below threshold, even with the 25 percent overshoot. The overshoot fraction was increased from 25 percent to 100 percent. Figure 5.15 shows the radiation profile for the 25 percent and 100 percent overshoot. The increase in overshoot succeeded in exciting the mode, although the resulting sound, dominated by the large *sforzando* at the onset, had a rather unique quality to it, almost like a cross between a pluck and a note blown across the mouth of a bottle. Figure 5.16 shows the transition from attack to steady-state. The harmonics present in the attack die out leaving behind an extremely smooth waveform, as shown in Figure 5.17, along with the corresponding spectrum. The fundamental is quite dominant here, which explains the flutelike quality in the tone.

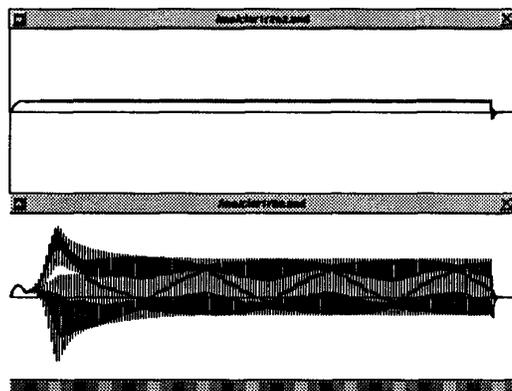


Figure 5.15: Pressure profiles for underblown clarinet simulation with attack overshoot of 25% (top) and 100% (bottom)

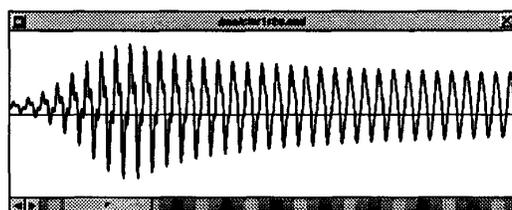


Figure 5.16: Transition from attack to steady-state for underblown clarinet simulation

In summary, the preceding experiment showed that, at least for the ClariNeXT, the attack affected the steady state only in the case of a marginal mode, where the attack could serve to set up an oscillation that could not have otherwise been excited. It must be re-emphasized that the perceived quality of a note is highly dependent on the transient startup phase, regardless of the steady state. These results therefore in no way imply that the attack parameters have no affect on the overall perceived tone, or that the simulation is deficient in this regard. They do, in fact, reinforce the notion of the psychoacoustic importance of the attack, simply because they do decouple the measurable attack from the steady state.

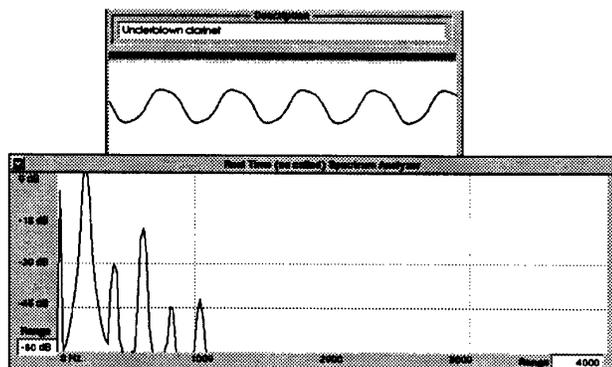


Figure 5.17: Steady-state pressure waveform and output spectrum for underblown clarinet

5.2.3 Influence of the Reed Table Parameters

A series of experiments was performed to determine the importance of the reed table parameters relating to stiffness and tip offset in clarinet operation. The reed table used was the piecewise linear one described in Chapter 4.

Reed stiffness was varied by changing the breakpoint, and the resulting slope of the operating region. Steepening the slope by moving the breakpoint to the right implied increasing pressure sensitivity, and softening the reed. Moving the breakpoint to the left reduced pressure sensitivity and stiffened the reed. In fact, the results were extremely linear. As could be expected, stiffening or softening the reed raised or lowered the output pressure accordingly. However, if the input pressure amplitude was adjusted to compensate for the change in reed stiffness, the waveform produced was almost identical to the original waveform; reed parameter changes were therefore equivalent to input pressure changes.

Figure 5.18 shows the input/output pressures for three different reed stiffnesses: the “nominal” reed, with breakpoint at 255, repeated from Figure 5.9 above; a “soft” reed, with breakpoint at 275; and a “hard” reed with breakpoint at 235 (recall that these numbers are in pressure units of the normalized reed table, with range 0-1024). Note that the three plots show the three different threshold pressures, but converge in the blowing regime. This convergence is somewhat surprising. Certainly the lay serves as an equalizer on the reed closure side, but the softer reed would still be expected to deflect more on the reed opening

side, allowing stronger compressions to develop. If the coupling from the reed is stronger on the closing part of the period, where under beating conditions, all reeds deflect the same amount, then the convergence, based only on the absolute maximum of pressure would be explained. However, the waveforms were biased on the positive, rather than the negative side, which belies this theory.

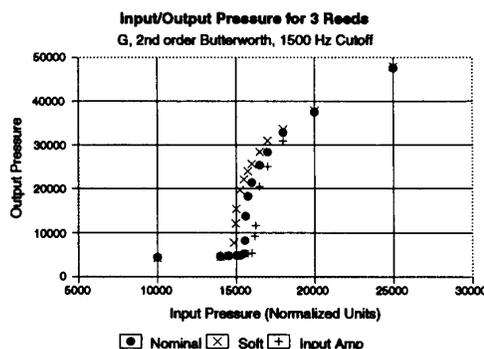


Figure 5.18: Input/Output pressure plots for three stiffnesses of reed in a linear reed table (output pressures scaled down by $4 \cdot 1.3$)

The results of changing the tip offset, or operating point, were also quite predictable. Closing the tip increased the intensity of beating for lower pressures, and opening it decreased the beating. The consequent harmonic generation changed accordingly to match the non-beating/beating regimes. As could be expected, the threshold blowing pressures also changed to match the new conditions. Because of this, the acceptable offset had a finite range which could support a tone. The nominal value was 700, and the table breakpoint at 255 (Recall from Chapter 4 that the total range of the table was 0-1024, representing a quantized soundfile-normalized pressure range of about 60,000). At an offset of 458, the reed was blown shut at the default input pressure level. At an offset of 713, the input pressure level dropped below threshold, so that a note could not be sustained.

In summary, variations in the reed table produced realistic, predictable results. For a linear reed table, the same effects of varying the reed parameters could be obtained by changing the input pressure, which is much easier and more intuitive. In a performance situation, it would probably be much more straightforward to “stiffen” or “soften” the reed by changing

a gain on the input pressure rather than by trying to change the reed table itself. A gain on the output would compensate for any consequent output volume changes.

It must be noted that the reed table does not have to be linear. Julius Smith has experimented with an effective nonlinear stiffness, with good results in terms of tone [63]. If a nonlinear reed table is used, the equivalence between reed parameter changes and input pressure changes no longer holds. In addition, a nonlinear reed table might result in the pitch/amplitude dependence which could not be replicated here.

5.3 Scattering Junction Experiments

This section deals with two implementations of a scattering junction — the first to simulate a register hole, and the second, to simulate a bore perturbation. The register hole was included as a new musical feature, and the bore perturbation as a research tool to explore acoustic behavior further.

5.3.1 Implementation of a Register Key

A register key was implemented following the procedures outlined in Chapters 3 and 4. It worked extremely well, without any additional variation in parameters. The twelfth was obtained cleanly and easily, and in fact was more accurate than the twelfth provided by shortening the delay line by the appropriate amount, which turned out to be a semitone flat if a Butterworth filter was used for the bell. This latter case will be discussed in a later section.

Figure 5.19 shows the bore resonances for a diameter ratio between hole and bore of 0.1. For these notes, the delay line length was set for a low C, so that the played note, a twelfth higher, was a high G'. Note that the first peak is suppressed, but not by much. The output spectrum, however, illustrated in Figure 5.20, shows indisputably that the playing pitch was based on the third harmonic; the first harmonic is completely absent, and all playing harmonics are multiples of the third. Note that the number of harmonics is quite low; recall that when harmonics are widely spaced, there will be fewer of them below cutoff. The resulting waveform is rather smooth, as shown in Figure 5.21. Also shown in Figure 5.21 is the waveform of the pressure escaping through the register hole. Notice how the highpass

nature of the register hole results in a much more triangular waveform than the waveform from the bell. The register hole also sounded a bit more “clarinety”, which is not surprising, since the real clarinet also exhibited a more triangular waveform.

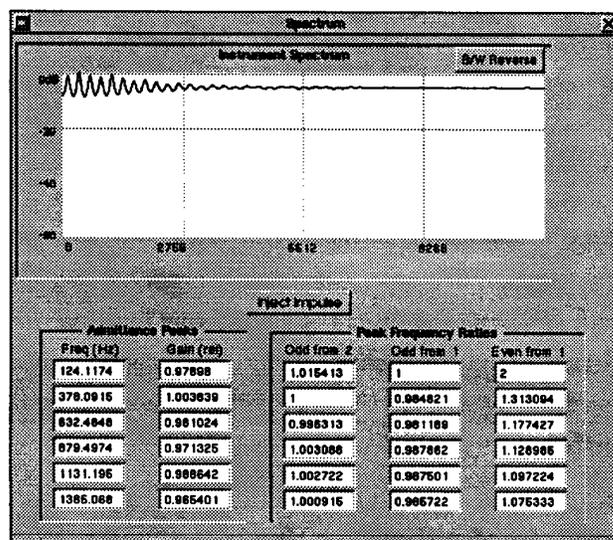


Figure 5.19: Bore resonances for a register hole G' with diameter ratio 0.1

The size of the register hole did not appear critical in achieving the jump, as long as enough input energy was supplied. Since the register hole represents a pure dissipation, its use naturally raises the threshold blowing pressure. For an input pressure of 16700, the acceptable range for the diameter ratio, α , was 0.04 to 0.15 for a base note of low C. A value of 0.02 actually achieved the jump as well, but with a rough, multiphonic sounding startup transient. Raising the base note to low A decreased the upper limit to 0.08, perhaps reflecting the frequency-dependent loss through the hole. Raising the input pressure to 18000 raised the upper limit of the hole size to 0.6.

One interesting effect of the register hole size was its influence on the bore resonance frequencies. Recall from Chapter 2 that Benade stated that a purely dissipative hole should serve primarily to attenuate, and not relocate, the fundamental resonance [9]. Figure 5.22 plots the locations of the first four peaks for a register hole to bore ratio ranging from 0 to 0.5. Figure 5.23 shows the percent change in peak location for the same range. The first

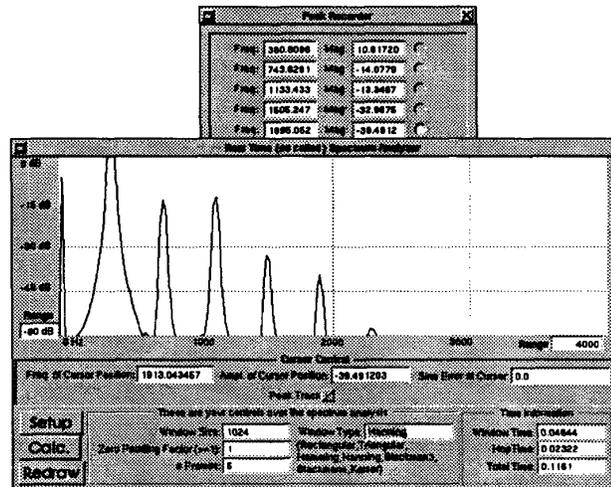


Figure 5.20: Output spectrum for a register hole G' with diameter ratio 0.1

peak was quite sensitive, moving by as much as 16 percent for a ratio of 0.5. The second peak, which will become the fundamental if the input pressure is high enough to sustain the note, moves by only 1 percent for the same range. In both cases, the peaks move downward in frequency. The third peak, on the other hand, moves upward, by as much as 2.6 percent. Since the first peak is not expected to be excited, the third peak should be unimportant in the final tone as well, since it is not a multiple of the second. Finally, the fourth peak exhibits small changes similar to that of the second. Thus, even without any reactance, the register hole has successfully displaced the first mode, largely independently of the other modes. Figure 5.24 shows the bore resonances for the extreme case where α is 0.5. Save for the second mode, the resonance pattern has completely deteriorated.

