Music 209
Advanced Topics in Computer Music
Lecture 1 – Introduction

2006-1-19

Professor David Wessel (with John Lazzaro)

(cnmat.berkeley.edu/~wessel, www.cs.berkeley.edu/~lazzaro)

Website: Coming Soon ...

Course Topic: Concatenative Synthesis

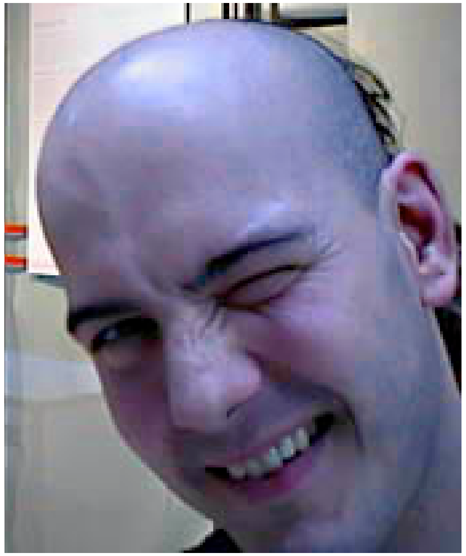

define: Concatenative synthesis methods stitch together segmented units of sound to produce the desired output. **(after Diemo Schwarz, IRCAM).**

Most synthesis methods that take a sample-based approach fit this definition, to some degree ... but some fit it more than others.

Let's begin with a commercially successful application that is "marginally" concatenative.

Grand piano

25,000 USD or more ...

Every one sounds different.

Must be tuned regularly.

To record well:
large room +
quality mics.

f_0 = 27.5 Hz

Interface looks uniform, but its design scales across 88 keys.

f_0 = 4100 Hz

Separate bridges for the lowest strings ...

and for all other strings.

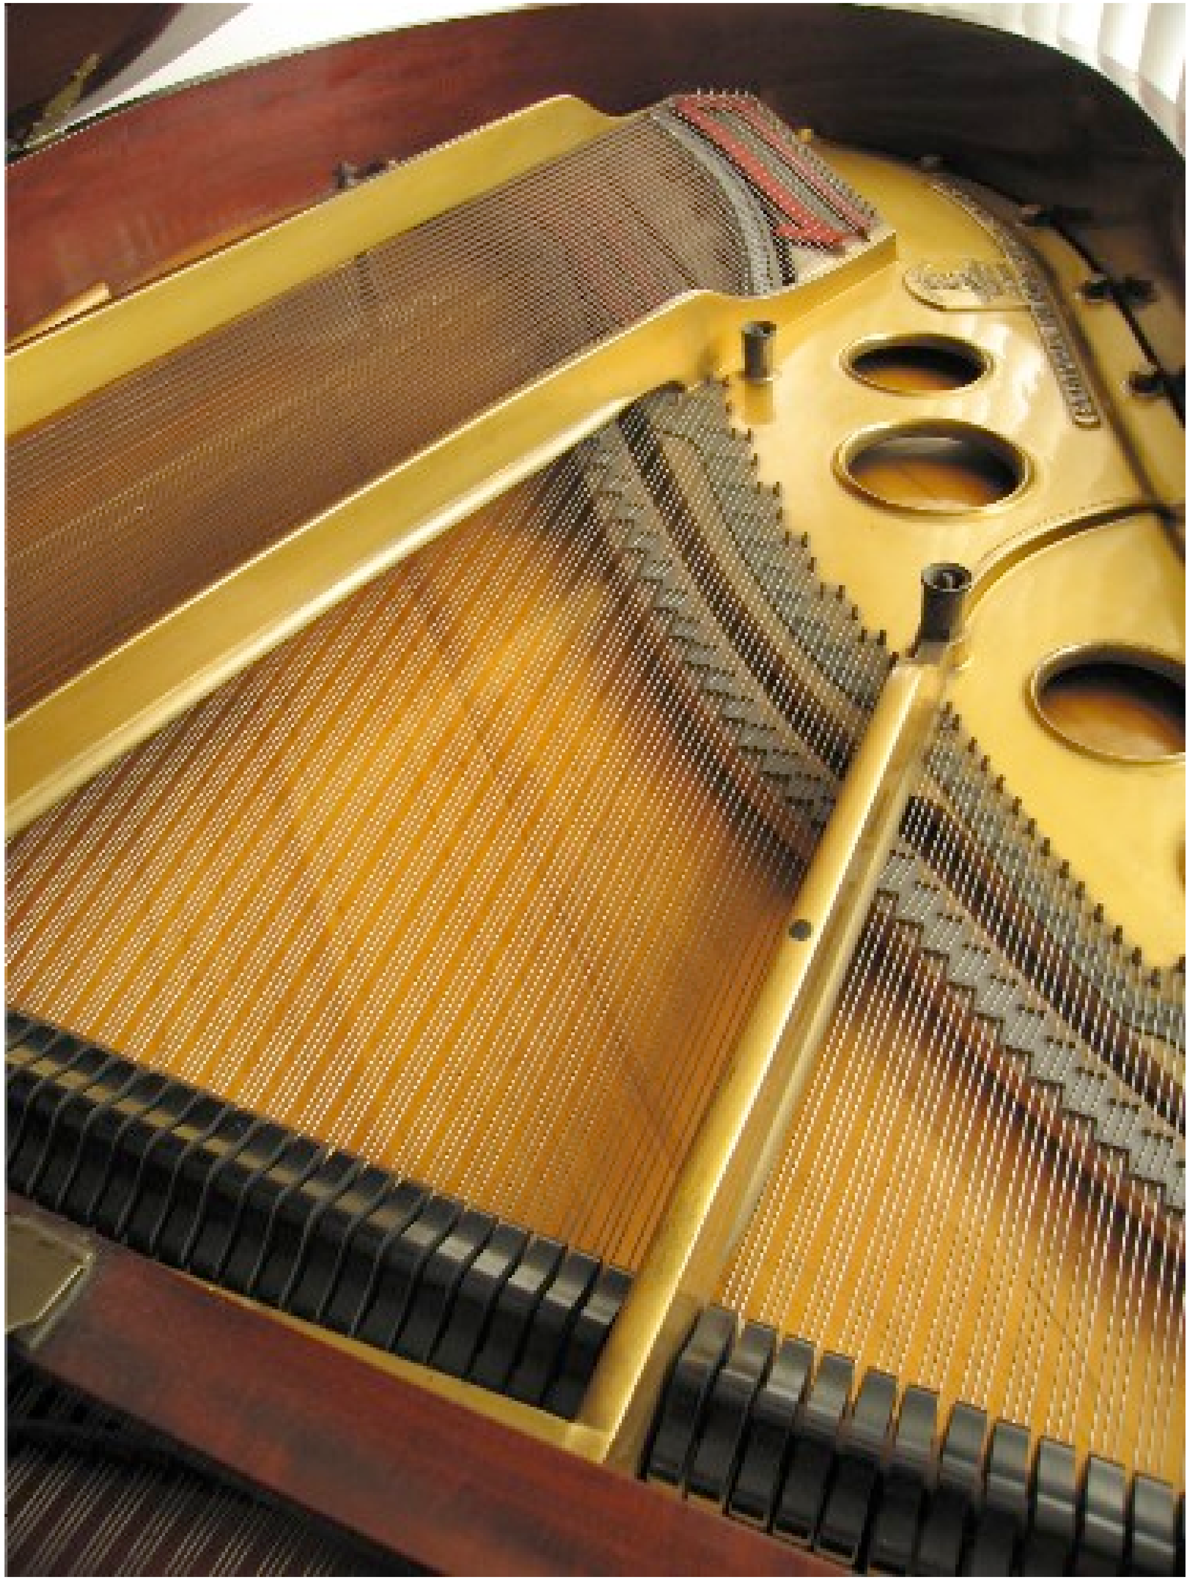

There's nothing remarkable about it. All one has to do is hit the right keys at the right time and the instrument plays itself.

-- Johann Sebastian Bach

Lowest notes: single wrapped string

Highest notes: 3 unwrapped strings (never precisely in tune)

Tuning pegs set string tension.
Total pressure of all 250 strings on harp = 20 tons!

Felt damper mutes string when its key is released.

