Music 209
Advanced Topics in Computer Music
Lecture 5 – Pitch Shifting

2006-2-16

Professor David Wessel (with John Lazzaro)

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

www.cs.berkeley.edu/~lazzaro/class/music209

UC Regents Spring 2006 © UCB

Today, we focus on **pitch shifting** techniques to change the note pitch of candidates.

Select candidate samples from db

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

Do the splice

#3(mod)

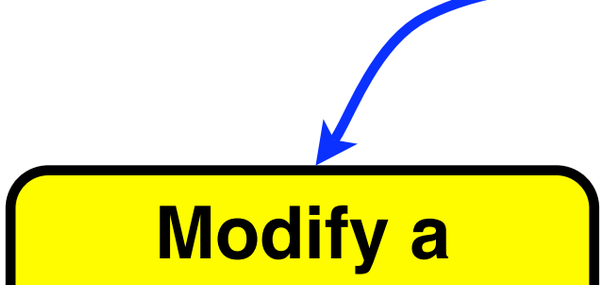

Y

Choose best match

#2

Do the splice

#2

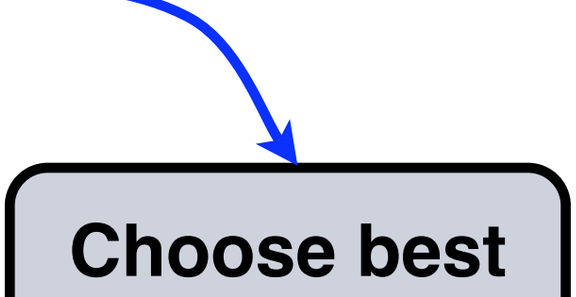

Candidates

#1

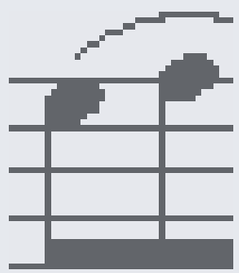

#2

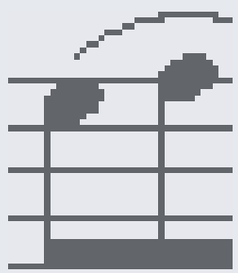

#3

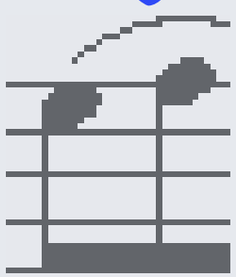

Topics for today ...

- **Pitch shifting: transposition of the pitch musical audio in a "transparent way".**
- **Time-domain pitch shifting algorithms.**
- **Project ideas on parade.**

UC Regents Spring 2006 © UCB

define: Pitch shifting

Play.

Pitch-shifted audio should sound like musician is playing a transposed score.

Original key is F

+5 semitones ...

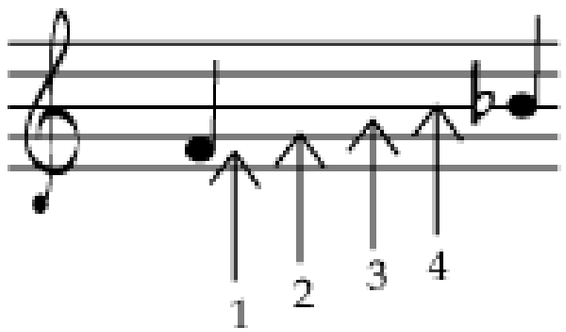

Move the key and every note up a perfect fourth.

New key is B flat

Play.

- Note and rest **durations** should be **unchanged.**
- **Timbre** should sound as it does when playing **transposed parts** on the instrument.
- Unpitched sounds should be unaffected.

UC Regents Spring 2006 © UCB

Music pitch mathematics ...

Pitch in Hz: The frequency of a sine wave whose pitch is heard to be the same as the played note.

Octaves: Pitches differ by an exact factor of two.

Equal temperament: Equal frequency multiple α between all adjacent piano keys.

$$440 \text{ Hz} = \alpha^{12} \times 220 \text{ Hz} \longrightarrow \alpha = \sqrt[12]{2} = 1.05946309 \ldots$$

Equal temperament is an engineering compromise ...

UC Regents Spring 2006 © UCB

Timbre and Pitch

Why are the timbres different?

Contributing factor: Height of bars differ, and evolve differently over time. "Spectral Shape"

