


Mathematics
Across the
Curriculum

COURSE GUIDE MUSIC AND COMPUTERS



A Theoretical and Historical Approach

Phil Burk
Larry Polansky
douglas repetto
Mary Roberts
Dan Rockmore

Your unique Access Code to password protected Web Resources can be found on a printed card in the plastic sleeve inside the back cover of your book.

MUSIC AND COMPUTERS

A Theoretical and Historical Approach

Course Guide

Phil Burk

SoftSynth.com

Larry Polansky

Dartmouth College

douglas repetto

Columbia University

Mary Roberts

Princeton University

Dan Rockmore

Dartmouth College



Key College Publishing

Innovators in Higher Education

www.keycollege.com

Phil Burk
75 Pleasant Lane
San Rafael, CA 94901

Larry Polansky
Department of Music
Dartmouth College
Hanover, NH 03755

douglas repetto
Computer Music Center
Columbia University
New York, NY 10027

Mary Roberts
Department of Music
Princeton University
Princeton, NJ 08544

Dan Rockmore
Department of Mathematics
Dartmouth College
Hanover, NH 03755

Key College Publishing was founded in 1999 as a division of Key Curriculum Press® in cooperation with Springer New York, Inc. We publish innovative texts and courseware for the undergraduate curriculum in mathematics and statistics as well as mathematics and statistics education. For more information, visit us at www.keycollege.com.

Key College Publishing
1150 65th Street
Emeryville, CA 94608
(510) 595-7000
info@keycollege.com
www.keycollege.com

Partial support for this work was provided by the National Science Foundation sponsored initiative: Mathematical Sciences and Their Applications Throughout the Curriculum under grant DUE-9552462 from the Division of Undergraduate Education.

© 2005 by Key College Publishing. All rights reserved.
Published by Key College Publishing, an imprint of Key Curriculum Press.

This work is protected by federal copyright law and international treaty. The book may not be translated or copied in whole or in part without the permission of the publisher, except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

®Key Curriculum Press is a registered trademark of Key Curriculum Press.

All other registered trademarks and trademarks in this book are the property of their respective holders.

Development Editor: Allyndreth Cassidy
Production Director: McKinley Williams
Production Project Manager: Ken Wischmeyer
Copyeditor: Erin Milnes
Text Designer: Ken Wischmeyer
Compositor: Interactive Composition Corporation
Art and Design Coordinator: Kavitha Becker
Printer: Data Reproductions

Editorial Director: Richard J. Bonacci
General Manager: Mike Simpson
Publisher: Steven Rasmussen

ISBN 1-930190-95-6
Printed in the United States of America
10 9 8 7 6 5 4 3 2 1 08 07 06 05 04

Contents

<i>Preface and Acknowledgments</i>	v
<i>How to Use the Course Guide and Web Site</i>	ix
<i>The Vision of Mathematics Across the Curriculum</i>	xi
<i>Introduction</i>	xiii
Part One: Abstracts, Questions, and Projects	
Chapter 1 The Digital Representation of Sound: Sound and Timbre	3
Chapter 2 The Digital Representation of Sound: Playing by the Numbers	9
Chapter 3 The Frequency Domain	19
Chapter 4 The Synthesis of Sound by Computer	27
Chapter 5 The Transformation of Sound by Computer	37
Part Two: Quizzes	
Quiz 1	47
Quiz 2	55
Part Three: Index of Applets, Soundfiles, Figures, Xtra Bits, and Tables	
Index 1 The Digital Representation of Sound: Sound and Timbre	65
Index 2 The Digital Representation of Sound: Playing by the Numbers	71
Index 3 The Frequency Domain	75
Index 4 The Synthesis of Sound by Computer	79
Index 5 The Transformation of Sound by Computer	85

Preface and Acknowledgments

This Web-based course is intended as a “user-friendly” introduction to music and computers for college and upper-division secondary students. By organizing a presentation around interactive visual and sonic examples, we hope to provide a resource and guide for those just beginning to explore the field of computer music, as well as for the more advanced computer composers who might benefit from a fresh insight. Our material can be used as a text for a one term or semester course, or as a module inside of a larger curriculum.

We have designed a flexible music synthesis learning environment that can accommodate a broad scope of curricular settings, including courses in contemporary music, composition, music technology, mathematics, computer science, and interdisciplinary studies. The many platform-independent interactive examples and audio aids stimulate the learner’s creative spirit while the multi-level approach (which contains plenty of optional tangents for in-depth mathematical studies) should cover all the bases for a comprehensive study.

Our personalized approach to generating most of the audio examples has been a homegrown one; we have worked to provide examples that can be easily replicated by laypersons who may have a minimal, yet working knowledge of a computer music language. We encourage all users of the Web-based text to take advantage of our sound-making modules and use and tweak them in any way they might imagine.

Our applets are written in Phil Burk’s JSyn and Java in another attempt to provide a platform-independent environment for designing interactive computer music examples. As instructors, we use a number of computer and music programs and platforms (SuperCollider, Csound, Max/MSP, Soundhack, JSyn, and others) but we were reluctant to make

any one of these a prerequisite for using this program. We of course encourage the instructors using this material to motivate their students to actually compose computer music, either by programming it themselves or by using higher-level software tools. The Web-based text is meant to be used in conjunction with that kind of hands-on activity. There is no substitute for doing it yourself!

One of the most difficult aspects of studying computer music is the rapid evolution and eventual extinction of software and hardware platforms—a familiar, yet frustrating aspect of the digital work environment. These developments are implicit for advancements in compositional tools, but difficulties crop up for keeping documentation current, especially with textbooks. The best standard texts (particularly those mentioned below by Roads and Dodge and Jerse) have elegantly avoided this problem by correctly focusing on ideas rather than software implementation, and we have followed suit.

Although we have created a great many of our examples in some of the most current, popular, and useful software (particularly those previously mentioned), we have no doubt that within a few years there will be a completely new set of standard tools. We also have no doubt that most of the fundamental ideas presented here will still be applicable, as they have been for some time now. We encourage students and instructors to find appropriate technology for their own compositions, and to avoid the kind of techno-consumer partisanship that so often inhibits real artistic and intellectual growth.

The important and advanced texts in the field of computer music, notably *Computer Music* by Charles Dodge and Thomas Jerse and *Computer Music Tutorial* by Curtis Roads, present comprehensive and in-depth studies of computer-generated sounds and related disciplines. We hope that our shorter, less technical work will lead the interested student and instructor to these essential works. We want our book to serve as the doorway, rather than the room itself.

Work on this text was made possible by a National Science Foundation Grant for Mathematics Across the Curriculum (MATC) at Dartmouth College. Many people at Dartmouth College have made invaluable contributions to the development of this book, including: Marcia Groszek and Dorothy Wallace of the Mathematics Department; Jon Appleton, Director of the Graduate Program in Electroacoustic Music and Co-Director of the Bregman Electronic Music Studio; and Claude

ents to
elves
to be

the

tion
arly
tly
tware

ie of

ere
bt that
e, as
ructors
void
l

sic,

nd
nes.
ed
to

at

g:
ent;
c Music
ude

Poux and Kim V. Rheinlander of the MATC Program. We would also like to thank our colleagues and the graduate students at the Bregman Studio for contributing in innumerable ways to this work: Charles Dodge, Dee Copley, Colby Leider, Lesley Stone, and Matthew Smith.

We would like to thank the staff at Key College Publishing for their work on the Web-based text and this *Course Guide*. Particularly, we would like to thank the reviewers solicited by Key to read the manuscript: James Bohn, University of Massachusetts-Dartmouth; Tim Koozin, University of Houston; Gary Lee Nelson, TIMARA Department, Oberlin College; and Eleanor Trawick, Ball State University.

Over the past several years, author Larry Polansky has taught an interdisciplinary computer music course at Dartmouth College. We owe a great deal of appreciation to the folks who have helped co-teach this interdisciplinary class: Dennis Healy and Charles Owen, both of the Dartmouth Mathematics and Computer Science Department. These indispensable and valued colleagues have made deep contributions to the collection of ideas presented here.

How to Use the Course Guide and Web Site

The *Course Guide* is your key to the online content and portion of this course. If you are visiting *Music and Computers* for the first time, you will need to register for the course with Key College Publishing at www.keycollege.com/online. There you will be prompted to enter the single-use registration access code printed on the card in the back of your *Course Guide*. After you enter the access code, it will expire, and you will be asked to create the unique username and password that you will use to access the online material for the duration of the course.

Before you begin, you will need to download the JSyn browser plugin to use the applets sprinkled throughout the online material. A link is provided on the *Music and Computers* Web site to lead you to the correct location to download this plugin. It provides real-time audio synthesis for Java applets and allows you to run interactive music and audio applets from your browser.