Surprisingly, the peak locations were not particularly sensitive to the location of the register hole. The first peak in particular barely moved over the working range, although the other peaks moved somewhat more, but still less than 1 percent. However, the amplitudes of the peaks were somewhat sensitive, and the resulting strength of the played note as well. Figure 5.25 shows the output amplitude as a function of the register hole position for a base note of low C. The optimum point in terms of this amplitude was not at precisely one third, or 0.333, as cited by Benade [9], but rather at 0.373. The total acceptable range for

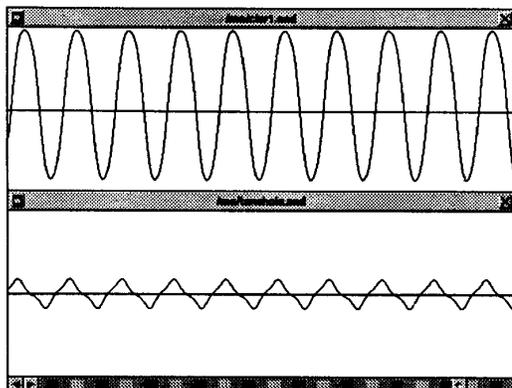


Figure 5.21: Waveform of a register hole G' with diameter ratio 0.1

this note, with a value of 0.1 for α , was from 0.293 to 0.413.

Finally, the register hole note can be compared with the equivalent shortened delay line note. The first aspect of this comparison - that the register hole produced a much better twelfth than the delay line truncation, was already mentioned. Figure 5.26 shows the input and output pressures for a high G achieved through register key, and through delay line length. For comparison, the plot for the low G is repeated. Note that the threshold blowing pressure is higher, and the transition steeper for the register hole version. Also, the linear part of the graph has a greater slope for the delay line than for the register hole, as could be expected because of the added dissipation from the hole. The bore resonances and spectrum of the delay line note are shown in Figures 5.27 and 5.28; the higher harmonics are a bit more pronounced than they had been for the register version presented earlier.

In summary, the register hole worked almost surprisingly well. It was not particularly sensitive to variations in its size and location, so long as the input pressure was high enough to overcome the extra dissipation. It produced an accurate interval of a twelfth, with a clean attack.

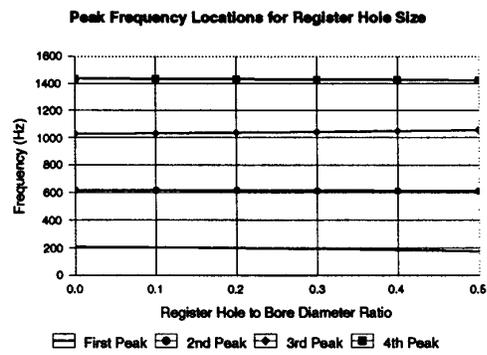


Figure 5.22: Bore resonant peak locations vs register hole diameter to bore diameter ratio (α)

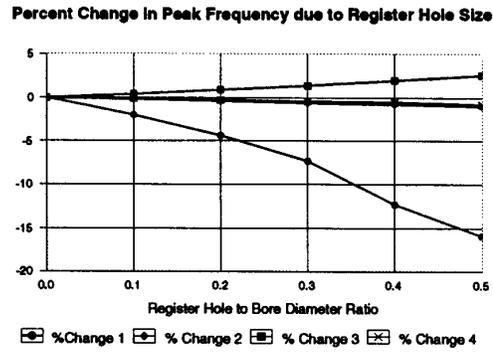


Figure 5.23: Percent change in bore resonant peak locations vs register hole diameter to bore diameter ratio (α)

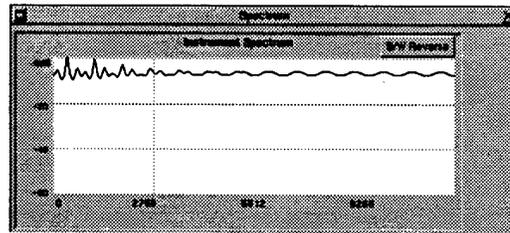


Figure 5.24: Bore resonances for a register hole G' with diameter ratio 0.5

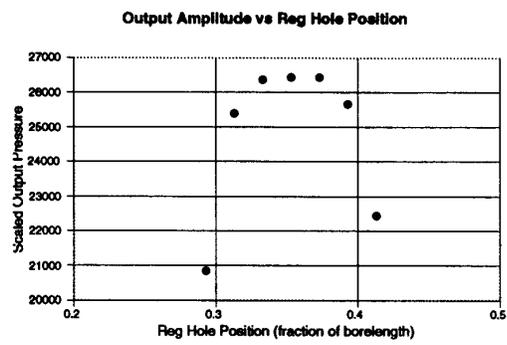


Figure 5.25: Output amplitude as a function of register hole position

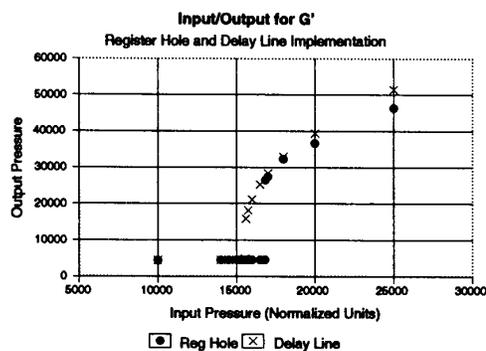


Figure 5.26: Input/Output relationship for a G' achieved through register hole and through delay line truncation (Output pressure scaled down by 4×1.3)

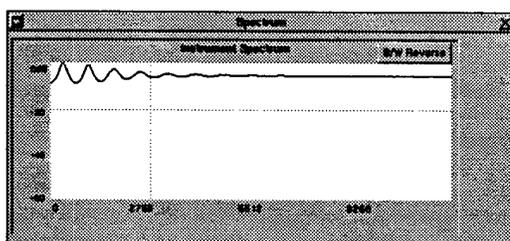


Figure 5.27: Bore resonances for a high G' achieved through delay line truncation

5.3.2 Experiments with Bore Perturbation

One of the difficulties in comparing the simulation results with published experimental data, as well as musical experience, is that the simulation worked too well. With no tonehole lattice, the resonances were extremely well aligned for all notes, and many of the artifacts of marginal modes could not be duplicated.

For this reason, a bore perturbation in the form of a lossless 2-port Kelly-Lochbaum junction was included. It served to misalign the resonances enough so that real instrument phenomena could be explored.

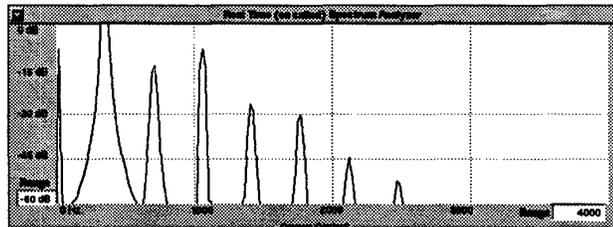


Figure 5.28: Output spectrum of a high G' achieved through delay line truncation

The major subject of interest was pitch dependency on input amplitude. The pitch changes which occur because of the increasing excitation of a misaligned mode as input amplitude increases, described by Worman [76], was discussed in Chapter 2. Figure 5.29 shows a typical resonance spectrum for the perturbed bore. It can be seen that the modes are no longer well aligned. Theoretically, the pitch should vary as the higher harmonics develop. This did not happen. Many parameter variations were tried, but in no cases did the pitch change noticeably with input amplitude, except for cases when a mode transition occurred. It is possible that the pitch changed slightly — this was difficult to measure with the resolution of the output spectrum display — but aurally, there was nothing like the pitch change one gets on a real instrument without adjusting the embouchure. The simple explanation of a compromise between harmonic locations did not hold. This indicates two conclusions: 1) the pitch dependency due to cooperating harmonics as input amplitude fluctuates is more complicated than a simple optimization between peak sizes and locations; and 2) the simulation model is still missing whatever causes this dependency to occur. It is possible that the conical bore contributes to the pitch deviation; Cook reported a distinct pitch fluctuation when he applied input amplitude vibrato to his waveguide saxophone in [18].

Although the bore perturbation was not intended as a musical feature, it actually produced some fairly interesting tones, including multiphonics, mode transitions, and metallic timbres that could be musically useful. For the case of the bore resonances illustrated in Figure 5.29 the result was a sound which started multiphonic, with a base note and an extremely flat twelfth, and eventually made the full transition to the higher note. This mode transition was amplitude dependent; for some levels, the lower note became dominant. For some, the higher note sounded. Figure 5.30 shows the spectrum as it develops from the multiphonic

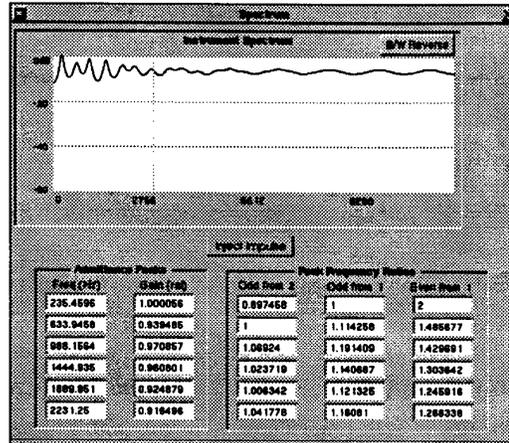


Figure 5.29: Bore Resonances for a perturbed bore

stage to the pure tone. The multiphonic spectrum, in the lower part of the figure, has many peaks, but even the two largest, the second and fourth, are not actually aligned. This is what causes the multiphonic tone to be perceived; if the modes were aligned, one would hear only one note, with its timbre influenced by the presence of the higher harmonic. Figure 5.31 shows the multiphonic portion of the wave form.

The implementation of the bore perturbation would be improved by a better model, particularly one which included the mass reactance at the bell. This would improve its potential for future investigation of resonance and mode transition phenomena. However, even as is, it presents a possible source of alternative musical sound.