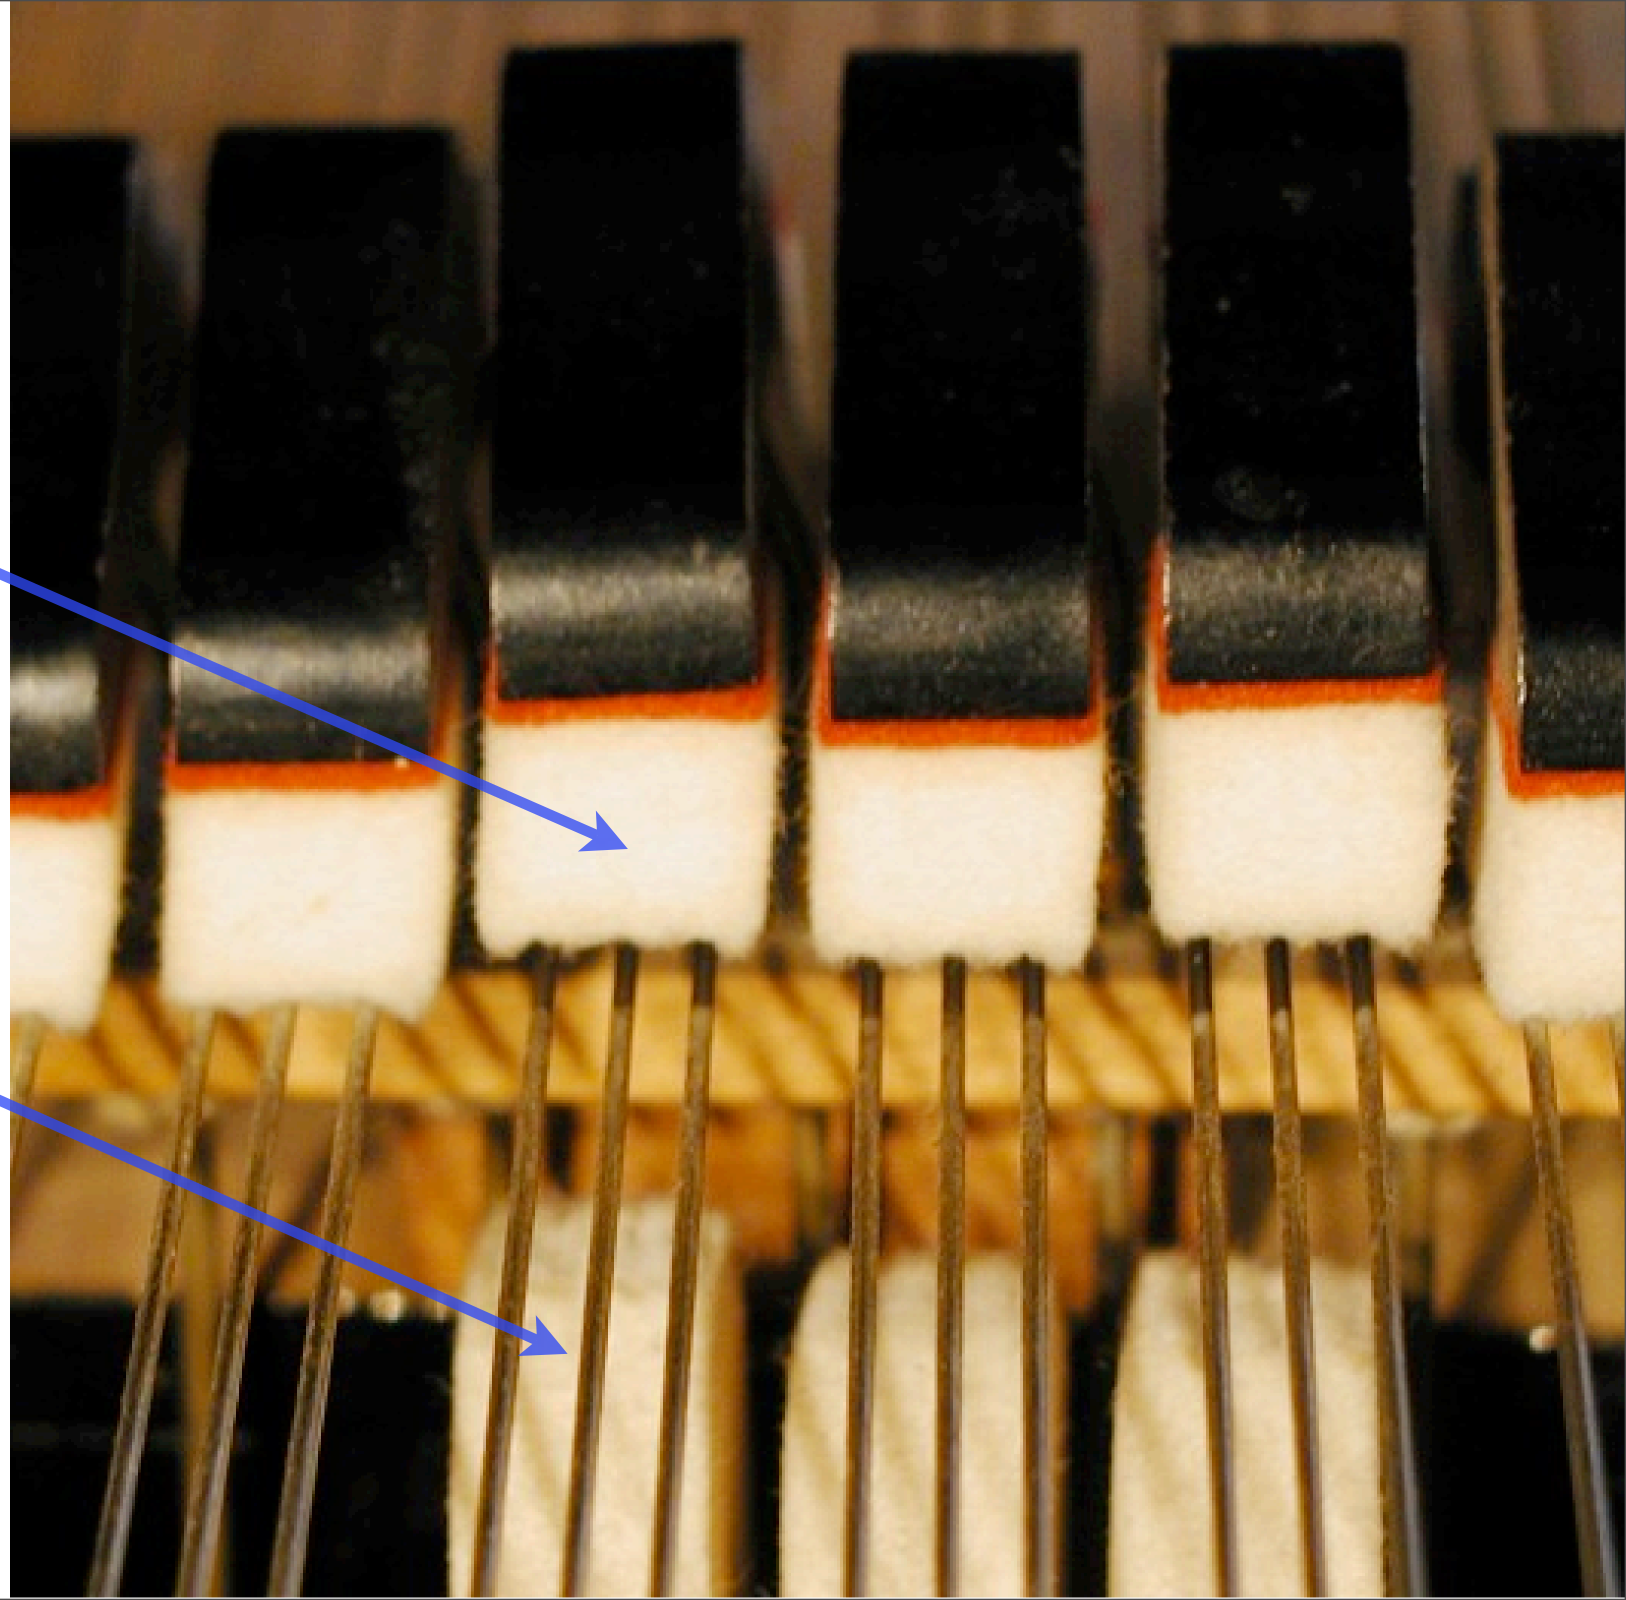

Felt hammer strikes strings when its key is depressed

Sustain pedal lifts all 88 dampers until pedal is released.

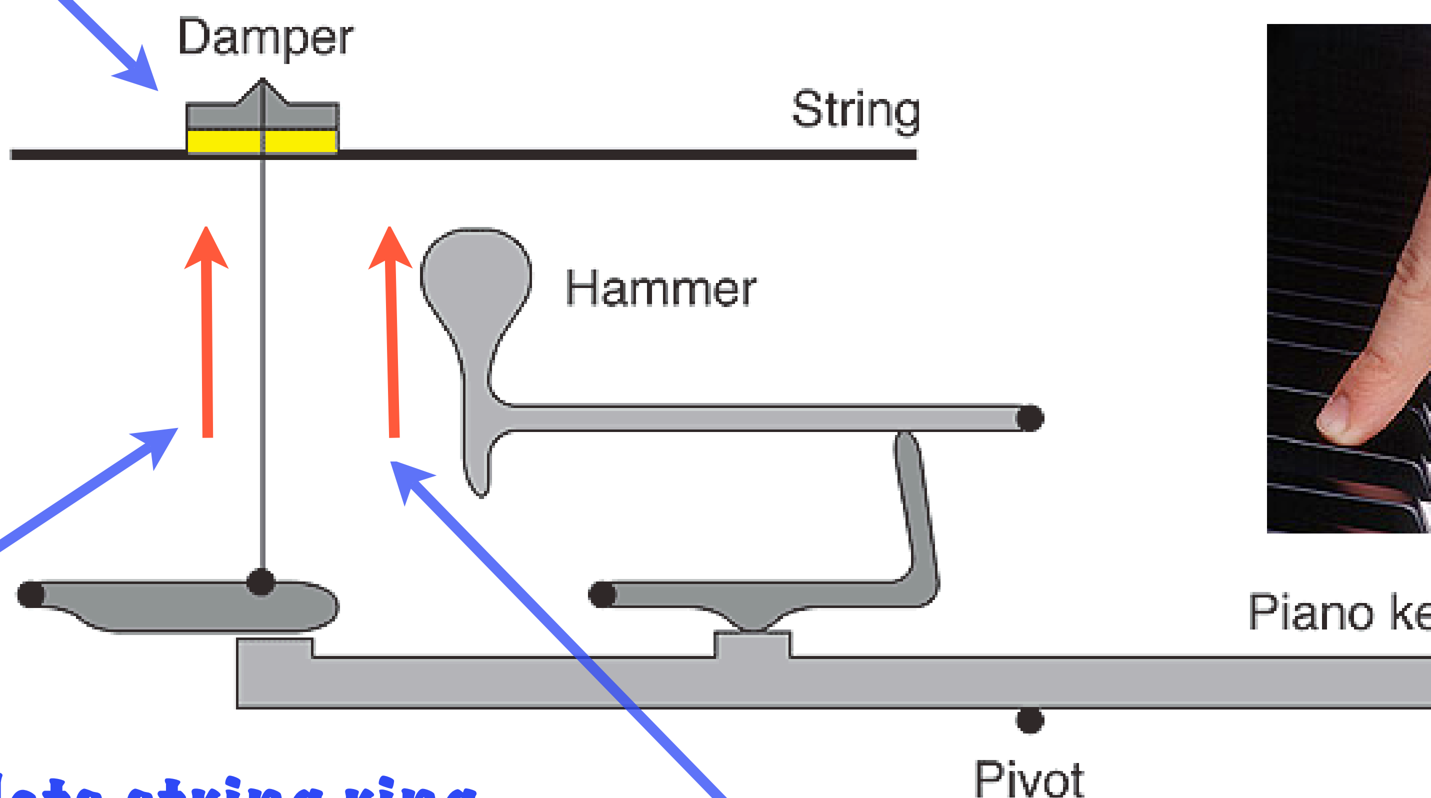

Damper lets string ring free until key is released.

Hammer velocity a function of key velocity.

"Escapement" -- hammer design permits quick repeats.

Figures from Sound on Sound Oct 02 (Synth Secrets)

The physics behind the unique piano attack sound (listen)