Same pitch, different timbre

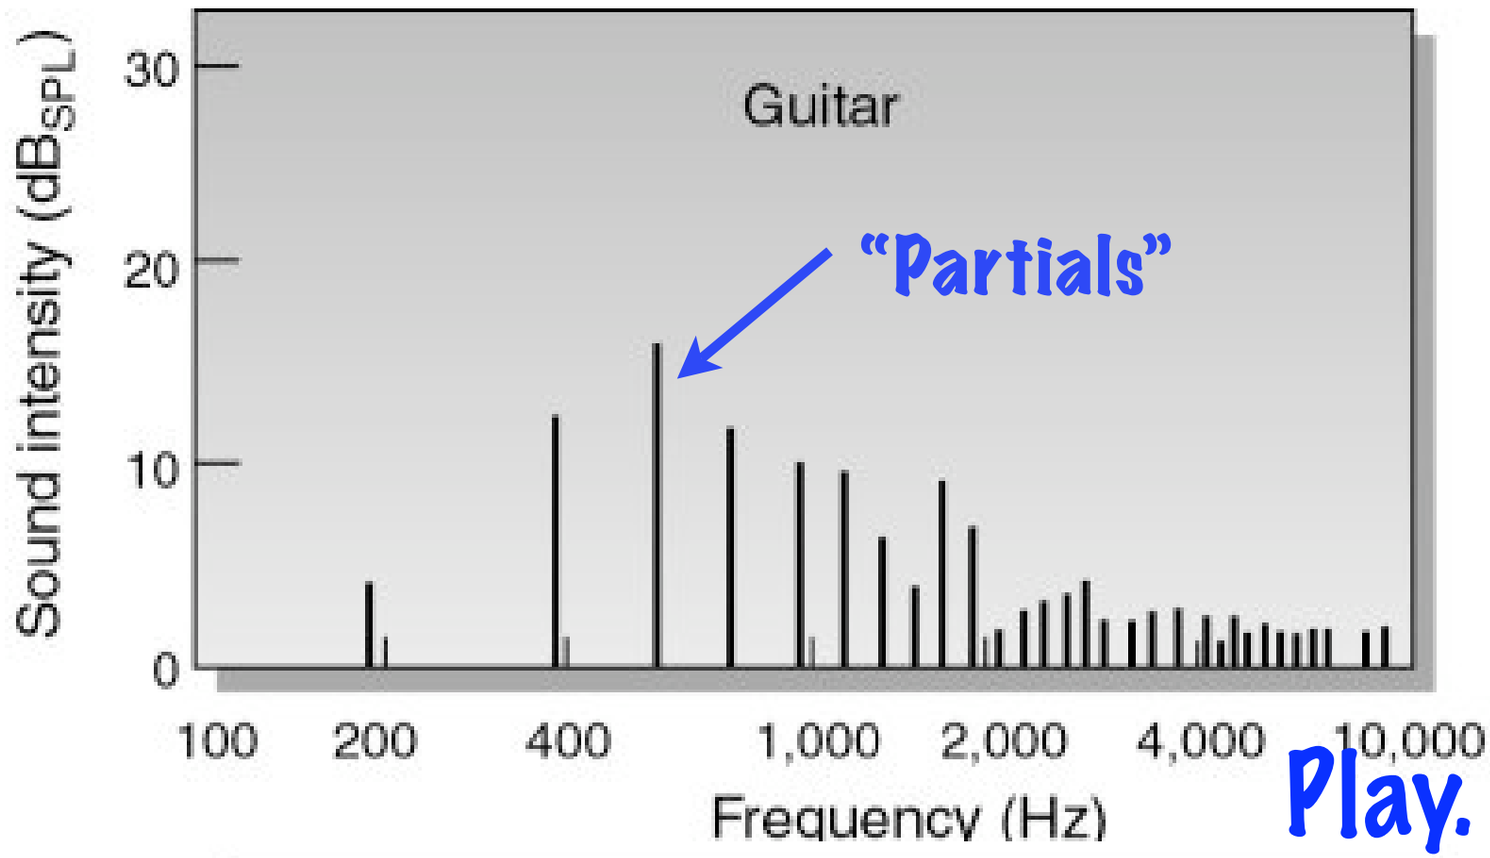

Play.

Why are both sounds pitched?

Why is the pitch the same?

Frequency placement of bars share a common structure & placement...