5.4 Implementation of the Bell

This section investigates issues of bell implementation. As has been discussed, the bell model consists of two portions — the lowpass reflection filter, and the highpass transmission filter. Because the reflection filter has a strong influence on the fundamental internal wave behavior, while the transmission filter is merely an operation on the output, much more attention was paid to the reflection filter. From a musical point of view, of course, the

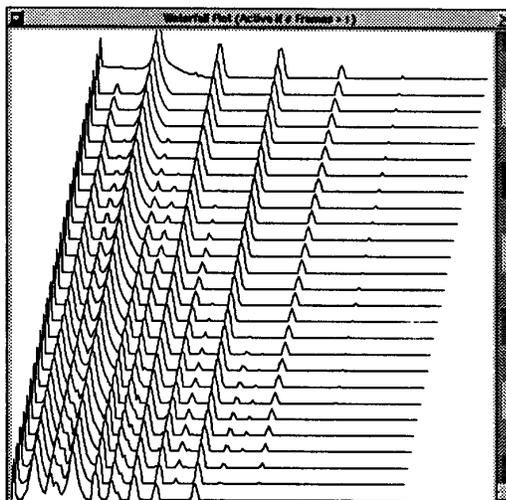


Figure 5.30: Mode transition from a multiphonic tone (bottom) to a pure high tone (top), achieved with a Kelly-Lochbaum bore perturbation

transmission filter is very important, since through it, the actual sound radiates.

5.4.1 Bell Reflection Filter

For the reflection filter, two forms were offered: a simple averaging filter and a Butterworth filter. The Butterworth filter allowed the choice of second, fourth, and sixth order, as well as a number of optional cutoff frequencies. Figure 5.32 shows typical bore resonances for

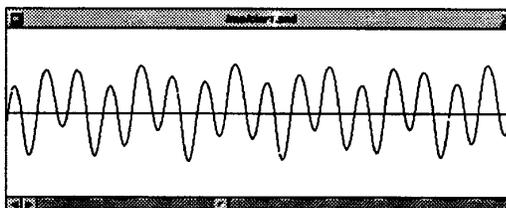


Figure 5.31: Multiphonic waveform, achieved with a Kelly-Lochbaum bore perturbation

the averager and the three orders of Butterworth. The resulting output spectra are shown in Figures 5.33 and 5.34. These were taken for an input pressure of 17000, slightly higher than the default of 16700, to sharpen the attack for the sixth order Butterworth.

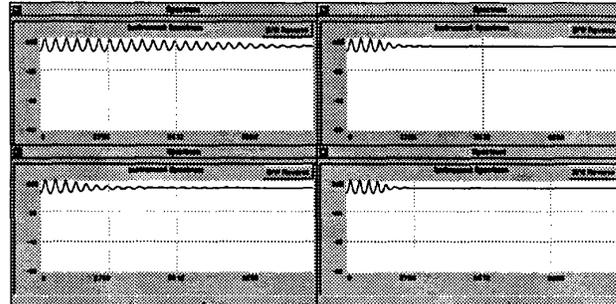


Figure 5.32: Bore resonances for four reflection filters: averager, and 2nd, 4th, and 6th order Butterworths with 1500 Hz cutoff

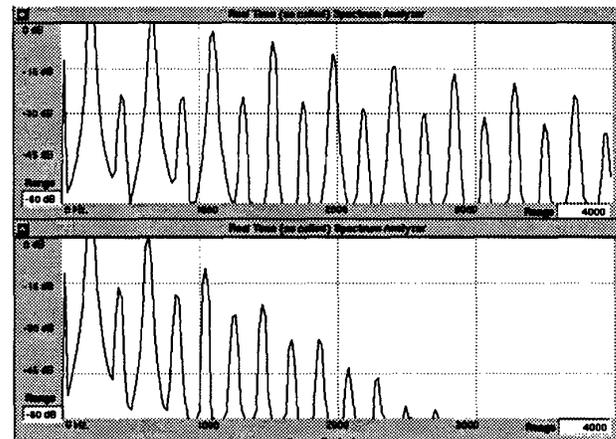


Figure 5.33: Output spectra for averaging reflection filter and 2nd order Butterworth filter reflection filter with 1500 Hz cutoff

The advantage of the averaging filter was its efficiency: it could be implemented with a single 1-bit right shift and add. The Butterworth filter requires more computation, but is much closer to the physical case, representing in a simplistic manner the radiating piston of

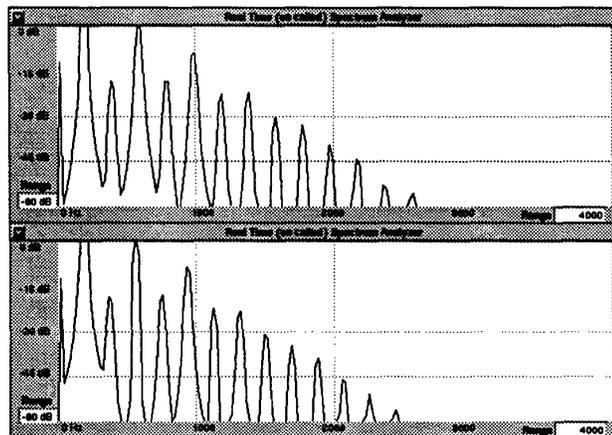


Figure 5.34: Output spectra for 4th and 6th order Butterworth reflection filters with 1500 Hz cutoff

air at the end of the bore. The Butterworth filter also allows direct control over the cutoff frequency, which is a key element in determining tone [9]. The averager, on the other hand, being a first order one-zero filter, does not impose as steep a cutoff, and in addition, does not allow direct control of cutoff frequency. One does have some measure of control through the radius of the zero, which can be varied along the real axis of the z -plane.

The results of the experiments can be summarized as follows:

- The averaging filter produced a very bright, but somewhat clarinet-like sound. The added high harmonics probably contributed much to the very square waveform the simulation produced.
- The Butterworth filter produced a somewhat more muffled tone. This would be expected from all of Benade's work on the importance of cutoff frequency [9], and from general experience in the influence of higher harmonics on timbre. However, the Butterworth filter also allowed much more flexibility in defining the tone, since it offered the parameters of both cutoff frequency and order. Figure 5.35 shows the output waveforms for the second order and fourth order Butterworth filters. As the filter order increases, and the resulting cutoff slope steepens, the highest harmonics

are more attenuated. As the figure shows, this results in an easing of the squareness of the waveform, and, one would expect, a darkening of the tone. It was actually difficult to compare the aural characteristics of the tones resulting from the different orders of Butterworth filter because of the pitch differences. A slightly sharper note tended to sound brighter when compared to a flatter note, even if the flatter note had more pronounced harmonics.



Figure 5.35: Output waveforms for 2nd order (top) and 4th order (bottom) Butterworth reflection filters with 1500 Hz cutoff frequencies

- The Butterworth filter implied a phase delay that translated into a flattening of the pitch; each two orders of filter resulted into approximately a semitone drop for the low A. In addition, the relative tuning for a given filter was corrupted. The interval between C and C', for example was somewhat less than an octave. This could be expected, because the phase delay amounts essentially to the lengthening of the bore, and is a smaller fraction of the total borelength for low notes than for high notes. The averaging filter always yielded the desired note in this range. However, this would not be the case for very high notes or very low sampling rates.

This aberration in pitch introduces a musical problem. If the Butterworth filter is to be used, the tuning defects it causes must be fixed. One obvious option for tuning adjustment would be minute changes in the sampling rate; however, this is not currently possible on the NeXT system, which allows only 22K and 44K sampling. A coarse correction could

be effected by including the phase delay of the bell in the calculation of the length of the waveguide. However, as was previously noted, the length of the delay line itself is by necessity heavily quantized, and in that, is part of the problem anyway. In addition, this solution offsets only one source of tuning deviation. Finer, and more general, tuning could be achieved by appending an allpass filter, such as that described in [31], to correct for any unwanted delays, be they the result of the bell or of other elements in the model [20] [69].

The adjustments described above could be precalculated for each available note, and accessed by the program through a lookup table or parameter list. A more dynamic approach, which would require some work, would be to implement a feedback control loop that adjusts parameters to bring the played pitch to the desired pitch. This approach is inviting only if one wants a fixed scale. In fact, experienced musicians often vary the playing frequency of a given note to fit the musical and harmonic requirements. The same note value will be played differently if it forms part of an interval of a third than if it completes a fourth. A good control system would produce some user-specified equilibrium point, perhaps even-tempered, perhaps meantone, about which the player can bend the pitch to his needs.

The idea of using a feedback control system to regulate instrument output extends beyond pitch. The simulation provides many variables — reed parameters, embouchure, the tuning adjustments described above, bore perturbation, etc. — and many aspects of output quality — loudness, pitch, spectral shape. In a very sophisticated control scheme, the user might define a cost function weighting the various contributors and outcomes. The combination of this and the definition of a set of gestures interpretable by the system could result in considerable artistic freedom. However, it would also require considerable computational time, and is probably not be feasible with the hardware available today.

- The Butterworth filter, particularly the fourth and sixth order filters, tended to interact with the rest of the instrument to produce more transitional modes, particularly at threshold blowing pressures.
 - The threshold blowing pressure required for a clean attack was somewhat dependent on the bell filter used. A sixth order filter required a slightly higher attack pressure.
 - Changes in the cutoff frequency should have produced a brighter or darker tone, according to Benade [9]. However, the cutoff frequency also affected the pitch, because the filter delay increases as the cutoff frequency decreases. A slightly sharper
-

tone, as mentioned above, does also sound a bit brighter. Figure 5.36 and Figure 5.37 shows the spectra for a second order Butterworth with cutoff frequencies at 1000 Hz, 1450 Hz, 1550 Hz, and 2000 Hz.

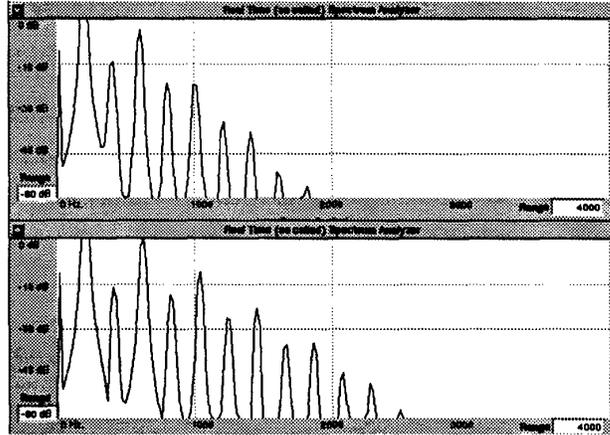


Figure 5.36: Output spectra for 2nd order Butterworth filter reflection filters with 1000 Hz and 1450 Hz cutoffs

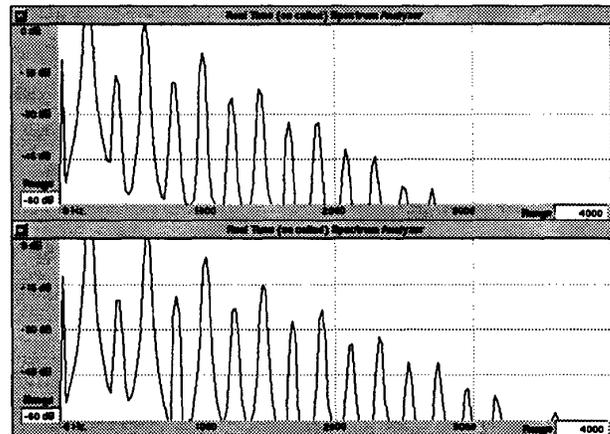


Figure 5.37: Output spectra for 2nd order Butterworth filter reflection filters with 1550 Hz and 2000 Hz cutoffs

5.4.2 Transmission Filter

The transmission filter, which ideally should radiate all the unreflected energy, is generally modelled as a highpass filter [64]. Recall that the wave energy which escapes is but a small portion of the internal wave energy and that the input to the highpass filter is very much defined by the upstream reed/bore coupling which is intensified by the action of the lowpass filter [9]. To be correct, the transmission filter should be complementary to the reflection filter. For the sixth order Butterworth filter, $H(z)$, this implies having an associated complementary sixth order highpass filter, $1-H(z)$. Because the work performed here focused more on the mechanics of the internal interactions than the actual sound, except in a relative sense, little computational time was allocated to the transmission filter. Instead, the simple one-pole filter used by Cook in his clarinet was implemented.