The JSyn Web site is maintained by author Phil Burk and gives you all the information you need to get the plugin up and running. If you have any problems downloading and installing the plugin, please contact the helpful folks that maintain the plugin site.

All of the content for this course can be found on the Web site. The *Course Guide*, which is divided into three parts, has been designed to help you in your studies and the use of the Web-based text. The first part will be useful as you prepare for tests and review what you have read online. Each section of the text is summarized in an abstract, and you are asked questions about the section. At the end of each chapter in the *Course Guide* are in-depth projects that your instructor may assign to you or that you may wish to research on your own.


The second part of the *Course Guide* contains two quizzes. The first quiz covers the material in the first two chapters of the online material. The second quiz covers the last three chapters. Try answering these questions without referring to the online material.

The final part of the *Course Guide* is an index of all applets, figures, soundfiles, tables and Xtra bits from the online material. This is useful if you need to quickly find an applet or a sound file. Each item also has a short description to remind you of its function.


Because this text is probably different from texts that you have read in the past, we encourage you to write notes in the *Course Guide*. Keep track of important concepts and terms and attempt to answer the questions that relate to each section.

The Web-based text is organized around numerous examples, and the narrative serves as a connecting thread. Students and instructors should use the individual topics as starting points for discussions; possible themes of departure could be musical, technological, or philosophical. You are encouraged to construct your own digital musical instruments using any available computer music system. Because creating and constructing are the most powerful learning strategies, you are expected to apply what you learn by working through the short compositional/ technological exercises (in the form of applets) included with each chapter.


While working through the online material, you will encounter three icons:



Applets for exploring interactive sound-making modules and images.



Xtra bits for learning more about a particular topic.



Sound files for hearing musical examples of the material described.

Although the Web-based material is best used sequentially, the first two chapters, the third chapter, and the last two chapters could function

ie first
terial.
se
gures;
seful if
has a
ve read
Keep

, and the
should
le
nical.
nents
l
pected
nal/
chapter.
er three
odules

ic.
aterial
he first
unction

as three independent units, which might be used in parallel. That is, teaching acoustics and basic computer ideas might (and ought to) be taught while teaching basic synthesis techniques, which are closely derived from computational and psychoacoustic precepts. Similarly, the third chapter, which covers the basics of DSP and Fourier analysis, is not necessarily a prerequisite to making computer music, or even to understanding most of the material in the later chapters. However, taught in parallel, this important mathematically-oriented material will greatly enhance the student's understanding of computer music and its tools.

The Vision of Mathematics Across the Curriculum

Dear Reader,

In 1994, Dartmouth College received a generous grant from the National Science Foundation to integrate mathematics throughout the undergraduate college curriculum in a five-year project, Mathematics Across the Curriculum (MATC). The project has involved over 40 faculty members from Dartmouth and various other colleges and universities representing departments of biology, chemistry, music, drama, English, art history, computer science, physics, earth science, economics, engineering, medicine, mathematics, and Spanish, producing lesson plans, short books, videotapes, and a Web site with images and text. The series of volumes published by Key College Publishing represents some of the best of the MATC collection.

These materials will make it easier for students to become more quantitatively literate as they tackle complex, real-world problems that must be approached through the door of mathematics. We hope that you, the reader, will appreciate our efforts to place the mathematics in this book completely in the context of your field of interest. Our goal is to help you see that applied mathematics is a powerful form of inquiry, and ever so much richer than mere "word problems." We trust that you will like this approach and want to explore some of the other volumes in the series.

Sincerely,

Dorothy Wallace

Professor of Mathematics

Principal Investigator: Mathematics Across the Curriculum Project
Dartmouth College

Introduction

In the early 1960s, a way of efficiently turning sound into bits and bytes was discovered. These analog-to-digital converters, capable of turning one second's worth of sound into 300,000 numbers (on the fly!) made it possible to transform the acoustic waves that we hear as sound—conversations, car horns, your mother-in-law's voice, and even music—into long sequences of numbers, which were then stored in computer memories.

From there the numbers could be turned back into the original sounds, but this was not a big advance because the ability to simply record a sound had been used for quite some time. The major advance was more than that: Now sound could be manipulated very easily, just as a chunk of data. The creative potential for musical composition and sound generation empowered a revolution in the world of music. That revolution, electroacoustic music, engendered a wonderful synthesis of music, mathematics, and computing.

This is a complete course that introduces the reader to the mathematical, physical, and computer science challenges and achievements that have made digital audio and electroacoustic music possible.

PART ONE

ABSTRACTS

QUESTIONS

PROJECTS

CHAPTER ONE

The Digital Representation of Sound Sound and Timbre

SECTION 1.1

What Is Sound?

Abstract

Sound is a complex phenomenon involving physics and perception. A mathematical function is an input/output machine and can carry information about the physical characteristics of sound, including amplitude, frequency, and timbre. Sound sensation and perception is the end result of a complicated process, including the physics of sound and the physiology of our ears. The relationship between a physical sound and its physiological interpretation is incomplete, due in part to an imprecise descriptive vocabulary.

Questions

1. What two values are plotted to make the graphical “waveform” of a sound?
2. What is a transducer, and which part of the ear functions as one?

3. What is the basic perceptual correlate of a sound's amplitude? Of its frequency?

SECTION 1.2

Amplitude

Abstract

Acoustic and cognitive terminologies describing sound are related, but not identical. The perceived loudness and pitch of a sound, though independent, have a subtle relationship. The mathematical concept of amplitude is related to the mechanical concept of intensity. The brain perceives changes in pitch and intensity primarily in terms of ratios.

Questions

1. Are perceived pitch and loudness completely independent of one another? Completely dependent? Somewhere in between? Explain.
2. What term describes the effect of a sound's amplitude on a particular medium?
3. Are changes in a sound's perceived pitch linear or logarithmic? How does this affect our perception of those changes?

Frequency, Pitch, and Intervals

Abstract

The frequency of a sound is a measurement of how often a given event repeats in time and can be defined in terms of the waveform's period and wavelength. The perception of pitches is almost universally based on logarithmic intervals between frequencies. Equal-loudness contours (such as the Fletcher-Munson curves) indicate how much intensity is needed at a certain frequency to produce the same perceived loudness as a tone at a different frequency and can be used to explain the "functional" relationship of perceived loudness to a sound's pitch.

Questions

1. Give the wavelength and period of a 2.4 kHz tone (assuming standard conditions).
2. What is the mathematical relationship between octaves?
3. Listening to a 50 Hz tone, a 5 kHz tone, and a 20 kHz tone all sounding at the same intensity, which tone will be perceived to be the loudest? The softest?

Timbre

Abstract

Timbre is the term used to approximate the “quality” or color of a sound, although its scientific meaning is less than clear. The envelope and spectra of a sound are thought to be the most significant components of timbre. Every periodic waveform can be expressed as a sum of sinusoids whose frequencies are integer multiples of the fundamental and whose amplitudes are variable. Fourier analysis is used to explain sounds and other wave phenomena.

Questions

1. What two components of a sound are thought to most influence its timbre?
2. What are the three basic sections of an amplitude envelope?
3. You have just performed the Fourier analysis on a trumpet tone. What information does the series of sine waves you end up with give you about that sound? Does this information alone define the sound's timbre? Why or why not?

PROJECTS

1. Write a simple sonification program of the digits of π , but be creative! Consider how your mapping algorithm can translate the input into changes not only in the pitch, but also in the rhythm, envelope, and even timbre of the tones.
2. Take a short piece of music (maybe “Happy Birthday” or “Mary Had a Little Lamb”) and code it so that the loudness of each note is determined by its pitch. Why does this sound so unnatural?
3. Find out how your ears compare to your friends’. Write a program to sweep a sine wave from 0 Hz to, say, 25 kHz, and see where the limits of your hearing are. Test your friends and your professors. Do you see any consistent differences due to age, listening habits, or any other factor?
4. Research the ways different animals—mammals, insects, anything—make use of sound. What are some common uses? Uncommon?
5. Collect a series of instrumental tones, and make a copy of each with its attack removed. Test how well your friends can identify the set without attacks, and then the set with attacks. What could explain any difference in performance?

CHAPTER TWO

The Digital Representation of Sound Playing by the Numbers

SECTION 2.1

The Digital Representation of Sound

Abstract

In order to store a finite approximation of a continuous real-world function, values are recorded at discrete time intervals, a process known as sampling. Analog to digital and digital to analog conversion are the two essential processes required for the recording and playback of sampled material, respectively.

Questions

1. Why is it necessary to sample sounds we wish to record in the real world?

2. Explain, roughly, what the digital to analog converter in a CD player does to help turn the sampled data on a disc into sound from a pair of speakers.