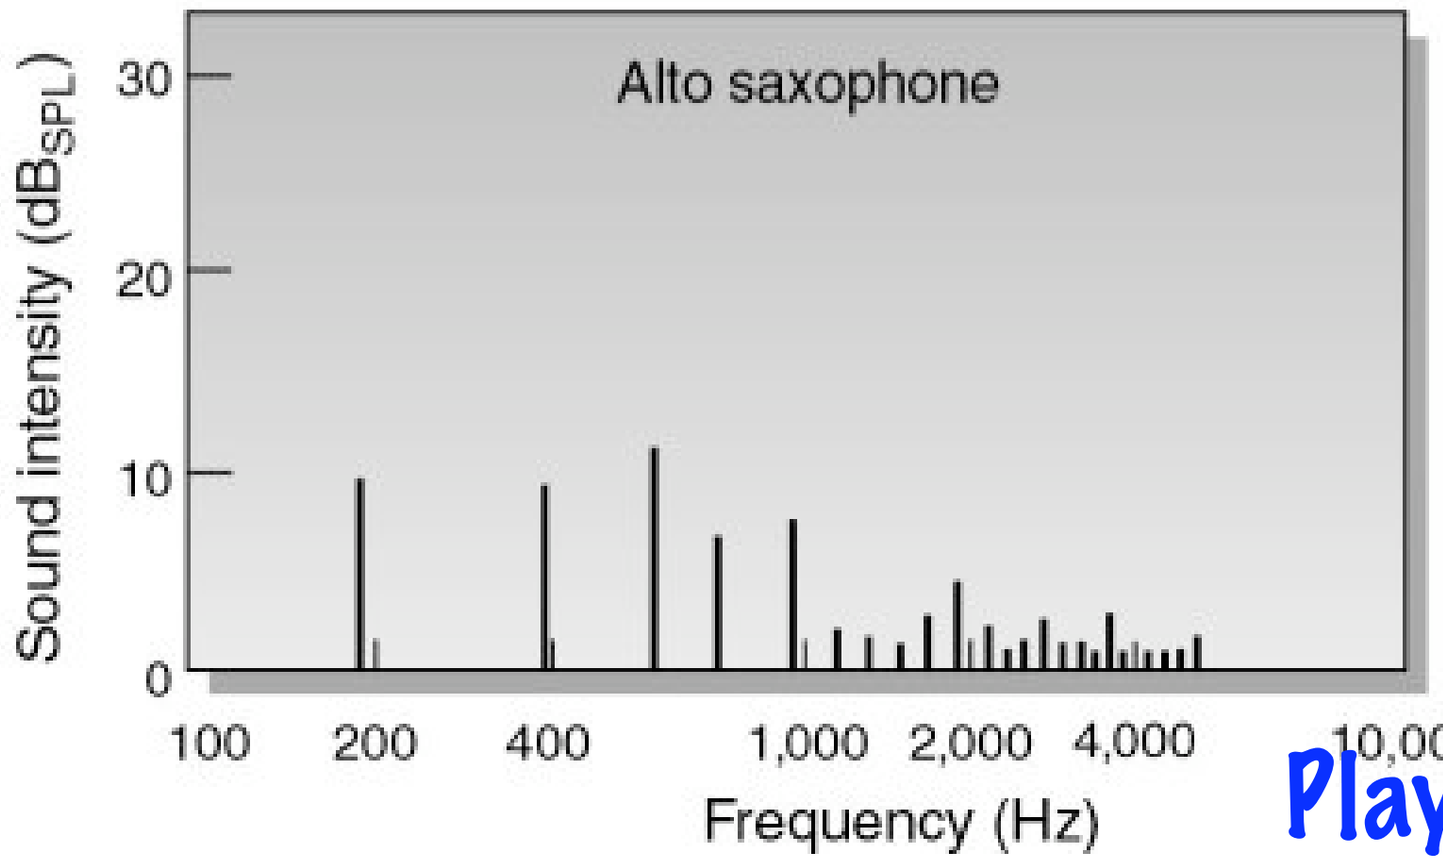

Play.

EXAMPLE BY KENNETH STEELE, APPALACHIAN STATE .

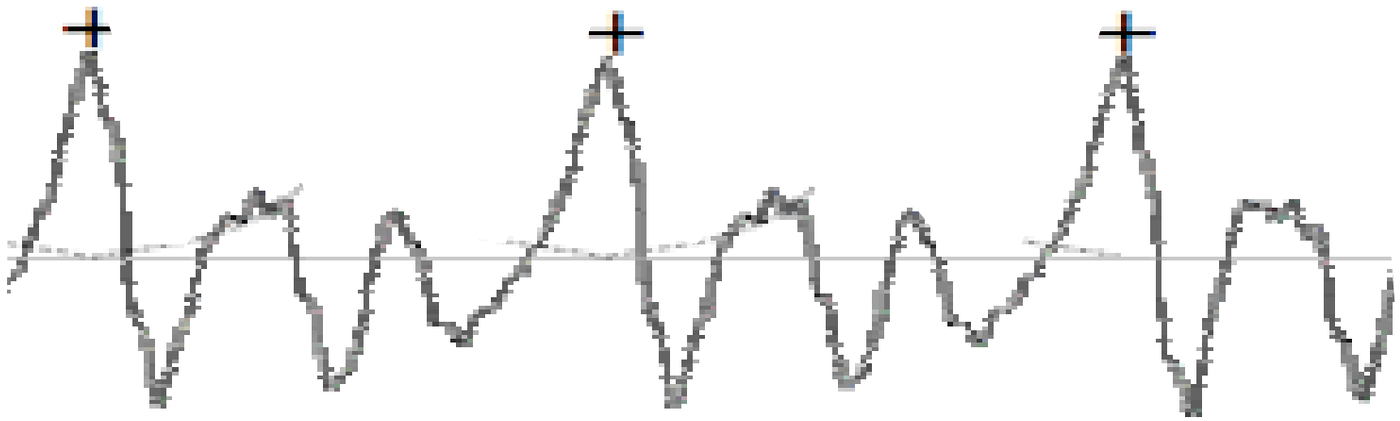

Summed waveform repeats at pitch frequency.