Since the simulation results have been compared to a clarinet, it is worthwhile to at least touch upon the nature of the transmission filter. The waveform emanating from the register hole discussed earlier, which was, by the nature of the scattering junction, highpass filtered, had a nice tone to it, and as was shown, a much different shape. On the other hand, Cook's experimental clarinet waveforms, which were measured internally and discussed in Chapter 2, were reminiscent of the actual clarinet waveform, suggesting that the internal and external sounds should be very close. Of course, the only real answer would be one obtained by measuring the internal and external spectra of a single instrument.

Chapter 3 suggested a different implementation of the sound radiation, which involved radiating the total pressure at the node, which, recall, equates to the residual left from the lowpass reflection filter. This implementation, suggested for this simulation by Cook [20], is much more efficient than the explicit highpass filter, and is inherently complementary.

Figure 5.38 shows the waveforms for: the default filter parameters; the internal waveform, obtained by setting the filter radius to 0; and the implicit highpass filter discussed here. Note that the waveforms for the one-pole filter and for the internal spectrum are almost indistinguishable. The waveform for the implicit highpass filter, on the other hand, is profoundly different. Although still symmetric, there is no longer the character of a square wave. As could be expected from the boosted higher harmonics, the timbre of the note is quite a bit brighter. The internal waveform sounds muffled in comparison. Figure 5.39 shows the waveforms for the second register G, using the register hole. Here, the implicit

filter imparts a more triangular look, but has a less dramatic effect. Figure 5.40 shows the output spectra of the first and third waveforms for the lower note (the internal pressure spectrum was almost identical to the one-pole radiated spectrum, and is not shown). The fundamental harmonic is about 8 dB stronger for the one-pole filter, but the other harmonics are boosted for the implicit filter. The bore resonances for these two cases are, of course, identical. Changing the transmission filter parameters has no effect on the internal behavior.

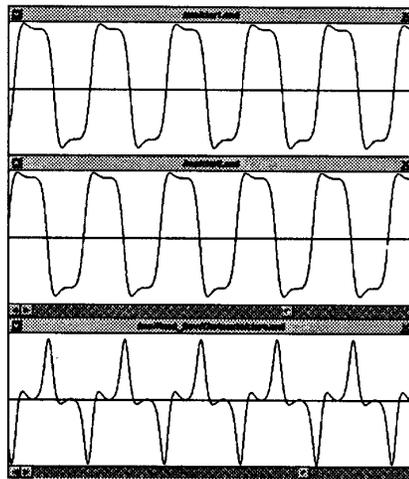


Figure 5.38: Waveforms for a low A: 1) radiated pressure using one-pole highpass filter with radius of 0.9; 2) Internal pressure waveform of the right-going wave; 3) radiated pressure using implicit highpass filter.

In [70], Stewart and Strong, with their model, produced waveforms somewhat similar to these. They defined a low blowing pressure, for non-beating reed, and a high blowing pressure, for beating reed, and reproduced both mouthpiece and radiated pressures. Their low blowing pressure waveforms had few harmonics, and matched in form the fundamental-dominated waveforms our simulation produced for the register hole notes. For the low pressure case, their internal and external waveforms were extremely similar. Although their high pressure mouthpiece pressure waveform did not have the definition ours did, at least in terms of the internal spectrum discussed, the radiated pressure looked much like the radiated pressure waveform depicted for the implicit highpass filter in Figure 5.38. Although one might conclude from Stewart's waveform that the very different shape of the

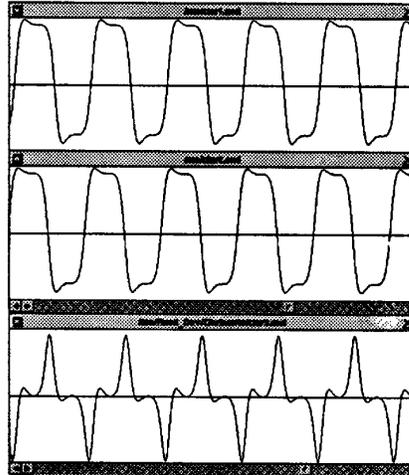


Figure 5.39: Waveforms for a high G , with register hole, using one-pole and implicit transmission filters

high-pressure example was due to the beating of the reed, our corresponding case resulted from a nonbeating reed; it is the presence of the harmonics in the internal spectrum due to the high input pressure, and not the beating reed, which are amplified by the output filter. It is this which results in the dramatic difference in waveform from less harmonically complex internal pressure fluctuations.

The work on the effects of the bell has really just begun. The experiment with the Glastonbury Pipe in Chapter 2 indicated the importance of the bell in defining midrange harmonic formants, and this cannot be modelled with the simplistic approach taken here. The filter parameters could be modified to duplicate the midrange boosting which occurs in the frequency domain. This would require increasing the order and complexity of the filter. A more physical approach, consistent with the philosophy of physical modeling, would be to model the mass of air at the end of the bore more accurately, as, for example, Stewart and Strong did with their model of a radiating piston [70]. For such a model to be musically useful, however, it would have to be very good. Otherwise, all the advantages of having the model, in this case, are lost, and a frequency-domain based filter will give better results. There is a definite trade here. In the first approach, experimental data can be analyzed with, for example, linear prediction techniques to produce a filter that gives realistic

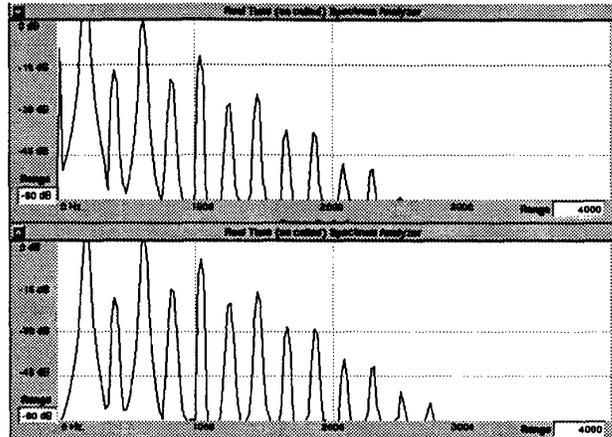


Figure 5.40: Output Spectra for a Low A, using one-pole and implicit transmission filters

effects, but is still based on non-physical specifications. In the second approach, the filter parameters are based on the physical model of the bell, but they must be implemented very accurately or the spectral results so easily obtained in the first approach will not be obtained.

5.5 Implementation of a Reed Model

The previous sections all used the static reed table proposed by Smith in [63]. In this series of experiments, the dynamic models discussed in Chapter 3 were investigated. The issues of interest were:

- Influence on the transient and steady-state portions of the waveform;
- Influence of the reed resonance on steady-state timbre and note stabilization;
- Influence of the reed on mode transitions
- Influence of the elasticity of collision against the lay.

The first point covers basically the transient effects of the dynamic overshoot by the reed,

which cannot be duplicated by the static reed table. The second deals with the interaction of reed resonance with the regime of oscillation, as discussed by Thompson in [72]. The third examines a phenomenon commonly known as “squeaking,” a musically undesirable event well known to beginners. Finally, the last treats one particular aspect of reed behavior for which contradicting experimental results have been published; Stewart and Strong claim total inelasticity [70] and Hirschberg claims elasticity [27] in the collision.

5.5.1 Influence of the Time-Domain Properties of the Reed Model

The first issue investigated was the time-domain effects of the reed model, that is, the effects of dynamic overshoot on the instrument timbre. Figure 5.41 show the waveforms associated with the reed table and the two reed models (recall that one reed model accounts for the different treatment of right and left going waves in the look-up table approach). The reed model was designed to be essentially equivalent to the reed table in the non-beating case. For the first reed model, this equivalence worked out to be accurate, at least in terms of maximum output pressure, which is governed in part by the fluctuations of the reed in response to bore pressure, and is therefore very sensitive to reed parameters. An input pressure of 16820 matched the reed table output with the reed model output for the default input pressure of 16700. The input pressures resulting in the waveforms in Figure 5.41 were set to equalize the unnormalized output pressure. They were not, however, normalized as was usually done to the soundfile, since the resultant stretching made it harder to compare the overshoots. The second reed model, which is more realistic, was not as closely matched in response; it required an input pressure of 17110 for a matched output. This was entirely expected, because the first reed model was designed specifically to be compatible with the reed table, while the second reed model was not.

The waveforms displayed in Figure 5.41 shows some definite differences among the implementations. The reed models show more overshoot at the onset of the peaks, as could be expected from the momentum of the mass in the model. For the second reed model, these overshoots were clearly defined, as were the transitions out of the overshoot phase. While the overshoots were also well defined for the second reed model, the transitions were much smoother. These differences were also evident in the reed position plots, which are not presented here because they provide no additional information. The first reed model produced the harshest tone, while the second produced the most mellow. The second tone

was also slightly flat, as would be expected because the second reed model involves more phase lag. In a direct comparison between tones of slightly different pitch, the flatter tone does tend to sound a bit darker, regardless of the harmonic structure. The spectra, shown in Figure 5.42, yield more information about the differences. The first reed model and the reed table are quite close out to about 1800 Hz, at which point the harmonics of the reed model do not drop off as quickly. This could be expected from the comparison of waveforms and the tonal differences. The second reed model spectrum was less predictable; the even harmonics, particularly the first two, were boosted while the flanking odd harmonics dropped a bit. In addition, the harmonics above 1800 Hz dropped more rapidly than in the other two implementations.