3. What pure tone is the basic building block of all sounds?

SECTION 2.2

Analog Versus Digital

Abstract

There is an important distinction between the continuous nature of physical information and the discrete nature of digital information. A graphical display of the quantization of digital data aids in the comparison of analog and digital waveforms, underscoring the necessarily approximate nature of all sampled data.

Questions

1. What is the main difference between analog and digital representations of information?
2. What is quantization, and why is it a necessary byproduct of sampling?
3. What would be an advantage of a high sampling rate? A disadvantage?

Sampling Theory

Abstract

The Nyquist theorem states that the sampling rate must be twice the highest frequency sampled to provide an accurate recording.

Undersampling occurs when frequencies higher than the Nyquist rate are sampled, causing unwanted aliasing. Anti-aliasing filters are necessary in data conversion to prevent these unwanted frequencies from entering the system.

Questions

1. Bats can produce and hear frequencies from 10 Hz to 100 kHz. According to the Nyquist theorem, what sampling rate would be necessary to accurately record an a cappella bat group?

2. What problems can undersampling cause in a recording?

3. Considering the way humans do—or don't—perceive frequencies above 20 kHz (as well as the fact that hearing is also a form of sampling), can you conclude that the human auditory system involves band-pass filtering? Explain.

Binary Numbers

Abstract

Number systems are built on the concept of bases and are ultimately equivalent representations of a given quantity. In the binary world, each bit of a number signifies either the presence (1) or absence (0) of successive powers of 2. A binary number of n bits can represent 2^n values.

Questions

1. What is the decimal value of the binary number 1001011? What is the binary value of the decimal number 120?
2. Keeping in mind the basic idea of sampling, why does it make sense that digital systems operate with binary numbers?
3. A CD contains audio data with a sampling rate of 44.1 kHz and 16 bits of information per sample. How many bits make up a 5-second segment of music on a CD?

SECTION 2.5

Bit Width

Abstract

The digitization of sound carries an inherent tradeoff between accuracy in representation and storage requirements. Bit width (or depth) refers to the amount of data, in bits, recorded at each sample point. The resolution of a sample depends on the resolution available in the number of bits used to record it.

Questions

1. Why does a larger bit width produce a more accurate sample? What is the downside of using large bit widths?
2. How many bits are in a nibble? In a byte?
3. How does bit width affect the perceived timbre of a recording?

SECTION 2.6

Digital Copying

Abstract

Because digital recordings are simply lists of numbers, reproductions are generally perfect—in fact, essentially identical. Analog reproductions,

however, introduce noise into successive generations, gradually degrading the quality of the copies. The difficulty in determining the original source of a digital recording has led to serious legal questions about intellectual property and copyright.

Questions

1. Why can digital copying claim to be free of information loss, unlike analog copying?
2. What details are lost in copies of analog recordings? What errors are introduced?
3. Why do you think the rise of digital recording has posed a much bigger problem for copyright issues than analog recording?

SECTION 2.7

Storage Concerns: The Size of Sound

Abstract

Digital information is written to a CD as a series of long and short pits, which represent the individual bits of the data and are read by light reflected from a laser. In general, CD quality sound requires about 10 MB per minute.

degrading
nal source
tellectual

unlike

rors are

uch bigger

ort pits,
light
out 10 MB

Questions

1. How is digital information encoded on a CD?
2. Roughly speaking, how much storage space is required for one minute of CD quality sound?
3. How much storage space is required to store one minute of stereo sound sampled at 192 kHz with a 64-bit resolution?

SECTION 2.8

Compression

Abstract

Data compression is intimately related to the transmission of digital information and involves techniques that effectively shorten the string of symbols required to represent a sound without compromising its quality. Three common techniques are elimination of redundancy, perceptual encoding (μ -law), and prediction-based encoding. A tradeoff generally exists between computation time and effectiveness of compression.

Questions

1. What feature of human sound perception does μ -law encoding exploit?

2. Which component of sound are prediction algorithms typically based upon?
3. Huffman coding, which is widely used as a final layer of compression in the transmission of all digital data, involves ordering strings of bits so that the common strings can be represented more compactly than uncommon strings. Which encoding technique would describe this algorithm?

PROJECTS

1. Using a cassette-to-cassette recorder, record a short piece of music back and forth several times so that the each copy is of a previous copy. How quickly do you notice a significant degradation in sound quality? Why is this happening?
2. Having just purchased a 60-gigabyte hard drive, how many minutes of CD quality sound (44.1 kHz, 16-bit) will you be able to store on the entire disk?

lly

3. How does MP3 encoding work? With a little experimentation of your own, determine what compression ratios the MP3 codec typically produces.

pression in
of bits so
, than
be this

4. Do a bit of research into how the recording industry is handling the digital copying explosion. How do you think this phenomenon will ultimately impact the industry?

music back
: copy. How
dity? Why

5. The current trend in audio engineering is toward recording at 192 kHz with 64-bit resolution. Why does it seem like a waste of space to use such a high sampling rate and bit width? Search the Internet to find what the purported advantages of extremely high-resolution sampling are. (Try www.digidesign.com for starters.)

minutes of
on the

CHAPTER THREE

The Frequency Domain

SECTION 3.1

The Frequency Domain

Abstract

Changes in a sound's amplitude can be described by a peak envelope and an average value envelope, both of which exist in the time domain. A sonogram incorporates frequency, time, and amplitude into a single graph. A visualization of a sound's frequency content better represents its timbral qualities than a time-domain plot.

Questions

1. What is the difference between the transient and steady state portions of a sound?
2. What are two common enveloping techniques used to describe the amplitude of a sound? How do they differ?

3. What information does a sonogram provide that a time-domain plot does not?

SECTION 3.2

Phasors

Abstract

Phasors are useful mathematical tools for describing the basic sinusoids that make up all sounds. The Fourier transform supplies the sinusoids that constitute a sound's frequency content. Waterfall plots show how a sound's spectrum changes over time.

Questions

1. In terms of phasors, how is a sine wave related to a cosine wave?
2. How are partials mathematically related to the fundamental of a periodic tone?
3. What is the result of adding two identical sinusoids of opposite phase?

Fourier and the Sum of Sines

Abstract

Fourier series are infinite sums of sines and cosines with coefficients that describe the properties of each partial in a sound. Fourier analysis determines the amplitude, frequency, and phase of each of the partials, and Fourier synthesis re-creates a sound from the analyzed sinusoids. A sound's spectrum can be modified directly between these two steps by filtering.

Questions

1. What does the DC term in a Fourier series tell you about a sound?
2. With respect to the audible spectrum, where is most of the energy in a typical sound wave found?
3. Which frequencies are in the passband of a high-pass filter? Of a low-pass filter?

SECTION 3.4

The DFT, FFT, and IFFT

Abstract

The FFT is an optimized computer algorithm for quickly computing the mathematical discrete Fourier transform of a signal. The FFT involves windowing, periodicizing, and transforming the signal into magnitude/phase pairs for given frequency bands in the spectrum.

Questions

1. What assumption is made in computing the Fourier transform of small sections of an aperiodic waveform?
2. Why is it necessary to break a signal's spectrum into discrete frequency bands when computing its FFT?
3. How does the resolution of the FFT change over the frequency spectrum?

SECTION 3.5

Problems with the FFT/IFFT

Abstract

Because of the mismatch between the FFT's linear analysis of frequency data and the logarithmic nature of human auditory perception, the FFT

provides poor frequency resolution at the low end of the spectrum and extreme resolution at the high end. Consequently, a compromise is necessary between time and frequency resolution, which, depending upon the FFT size, smears information in one domain or the other.

Questions

1. In what part of the spectrum is most of an FFT analysis “wasted”? Why?
2. What information is lost by expanding the frame size of an FFT?
3. What possible solutions or workarounds to the time/frequency tradeoff of the FFT can you propose?

SECTION 3.6

Some Alternatives to the FFT

Abstract

Wavelet analysis, which has dynamic bandwidth and makes use of non-sinusoidal waveforms, can provide a more accurate transform but is significantly more complicated to implement. McAulay-Quatieri (MQ) analysis employs frequency tracking to map the trajectories of spectral components over time. MQ analysis provides more perceptually significant information than the standard FFT.

Questions

1. How does wavelet analysis offer a solution to the time/frequency tradeoff inherent in the FFT algorithm?
2. What additional information does MQ analysis provide about the sinusoids extracted in the standard FFT algorithm?
3. What features of human psychoacoustics does MQ analysis take advantage of?