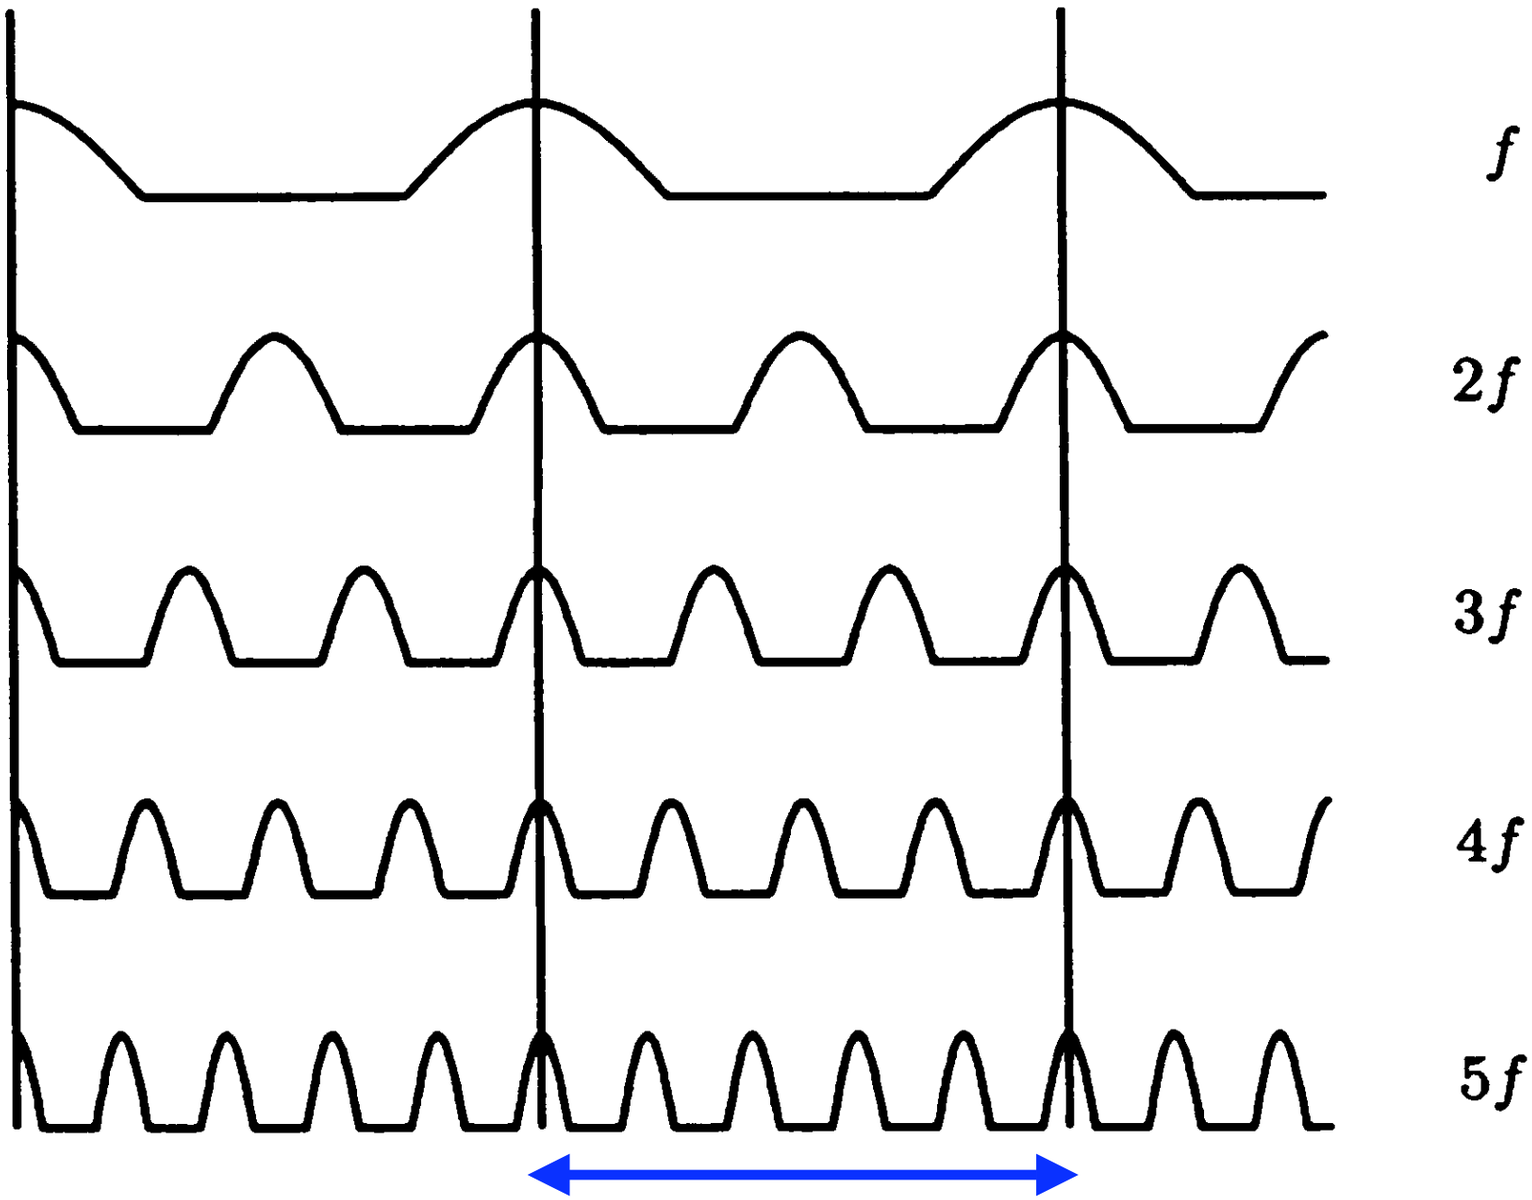

Frequencies of partials are integer multiples of an underlying fundamental.

Pitch Period = 1/(Pitch Frequency)

Caveats

First partial not necessary to detect pitch - A and B → are heard with same pitch.