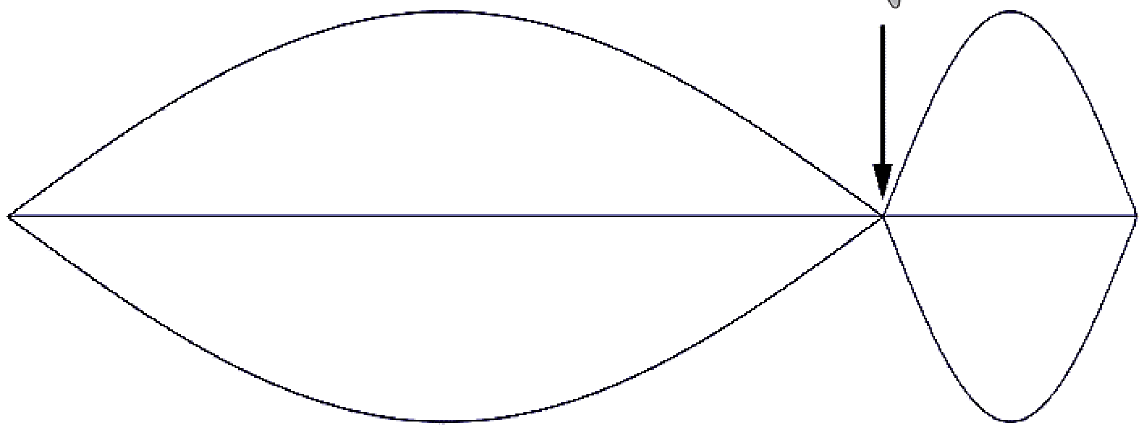

Hammer impact: **Inharmonicity.**

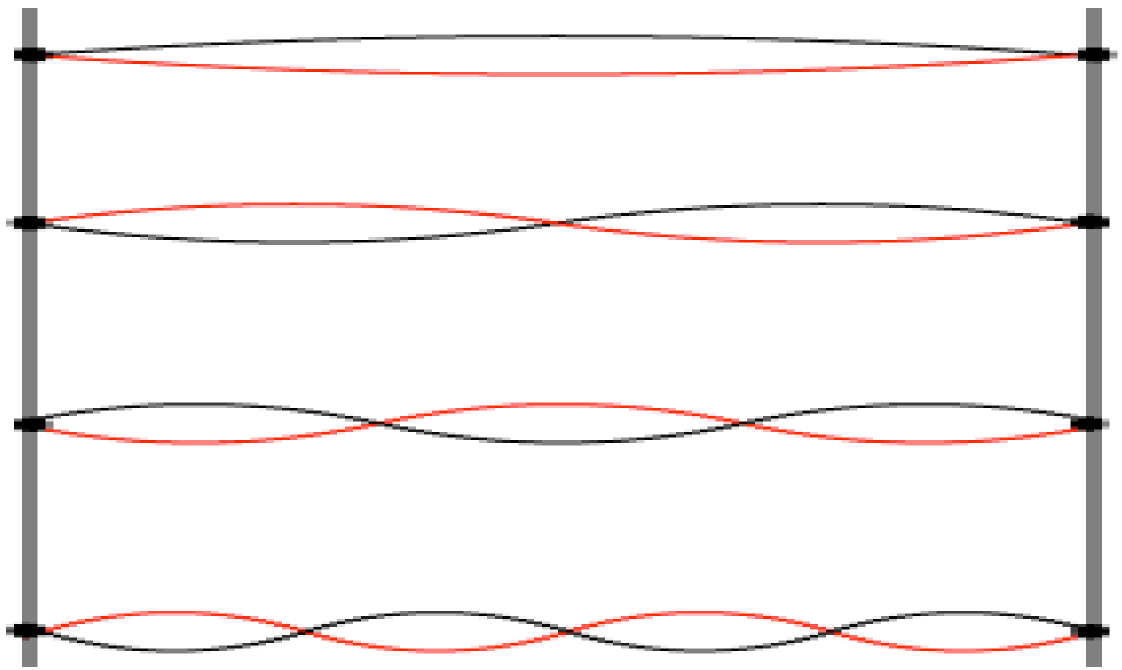

After a few hundred ms: **A harmonic spectrum ...**

Figures from Sound on Sound Oct 02 (Synth Secrets)

Three strings are never perfectly in tune. Energy exchange gives sustain a unique sound (listen).

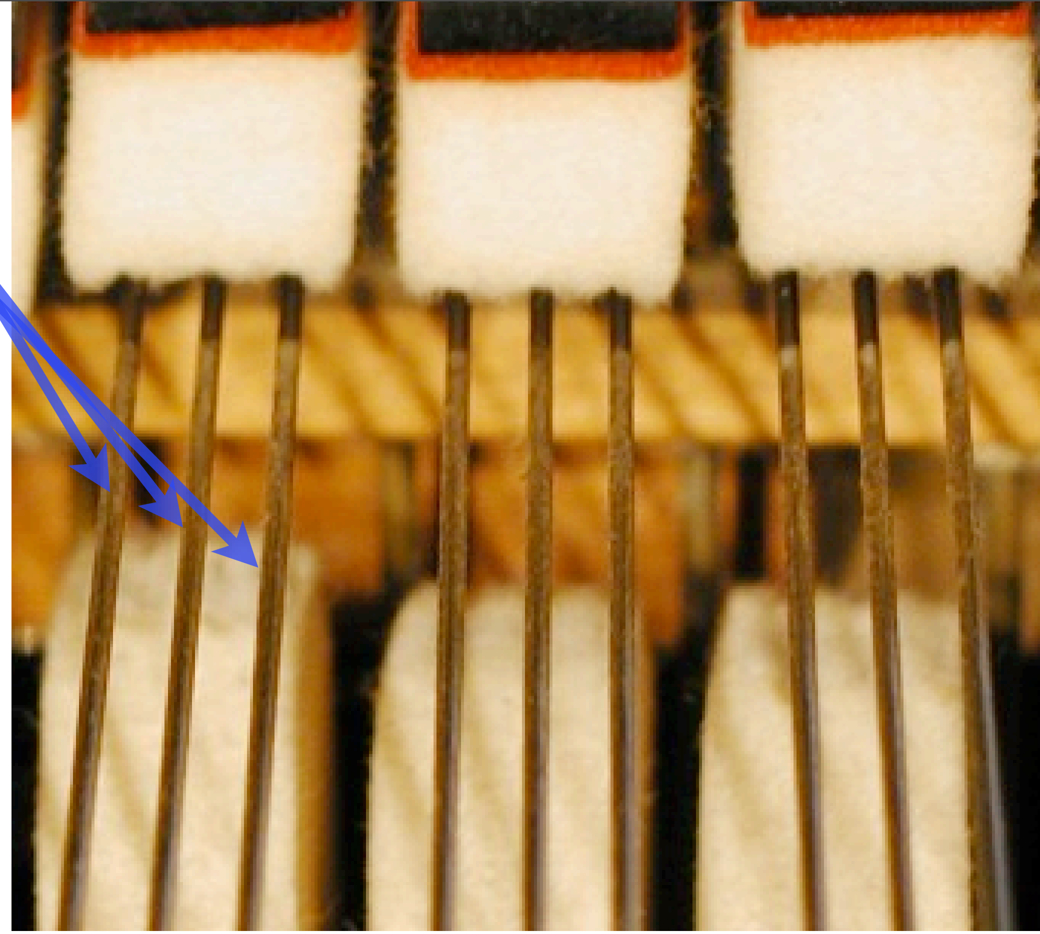

If pedal is down, strings for all other keys vibrate: sympathetic resonance.

Energy exchange between the soundboard and the strings also shapes the evolution of the sound.

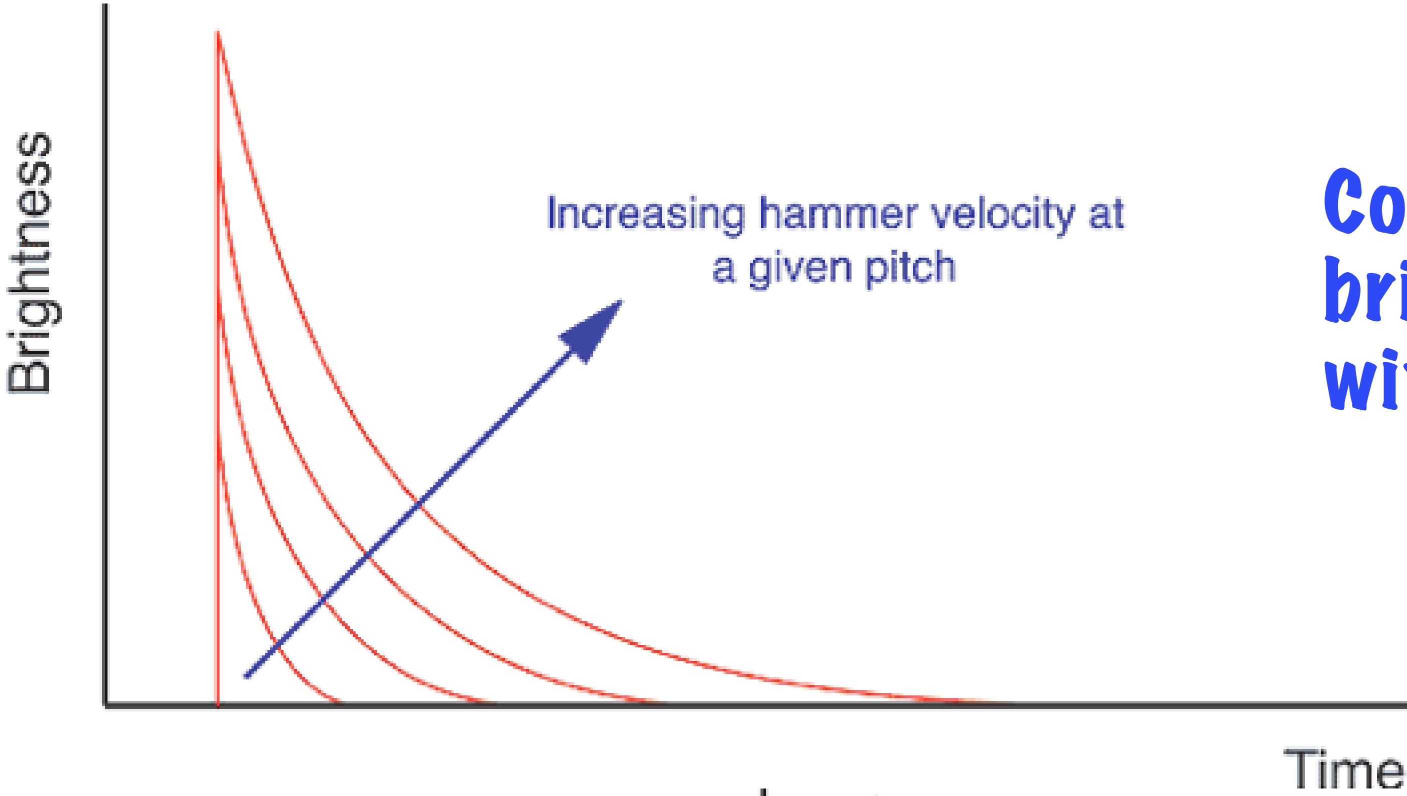

Complex changes in brightness and decay with hammer velocity.

Decay time for high velocities: 40 seconds for lowest notes, 10 seconds for mid-range, 3 for highest notes.

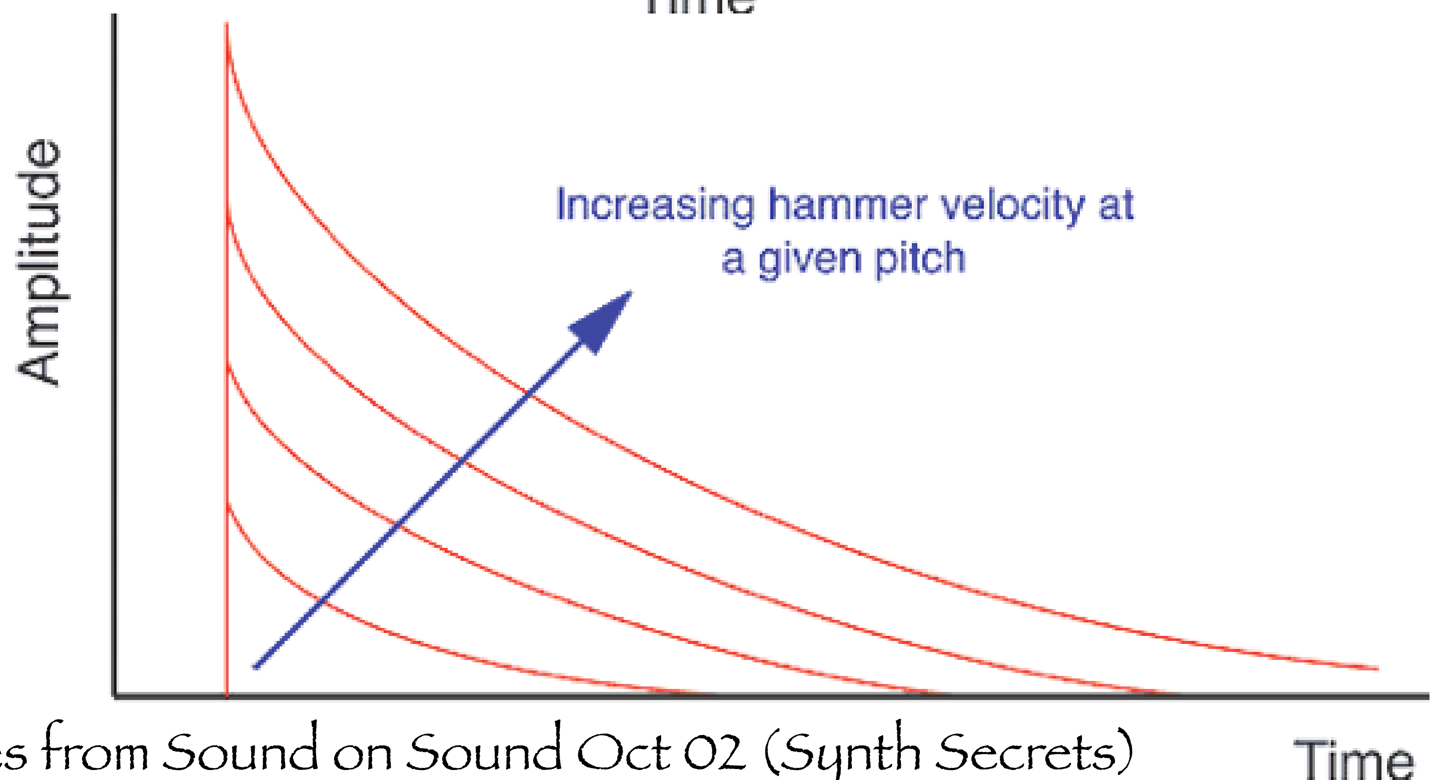

Figures from Sound on Sound Oct 02 (Synth Secrets)

Brightness and decay changes across the keyboard.