PROJECTS

1. Audio engineers frequently need to deal with the problem of unwanted phase filtering effects (also known as comb filtering) when recording one source with multiple microphones. Comb filtering arises when two microphones record the same sound source at slightly different times due to the unequal distances they are placed from the source. Based on your knowledge of sinusoidal phase and vector addition, what sort of filtering would you expect as a result of adding two identical recordings that are only slightly out of sync with one another? Try using two microphones to record the same source, and experiment with their relative distances to see how the mixed sound is affected.

icy

2. Research wavelet analysis on the Internet. Does this technology seem like a viable alternative to the FFT? What are some of its drawbacks?

the

3. What is the bin width for a 1,024-point DFT of a sound sampled at 44.1 kHz? For a 512-point DFT of a sound sampled at 22 kHz? How should the window size be adjusted to provide better time resolution?

ike

4. Use a program with spectral analysis features (such as SoundHack) to examine the spectral content of several sounds. Play with the processors and effects that make use of the FFT to see how the different options affect the output of the algorithm.

inwanted
ording
when two
nt times
Based on
t sort of
recordings
two
t their

5. Additive synthesis combines pure tones to create more complex waveforms. Using any standard music language, try building new sounds from a set of sines. How should the sines be related to one another to be harmonic? Experiment by adjusting the ratios of the partial frequencies and individual amplitudes.

CHAPTER FOUR

The Synthesis of Sound by Computer

SECTION 4.1

Introduction to Sound Synthesis

Abstract

Computer sound synthesis involves generating audio data internally rather than recording it from the real world. At the simplest level, a computer generates a list of sample values given a particular function (the waveform) and a series of time points as input. More generally, filters are used to process and transform existing audio data.

Questions

1. What is the basic difference between one color of noise and another?

2. How does the input to a function differ from the input to a filter in computer music?

3. Given a playback rate of 44.1 kHz, what is the highest-frequency sine wave you could generate without aliasing?

SECTION 4.2

Additive Synthesis

Abstract

Additive synthesis involves combining sine wave partials to create more complex tones. Although it is conceptually straightforward, the process is computationally expensive and is better suited to steady state sounds, which are less interesting than dynamic sounds in a psychoacoustic sense. Shepard tones take advantage of the human perception of pitch "chroma," by creating apparently infinite glissandi effects.

Questions

1. How is a square wave constructed from the addition of sine waves?
2. What aspect of natural sound is perceptually most significant to the brain?
3. What feature of human pitch perception do Shepard tones exploit?

Filters

Abstract

Subtractive synthesis relies on filters to “sculpt” a signal, removing (and possibly adding back) portions of its spectrum. The four basic types of filter are high-pass, low-pass, band-pass and band-reject. FIR (finite impulse response) filters always output less than their input, whereas IIR (infinite impulse response) filters can output more, since they feed the processed signal back into the filter. Feedback, however, makes these filters prone to “blowing up.”

Questions

1. Which frequencies are in the stopband of a high-pass filter? Of a low-pass filter? Of a narrow band-pass filter centered at 10 kHz?
2. What is the significance of the value .707 in terms of a filter's cutoff frequency?
3. What is the essential difference between FIR and IIR filters? Why are IIR filters more likely to “blow up”?

Formant Synthesis

Abstract

Formant synthesis involves modeling the resonant physical structures of a system (for example, a wind instrument or vocal tract). Formants are fixed frequency peaks that provide significant information about an instrument's timbre. By manipulating formants, synthetic vocal sounds (primarily vowels) can be accurately produced.

Questions

1. How are the formants affected when a trumpet player plays an A followed by a C?
2. Why do instruments with similar shapes share similar timbres?
3. Why is formant synthesis an incomplete approach to speech synthesis?

Amplitude Modulation

Abstract

Modulation is a powerful synthesis technique capable of producing and controlling a wide range of tones. A basic amplitude modulated (AM)

signal is created by connecting the output of one oscillator (the modulator) to the amplitude port of another oscillator (the carrier). Low frequency amplitude modulation produces tremolo, whereas higher modulation frequencies change the spectrum of the carrier.

Questions

1. What is the effect of low frequency modulation on the output of an FM signal?
2. What is a unit generator? Give a few examples.
3. What do the sidebands of a modulated signal represent?

SECTION 4.6

Waveshaping

Abstract

Waveshaping requires a transfer function that distorts an input signal (often a simple sine wave) to create a more complex tone. Chebyshev polynomials are popular transfer functions, since they allow for direct control over the harmonic content of the output signal. A common idiom in computer music is to create a wave table to store pre-computed values of a complex transfer function and then do a "look up" in the table according to the current value of the input signal.

Questions

1. How many harmonics can be present in a signal distorted by a fifth-order polynomial? What is the highest-frequency signal that could be passed through this transfer function at a sampling rate of 44.1 kHz without aliasing?
2. Why are Chebyshev polynomials considered frequency multipliers?
3. Who is Don Buchla, and what does he do?

SECTION 4.7

FM Synthesis

Abstract

Frequency modulation synthesis uses the output of one oscillator (the modulator) to control the frequency of another (the carrier). When audible frequencies are used for the modulator, the sidebands are introduced into the carrier's output signal. By controlling these sidebands, or spectra, FM synthesizers are capable of efficiently producing a wide range of timbres.

Questions

1. What parameter of a basic FM oscillator setup controls the depth of modulation?

2. What effect is produced by subsonic modulation frequencies (those lower than 30 Hz)?

3. What determines the frequencies of the partials in an FM spectrum?

SECTION 4.8

Granular Synthesis

Abstract

Granular synthesis involves the addition of a number of grains of sound, which are short bursts of sound derived from any waveform, not simply sinusoids. Manipulating the amplitude envelopes, frequencies, and density of a cloud of grains can produce unique sonic results.

Questions

1. How is granular synthesis similar to additive synthesis? How is it different?

2. What is a typical length of a single grain?

3. What two composers are considered the pioneers of granular synthesis?

Physical Modeling

Abstract

The aim of physical modeling synthesis is to understand how a sound is generated in the real world and to re-create the process computationally. The Karplus-Strong algorithm is a simple example of this technique, in which the repeated averaging of digital noise is used to simulate a plucked string. In general, the parameters of physical models correspond closely with the physical attributes of a sound source.

Questions

1. How does the harmonic content of a plucked string change over time?
2. How does averaging random noise simulate this change?
3. What are the computational correlates of the pitch, initial energy, and length of vibration of a plucked string?

PROJECTS

1. Experiment with different partial ratios in the creation of an additive synthesis tone. Do any patterns become apparent?

2. Use subtractive synthesis techniques to create interesting timbres from different colors of noise. How does this technique compare to additive synthesis?

3. Using the code provided in Section 4.7 as a starting point, experiment with the creation of FM tones in Csound. Try to understand how the parameters of the foscil generator change the timbre of the sound.

4. Write a program for waveshaping a simple tone, and compose a short piece with your software.

5. Research the work done on physical models of your favorite instrument. How closely do the designs relate to the actual instrument?

CHAPTER FIVE

The Transformation of Sound by Computer

SECTION 5.1

Introduction to the Transformation of Sound by Computer

Abstract

Computers offer a powerful platform for editing and transforming pre-existing sounds. One obvious technique involves the rearrangement of sonic information, which digital software makes possible at the sample level. Samplers and drum machines, which are often incorporated into computer software, provide tools for working with samples. Digital-audio workstations (DAWs) have made quality recording and editing capabilities widely available, revolutionizing the production of music.

Questions

1. How is the term *sampling* used differently from its traditional meaning in regard to certain compositional styles?
2. What company produced the first DAW?

3. How are drum machines and samplers related?

SECTION 5.2

Reverb

Abstract

Reverberation is an extremely popular, versatile computer music effect. One approach to producing reverb uses physical modeling to simulate the delayed copies of a sound as it bounces around a space. With increases in processing power, producing reverb through convolution has become more popular. Convolution involves the cross-multiplication of samples in the time domain or, equivalently, the point wise multiplication of Fourier coefficients in the frequency domain.

Questions

1. Explain why convolution is considered a “running average.”
2. What is meant by the number of “taps” for a convolution filter?
3. What are the components of an all-pass filter? How are these filters used to model reverberation?

Localization/Spatialization

Abstract

Spatial location of sound in human perception depends on three primary factors: the relative loudness, time difference, and frequency content of a sound at each ear. Head-related transfer functions (HRTFs) describe the way the human head and shoulders filter frequencies at the ears. A binaural dummy head records sounds much the way they would be perceived by a real person. Interestingly, although humans can accurately and quickly locate a sound in the horizontal plane, they have difficulty perceiving sound sources in the vertical plane.

Questions

1. Why is relative loudness alone not a very realistic simulation of sound location?
2. What is the overall filtering effect of the human head on sound?
3. What is one possible explanation for why humans lack accurate sound localization in the vertical plane?