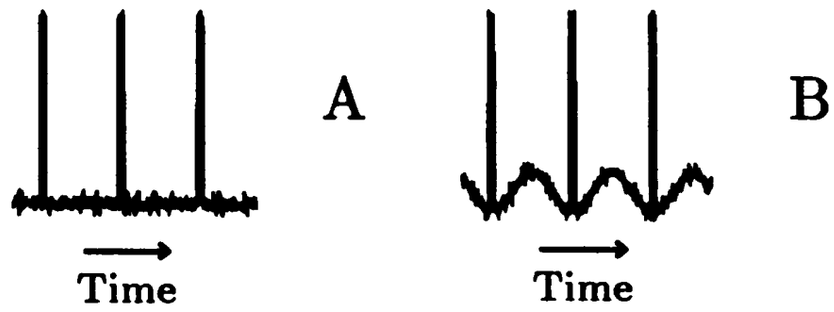

Relative phases of partials need not be aligned - any phase relation yields a strong pitch.

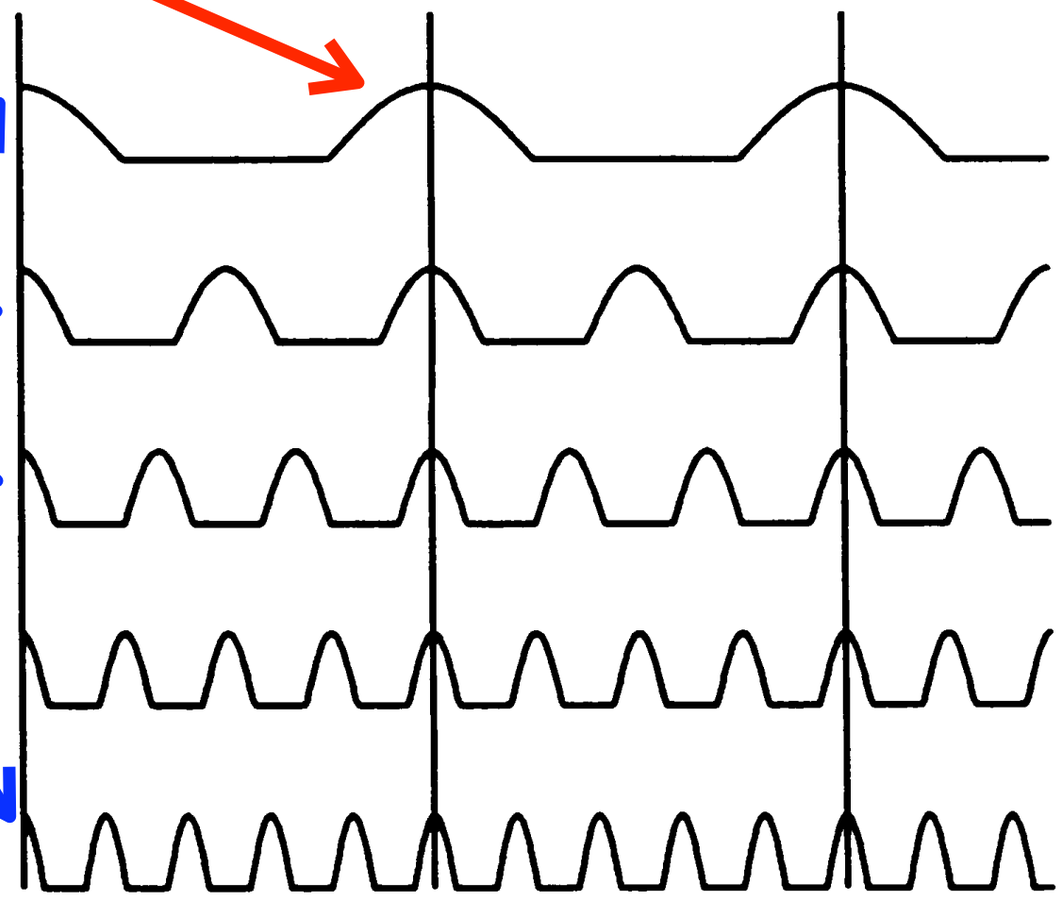

Thus ... repeating shape may be subtle to detect directly.

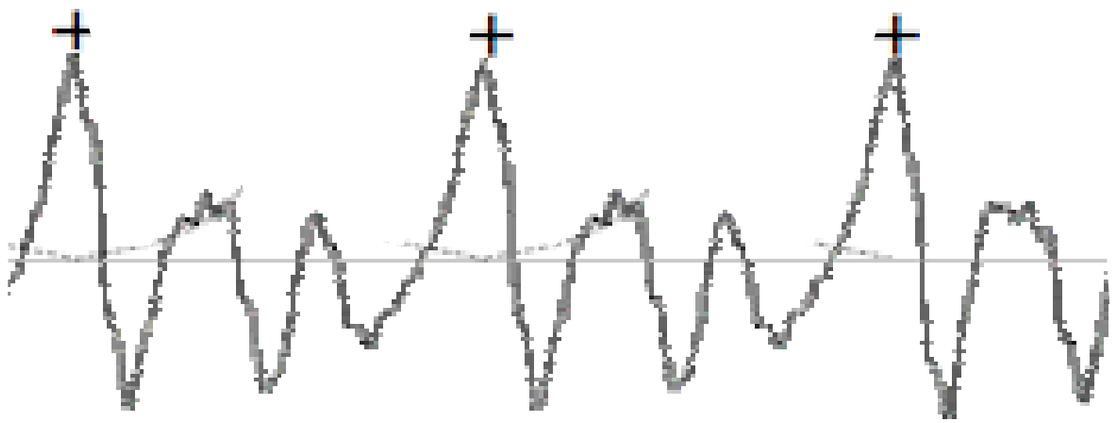

Sounds whose partials are not quite integer - related still yield a sense of pitch -