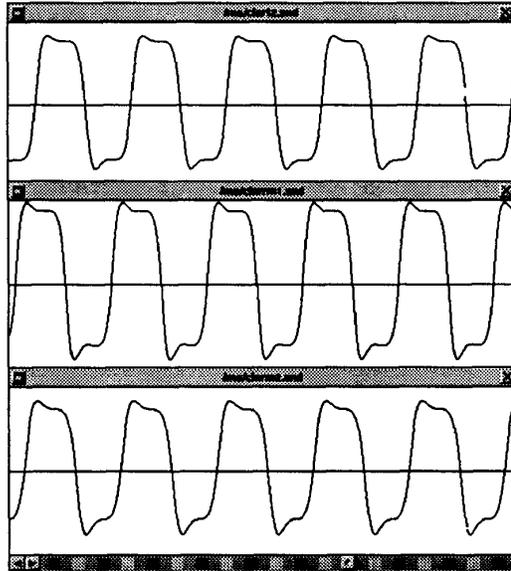


Figure 5.41: Nonbeating waveforms for three reed implementations: reed table, reed model 1, and reed model 2

The previous examples showed that the reed model did make a difference, if a second or third order difference, in the tone. For the beating reed, the results were not quite as sensitive, as Figure 5.43. The input normalization required took a different direction. Here, the second reed model gave the strongest output for a set input pressure. To match the

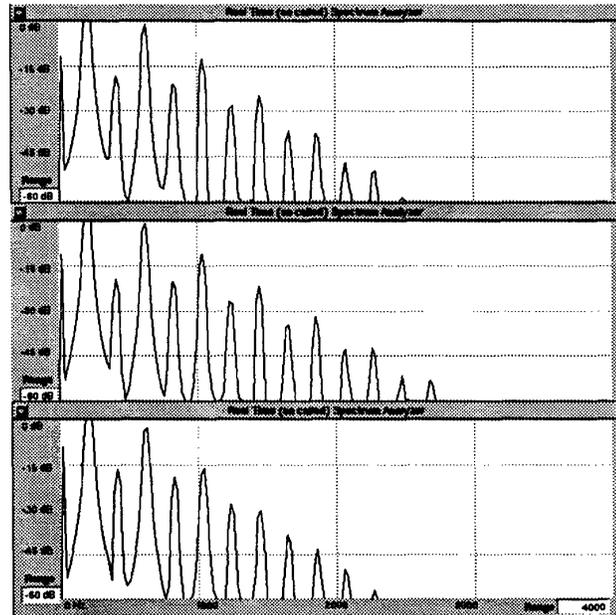


Figure 5.42: Nonbeating output spectra for three reed implementations: reed table, reed model 1, and reed model 2

first reed model at 20000, the reed table input pressure had to be raised to 20353, while the second reed model pressure was lowered to 19859. This was a less intuitive result, since it would seem that beating would intensify the dynamic interactions. However, it would appear that, at least with inelastic collisions, the act of beating in itself, whether the model was static or dynamic, dominated the results. The effects of elasticity in the collision, which is direction-dependent and could be implemented only with the dynamic models, will be discussed presently.

5.5.2 Influence of the Reed Resonance on Stable Modes

To reiterate from earlier discussions, the reed resonance is believed to play an important role in the stabilization of upper register notes, for which the cutoff frequency and more widely spaced harmonics result in a sparser regime of oscillation than for lower register

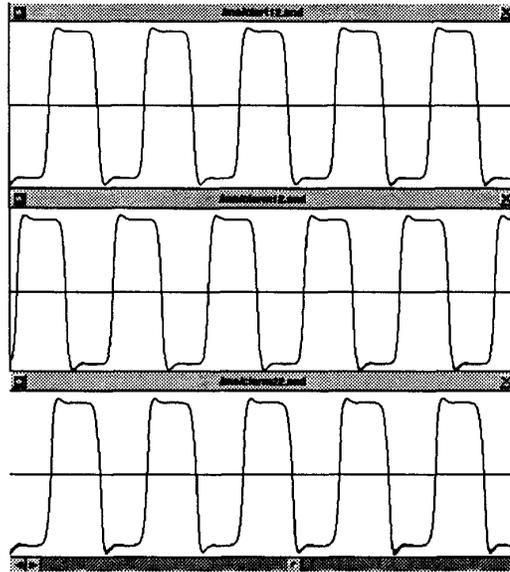


Figure 5.43: Beating waveforms for three reed implementations: reed table, reed model 1, and reed model 2

notes [72]. A number of experiments were performed on the effects of reed resonance. The following variations were performed:

- Varying the reed resonance through the embouchure damping to ascertain its effect on the output strength;
- Varying the reed resonance through the base reed resonance;
- Aligning the reed resonance to the first bore resonance, rather than the ideal waveguide designed playing frequency;
- For chalumeau notes, a bore perturbation was introduced to misalign the resonances so that the stabilizing effects of the resonance could be ascertained. Clarion notes with the register key could not be tested in this manner because the simulation allows for only one scattering junction, and the register key and the bore perturbation cannot both be selected simultaneously. However, the delay line was truncated to

produce the clarion register, and the bore perturbed.

None of these variations had any appreciable effect. The output amplitude did increase somewhat linearly as the damped resonance was increased, and the damping lowered, but there were no local maxima observed where the reed resonance matched a harmonic. Again, the simulation was limited by its own inherent stability. Even with the misaligned modes, no reed resonance advantages could be measured. The stabilization which the reed participates in is most likely in situations where undesired mode transitions are possible, and this simply wasn't the case here, for the most part. The reed resonance at this stage in the development contributed little to either musical potential or acoustic understanding. However, a more comprehensive study would be valuable.

Recall that the reed model panel was carefully set up to automatically align the reed resonance with desired playing note. In fact, this procedure was not entirely advantageous because the actual playing frequency, by virtue of the phase lag induced by the bell discussed earlier, was not necessarily anywhere near the desired note. The tuning filter, described above, would reduce this problem by bringing the desired and played notes closer together.

5.5.3 Influence of the Reed Model on Mode Transitions

It turned out that for reproducing musically undesirable events, the reed model could prove extremely useful. In particular, the reed model could duplicate, in various ways, the phenomenon of "squeaking".

The first form of squeaking discovered was a register change that occurred near threshold blowing pressure. This was very dependent on the bell filter used, as well as the note played. Figure 5.44 shows the spectra for a low A , obtained with the 2nd, 4th, and 6th order Butterworth filters. The 2nd and 4th order filters interacted with the reed to cause a jump to the second register; the 6th order filter resulted in a jump to the third register! If unexpected, this could certainly be perceived as a squeak. For the high A , both 2nd and 4th order filters produced the proper note, but the 6th order filter produced a 2nd register jump.

Raising the input pressure or increasing the damping both served to halt the squeaking, although the damped tone was still marginal. One is tempted to conclude that the cause of

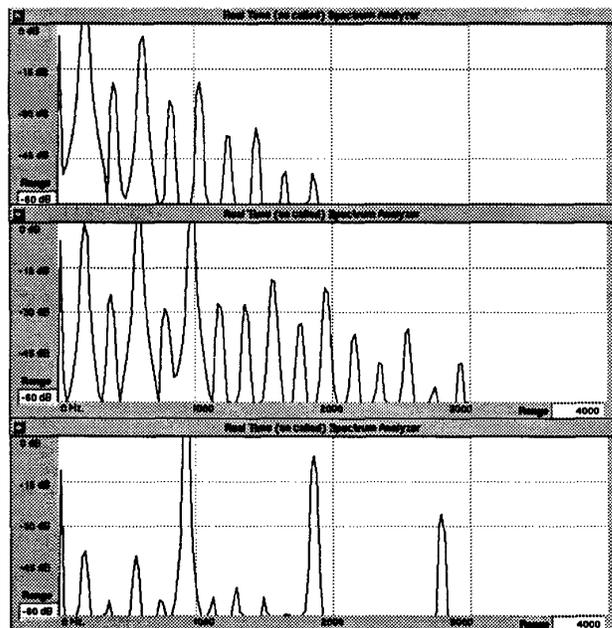


Figure 5.44: Register shifts for the reed model with 2nd, 4th, and 6th order bell reflection filters

squeaking is insufficient damping on the part of the player. This was suggested by Wilson and Beavers in [75], who studied the effects of damping on reed resonance effects, although in their case, the squeak was a reed resonance oscillation and not a bore register. In fact, a reed teacher will probably tell you (when asked) that squeaking is due to insufficient support, i.e., not blowing hard enough [10][40]. The musical “problem” duplicated with the simulation was one of insufficient input pressure, not insufficient damping; an experienced musician would not blow at threshold blowing pressure, and is much less likely to squeak. Wilson did note the effects of raising pressure, and also hypothesized that the threshold blowing pressure for the higher mode was lower than that for the lower mode, and so more likely to be excited at lower blowing pressures; his conclusions, however, focused on the damping. Benade discussed the octave “sneak” in [9] in the saxophone and taragato, the result of a weakened regime of oscillation in which the second mode could take over on its own. This is unlikely to be the case here, however, as Benade attributed the problem to a

defective instrument, and the simulation is highly idealized. Benade did not describe any such jumps to the third register. In fact, a “sneak” does not have quite the aural effect of a “squeak”; a third register jump is bound to be more jarring than a second register jump.

Mode transitions were not merely the result of blowing pressure that was too low. A register shift was also obtained with an input pressure of ranging from 23600 to 26000, which is relatively high. Interestingly, this register shift was obtained with both the first reed model *and* the reed table, although the input range for the reed table was much more narrow. In both cases, the second order Butterworth filter was used. The second reed model did not produce such an effect, although it probably would have with some exploration. Experience has demonstrated that overly enthusiastic blowing can also result in squeaking, although probably at a higher frequency than the second register attained here.

Finally, the most commonly postulated source of squeaking is the reed resonance itself. This musical feature was also available from the simulation. One example is shown in the spectrum of Figure 5.45. Here, the reed damping was set to 0.39, resulting in a damped frequency of 2413 Hz. The base frequency was 2620 Hz. Figure 5.45 shows a clear dominant mode at 2650 Hz, which is higher than both the natural and the damped reed frequencies. However, 2650 is almost precisely 13 times the frequency of the first peak in the spectrum, which lies at approximately 204 Hz. This is the location of the fundamental peak for the same case with the reed table, as has been shown in the sections above. The irony is that here, the fundamental mode has apparently helped to stabilize the reed resonance, whereas the reed resonance is supposed to stabilize the fundamental (or really, the third harmonic). The resonance was produced with the “nominal” blowing pressure of 16700; however, increasing the pressure did eventually eliminate the squeak (with some very interesting transitional tones on the way up).

The reed resonance demonstrated above was quite clean. In fact, the damping ratio that would allow such a clean tone was extremely narrow; and variation up or down would result in a harsh multiphonic tone with many peaks, sometimes focusing onto one pitch by the end of the note. A damping of 0.425 and higher produced the fundamental clearly, with no reed resonance present at all.

In summary, the reed played a clear role in certain mode transitions in the instrument. In one case of mode transition, a single-register shift was obtained for a band of moderately

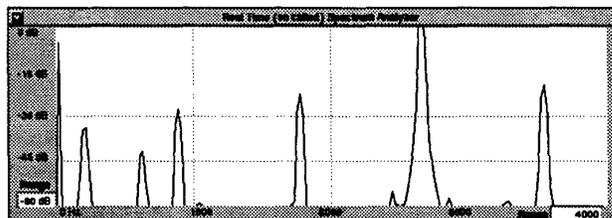


Figure 5.45: Reed resonance mode on ClariNeXT model

hard input pressure that could be obtained with both reed table and reed model. This could be meaningful in light of the discussion about overblowing earlier. There, it was postulated that either the corruption of the equivalent volume or the ability of the reed to support the slower oscillation given a reduced tip displacement aided in overblown register shift. These results suggest that the right conditions can cause a register shift even without an narrowed tip opening (in fact, the mean tip opening is larger for harder blowing because the reed will open further, while being constantly limited by the lay on the closure side).

The shift to the third register was also significant, because it suggests a new explanation for reed squeaking that occurs only at very low breath pressure, where beginners tend to tread. The reed resonance was also pressure dependent, but occurred at higher pressures where one would expect the tone to be stable. It would be instructive to try to accurately measure the squeaks of an instrument, if one could find a willing subject; a beginner would be unpredictable, and an experienced player unlikely to subject his instrument to the indignity of such a sound.