Introduction to Spectral Manipulation

Abstract

The phase vocoder is made up of a set of techniques for deconstructing a sound, manipulating it in the spectral domain, and resynthesizing it. This process relies heavily on Fourier analysis and the FFT in particular. One powerful application of the phase vocoder is changing the pitch of a sound without changing its duration, and vice versa.

Questions

1. Why is duration directly related to frequency in standard playback of recorded sounds?
2. What artifacts does a Fourier analysis produce when smoothing windows are not used?
3. How would a sound change if it were resynthesized using a larger frame size?

More on Convolution

Abstract

The phase vocoder can also be used to efficiently implement convolution by taking advantage of its equivalence to spectral multiplication. Beyond using the impulse response of reverberant spaces to produce reverb, convolution can be used with any two arbitrary samples for cross-synthesis. The results of such convolutions are often physical impossibilities, for example, a speaking guitar chord (composed of a singer's voice and a guitar chord).

Questions

1. How does the phase vocoder facilitate an efficient implementation of convolution?
2. What is the general definition of cross-synthesis? Give an example of its use.
3. What is the effect of convolving a sound with white noise? Can you explain this result?

Morphing

Abstract

Sonic morphing can be achieved in a variety of ways: in the time domain, with replacement and interpolation morphing; and in the frequency domain, with feature morphing. Time-domain morphing is a sort of “smart” cross-fade between two sounds and works with sample values rather than frequency content. Feature morphing, on the other hand, takes advantage of a salient feature of two spectra (for example, their centroids) and morphs one into the other over time.

Questions

1. What type of morph can be described as “morphing somewhat all of the time”? What type of morph can be described as “morphing completely some of the time”?
2. What perceptual characteristic of a sound does its spectral centroid measure?
3. How is a centroid defined mathematically?

Graphical Manipulation of Sound

Abstract

Graphical manipulation of sound is a developing field in which techniques are designed for visually representing and editing sonic material. Prior to the digital era, several audio researchers designed systems for transforming visual input into sonic output, notably Xenakis's UPIC. With computer technology, many more advanced systems have become available, such as AudioSculpt and SoundHack, which provide the user direct access to a sound's spectrum in a visual environment.

Questions

1. What does the term *synaesthetic* mean? How does it apply to the way humans intuitively describe sounds?
2. Why are graphical representations of sound often ambiguous?
3. What form of synthesis did Hugh LeCaine's Spectrogram employ? How was this device similar to Xenakis's UPIC system?

PROJECTS

1. Try "shooting" your own reverb. Record the sound of a balloon popping in a space you'd like to capture (make sure you've got permission first if you need it!) and then, using software such as SonicFoundry's Sound

Forge, generate the reverb characteristic of the room. Now try convolving this sound with another sample. Can you “hear” the room?

2. As described in Section 5.3, time delays between stereo channels can produce localization effects. Yet, when the delay becomes long enough, the illusion disappears and we simply hear two copies of the sound at distinct times. Experiment with a stereo sound file to see how long you can delay one channel before the localization effect disappears. Do your friends perceive this “critical duration” the same way?

3. Compose a piece based on Steve Reich’s instructions for “Slow Motion Sound” and his additional comments.

4. Use a convolution patch (available in current versions of Sound Forge and SoundHack) to experiment with cross-synthesis. Combine many different sounds and note the results. Can you describe any general rules for achieving particular results with cross-synthesis?

5. Find a short QuickTime movie and save its audio track as a separate sound file. Using SoundHack’s QT-coder, transform the audio track into a QuickTime movie, and transform the original movie data into audio. Does this new “inverted” movie bear any relation to the original?

room?

els can
g enough,
ound at
long you
s. Do your

ow Motion

ound Forge
bine many
general

a separate
dio track line
a into audio
riginal?

PART TWO

QUIZZES



Quiz One

A. Sound as a Function

1. Describe a way in which sound can be viewed as a mathematical function. What are the possibilities for the axes? Why should sound be considered a function?
2. What is meant by the envelope of a sound? Sketch an example. Label the parts of the typical sound event.
3. What function (to a very good approximation) describes the sound that a tuning fork generates? Include a brief justification as to why this might be a reasonable model.
4. What is a periodic function?

5. For the sinusoidal function $f(t) = 5 \sin(10\pi t + \pi/6)$,
- Identify and define its amplitude, frequency, phase, and period.

- Sketch the graph of the function.

6. Suppose that the fundamental frequency of a periodic function is 200 Hz.

- What is its period?

- What is the frequency of its second partial?

c. What is the amplitude of its second partial?

7. The graphical representation (in terms of amplitude) of the sound generated by playing, say, a violin note, can be broken up into three parts.

a. Sketch a likely amplitude envelope for such a sound.

b. Name the three primary parts and describe them in terms of their frequency content.

8. What frequency is at one octave above 200 Hz?

9. What is the approximate range of frequencies of the human voice? Of a piano? Of human hearing? What do we call frequencies below this range? Above this range? How do we as humans organize that frequency range for the purposes of pitch perception?

10. We now know that normal conversation occurs at about 60 decibels. What number of decibels correspond to a sound twice as loud?

11. Describe how we get from amplitude to loudness, both from a perceptual and a mathematical perspective. Provide common units and include "intensity" as an intermediary measure.

B. The Digital Representation of Sound

12. What are two potential problems in trying to represent on a computer the analog or continuous functions that describe sound?

13. What is quantization?

14. What is bit width?

15. Using 6 bits:

a. What is your age in binary? (Give it in decimal as well, so that your binary response can be checked.)

b. What is 57 in binary?

c. What is the highest number you can represent?

d. What is the most significant bit? The least significant?

16. In order to represent a sound that contains only frequencies up to 10,000 Hz, how fast do you need to sample it? What is the name of this sampling rate?

17. What is aliasing? Explain with a picture.

18. What is a high-pass filter? Sketch its graph. Label all relevant portions of the picture.

C. Acoustics and Cognition

19. What are the psychoacoustic correlates of amplitude? Frequency?

20. What is timbre?

21. Explain, succinctly and accurately, what Fletcher-Munson curves describe.

D. Storage

22. Give an example of a compression technique, and explain a potential (sonic) drawback.

23. If you sample a stereo signal at 20,000 Hz for 1,000 seconds and use 2 bytes to store each sample, how much memory will you need?

to
e of

ortions

E. Exposition


Your roommate hears that you are taking a course called "Computers and Music" and asks you how it is that mathematics and computers can have anything to do with music. You are suitably shocked. What do you tell your naïve friend? (Write one to two paragraphs.)

Quiz Two

A. The Frequency Domain

1. What is the psychoacoustic correlate of frequency?

2. The following figure should look familiar: it is a time/frequency plot of 1.2 seconds of sound taken from the Web text. Time is indicated on the horizontal axis and frequency on the vertical axis.



- a. Explain in words what the graph represents (i.e., indicate the meaning of the darker and lighter shades of grey, the horizontal lines of varying widths, and so forth).
- b. Sketch the spectral envelope (i.e., the frequency content of the sound) at 0.1 second.
- c. Sketch the way in which the amplitude of the partial at 2,000 Hz changes over the 1.2 seconds of sound.
3. In honor of the change of seasons, your artist roommate gives you as a gift a snippet of Wagner's "Spring Song" from *The Valkyrie*, laminated on a piece of redwood. You are suitably touched—and in return you take a moment to explain how musical notation is also a time/frequency representation. Write a paragraph discussing this theme.

ean-
of

4. Explain what a phasor is and how it represents a sinusoid. In particular, explain the characteristics of a phasor that is meant to represent the sinusoid

$$s(t) = 12 \sin(6\pi t + \pi/4)$$

B. Fourier Analysis

5. What is Fourier analysis?

6. What is a Fourier series?

7. What is the Fourier transform?

8. What is Fourier synthesis?

9. What is the inverse Fourier transform?

C. The FFT

10. What is an algorithm?

z

l as a
ited
ou

e.

11. What is the fast Fourier transform (FFT)?

12. What is the input to the FFT?

13. What is the output of the FFT?


14. For what particular application was the FFT originally invented?

15. What is meant by the complexity of an algorithm?

16. In what manner did the FFT change the face of electroacoustic music?

17. If the FFT takes as input 2^{11} data points, roughly how many operations are required to complete the FFT? Roughly how many operations would be required if we simply computed the answer directly (i.e., without using the FFT)?