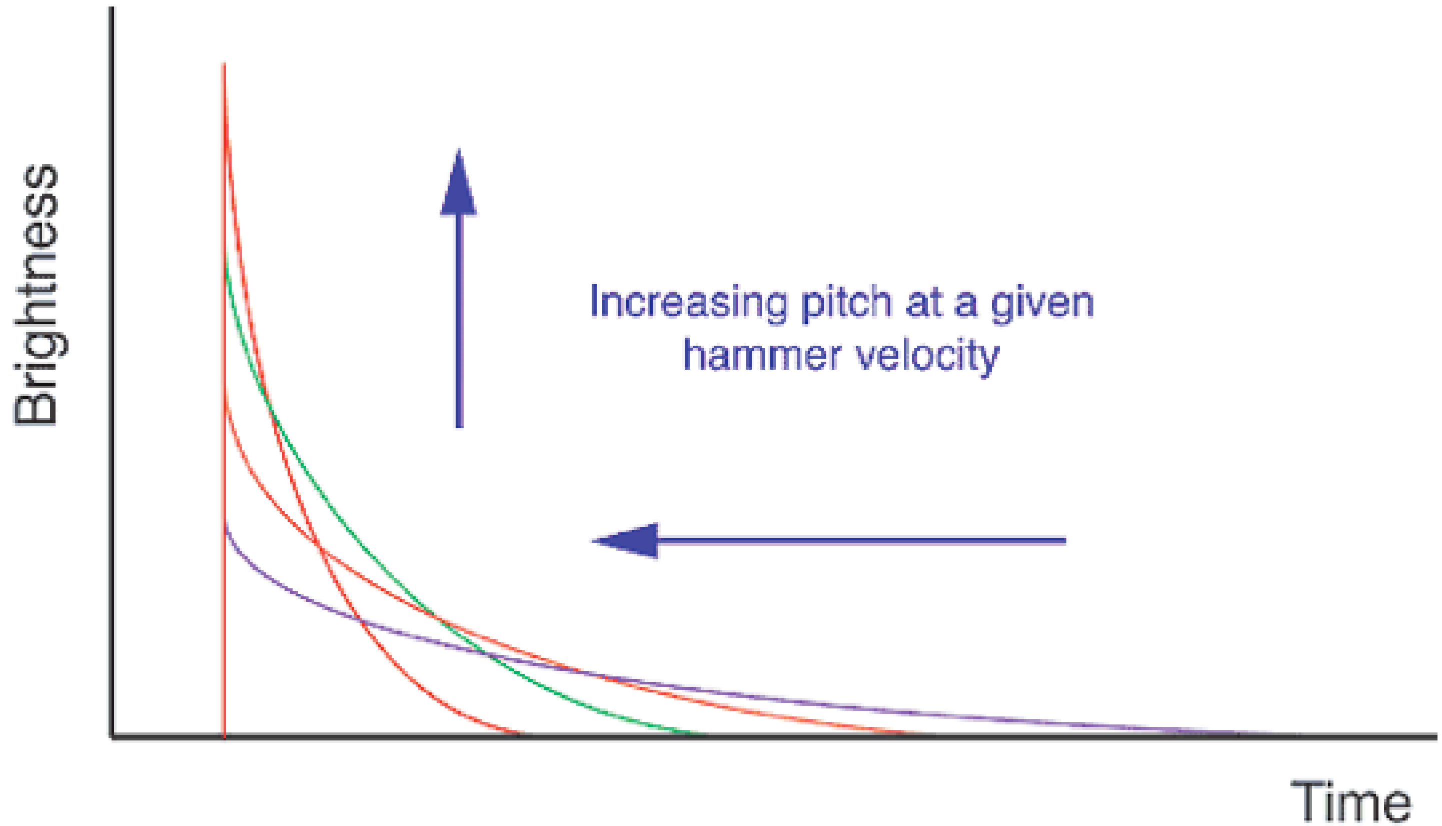

The way the mechanical noise of the strike fuses with the string sound also changes across the keyboard.

Figures from Sound on Sound Oct 02 (Synth Secrets)

Where you listen has a big impact on what the piano sounds like. Solo classical recordings often choose a good room, and find the "sweet spot" where the room sound meets the piano sound.

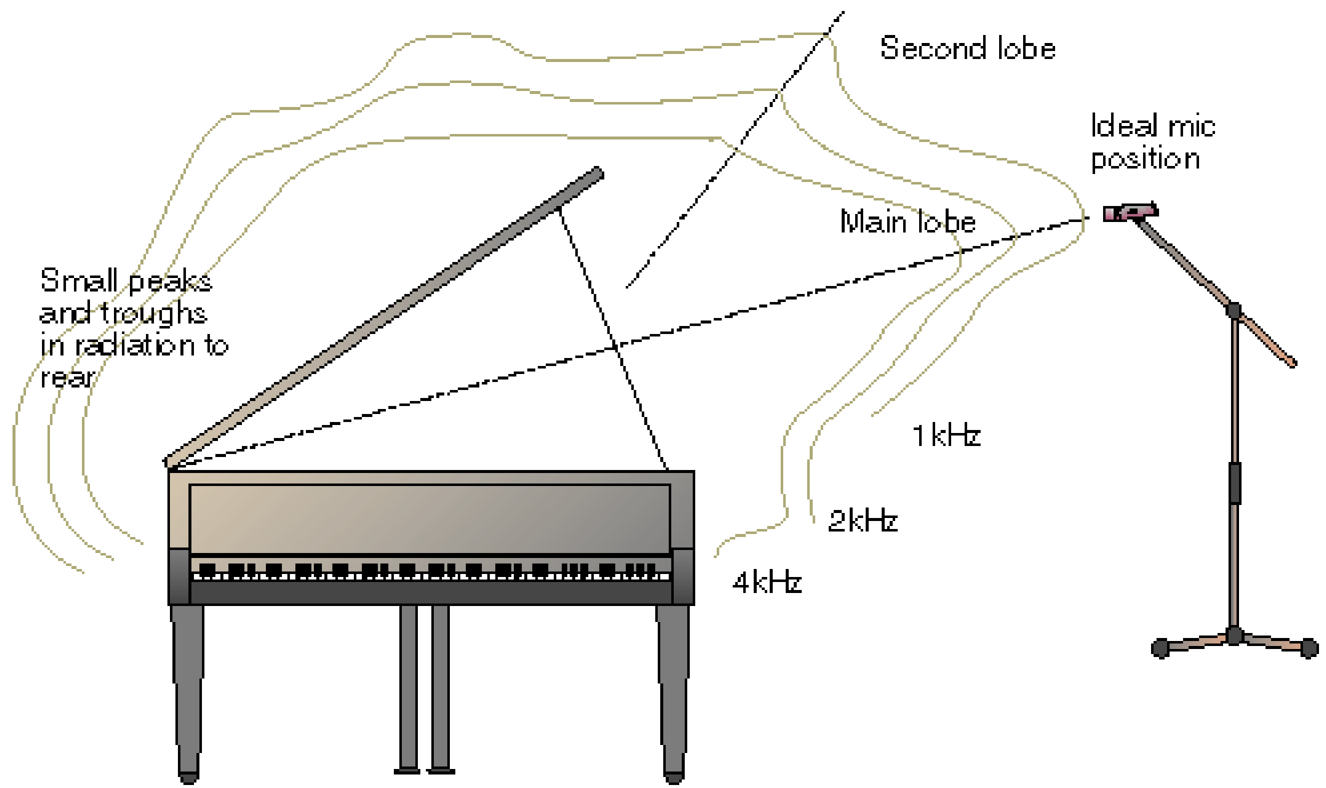

Figures from Sound on Sound May 99 (Piano Principles)

Pop recording techniques often use a close-in sound.

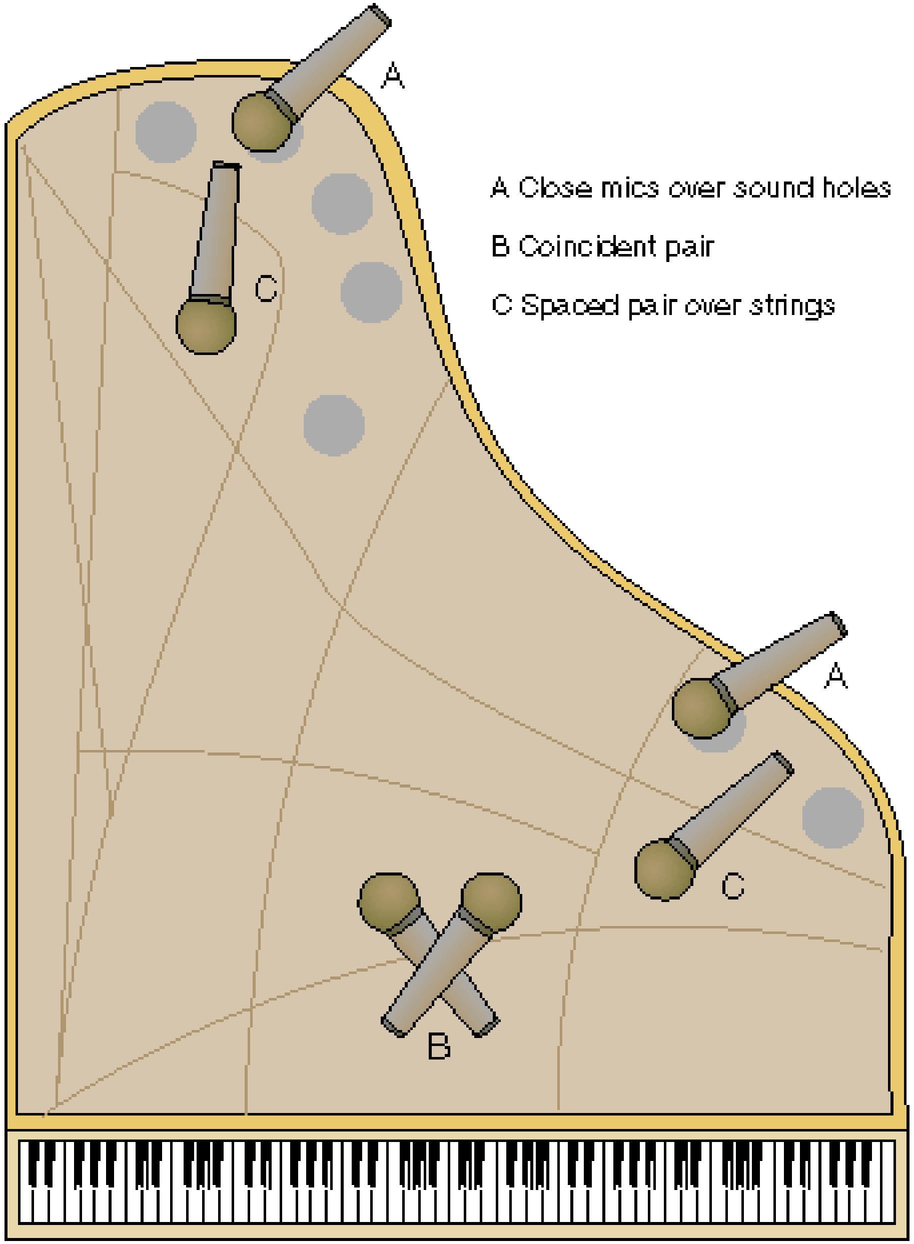

Understanding the physics of pianos well enough to model them from first principles is beyond our knowledge, not to mention compute cycles.