5.5.4 Influence of Elastic Collisions

The dynamic reed model included a boundary condition to allow any degree of elasticity in the collision. The purpose was to better understand the effects of elastic reed collision on the waveform. In addition, Hirschberg had described a link between the sought-after pitch/amplitude dependency and a beating elastic collision [27], and this was to be tested.

Figure 5.46 shows the waveforms and reed position plots of the simulation with inelastic and elastic collisions at the point that beating begins. There is no discernible difference

between the two. This is not really surprising when one considers the velocity of the reed as it strikes the lay. Since the reed is very close to turning around on its own at this point, its velocity is very low, and the impetus added to its return trip is insignificant compared with the pressure forces acting on it.

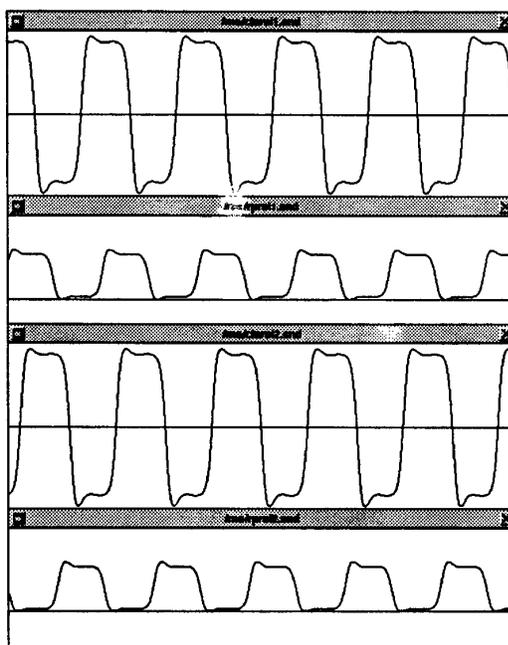


Figure 5.46: Pressure waveform and reed position plot for inelastic and elastic collisions under just-beating conditions

A more extreme case is shown in Figure 5.47. Here, an input amplitude of 19000 assured hard collisions, as the reed position plots indicate. Now there is a noticeable difference between the two, because the reed chatters against the lay in the elastic case until the bore pressure releases it. Except for this chattering, however, the waveforms do not differ. The initial overshoot on the upper peak is no larger for the elastic case than the inelastic case. Again, the real driver is the pressure driving the reed against the lay. In the final beat, the pressure wave, which is not completely square, does not have as much force as for the initial beat, and the final velocity is not high enough to provide much bounce. The spectra in Figure 5.48 verify this. The low frequency harmonics all match well, and it is only in

the last quarter of the band shown that the chattering is evident for the elastic example. Aurally, the chattering is very apparent, and does change the perceived tone, even though the chattering itself sounds like chattering and not like some multiphonic.

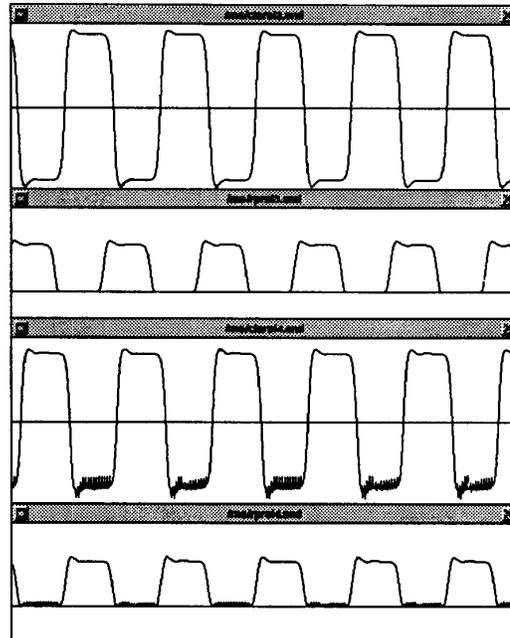


Figure 5.47: Pressure waveform and reed position plot for inelastic and elastic collisions under strong beating conditions

Thus, elastic collision appears to affect only behavior at the lay, and does not provide enough momentum to influence the rest of the waveform; the reed motion is still dominated by the force of the standing pressure wave in the bore. With a more triangular pressure waveform, as was evidenced by the real clarinet, there might be more opportunity for the elasticity to make a difference. In addition, if the reed were modelled as more of a nonlinear spring near the lay, as suggested by Hirschberg [30], the model might be more effective.

The other elastic experiment was on pitch/amplitude dependency. Once again, no dependency was found. With a purely elastic attack, and the offset set as far as 11000, the pitch remained constant.

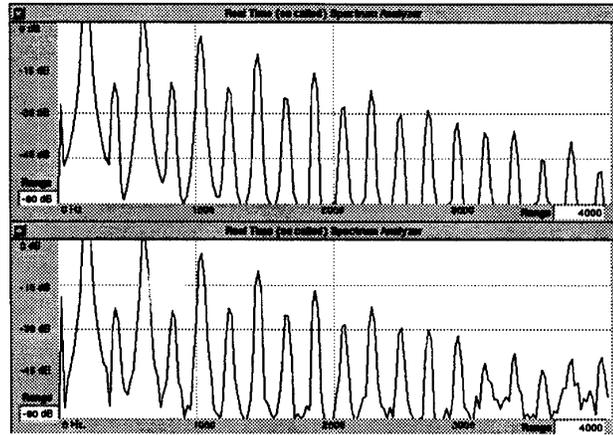


Figure 5.48: Output spectra for inelastic and elastic collisions under strong beating conditions

5.6 Implementation of Hydrodynamics (“Bernoulli Effect”)

The previous discussion of elastic collision was an example of an effect which could be implemented much more easily with a dynamic reed model than in the precomputed reed table (although a velocity-dependent pair of reed tables, one for closing and one for opening, could be designed for direction-dependent phenomena). In this section, the implementation of a hydrodynamic behavior model in both reed table and reed model, as discussed in Chapter 3, is investigated.

From the nature of the Bernoulli implementation, we can expect an effectively softer reed, because it will close at lower pressures. Beating will consequently occur more easily. It must be noted again that the assumptions on which this model is based are suspect [27]. On the other hand, it is an accepted truism, be it true or not, that one of the chief characteristics of the double reed is its hydrodynamic behavior, and in particular, the Bernoulli effect. For this reason, it was worth trying.

First, the reed table form of the Bernoulli model was simulated. Recall that including the Bernoulli effect in the reed table would result in a translated breakpoint, because the reed will tend to close for lower pressure differences, and a steeper slope at closure, similar to

a nonlinear spring. A typical such reed table was illustrated in Figure 4.8. Recall that the result of the effect is a “sucking” of the reed toward closure as it nears the lay (or the other reed). Figure 5.49 shows the result. All three waveforms resulted from an input pressure of 18201, which normally produces strong beating. The first waveform includes a Bernoulli force of strength 250. The second waveform shows the waveform without the Bernoulli force. For the third waveform, the linear reed table was adjusted so that the new breakpoint coincided with the Bernoulli breakpoint from the first waveform — this was to eliminate differences due to the change in effective stiffness.

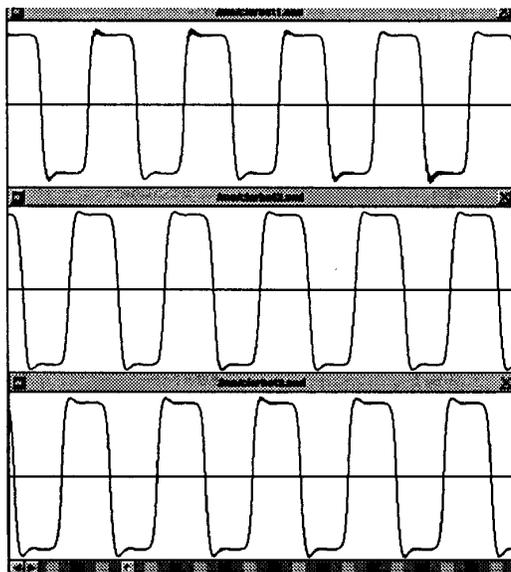


Figure 5.49: Bernoulli effect: waveform 1 has Bernoulli-type force of level 250; waveform 2 represents the same reed table with no Bernoulli force; waveform 3 represents the reed table with adjusted breakpoint

The three waveforms do not look remarkably different. They do, however, sound considerably different. The Bernoulli waveform has a distinct, almost nasal timbre. In fact, the spectra of the first and third cases, presented in Figure 5.50 show a marked difference. The odd harmonics match fairly well; the even harmonics are, however, much more prevalent for the Bernoulli spectrum. This effect was obtained even with much lower Bernoulli levels.

It might be mentioned as an aside that, while the Bernoulli-adapted reed table looks and acts much like a non-linear spring, there was still no noticeable amplitude/frequency dependence.

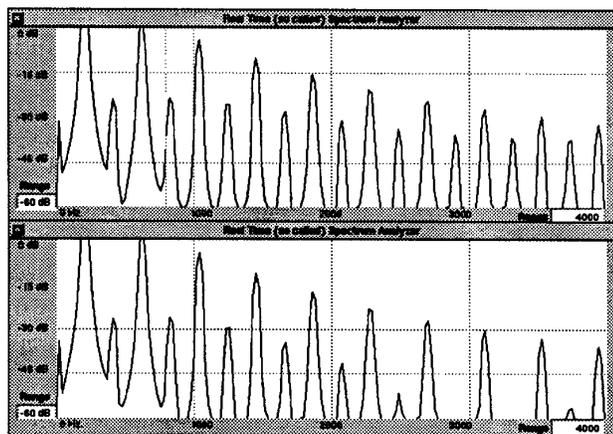


Figure 5.50: Bernoulli effect: Top spectrum with, Bottom spectrum without

The reed model was similarly affected, except as with the elastic collisions, there was much activity during the period where the reed touched the lay. Figure 5.51 compares the waveforms with and without the Bernoulli model for the default input pressure of 16700. Whereas before, the two waves were almost indistinguishable in their time-varying form, the two waves here are quite distinct by virtue of the ragged beating portion. Figure 5.52 shows the spectra. Although it is likely that the even harmonics might have had the prevalence they had with the reed table, the ragged portion of the waveform significantly disrupts the higher harmonics, dominating the sound to some extent.