18. The following is a "waterfall plot" taken from the Web text. Please explain it.



19. Suppose that the output of an FFT on 16 samples looked like this:

(1, 2, 3, 4, 5, 22, 3, 45, 50, 9, 11, 12, 15, 22, 100, 10000)

a. Explain what these numbers mean.

b. How could we transform this vector of numbers in order to effect a low-pass filter?

c. How could we transform this vector of numbers in order to effect a high-pass filter?

20. Suppose that we are sampling at the Nyquist rate, with a frame size of 1,024 samples.

a. How many frequency bins will we get?

b. What is the width of each of the frequency bins?

21. Briefly explain the tradeoffs between resolution in time and frequency when considering whether you should increase the frame size.

22. What is time smearing?

D. An Odd and an End

23. Here is a complex number:

$$4 + 3i$$

a. What is its real part?

of

ncy

- b. What is its imaginary part?

 - c. Represent this complex number as a point in the complex plane.
(Don't forget to label the axes!)

 - d. What is its magnitude? Give a geometric interpretation.

 - e. What is its phase? (It is sufficient to give a formula for computing it.)
24. Sinusoids are but one way of decomposing a sound. Can you name one other?

PART THREE

APPLETS

SOUNDFILES

FIGURES

XTRA BITS

TABLES

SECTION 1.1

What Is Sound?

Applet 1.1 Hearing a sine wave oscillator

Figure 1.1 Graph: Time-domain plot of a waveform zoomed out many times

Figure 1.2 Graph: Close example of the previous waveform zoomed in many times

Xtra bit 1.1 Sonification

Figure 1 Photo: Charles Dodge

Soundfile 1 Excerpt from "Earth's Magnetic Field," by Charles Dodge

Soundfile 2 Sonification of the raw data on a hard drive ("I am a Nerd," by Phil Stone)

Figure 2 Graph: Discretization of space (Edward Childs's fluid dynamics sonification)

Soundfile 3 CFDSound2, by Edward Childs

Soundfile 4 CFDSound5, by Edward Childs

Figure 1.3 Graph: Temperature as a linear function

Soundfile 1.1 Synthetic voice

Figure 1.4 Graph: Synthetic voice waveform for Soundfile 1.1

Soundfile 1.2 Bird sounds

Figure 1.5 Graph: Bird sounds waveform

Soundfile 1.3 Gamelan

Figure 1.6 Graph: Gamelan waveform

Soundfile 1.4 Tuba sounds (example of "low")

Soundfile 1.5 More bird sounds (example of "high")

Figure 1.7 Diagram: The inner ear

Soundfile 1.6 Trumpet sound

Soundfile 1.7 Loud sound that gets softer

Soundfile 1.8 Whooshing noise

Soundfile 1.9 Chirping

Soundfile 1.10 Phat groove

SECTION 1.2

Amplitude and Loudness

Xtra bit 1.2 Another sonic universe

Figure 1.8 Chart: Psychoacoustic correlates

Applet 1.2 Changing the amplitude of a sound

Soundfile 1.11 Two sine waves played in sequence with different amplitude

Figure 1.9 Graph: Two sine waves with the same frequency and different amplitude

Soundfile 1.12 Two sounds played in sequence with different waveforms

Figure 1.10 Graph: A sine wave and a saw tooth wave with the same frequency and amplitude

Figure 1.11 Graph: An envelope of a sound wave

Figure 1.12 Graph: Two sine waves starting at different phase points

Soundfile 1.13 A sine wave swept over the frequency range 0–20,000 Hz (sometimes called a *chirp* in acoustic discussions)

Figure 1.13 Graph: Bubba & U (the difference between ratiometric change and fixed arithmetic change)

Xtra bit 1.3 More on decibels

Table 1.1 Average decibel levels for environmental sounds

Applet 1.3 Amplitude versus decibel level

SECTION 1.3

Frequency, Pitch, and Intervals

Xtra bit 1.4 Tapping a frequency

Applet 1.4 Hearing frequency and amplitude

Soundfile 1.14 Two sine waves with different frequencies

Figure 1.14 Graph: Two sine waves with the same amplitude, but with frequencies differing by one octave (a ratio of 2:1)

Applet 1.5 When ticks become tones

Table 1.2 Some frequency ranges (lowest and highest frequencies of a piano, female and male speech, a compact disc, and the human hearing)

Soundfile 1.15 Low gong (a Javanese gong sound gives an example of very long wavelengths)

Figure 1.15 Photo: Central Javanese gongs

Xtra bit 1.5 Animals and frequency (a comparison of animal and human hearing)

Applet 1.6 Transpose using multiply and add

Applet 1.7 Octave quiz

Figure 1.16 Graph: Comparison of logarithmically increasing frequencies (doubling the frequency to move up an octave) and linearly increasing frequencies

Applet 1.8 Example of Fletcher-Munson curves

Figure 1.17 Graph: Fletcher-Munson curves (equal-loudness contours)

SECTION 1.4

Timbre

Applet 1.9 Draw a waveform (The timbre of a tone is related to the shape of the repeating waveform.)

Figure 1.18 Graph: Attack, sustain, and decay of a trumpet tone: waveform and simplified graph of the envelope

Soundfile 1.16 Tuning fork at 256 Hz

Figure 1.19 Photo: Tuning fork

Soundfile 1.17 Sine wave at 256 Hz

Xtra bit 1.6 The mathematics of pure tones

Figure 1 Illustration: Triangle wave

Figure 2 Illustration: Sine wave

Figure 3 Graph: Lots of sine waves (same amplitude, different frequency)

Figure 4 Graph: Lots of sine waves (same amplitude, same frequency, different phases)

Figure 5 Graph: Single sine wave (result of adding up out-of-phase example shown in Figure 4)

Figure 1.20 Animation: Click on each tuning fork to listen to the different pitches and view the audiograms

Soundfile 1.18 Warren Burt's piece for tuning forks, "Improvisation in Two Ancient Greek Modes"

Figure 1.21 Animation: Sine and cosine waves

Figure 1.22 Graphs: Spectra of a sawtooth and a square wave (relative amplitudes of sinusoidal components of simple waveforms)

Soundfiles 1.19–1.30 Sound examples of clarinet, flute, piano, trombone, violin, and voice (one standard sample and one with the attack lopped off)

Figures 1.23–1.28 Photos: Clarinet, flute, piano, trombone, violin, and singer

ape

ncy,

e

The Digital Representation of Sound Playing by the Numbers

SECTION 2.1

Digital Representation of Sound

Figure 2.1 Graphs: Genetic algorithms (GA) (an example of a “continuous” or an “analog” function)

Figure 2.2 Illustration: Recording and playback of sounds through an ADC/DAC (analog to digital converter/digital to analog converter)

Figure 2.3 Graphs: Two different graphical representations of sound

SECTION 2.2

Analog Versus Digital

Figure 2.4 Animation: The path of a bouncing ball

Figure 2.5 Animation: The same path of a bouncing ball, now sampled by blinking

Applet 2.1 Sampled fader

Figure 2.6 Graph: An analog waveform and its digital representation. The analog waveform has smooth and continuous changes. The digital version of the same waveform has a staircase type look. The black squares are the actual samples taken by the computer.

SECTION 2.3

Sampling Theory

Xtra bit 2.1 A free sample: A moving tonsorial tale by Dan Rockmore

Figure 2.7 Graph: Undersampling and aliasing of a sine wave

Applet 2.2 Oscillators: Band-limited and not band-limited

Soundfile 2.1 Song excerpt sampled at 1,024 Hz

Soundfile 2.2 Same song excerpt sampled at the standard 44,100 samples per second

Figure 2.8 Graph: An undersampled waveform (512 Hz)

Figure 2.9 Graph: Same waveform sampled at 44,100 Hz

Applet 2.3 Play a sample with the mouse (scrubber applet)

Figure 2.10 Graph: Frequency domain graph of foldover aliasing

Soundfile 2.3 Chirping: A 10-second sound file sweeping a sine wave from 0 Hz to 44,100 Hz

Figure 2.11 Graph: An anti-aliasing low-pass filter

SECTION 2.4

Binary Numbers

Table 2.1 Number base chart: 0 to 15 in binary, octal, decimal, and hexadecimal

Figure 2.12 Diagram: A 4-bit binary number (called a nibble)

Table 2.2 Bits and their values

Applet 2.4 Hearing a binary counter

Table 2.3 The number of numbers an n -bit sequence of digits can represent (8, 16, 24, and 32 bits)

SECTION 2.5

Bit Width

Figure 2.13 Illustration: Staircase wave (representation of a 3-bit sound file)

Figure 2.14 Illustration: Staircase wave (representation of a 6-bit sound file)

Table 2.4 Convenient units for dealing with bits (nibble to terabyte)

Table 2.5 Examples of a sound file at different sampling rates (44,100; 22,050; 11,025; and 5,512.5) and bit widths (8 and 16)