And so, the best synthesized pianos have always used recordings of pianos. We now describe a typical product.

Black Grand. Sample library for a Steinway D Hamburg.

3 DVD-ROMs, $169. Meant to be used with software samplers in Digital Audio Workstations (DAWs)

DAWs support real-time (5 ms latency) control from a MIDI keyboard

12672 stereo recordings: 144 @ 88 keys

* 16 strikes at different velocities, with the sustain pedal up, held until note fades to silence.
* 16 strikes at different velocities, with the sustain pedal down, held until note fades to silence.
* For pedal up, 16 recordings of the release sound after key is lifted.

144? Three sets of coincident microphones: close, medium, far - in Orebro Hall (Sweden).

18 GB: 9 hours of stereo audio (24 bit, 96 kHz).

Why do we consider this concatenative?

We're linearly mixing the sounds of different notes, not stitching them!

Recall: Concatenative synthesis methods stitch together segmented units of sound to produce the desired output.

Answer:

The release sample is stitched onto the strike sample when the player lifts a key.

Pedaling may also require stitching.

If stitches sound artificial, illusion is lost.

Black Grand Sample Library Demo

- Close-up microphones — Playing
- Medium range microphones — Playing
- Ambient microphones — Playing

History ...

1984: Kurzweil 250, first commercial sampled piano - on Stevie Wonder's request. 512 KB piano sample set in ROM.

Sounds good to this day ... many technology and music tricks used to handle small ROM. We will cover these tricks in future lectures.

Current Status ...

Sampled piano sales dwarf acoustic piano sales.

- As memory storage became cheaper since 1984, "tricks" have been gradually replaced with memory.
- **Biggest challenges:** Choosing the **right** piano to sample, keeping it **in tune** throughout the project, making a **perfect recording**, clean editing ...
- **Pedal-up sympathetic resonances:** Requires (???) an algorithmic approach: too many combinations to sample.

Features (ex: unusual tunings)

Why pianos were relatively easy ...

... compared to a violin.

	Piano	Violin
Articulations	One	Many
Expression During Sustain	No	Yes
Legato and Portamento	No	Yes

Playing

But recent progress has been made ...

SYNFUL
ORCHESTRA

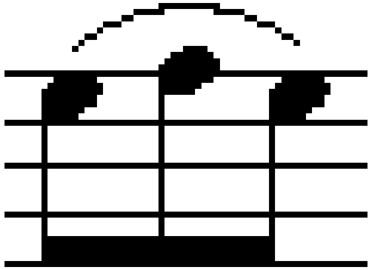

	Piano	Violin
Legato and Portamento	No	Yes

“Online” stitching of a legato run from 3 samples in a library.

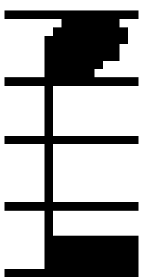

splice to

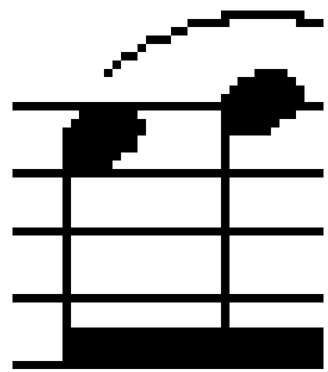

splice to

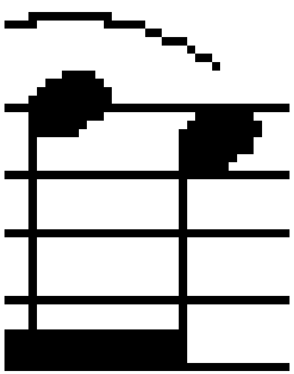

Sample #1:
isolated E

Sample #2:
E to F interval
played legato

Sample #3:
F to E interval
played legato

	Piano	Violin
Legato and Portamento	No	Yes

Performance Libraries

Scale runs, legato intervals and repeated notes, grace notes, glissandos.

Performance libraries for the Vienna Symphonic Library Solo Violin

Performances

ARTICULATION GROUP	EXISITNG CONTENT	NEW CONTENT
Interval performances	Legato Portamento forte Détaché Marcato forte Spiccato	Legato on the same string Portamento piano Marcato piano Legato with progressive vibrato "Zigane" (gipsy) style
Fast interval performances		Legato Marcato Spiccato Harsh
Multi interval performances		Performance trills, legato
Repetition performances	Legato Portato medium and fast Staccato Bow vibrato Normal and dynamics	Portato slow Spiccato Harsh
Fast repetitions	9 repetitions, 16ths from 150 to 200 BPM Normal and dynamics	Ricochet repetitions 150 to 210 BPM 3 rebounds accelerando, ritardando
Scale runs	Octave runs, legato major, minor, chromatic and whole tone Spiccato major	Spiccato minor, chromatic and whole tone
Grace notes	Minor and major 2nd, up and down	Slow grace notes "Zigane" (gipsy) style, minor 2nd to major 3rd
Glissandos	Octave glissandos, medium speed, up and down. Performance glissandos on the G, D, A, and E string	Octave glissandos, fast

	Piano	Violin
Articulations	One	Many

Solution: Sample each note with each articulation.

Articulation libraries for the Vienna Symphonic Library Solo Violin

Short notes	Staccato, short détaché, long détaché with and without vibrato 2 variations (downstroke and upstroke)
Long notes	Sustained without, with medium and progressive vibrato, Espressivo Marcato
Dynamics	Medium crescendo and diminuendo with and without vibrato, various durations Strong crescendo and diminuendo with vibrato, various durations
Tremolo + Trills	Fortepiano, sforzato, sforzatissimo Crescendo-diminuendo with vibrato, 2 durations Tremolo sustained Crescendo and diminuendo (2 durations). Half and whole tone trills Constant dynamics, crescendo and diminuendo Constant speed and accelerando.
Pizzicato + col legno	Pizzicato normal, secco, snap (Bartók pizzicato) Col legno 2 variations each
Ponticello	Staccato, détaché, sustained, sforzato, sforzatissimo, tremolo, repetitions
Sul tasto	Staccato, détaché, sustained, sforzato, tremolo
Harmonics	Artificial harmonics: Staccato, sustained, repetitions, glissandi

Duplicates for expression during sustain ...

Tradeoffs: Memory size vs algorithms

Vienna approach: systematic recording of a scored database.

82 GB for the solo violin library.

5 DVDs, $950 USD

SYNFUL ORCHESTRA

32 MB for the solo violin library.

How? Database of longer phrases. Use signal processing to find closest phrase, modify to suit.

Control: How to "Play" an 82 GB library?

One answer: Keyswitches

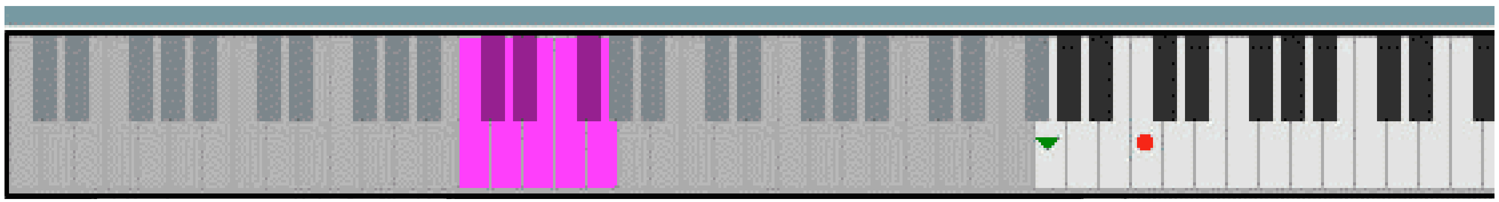

Grey keys: Map to control commands

Sample mapping

C1	staccato
C#1	short détaché
D1	long détaché w/vibrato
D#1	sustained note
E1	sforzando w/vibrato
F1	Piano-forte-piano w/vibrato, 4 sec. long
F#1	Tremolo
G1	Pizzicato w/vibrato

White keys: Sound notes as normal

Another Approach: Wheels + Pedals

1. **Mod Wheel**
Controls Volume
and Expression

2. **Note Velocity**
Contols Attacks
and Accentuation

3. **Sustain Pedal**
Controls Legato
for Smooth Playing

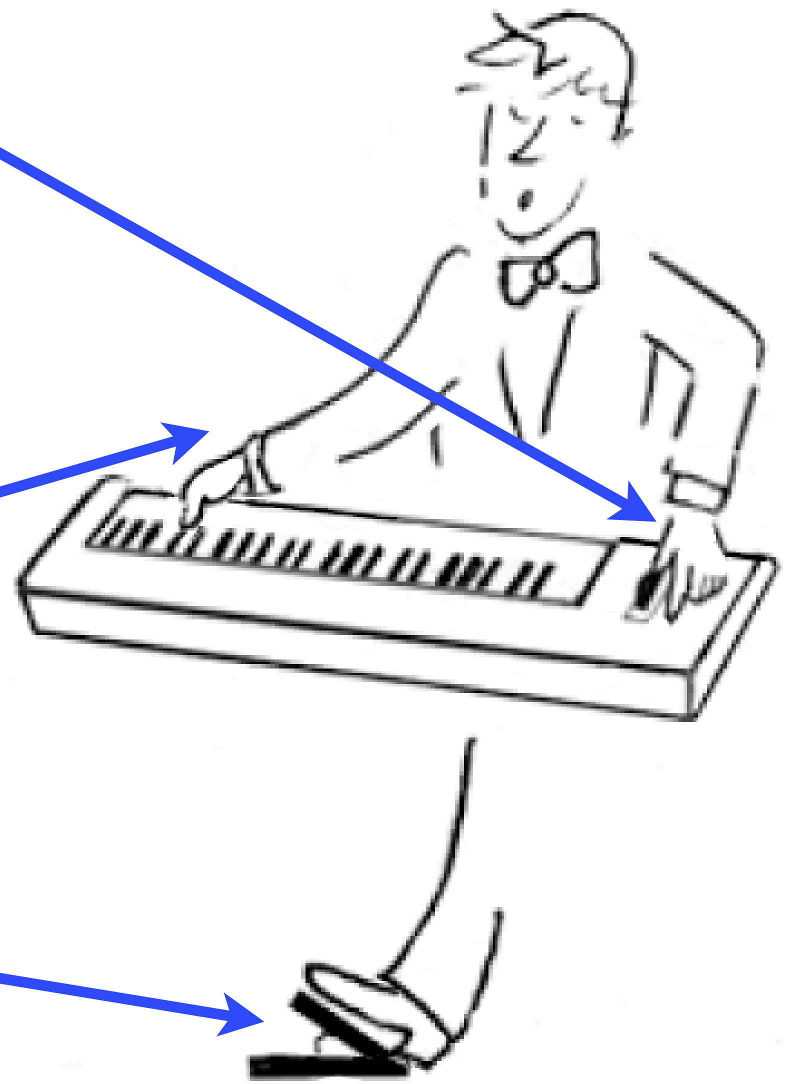

Infer Performance Intent from Playing ...