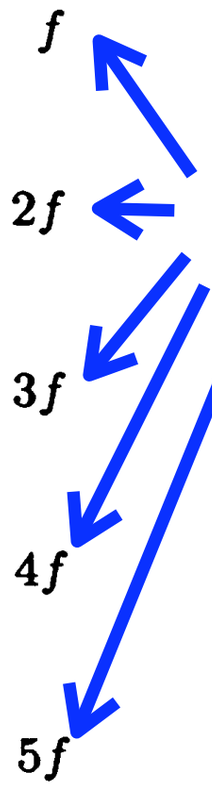

Bells ...

Lowest partials are exact integers, but higher partials are quasi-harmonic (4.2, 5.4, 6.8).

We still hear the bells as having a definite pitch.

Play.

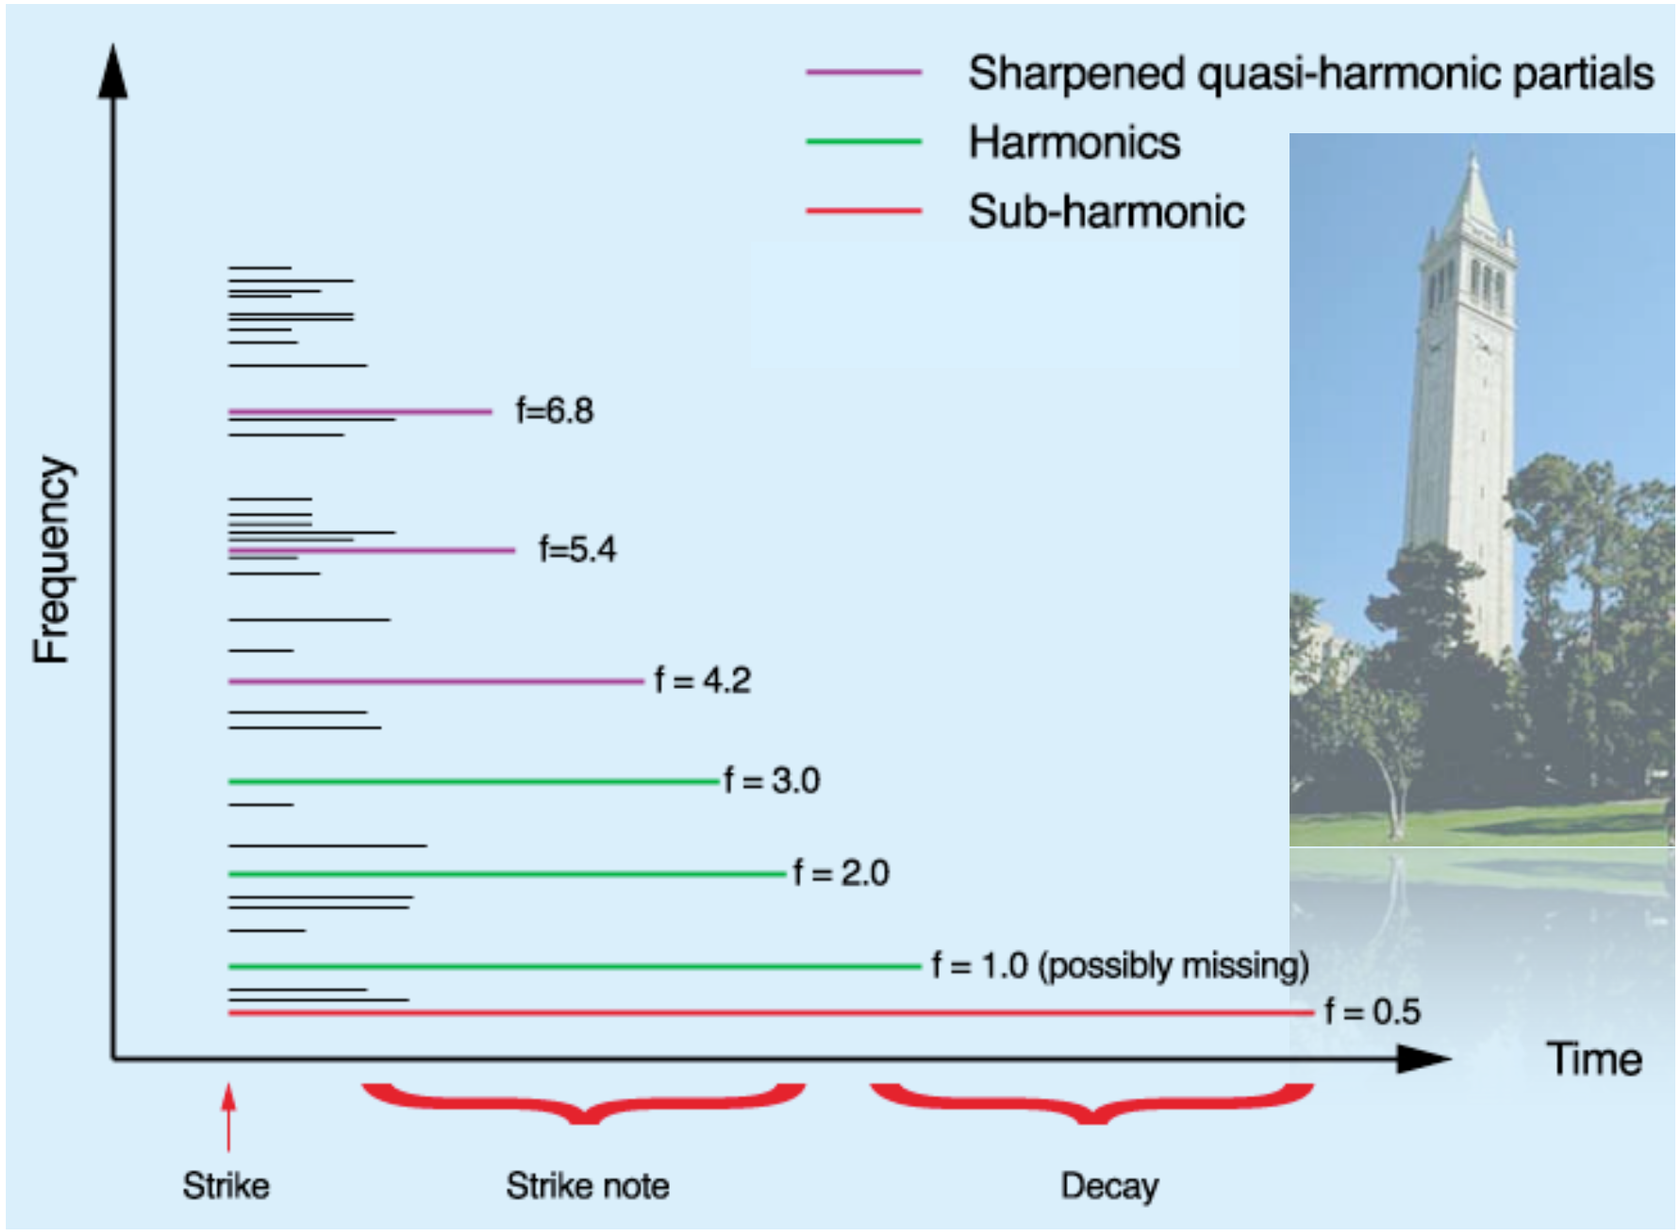

FROM SOUND ON SOUND, SYNTH SECRETS, AUG 02

Return to today's topic: **Pitch shifting**

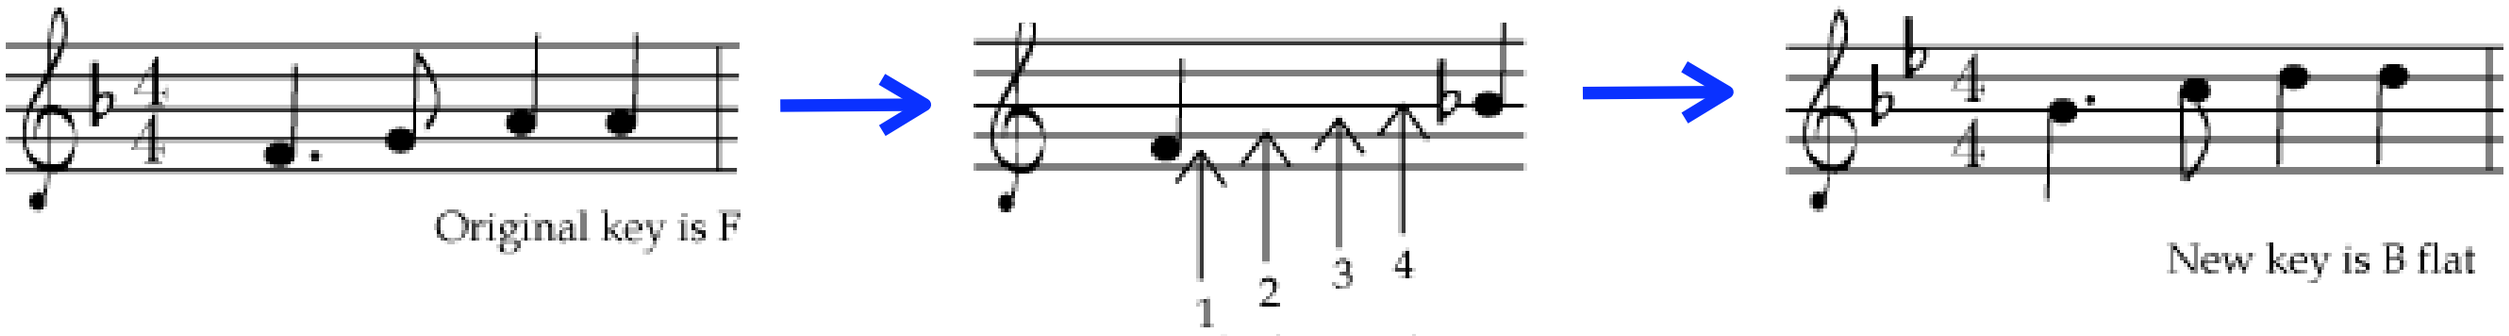

+5 semitones = $(\sqrt[12]{2})^{5}$ = 1.33483985