SECTION 2.6

Digital Copying

Figure 2.15 Illustration: The degradation of the word *analog* after a couple of repetitions

Xtra bit 2.2 The peanut butter conundrum (the problem of cyber-piracy)

Applet 2.5 Make a noisy copy

Xtra bit 2.3 Errors in digital copying: Parity

Soundfile 2.4 A cheap imitation of a great piece by Alvin Lucier (example of the degradation of a recorded spoken fragment)

Soundfile 2.5 "Pretender," from John Oswald's *Plunderphonics*

Figures 2.16–2.19 Illustration: Graphic (excerpt) prepared by BMI (Broadcast Music Inc.) to illustrate the problem of copyright laws and new technologies

Soundfile 2.6a and 2.6b Music excerpt in original form and degraded after numerous copies

Soundfile 2.7 Excerpt from "Singing in the Style of the Voice of the Poet," by David Mahler

Figure 2.20 Photo: David Mahler

Xtra bit 2.4 Digital watermarking

SECTION 2.7

Storage Concerns: The Size of Sound

Figure 2.21 Photo: Standard CD pits under high magnification

Figure 2.22 Diagram: Three layers of a CD

Figure 2.23 Chart: From bits to CD (calculations for converting minutes of audio to megabytes)

Xtra bit 2.5 Hard drives

SECTION 2.8

Compression

Figure 2.24 Illustration: Perception-based encoding (weather data and the conclusion we draw from it)

Xtra bit 2.6 MP3

Soundfile 2.8–2.10 Three music excerpts compressed into the MP3 format at 128 Kbps, 64 Kbps, and 32 Kbps, respectively

Table 2.6 Entries from a typical μ -law table

Xtra bit 2.7 Delta modulation

SECTION 3.1

Frequency Domain

Figure 3.1 Graph: Time-domain representation of a waveform

Soundfile 3.1 Monochord sound

Figure 3.2 Graph: Signal, average signal envelope, and peak signal envelope for a monochord sound

Soundfile 3.2 Trumpet sound

Figure 3.3 Graph: Signal, average signal envelope, and peak signal envelope for a trumpet sound

Figure 3.4 Graph: Sonogram for sound from Figure 3.1

Soundfile 3.3 Mystery sound

Xtra bit 3.1 MatLab code to plot amplitude envelopes

Soundfile 3.4 Song of the hooded warbler

Figure 3.5 Photo: "Old-fashioned" sonogram of the song of the hooded warbler

Figure 3.6 Photo: Phonophotographic image of recording of "Swing Low, Sweet Chariot"

SECTION 3.2

Phasors

Figure 3.7 Animation: Sine waves and phasors

Figure 3.8 Diagram: Phase as angle measure

Figure 3.9 Illustration: The relation between circular movement and a sine wave

Figure 3.10 Graph: Basic sinusoid

Applet 3.1 Sampling a phasor

Figure 3.11 Graph: Bigger sinusoid

Figure 3.12 Graph: Bigger, faster sinusoid

Figure 3.13 Graph: Phase-shifted sinusoid

Figure 3.14 Graph: Sinusoid with 90-degree phase shift (cosine)

Applet 3.2 Building a triangle wave partial by partial

Figure 3.15 Illustration of the following process: hardcore singer to waveform to computer to digital conversion to sine waves sum representation

Figure 3.16 Graphs: Fast Fourier transform

Figure 3.17 Graph: FFT plot of a gamelan instrument

Figure 3.18 Graph: Waterfall plot (3D graphical representation of a stereo sound)

Figure 3.19 Graph: Time-amplitude plot of aliasing, foldover

SECTION 3.3

Fourier and the Sum of Sines

Figure 3.20 Spectrograms of a sine wave, triangle wave, sawtooth wave, and pulse

Figure 3.21 Illustration: Adding two sine waves

Soundfile 3.5 Example of sine waves being added step by step

Figure 3.22 Graph: 2D and 3D views of an FFT analysis of a trumpet note

Soundfile 3.6 Trumpet note analyzed in Figure 3.22

Applet 3.3 Play with an FFT

Soundfile 3.7 Sampled file of bugs by David Dunn

Soundfile 3.8 High-pass filtered version of David Dunn's bugs

Soundfile 3.9 Low-pass filtered version of David Dunn's bugs

SECTION 3.4

The DFT, FFT, and IFFT

Xtra bit 3.2 History of the FFT

Figure 3.23 Graph and table: Spectral components

Figure 3.24 Graph: Periodic function, $f(t)$

Figure 3.25 Graph: The window function, $w(t)$

Figure 3.26 Graph: Both functions plotted together

Figure 3.27 Graph: Multiplication of $f(t)$ by $w(t)$

Figure 3.28 Graph: Periodic extension of the windowed function, $f(t) \times w(t)$, along the t -axis

Figure 3.29 Illustration: SoundHack's phase vocoder window

Xtra bit 3.3 The FFT: Real, imaginary, magnitude, phase

Figure 1 Graph: An imaginary number in a two-dimensional space

Figure 3.30 Illustration: Interface from SoundHack presenting sound as a histogram of frequencies (the frequency spectrum without a time axis)

SECTION 3.5

Problems with the FFT/IFFT

Figure 3.31 Chart: Frequency range, in Hz, of common musical instruments and voice

Soundfile 3.10 A sine wave swept from 50 Hz to 10 kHz, processed through an FFT analysis

Figure 3.32 Graph: Sweeping sine wave through an FFT analysis

Figure 3.33 Graphs: An analysis using an FFT size of 512 samples (giving us good time resolution) and another analysis of the same sound, using an FFT size of 2,048 samples (better frequency resolution, worse time resolution)

Soundfile 3.11 High-pitched male voice (sound file represented by the analysis in Figure 3.33)

Soundfile 3.12 Music excerpt (beat sound)

Soundfile 3.13 Same excerpt after being processed to provide accurate frequency resolution but inaccurate time resolution (rhythmic smearing)

Soundfile 3.14 Same excerpt processed in the other way—to provide accurate time resolution but bad frequency resolution (spectral smearing)

SECTION 3.6

Some Alternatives to the FFT

Figure 3.34 Illustration: Four common wavelet analysis waveforms (Daubechies, Coiflet, Haar, Symmlet)

Figure 3.35 Graph: Frequency tracks in Lemur Pro

Figure 3.36 Photo: MQ plot of 0'00" to 5'49" of Movement II from *Sud*, by Jean-Claude Risset

The Synthesis of Sound by Computer

SECTION 4.1

Introduction to Sound Synthesis

Applet 4.1 Basic types of noise

Xtra bit 4.1 Computer code for generating an array of a sine wave

Applet 1 Sawtooth faders

Xtra bit 4.2 Composer Robert Marsanyi's unique approach to synthesis

Xtra bit 4.3 Early computer music network ensembles

SECTION 4.2

Additive Synthesis

Figure 4.1 Illustration: Two waves joined by a plus sign

Figure 4.2 Photo: Large church organ

Applet 4.2 Hearing mixed sounds and clipping

Soundfile 4.1 Short excerpt from Kenneth Gaburo's composition *Lemon Drops*, a classic of electronic music made in the early 1960s

Figure 4.3 Graphs: Generating a square wave by means of additive synthesis of sine waves

Figure 4.4 Photo: Synclavier

Applet 4.3 Making a waveform using additive synthesis

Applet 4.4 Adding spectral envelopes to a number of partials

Xtra bit 4.4 Spectral formula of a waveform (contains formulas for triangle, square, and sawtooth waves)

Soundfiles 4.2–4.5 Speech reconstructed as sine waves and original speech

Figure 4.5 Diagram: A typical ADSR (attack, decay, sustain, release, steady-state modulation)

Figure 4.6 Graph: Bell-shaped spectral envelope

Soundfile 4.6 Chromatically ascending tones (demonstration of Shepard tones)

Figure 4.7 Illustration: Shepard tones represented by a keyboard and a chromatic scale

Soundfile 4.7 The spiraling Shepard tone effect

Soundfile 4.8 "For Ann (rising)," by James Tenney (1969)

SECTION 4.3

Filters

Soundfile 4.9 Telephone simulation

Soundfile 4.10 High-pass filtered noise

Soundfile 4.11 Low-pass filtered noise

Figure 4.8 Graphs: Four common filter types (low-pass, high-pass, band-pass, band-reject)

Applet 4.5 Hearing different filters

Applet 4.6 Comb filters

Figure 4.9 Diagrams: FIR and IIR filters

SECTION 4.4

Formant Synthesis

Figure 4.10 Graph: A trumpet playing two different notes, a perfect fourth apart (the formants stay in the same places)