Synful Orchestra Features

Features

- **Realistic slurs, tonguing, and bowing created automatically from standard MIDI.** No special key switches or phrasing tools.
- **Synful Pitch Wheel mode** for realistic portamento slurs and pitch slide effects. Midi look-ahead anticipates upcoming notes and begins note transitions early –like a real player.

VIENNA SYMPHONIC LIBRARY

Performance Detection

Interval Detection
Provides natural intervals and note transitions for legato, portamento, glissando, spiccato, marcato, détaché styles in real-time.

Repetition Detection
Lets you play natural repeated notes in legato, portato, staccato, spiccato styles at any speed!

Pattern and Trill Detection
Avoids same-sample occurrences, even within complex phrases and trills.

Speed Detection
The tempo of your performance automatically triggers the appropriately articulated samples and switches seamlessly between articulations.

For best results: Edit by hand.

Example: Piano, cello, violin. One string instrument is programmed into Synful, one is a real recording.

Playing

Example: Paganini's Capriccio #1, for Garritan Stradivari.

Playing

What does editing look like?

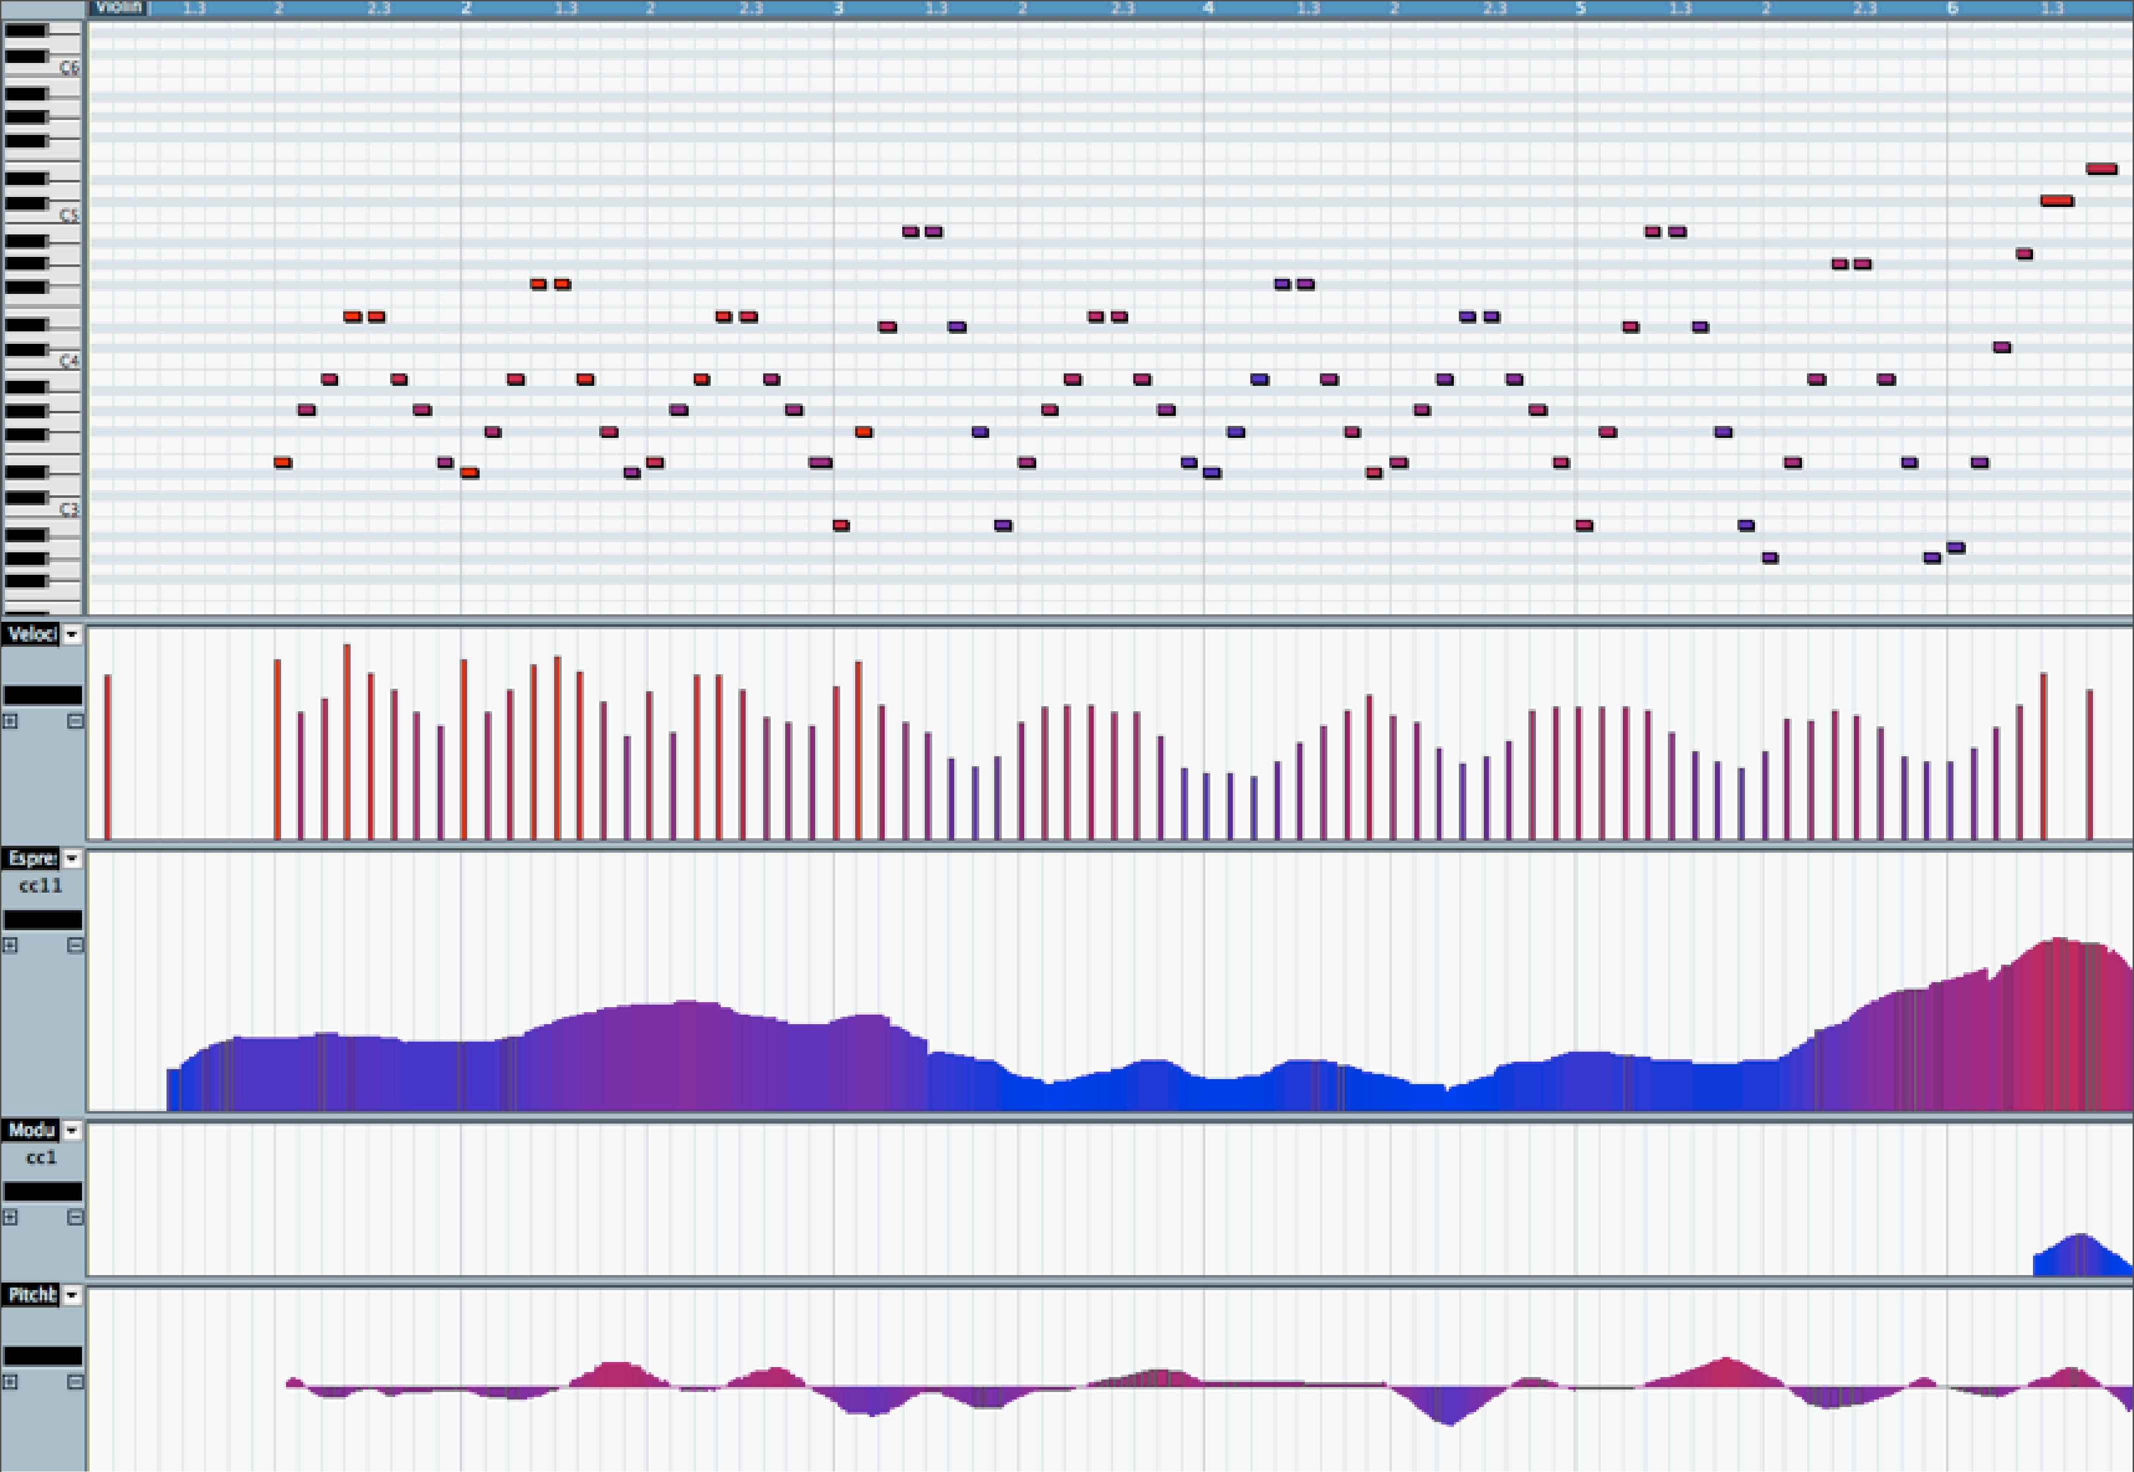

Fine tune note velocity + several controllers by hand.