Figure 5.53 shows two more extreme cases of the Bernoulli effect on the reed model, with the levels set at, respectively, 500 and 1000. Certainly part of the cause of the chatter is physical; as the reed is sucked toward the bore, it moves with increasing velocity. When it hits, cutting off the flow, the hydrodynamic force vanishes, leaving only the restoring force of the bore/mouth pressure to pull it back. For this reason, the most extreme effects are obtained at a moderately low input pressure when the beating is caused by the Bernoulli force alone. With such an arrangement, the bore pressure will always pull the reed back up as soon as the Bernoulli force is cut off. At this point, some possible computational problems

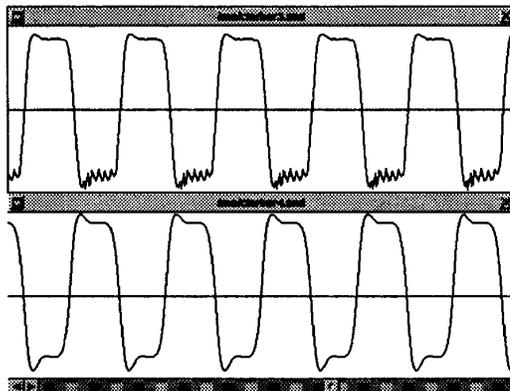


Figure 5.51: Bernoulli effect on reed model: waveform 1 has Bernoulli-type force of level 250; waveform 2 represents the same reed table with no Bernoulli force

arise. The reed has one sample time to rise up off the lay. Because the flow is no longer cut off, the next pressure calculation will once again include the Bernoulli force, which may then be sufficient to slam the reed back into the lay, where the force vanishes again. Thus, a limit cycle may be occurring because of the on-off nature of the Bernoulli model in this implementation. This is particularly apparent in the 1000-level force in Figure 5.53, but is potentially invasive for all models with lay discontinuities. The most obvious solution would be to multisample the reed, that is, run the reed at a higher rate than the rest of the simulation. This way, the reed will not have as much rise time before encountering the force again. It is quite possible that the elastic collision results suffered from the same type of cycling, due to the instantaneous velocity change that was part of the boundary condition.

In summary, the Bernoulli model did yield some interesting sound differences, although it did not produce a magical double reed. The reed table implementation worked quite nicely, and yielded a smooth spectrum with boosted even harmonics. The reed model gave more ragged results, although in some respects more physical results. Continued experimentation with the reed model would require multisampling it so that numerical artifacts can be reduced.

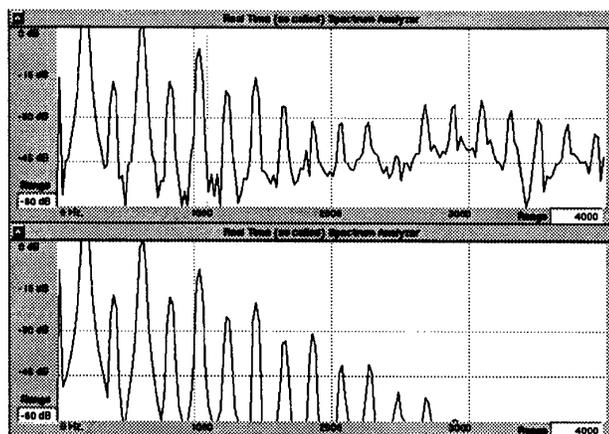


Figure 5.52: Bernoulli effect on reed model: spectrum 1 has Bernoulli-type force of level 250; spectrum 2 represents the same reed table with no Bernoulli force

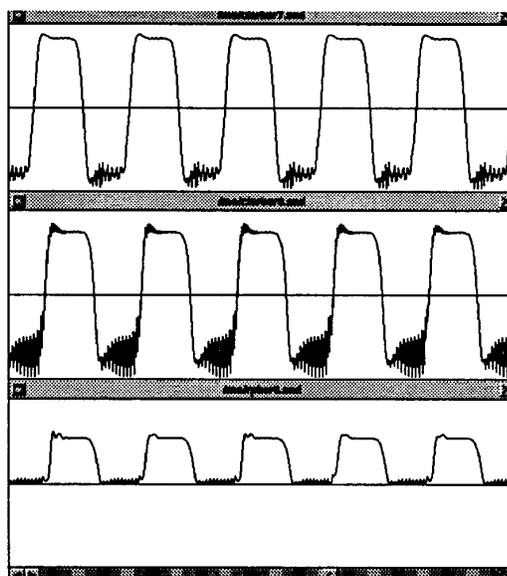


Figure 5.53: Bernoulli effect: Levels of 500 and 1000 on reed model

Chapter 6

Conclusions

The ClariNeXT simulation provided a convenient testbed for checking out acoustical concepts and models. Its object orientation allowed for easy incorporation of new enhancements, and its interface allowed for intuitive, almost real-time experimentation. It serves the purpose of both the earlier research simulations, in demonstrating acoustic phenomena for better understanding, and of a development environment for the creation of, eventually, a real-time instrument.

Some acoustical behavior was quite realistic - the generation of harmonics with amplitude, the register jump due to the register hole, the ratio of nonbeating regime to nonplaying regime, and mode transitions, all followed prediction from the literature. Other behavior could not be repeated. The basic dependency of pitch on amplitude which every musician learns to correct was not reproducible, at least to the ears of the author, by implementation of any of the three possible explanations included here: cooperation among misaligned modes; effective reed beating resonance; and nonlinear spring behavior. These implementations were somewhat crude however, and might have yielded more had they been refined. Also, the pressure waveform radiated was much more square than that from a real clarinet, and this may have masked other results.

From a musical standpoint, the register hole scattering junction and the variable cutoff Butterworth filter, particularly the 2nd order filter, were the most promising. These were relatively inexpensive, and could be efficiently included in a real-time model. The register hole simply worked, as did the filter, which allowed a more realistic cutoff frequency. The

mode transitions which could be induced by a lossless scattering junction were not as desirable for a basic clarinet, but did produce some interesting sounds. Given that the “clarinet” is just a starting point in providing a physical vocabulary of gestures to the musician, the method of mode misalignment might lead to a new variant of the instrument.

The incorporation of the reed dynamics was less successful from a musical standpoint, although it was certainly successful from a simulation standpoint. The use of the reed model afforded no real musical advantages over Smith’s reed table. It did exhibit some reed resonance behavior that the reed table could not, and was effective as a study tool, but at this point would probably not be worth the cost in a real-time simulation. Improvements in results might be had by a more refined reed model. At the very least, the reed could be run at a higher rate to avoid some of the high amplitude limit cycling during reed closure times. A trapezoidal rather than a rectangular integration scheme would implement a bilinear transform in the propagation of the dynamics, and might help as well. In addition, the reed itself might be modeled more carefully, with account taken of its cantilevered nature (as in Stewart’s model in [70]) and the nonlinear spring effects at the lay or junction with the other reed.

One of the greatest advantages of using the reed model was the ease with which additional nonlinear behavior could be incorporated. For example, while the elastic collisions did not prove very interesting, they were easy to incorporate, and the results quite believable. Multisampling the reed would improve the utility of the reed model, since any nonphysical limit cycling at discontinuities would be reduced.

Although one of the original objectives of this author was to design the definitive double reed model, it turned out that so little was actually understood about double reed behavior that there was little to base this definitive model on. On the other hand, experimental data, both from the author’s experiments and published results, indicate that the behavior unique to the double reed may yield only a second order effect (although certainly not in terms of the technique required to play it, as a double reed player would be quick to point out). The characteristic sound of the orchestral double reed seems to be, at the primary level, more due to the conical bore than the reed itself.

Similarly, it is probably not worthwhile at this point to try to develop a more accurate hydrodynamic model; the data simply isn’t available. The Bernoulli force model attempted

to implement the “essence” of this folkloric effect, and in this, was useful. Both reed table and reed model did yield some different timbres with this force included, although whether they displayed realistic behavior is debatable. The boosting of the even harmonics with the reed table was a significant result in controlling timbre.

In fact, the experiments with the Bernoulli force indicated that much could be done through manipulation of the reed table itself. Smith has indicated as much, and has reported tone improvement with nonlinear adjustments to the table [69]. Some features, such as nonlinear spring stiffness near beating, would be easy to precalculate in the reed table. Others, which are hysteretical or direction-dependent in nature, would be more difficult in the current formulation. It would be worthwhile to develop a tandem reed table — one for opening and one for closure — that could be used for implementation of these effects.

If the piecewise linear reed table is chosen for a real time instrument, a pair of gains on the input and output pressures would be functionally equivalent to varying the reed table parameters, except for the offset, which defines the operating point. Changing these gains to adjust stiffness would probably be more intuitive than trying to adjust the reed table itself.

Many of the mode transitional behavior, such as the reed resonance and the third-register shift, might have been duplicating real behavior, or might have occurred because of computational artifacts. The numerical issues affecting the computation should be investigated thoroughly to distinguish one from the other. If the register shift is truly a physical result, it is a significant one; all the work the author has come across has attributed the phenomenon of “squeaking” to reed resonance rather than register shifting. As mentioned above, the mode transition behavior, if properly understood, could add an important dimension to a real-time instrument.

The simulation could, of course, benefit from an addition of behavior models. The most tempting at this point is the efficient conical bore model which Smith is developing [68], which would add an entirely new dimension. A frequency-dependent lumped loss source is very important. No loss has been explicitly modelled thus far, save for sound radiation at bell, toneholes, and mouthpiece. Boundary layer effects were neglected, yet are known to be significant. It could well be that much of the deviation from realistic behavior was due to the deficient loss modeling. Finally, an allpass filter, as described in [31] must be

implemented for tuning if the simulation is to be at all musically practical.

Other potential improvements are not quite as pressing at this point, but still enticing. A better bell model could be designed, which better duplicates the effects of boosting midrange harmonics and improving perceived tone resonance, as well as merely imposing a cutoff frequency that affects primarily the higher frequency harmonics. Possible improvements to the reed table and reed model have already been discussed. Following along the development path of the register hole, a true tonehole, with both reactance and resistance, should be designed. Of course, a single tonehole leads to a tonehole lattice. Finally, a model of the vocal tract would prove interesting. These enhancements would require much experimentation with real instruments.

Even with the addition of the above models, the simulation still runs open-loop; any adjustments to the parameters are up to the user, and are fixed for the length of the run. Of course, the eventual goal is to have an abbreviated real-time version of the instrument, with an interface to a live performer, who would close the loop just as he would on a physical instrument. An ambitious undertaking would be to define a multivariable control system that would aid in the adjustment of the many parameters to achieve some tonal goal. The performer would select a cost function to his liking, and through this, and a battery of gestures understandable to the control system, he could concentrate his musical efforts on other aspects of the performance.

In summary, the ClariNeXT reed workbench turned out to be an effective tool for acoustic research and musical development. With continued development, it could play an important role in the developing field of physical modeling of musical instruments.

Appendix A

Default Simulation Parameters

All the results presented in Chapter 5 were, unless otherwise stated, based on the following set of default parameters, in “simulation units”.