Applet 4.7 Formants

Xtra bit 4.5 Resonances

Figure 1 Photo: Back of a hasapi, a two-stringed open-backed lute from the Batak of North Sumatra, Indonesia

Figure 4.11 Graph: Spectral plot of the voice, showing formants

Xtra bit 4.6 Human formant manipulations—throat singing, trumping

Soundfile 1 Tuvan throat singing (overtone singing)

Soundfile 2 Jaw harp

Figure 4.12 Photo: Composer Paul Lansky

Soundfiles 4.12 and 4.13 Two examples of Paul Lansky's work with human voice: "Notjustmoreidlechatter" (1988) and "idlechatterjunior" (1999)

Soundfile 4.14 "Trajectory of Her Voice," by Sarah Myers (1998)

Soundfile 4.15 Synthetic speech ("Fred" voice from the Macintosh computer)

Soundfile 4.16 Carter Sholz's one-minute piece "Mannagram," based on a reading by Australian sound poet Chris Mann

Soundfile 4.17 A tune on the trump, also called the jaw harp

SECTION 4.5

Amplitude Modulation

Applet 4.8 Low-frequency modulation

Figure 4.13 Diagram: A simple amplitude modulation scheme

Soundfile 4.18 A low-pass moving filter modulated by a sine wave

Soundfile 4.19 A high-pass moving filter modulated by a sine wave

Soundfile 4.20 A low-pass moving filter modulated by a sawtooth wave

Soundfile 4.21 A high-pass moving filter modulated by a sawtooth wave

Figure 4.14 Diagrams: James Tenney's "Phases"

SECTION 4.6

Waveshaping

Figure 4.15 Graph: A simple sine wave as input and its output after waveshaping

Applet 4.9 Changing the shape of a waveform

Applet 4.10 Waveshaping using tables and Chebyshev formulae

Soundfile 4.22 Experimental waveshaping: Excerpt from "Toyaji Patch," by Larry Polansky

SECTION 4.7

FM Synthesis

Applet 4.11 Frequency modulation

Figure 4.16 Photo: Yamaha DX-7 synthesizer

Figure 4.17 Diagram: A simple frequency modulation scheme

Soundfile 4.23 Vibrato sound (carrier: 500 Hz; modulator frequency: 1 Hz; modulation index: 100)

Soundfile 4.24 Vibrato sound (carrier: 500 Hz; modulator frequency: 1 Hz; modulation index: 500)

Figure 4.18 Graph: 3D plot of the first 15 Bessel functions

Soundfile 4.25 Bell-like sound (carrier: 100 Hz; modulator frequency: 280 Hz; FM index: 6.0 \rightarrow 0)

Soundfile 4.26 Bass clarinet-type sound (carrier: 250 Hz; modulator frequency: 175 Hz; FM index: 1.5 \rightarrow 0)

Soundfile 4.27 Trumpet-like sound (carrier: 700 Hz; modulator frequency: 700 Hz; FM index: 5.0 \rightarrow 0)

Soundfile 4.28 FM sound (carrier: 500 Hz; modulator frequency: 500 \rightarrow 5000 Hz; FM index: 10)

Figure 4.19 Illustration: Csound orchestra and score blueprint for a simple FM synthesis instrument

SECTION 4.8

Granular Synthesis

Applet 4.12 Granular synthesis

Figure 4.20 Graphs: Creation of a grain

Figure 4.21 Graphs: A granular synthesis "score"

Soundfile 4.29 Excerpt from "Implements of Actuation," by Mara Helmuth

Physical Modeling

Applet 4.13 Karplus-Strong plucked string algorithm

Applet 4.14 Mass/spring unit generator

Figure 4.22 Graphs: Applying the Karplus-Strong algorithm to a random waveform

Figure 4.23 Diagram: Schematic view of a computer software implementation of the basic Karplus-Strong algorithm

Soundfile 4.30 Electronic heavy metal version of "The Star Spangled Banner" by engineer Charlie Sullivan, using a software "super" guitar

Figure 4.24 Illustration: Part of the interface from Perry R. Cook's SPASM singing voice software

Soundfile 4.31 An example of Perry R. Cook's SPASM

SECTION 5.1

Introduction to the Transformation of Sound by Computer

Applet 5.1 Making a piece with samples

Soundfile 5.1 Excerpt from "Chef d'Oeuvre," by Jon Appleton

Soundfile 5.2 Excerpt from "BS Variation 061801," by Huk Don Phun

Soundfile 5.3 "Puzzels & Pagans," by Jane Dowe, 1997.

Figure 5.1 Illustration: A window from Argeiphontes Lyre, used for sound deconstruction

Soundfile 5.4 A random cut-up made using Argeiphontes Lyre

Figure 5.2 Photos: Herbert Brün and graphic for SAWDUST

Soundfile 5.5 "Dustiny" (1978), by Herbert Brün (an example of SAWDUST)

Soundfile 5.6 Drum machine loops

Applet 5.2 DrumBox

Figure 5.3 Photo: The Mark of the Unicorn (MOTU) 828 breakout box, part of a DAW system with FireWire interface

Figure 5.4 Photo: Motor Mix™, the only DAW mixer control surface to offer eight-bank switchable motorized faders

SECTION 5.2

Reverb

Soundfile 5.7 Changing the reverberant characteristics of a sound (a nail being hammered) over time

Figure 5.5 Diagram: Introduction of a direct sound into a reverberant space

Figure 5.6 Diagram: The creation of reverb within a reverberant space

Figure 5.7 Diagram: Signal delay

Figure 5.8 Diagram: Signal delay with all-pass filter

Applet 5.3 Flanges, delays, and reverb

Applet 5.4 Multitap delay

Figure 5.9 Graph: The impulse and the room's response

SECTION 5.3

Localization/Spatialization

Applet 5.5 Filter-based localization

Soundfile 5.8 Music spatialized using SoundHack

Figure 5.10 Diagram: The impact of sound direction on the spectral contents of a sound

Figure 5.11 Photo: Binaural dummy head recording system

SECTION 5.4

Introduction to Spectral Manipulation

Soundfile 5.9 Unmodified speech example

- Soundfile 5.10 Speech made twice as long with a phase vocoder
- Soundfile 5.11 Speech made half as long with a phase vocoder
- Soundfile 5.12 Speech transposed up an octave with a phase vocoder
- Soundfile 5.13 Speech transposed down an octave with a phase vocoder
- Soundfiles 5.14 and 5.15 Two examples of time-domain pitch shifting/speed changing using the speech from Soundfile 5.9
- Xtra bit 5.1 Steve Reich's "Slow Motion Sound"
- Soundfile 5.16 A SoundHack varispeed function of the speech from Soundfile 5.9
- Figure 5.12 Graphs: Windowing FFT frames to smooth high-frequency artifacts
- Figure 5.13 Graph: Overlapping windowed signals so that the original signal is reconstructed without the sharp edges
- Figure 5.14 Diagram: FFT bins and the process of pitch shifting
- Soundfile 5.17 "Study: Anna, the Long and the Short of It," by Larry Polansky

SECTION 5.5

More on Convolution

- Soundfile 5.18 A simple convolution example: A guitar power chord, a vocal excerpt, and the convolution of one with the other
- Soundfile 5.19 Noise
- Soundfile 5.20 Inhalation
- Soundfile 5.21 Noise and inhalation convolved

SECTION 5.6

Morphing

Figure 5.15 Animation: Several stages of a morphing image

Soundfile 5.22 Excerpt from "Morphing Piece," by Max Mathews

Figure 5.16 Graph: An amplitude cross-fade of a number of different data points

Soundfile 5.23 Sound example of the cross-fade in Figure 5.16 (from a pipe organ to a didgeridoo)

Figure 5.17 Graph: The centroid curve of two sounds over time (a violin tone and the rapidly changing voice of Chris Mann)

Soundfile 5.24 The rapidly changing voice of Chris Mann

Soundfile 5.25 The violin tone in Figure 5.17

SECTION 5.7

Graphical Manipulation of Sound

Applet 5.6 Color organ, drawing sound

Figure 5.18 Photos: Hugh LeCaine and his electronic musical instrument, the Spectrogram

Figure 5.19 Illustration: FFT data displayed as a sonogram in the computer music program AudioSculpt

Figure 5.20 Illustration: An original image created by the QT-coder in SoundHack and the image after alterations

Soundfiles 5.26 and 5.27 The two sounds related to the images in Figure 5.20

Figure 5.21 Illustration: Window from Chris Penrose's Hyperupic

Soundfile 5.28 "American Jingo," by Chris Penrose, created with Hyperupic

Figure 5.22 Illustration: Interface of squiggy, a real-time spectral manipulation tool created by douglas repetto