Music 209 goals ...

Understand how to implement the sort of note-level concatenative synthesis the piano and violin products use ...

But, note-level synthesis isn't the sole focus of the course ...

Playing

Wayne Shorter

Steely Dan:
Walter Becker
and Donald Fagan

"We don't write scores for our invited soloists"

Solos arise from conversations between composers and soloist during the recording process ...

Note-based concatenative synthesis is not appropriate, because the composer doesn't know the notes he or she wants to hear!

Concept: Phrase-based synthesis

750 MBs of "one-breath long" sax solo phrases in many different styles.

Sorted by key, scale, length, tempo, style, instrument ...

Produces audition phrases in a browser, and assemble a solo. The "conversation with the cyber soloist"

Each phrase can be fine-tuned using Celemony Melodyne technology (resynthesis editor)

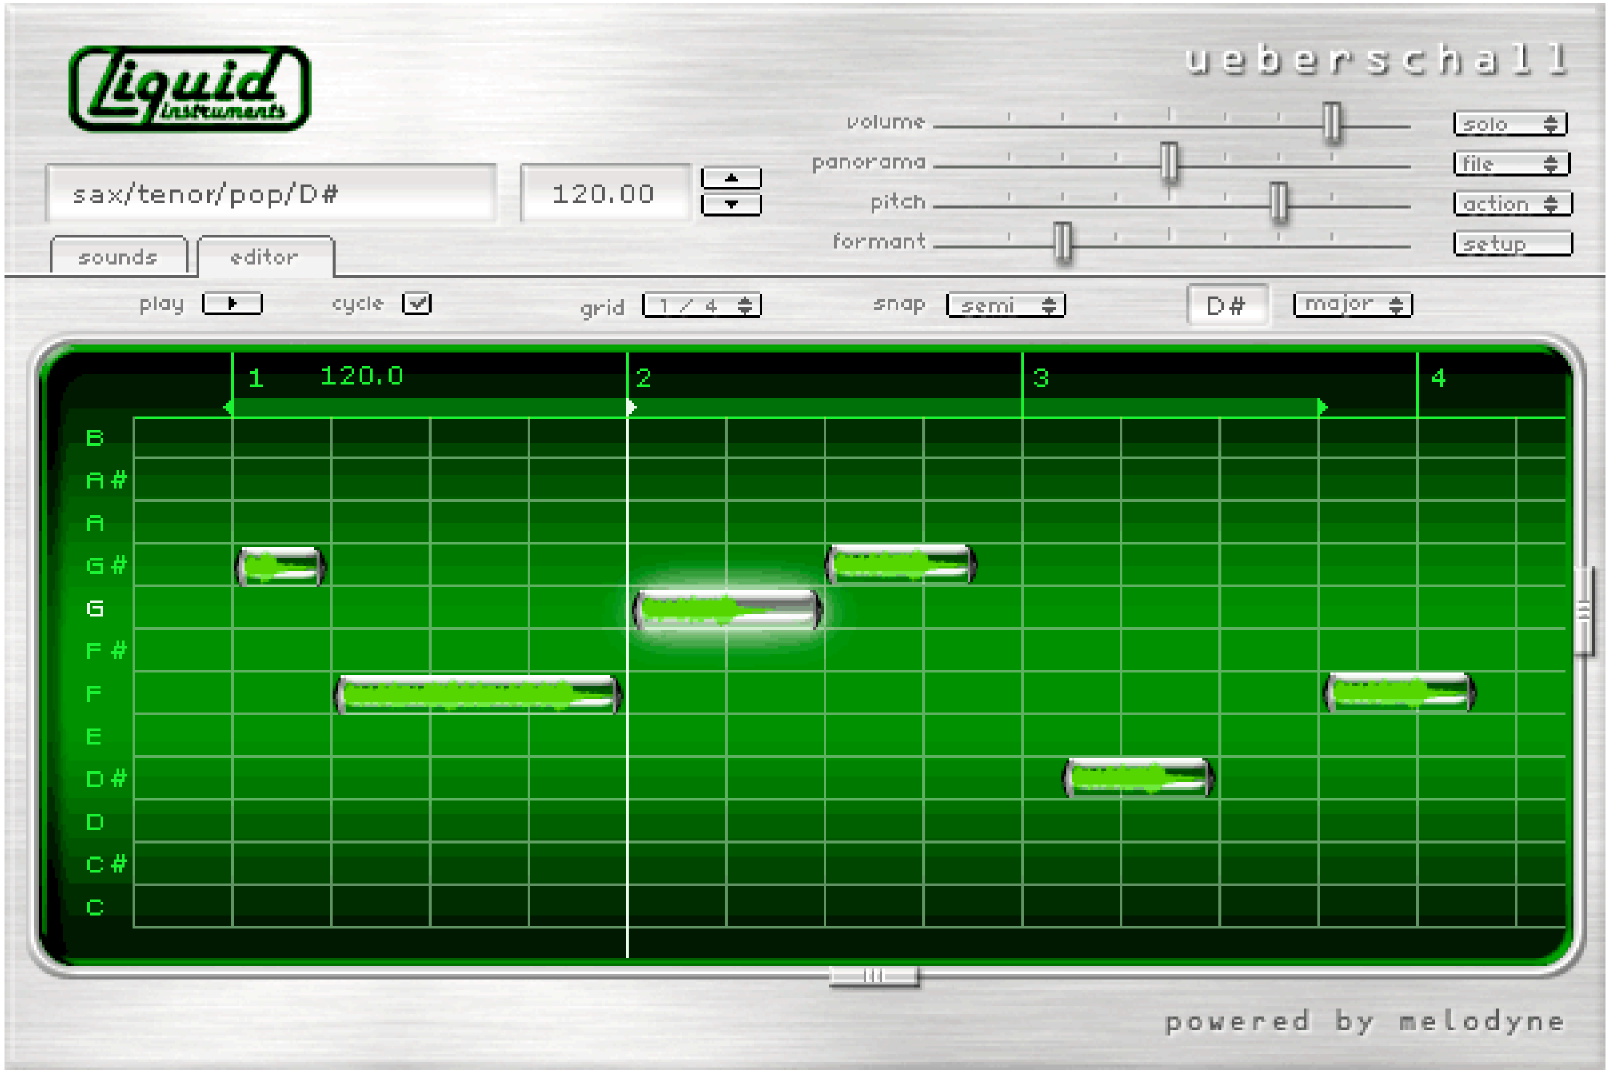

db funk ballad brazil

Course Outline

Week	Date	Title
1	Jan 19	**Overview.** Introduction to concatenative synthesis. Details on the course project.
2	Jan 26	**Splicing.** Evaluating the quality of a candidate concatenation, by comparing pitch, loudness, spectrum, and volume across the concatenation. Methods for performing splices in the time domain and the spectral domain (morphing).
3	Feb 2	**Speech Synthesis.** Guest lecture on concatenative speech synthesis. Concatenative music techniques that are patterned after speech synthesis (diphone synthesis).

4	Feb 9	**Time-Warping.** Concatenation of two rhythmic units may require the tempos to match. We describe techniques for matching tempos without otherwise altering units.
5	Feb 16	**Pitch-Shifting, Spectral Matching, Formant Matching.** Techniques for altering one sample to match another in preparation for a splice.

6	Feb 23	**Real-time Note Concatenation Control.** Techniques for real-time note-level control, such as a MIDI hardware controller.
7	Mar 1	Project Abstracts Due (via email, Mar 1 is not a lecture day). See link.
7	Mar 2	**Off-line Note Concatenation Control.** Techniques for synthesizing audio from score notation or guide audio.

8	Mar 9	**Phrase Concatenation Control.** Technology for assembling phrase-level units, as one does with Apple Loops in Garageband. Also, technology for concatenating completed compositions (example: DJ software) and reworking complete compositions (example: remixes).		[link]
9	Mar 16	**Automatic Classification.** Algorithms for speech/non-speech detection, and similar automatic classification techniques. Also, an introduction to the speech recognition pipeline. (David away)		[link]
10	Mar 23	**In-Class Student Progress Report Presentations (see link)**		[link]
11	Mar 30	No Class - Spring Break		

12	Apr 6	**Sound Separation.** State of the art in removing vocals from complete tracks, and related separation problems.
13	Apr 13	**Room Modeling and Spatialization.** Impulse responses of rooms and other interesting spaces. Concatenating spatialized samples.

14	Apr 20	**MetaData.** Approaches to labelling sound databases in a standardized way. Possible guest lecture by MuscleFish.	
15	Apr 27	**Computer Systems.** How disk and audio I/O work in audio today, and ideas for new APIs. May include final project presentations for CoreSample. (David away)	
16	May 4	**In-Class Student Final Project Presentations (see link)**	
17	May 8	**Final Project Report Due (see link)**	

Projects

Music 209 : Projects

The grade for this course is based on the completion of a single, semester-long project, whose topic is related to concatenative synthesis. The forms a project may take include (but are not limited to):

- Proposal for new and improved algorithms related to concatenative synthesis (example: a new spectral morphing algorithm). A good project of this type should include an audio demonstration of the algorithm that compares it to existing approaches.
- Software prototype of a complete concatenative system.
- Proposal for a new user-interface related to concatenative systems. A good project of this type should include storyboards for showing how the user-interface works, and a comparison with existing interfaces.
- A composition or performance that explores a novel use of concatenative synthesis.

A project may be a solo effort, or may be a collaboration between several students.

Key milestones for the project appear below.

Title	Due Date	Description	Percent of Grade
Project Abstract	March 1, 11:59 PM	A short (one or two page) description of the project. PDF or plain text format is fine -- please, no .doc files. Collaborative projects should include information on how the work will be split between team members. Email this abstract to the instructors (wessel [at] cnmat [dot] berkeley [dot] edu, lazzaro [at] eecs [dot] berkeley [dot] edu).	5 percent
Progress Report Presentation	March 23 in class	A 10-15 minute presentation to the class, describing the current status of the project. Group projects should share presentation duties between all members. Audio demos of work in progress is encouraged. Primary purpose of presentation is to solicit feedback from the audience.	15 percent

Final Presentation	May 4th or April 27 in class	A final presentation to the class describing your finished project. If your project includes a compositional aspect, the composition should be performed or played during the presentation. Software projects should include a live demo; user-interface projects should pitch the final storyboard of the UI. Nominal presentation time is 15 minutes, but can be longer if the project requires it (let us know in advance).	30 percent
Final Report	Monday May 8th, 11:59 PM	A final report describing the project: what you planned to do, how it turned out, lessons learned. Main body of the report should be a PDF file 10-20 pages in length, and should be emailed to the instructors by the due date. In addition to the main body of the report, you may also submit supporting materials (examples: audio files for completed compositions, storyboards for a UI, a Quicktime movie showing a demo of a software application, etc). Send support materials as	50 percent

You are free to propose a project topic of your own creation. Alternatively, you may choose one of the project ideas below (click on the link for a complete description).

Complete list to come ... here is one example.

Piano-related class project idea

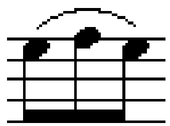

Legato: Play without attack.

Real pianos cannot (really) do this:
Every note starts with a hammer hit.

Project: Design a sample-based piano that detects legato MIDI playing style, and transitions from note to note by "splicing" to the new note at an equal energy point in the envelope.

Tools for the project ....

Developer Connection

Log In | Not a Member?