- Sampling Rate: 22050 Hz
- Waveguide Parameters

Base Frequency: C 262

Desired Note: A

Octave Drop: ON

Number of samples: 20000

- Scattering Parameters

Position: Scattering OFF

Diameter Ratio (when turned on): 0.1

Hole Position (when turned on): 0.3333

- Reed Parameters:

Reed Model Selection: Reed Table

Reed Table:

Breakpoints: 255 and 1024 normalized pressure units

Offset: 700

Scale Factor: 80

Reed Model:

Reed Resonance multiplier: 10

Open Position: -25840

Closed Position: 35600

Stiffness: 1

Tip Offset: 0

Collision Elasticity: 0 (inelastic)

Bernoulli Force: OFF

- Bell Parameters:

Reflection Filter Type: 2nd order Butterworth

Cutoff Frequency: 1500 hz

Gain: 1

Transmission Filter Type: one-pole filter

Transmission Filter Gain: 1.3

Transmission Filter Radius: 0.9

- Attack Parameters:

Note Length: 19000 samples

Steady State Amplitude: 16700

Attack Length: 2000 samples

Decay Length: 10 samples

Attack Overshoot: 0.25

Overshoot position in Attack: 0.3

Vibrato Level: 0

Vibrato Frequency: 6 Hz

Noise Level: 0

Bibliography

- [1] Y. Ando, "Structure and Timber Character of a Baroque Treble Recorder," paper presented at the 13th International Congress on Acoustics, Yugoslavia, 1989.
- [2] J. Backus, *The Acoustical Foundations of Music*, New York: W. W. Norton, 1969.
- [3] J. Backus, "The Effect of the Player's Vocal Tract on Woodwind Instrument Tone," *J. Acoust. Soc. Am.* 78(1): 17-20, July 1985.
- [4] J. Backus, "Input Impedance Curves for the Reed Woodwind instruments," *J. Acoust. Soc. Am.* 56(4):1266-1279, October 1974.
- [5] J. Backus, "Small-Vibration Theory of the Clarinet," *J. Acoust. Soc. Am* 35(1):305-312, March 1963.
- [6] A. Baines, *Woodwind Instruments and their History*, New York: W. W. Norton, 1957.
- [7] N. Bak and P. Domler, "The Relation between Blowing Pressure and Blowing Frequency in Clarinet Playing," *Acustica* 63:238, 1987.
- [8] A. H. Benade, "Equivalent Circuits for Conical Waveguides," *J. Acoust. Soc. Am*83(5):1764-1769, May 1988.
- [9] A. H. Benade, *Fundamentals of Musical Acoustics*. New York: Oxford University press, 1976.
- [10] M. Boenau, private discussions.
- [11] H. Bouasse, *Instruments à Vent*, Paris : Librairie Delagrave, 1929.
- [12] A. Carse, *Musical Wind Instruments*, New York: Da Capo Press, 1965.

- [13] C. D. Chafe, "Pulsed Noise in Self-Sustained Oscillations of Musical Instruments," Proc. of the IEEE Int. Conf. on Acoustics, Speech, and Signal Processing, A1.17, pp.1157-1160, Albuquerque, NM, 1990.
 - [14] J. M. Chowning in, *FM Theory & Applications*, Tokyo: Yamaha Music Foundation, 1986.
 - [15] J. M. Chowning, *Foundations of Computer Music*, ed. Roads and Strawn, Cambridge, MA: MIT Press, 1985.
 - [16] J. W. Coltman, "Mode Stretching and Harmonic Stretching in the Flute", J. Acoust Soc. Am. 88(5), Nov., 1990.
 - [17] P. R. Cook, *Identification and Control of Parameters in an Articulatory Vocal Tract Model, with Applications to the Synthesis of Singing*, Ph.D. Dissertation, Dept. of Electrical Engineering, Stanford Univ., 1991.
 - [18] P. R. Cook, "Implementation of Single Reed Instruments with Arbitrary Bore Shapes Using Digital Waveguide Filters," Stanford University Department of Music Technical Report Stan-M-50, May 1988.
 - [19] P. R. Cook, C. D. Chafe and J. O. Smith, "Pulsed Noise in Musical Systems, Techniques for Extraction, Analysis, and Visualization," Proc. of the International Computer Music Conference, pp.63-65, Glasgow, Scotland, 1990.
 - [20] P. R. Cook, private discussions.
 - [21] R. Dawson, personal discussions.
 - [22] The Diagram Group, *Musical Instruments of the World*, New York: Bantam, 1976.
 - [23] A. Fettweis, "Some Principles of Designing Filters Imitating Classical Filter Structures," *IEEE Trans. on Circ. Theory*, vol. CT-18, pp.314-316, March 1971.
 - [24] D. E. Hall, *Musical Acoustics: An Introduction*, Belmont, CA: Wadsworth, 1980.
 - [25] H. L. F. Helmholtz, *Sensations of Tone*, trans. A. J. Ellis, New York: Dover, 1954.
 - [26] A. Hirschberg, J. Gilbert, P. J. Wijnands "Flow through the Reed Channel of a Single Reed Music Instrument", *Colloque de Physique C2* Tome 51, n2, February 1990
-

- [27] A. Hirschberg, R. W. A. van de Laar, J. P. Marrou-Maurières, A. P. J. Wijnands, H. J. Dane, S. G. Kruijswijk, and A. J. M. Houtsma, "A Quasi-Stationary Model of Air Flow in the Reed Channel of Single-Reed Woodwind Instruments," *Acustica*, vol. 70, pp. 146-154, 1990.
 - [28] A. Hirschberg, P. J. Wijnands, J. V. Zon, "Some Fluid Dynamic Aspects of Single Reed Woodwinds," paper presented at the Acoustical Society of America meeting, November 24-30, 1990, in San Diego, CA.
 - [29] A. Hirschberg, "Some Fluid Dynamic Aspects of Speech", paper presented at the fourth Colloquium Signaalanalyse en Spraak, COLSAS 22-23 Oct 1990, at the Instituut voor Perceptie Onderzoek, Eindhoven, The Netherlands
 - [30] A. Hirschberg, private discussions and correspondence, 1990-1991.
 - [31] D. A. Jaffe and J. O. Smith, "Extensions of the Karplus-Strong Plucked-String Algorithm," *Computer Music Journal* 7(2):56-69, 1983.
 - [32] D. H. Keefe , "Acoustic streaming, Dimensional Analysis of Nonlinearities, and Tone Hole Mutual Interactions in Woodwinds." *J. Acoust.Soc.Am.* 73(5):1804-1820, May 1983.
 - [33] D. H. Keefe , "Experiments on the Single Woodwind Tone Hole," *J. Acous. Soc.Am.*72(3): 688-698, Sept. 1982.
 - [34] D. H. Keefe , "On Sound Production in Reed-Driven Wind Instruments," April 17 1989.
 - [35] D. H. Keefe , "Theory of the Single Woodwind Tone Hole," *J. Acoust.Doc. Am.*72(3): 676-687, Sept 1982.
 - [36] D. H. Keefe and M. Park , "Tutorial on Physical Models of Wind Instruments: II, Time Domain Simulations." July 1989.
 - [37] D. H. Keefe , *Woodwind Tone Hole Acoustics and the Spectrum Transformation Function*, Ph. D. Dissertation, Department of Physics, Case Western Reserve University, January 1981.
 - [38] R. W. A. Van de Laar, Unpublished paper, Technische Universiteit Eindhoven
-

- [39] E. Lindemann, "Routes to Chaos in a Non-linear Musical Instrument Model," Wave-Frame Corp: Boulder, Colorado.
 - [40] J. Linsenberg, private discussions, 1991.
 - [41] J. D. Markel and A. H. Gray, *Linear Prediction of Speech*, New York: Springer Verlag, 1976.
 - [42] *MATLAB*, South Natick, MA: The Mathworks, Inc., 1989
 - [43] M. E. McIntyre, R. T. Schumacher, and J. Woodhouse , "On the Oscillations of Musical Instruments" *J. Acoust.Doc.Am.* 74(5):1325-1345, November 1983.
 - [44] D. C. Miller, *The Science of Musical Sound*, New York: MacMillan, 1934 (first published in 1916).
 - [45] P. Monroe, "The Pipe and Tabor: an Introduction" Seattle, Washington: 1985.
 - [46] P. M. Morse, *Vibration and Sound*, Acoustical Society of America, 1936, 1976.
 - [47] P. M. Morse and K. U. Ingard, *Theoretical Acoustics*, Princeton, N.J.: Princeton University Press, 1968.
 - [48] H. Myers, "Observations: The *Mary Rose* 'shawm' ", *Early Music* 11(3), July 1983.
 - [49] H. Myers, *The Practical Acoustics of Early Woodwinds*, DMA Final Project, Stanford University Dept. of Music, 1980.
 - [50] H. Myers, Private discussions
 - [51] C. J. Nederveen, *Acoustical Aspects of Woodwind Instruments*, Amsterdam: Frits Knuf, 1969.
 - [52] T. Neuman, private discussions, 1990-1991.
 - [53] K. Ogata, *Modern Control Engineering*, Englewood Cliffs, NJ: Prentice-Hall, 1970
 - [54] F. Opolko and J. Wapnick, *McGill University Master Samples*, , McGill University Faculty of Music, Montreal, Quebec, 1987, 1988.
 - [55] G. R. Plitnick and W. J. Strong, "Numerical Method for Calculating Input Impedances of the Oboe," *J. Acous. Soc. Am* 65(3):816-825, March 1979.
-

- [56] Rayleigh, *Theory of Sound*, New York: Dover Publications, 1945 (first published in 1877).
 - [57] F. Rocaboy, "Proposed model for Reed Action in the Bassoon," *JCAS* 1(4):20-25, November 1989.
 - [58] R. T. Schumacher, "Ab Initio Calculations of the Oscillations of a Clarinet," *Acustica* 48(2):71-85, 1981.
 - [59] J. L. Shearer, A. T. Murphy, and H. H. Richardson, *Introduction to System Dynamics*, Reading, MA: Addison-Wesley, 1967.
 - [60] D. H. Smith, *Complete Instructions On Making renaissance Double Reeds* to be published by the University of Indiana Press
 - [61] D. H. Smith, *On Making renaissance Double Reeds Kings Trumpets and Shawms*: San Francisco, 1988.
 - [62] D. H. Smith, private discussions
 - [63] J. O. Smith, "Efficient Simulation of the Reed-Bore and Bow-String Mechanisms," *Proceedings of the 1986 International Computer Music Conference*. Computer Music Association, 1986.
 - [64] J. O. Smith, "Music Applications of Digital Waveguides," Music Department Technical Report STAN-M-39, 1987.
 - [65] J. O. Smith, MU420 Lecture Notes, Stanford University Music Department.
 - [66] J. O. Smith, "Techniques for Digital Filter Design and System Identification with Application to the Violin," PhD Dissertation, Stanford University Department of Electrical Engineering, 1985.
 - [67] J. O. Smith, "Waveguide Filter Tutorial," *Proc. of the International Computer Music Conference*, pp. 9-16, Urbana, IL, 1987.
 - [68] J. O. Smith, "Waveguide Simulation of Non-Cylindrical Acoustic Tubes," Presentation at the CCRMA Affiliates meeting, 9 May 1991.
 - [69] J. O. Smith, private discussions.
-

- [70] S. E. Stewart and W. J. Strong, "Functional Model of a Simplified Clarinet." *J. Acoust Soc Am* 68(1):109-120, July 1980.
- [71] S. C. Thompson, "The Effect of Reed Resonance on Woodwind Tone Production", *J. Acoust Soc Am* 66(5):1299-1307, Nov, 1979.
- [72] S. C. Thompson, *Reed Resonance Effects on Woodwind Nonlinear Feedback Oscillations*. Ph. D. Dissertation, Department of Physics, Case Western Reserve University, January 1983.
- [73] P. A. Tipler , *Physics* , Volume 1, New York: Worth Publishers, 1976.
- [74] F. M. White, *Fluid Mechanics*, New York: McGraw-Hill, 1979.
- [75] T. A. Wilson and G. S. Beavers, "Operating Modes of the Clarinet" *J. Acoust. Soc. Am.* 56(2):653-668, August 1974.
- [76] W. E. Worman, *Self-sustained Nonlinear Oscillations of Medium Amplitude in Clarinet-Like Systems*, Ph. D. Dissertation, Department of Physics, Case Western Reserve University, January 1971.
-