ADC Home >

Audio

Mac OS X delivers superior audio services, and with system-level services that streamline the development process for audio, it allows you to incorporate high-quality music, audio media and audio functionality into your applications. The audio system consolidates, integrates, and standardizes third-party audio and MIDI services and protocols, thereby shortening the hardware and software development cycle and streamlining device configuration. Mac OS X provides powerful system-level functionality, including native multichannel audio with plug-in support, native MIDI services, and driver support, as well as the services of Core Audio, including Audio Units and Hardware Application Layer (HAL) APIs, and MIDI services. With Mac OS X, you get the most comprehensive set of audio services ever provided for audio software and hardware developers. Read More...

Getting Started

A guided introduction and learning path for developers new to Audio.

Tools for the project ....

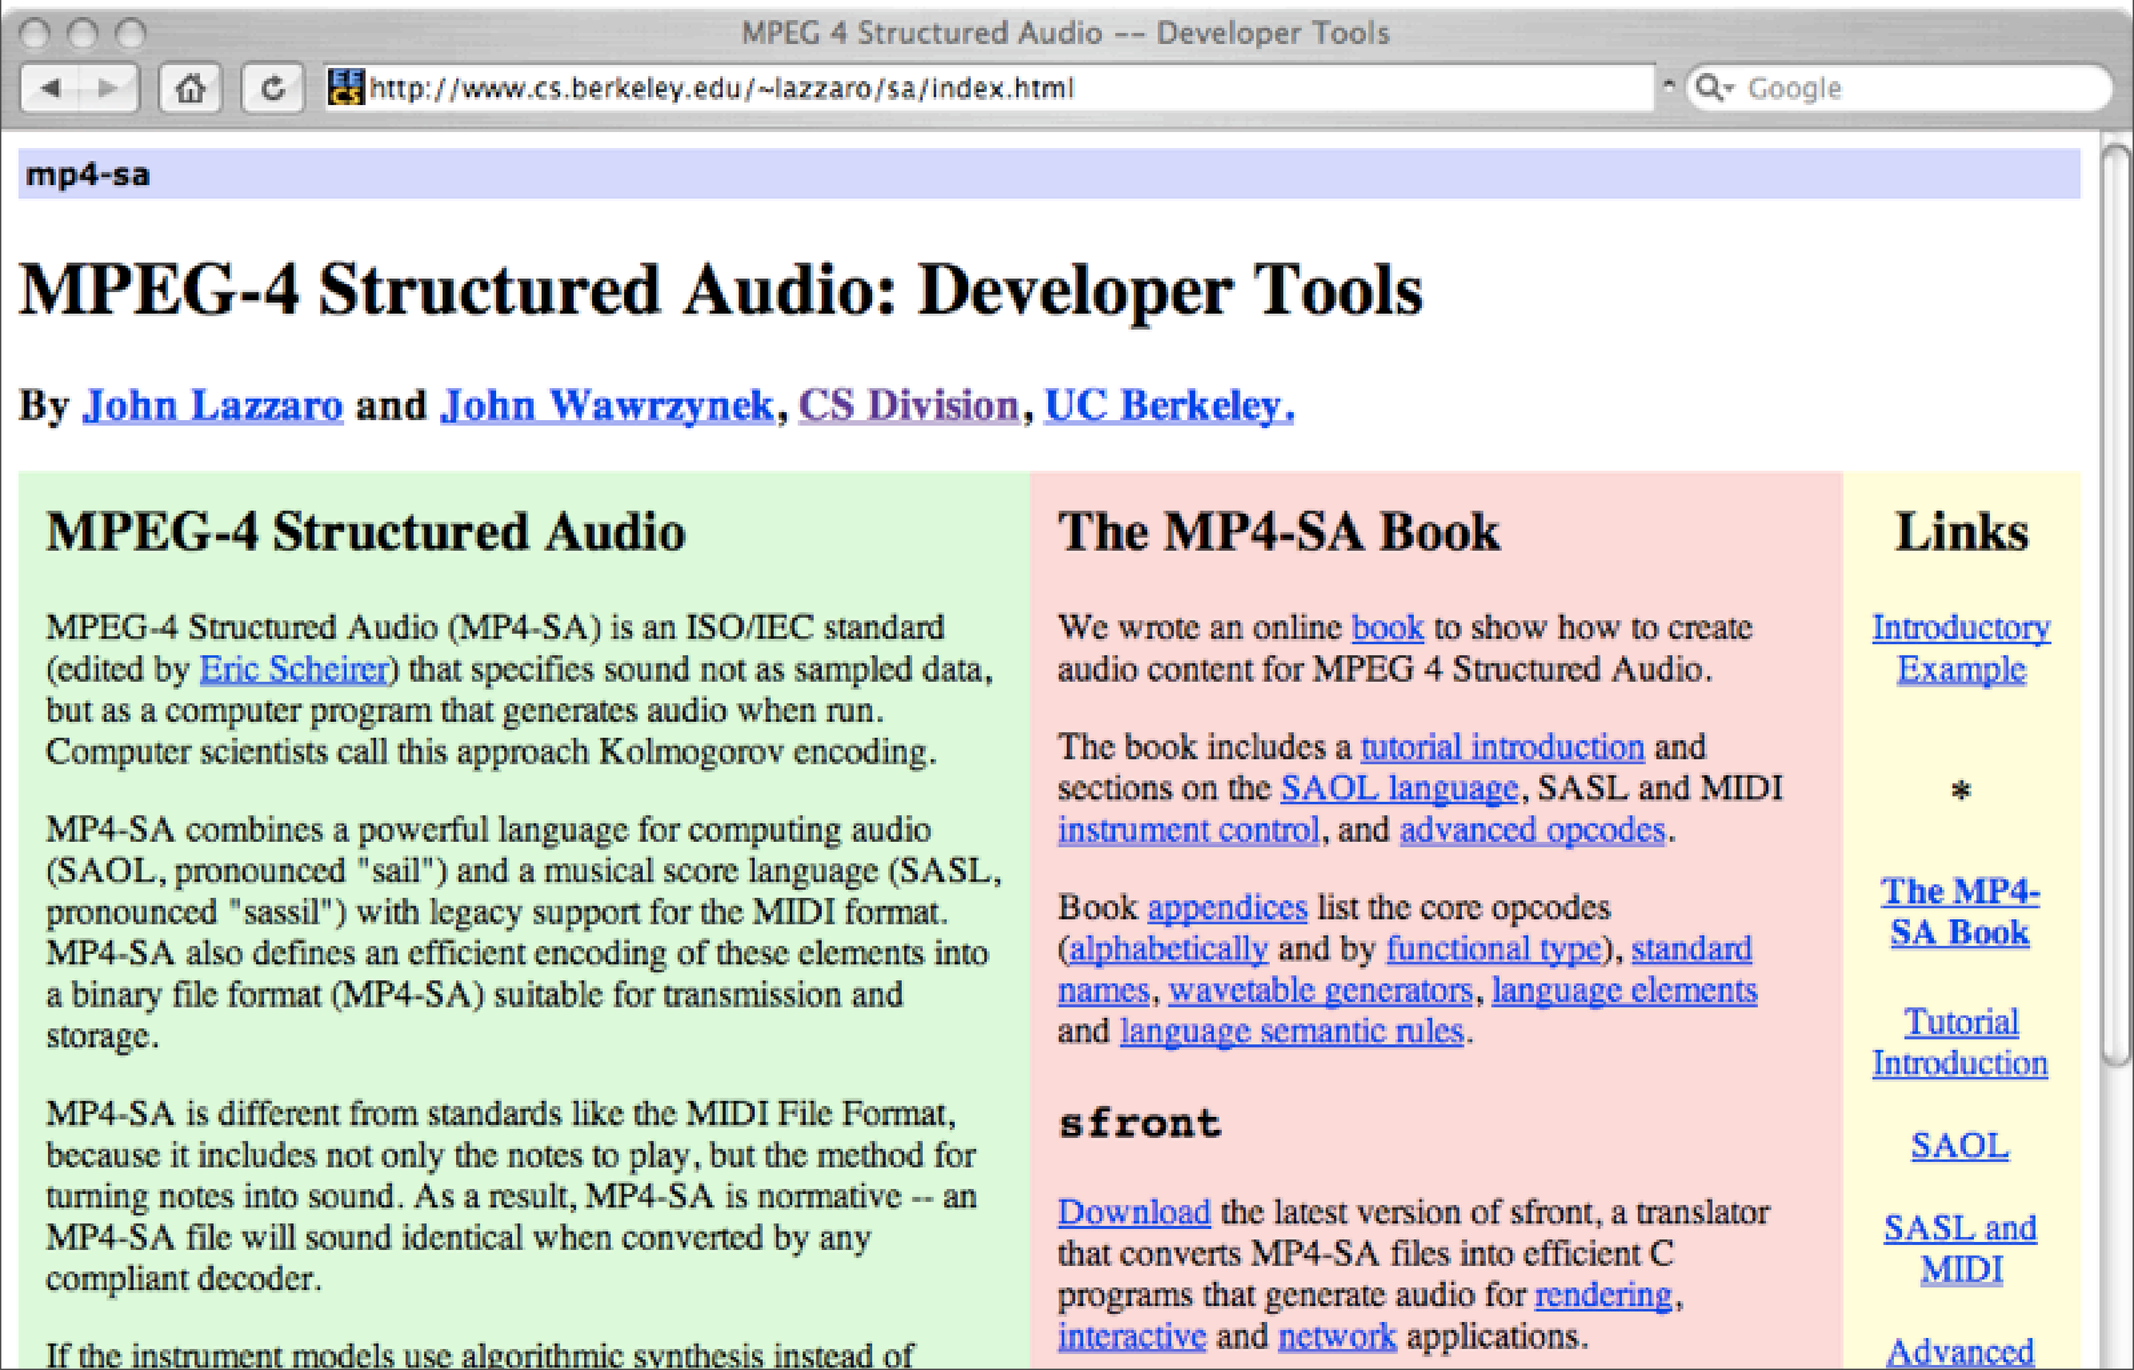

MPEG 4 Structured Audio -- Developer Tools

http://www.cs.berkeley.edu/~lazzaro/sa/index.html

Google

mp4-sa

MPEG-4 Structured Audio: Developer Tools

By John Lazzaro and John Wawrzynek, CS Division, UC Berkeley.

MPEG-4 Structured Audio

MPEG-4 Structured Audio (MP4-SA) is an ISO/IEC standard (edited by Eric Scheirer) that specifies sound not as sampled data, but as a computer program that generates audio when run. Computer scientists call this approach Kolmogorov encoding.

MP4-SA combines a powerful language for computing audio (SAOL, pronounced "sail") and a musical score language (SASL, pronounced "sassil") with legacy support for the MIDI format. MP4-SA also defines an efficient encoding of these elements into a binary file format (MP4-SA) suitable for transmission and storage.

MP4-SA is different from standards like the MIDI File Format, because it includes not only the notes to play, but the method for turning notes into sound. As a result, MP4-SA is normative -- an MP4-SA file will sound identical when converted by any compliant decoder.

If the instrument models use algorithmic synthesis instead of

The MP4-SA Book

We wrote an online book to show how to create audio content for MPEG 4 Structured Audio.

The book includes a tutorial introduction and sections on the SAOL language, SASL and MIDI instrument control, and advanced opcodes.

Book appendices list the core opcodes (alphabetically and by functional type), standard names, wavetable generators, language elements and language semantic rules.

sfront

Download the latest version of sfront, a translator that converts MP4-SA files into efficient C programs that generate audio for rendering, interactive and network applications.

Links

Introductory Example

*

The MP4-SA Book

Tutorial Introduction

SAOL

SASL and MIDI

Advanced

Tools for the project ....

After the break, David will describe other tools ...