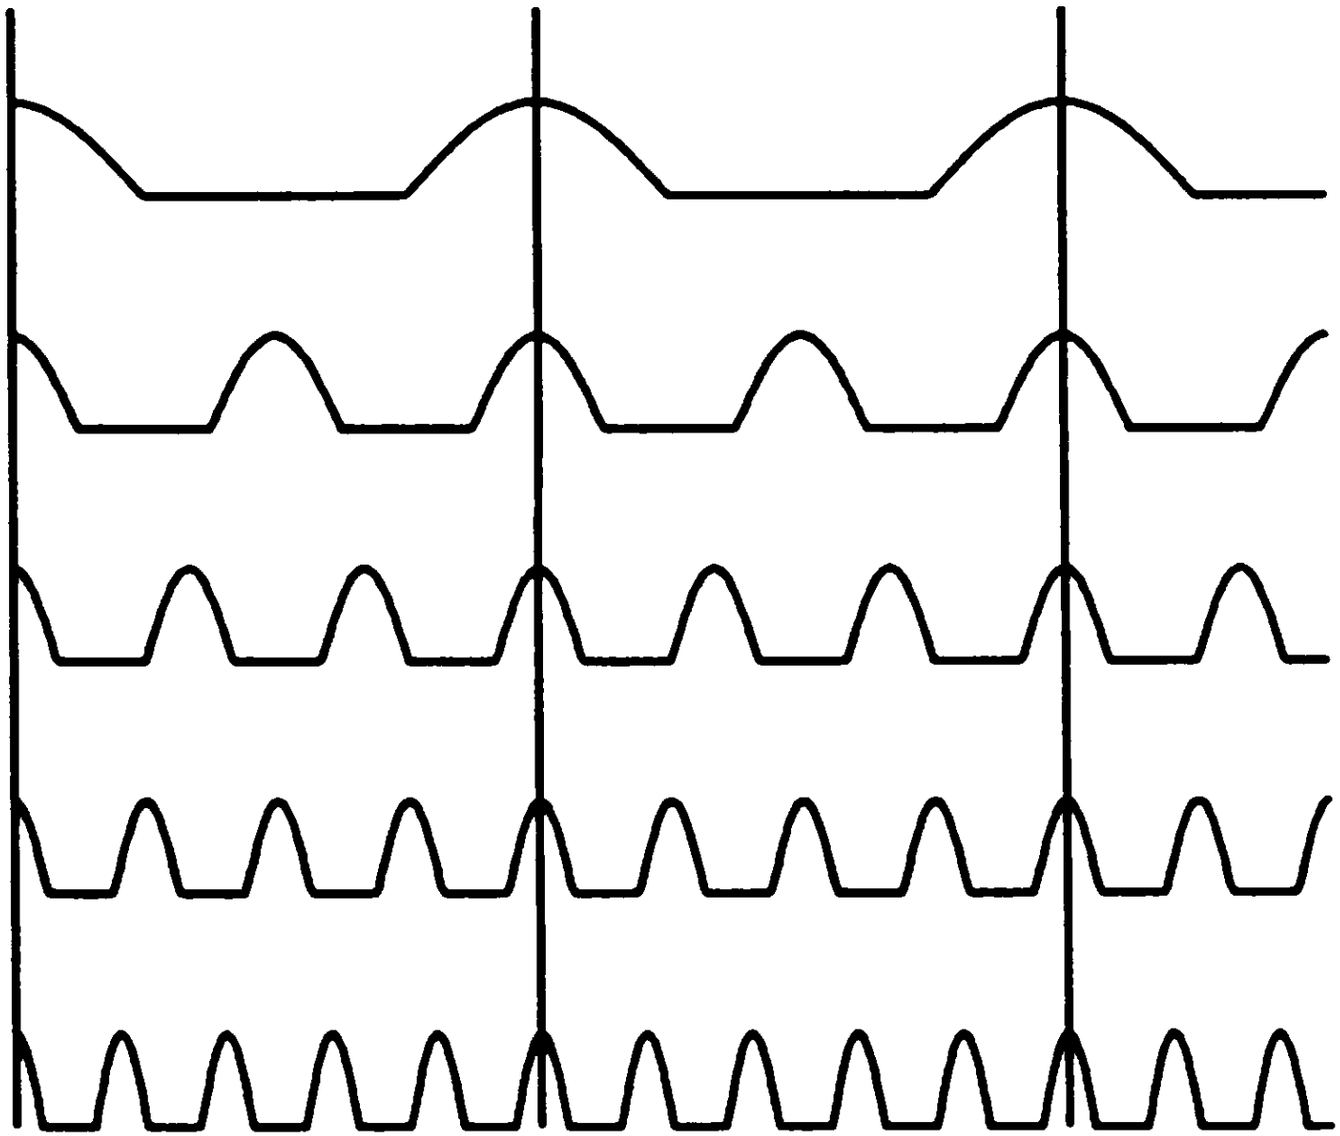

+5 semitones

1.33483985 f

2.66967971 f

4.00451956 f

5.33935942 f

6.67419927 f

UC Regents Spring 2006 © UCB

Possible strategy: resampling

Sample period: 1 / (44100 Hz)

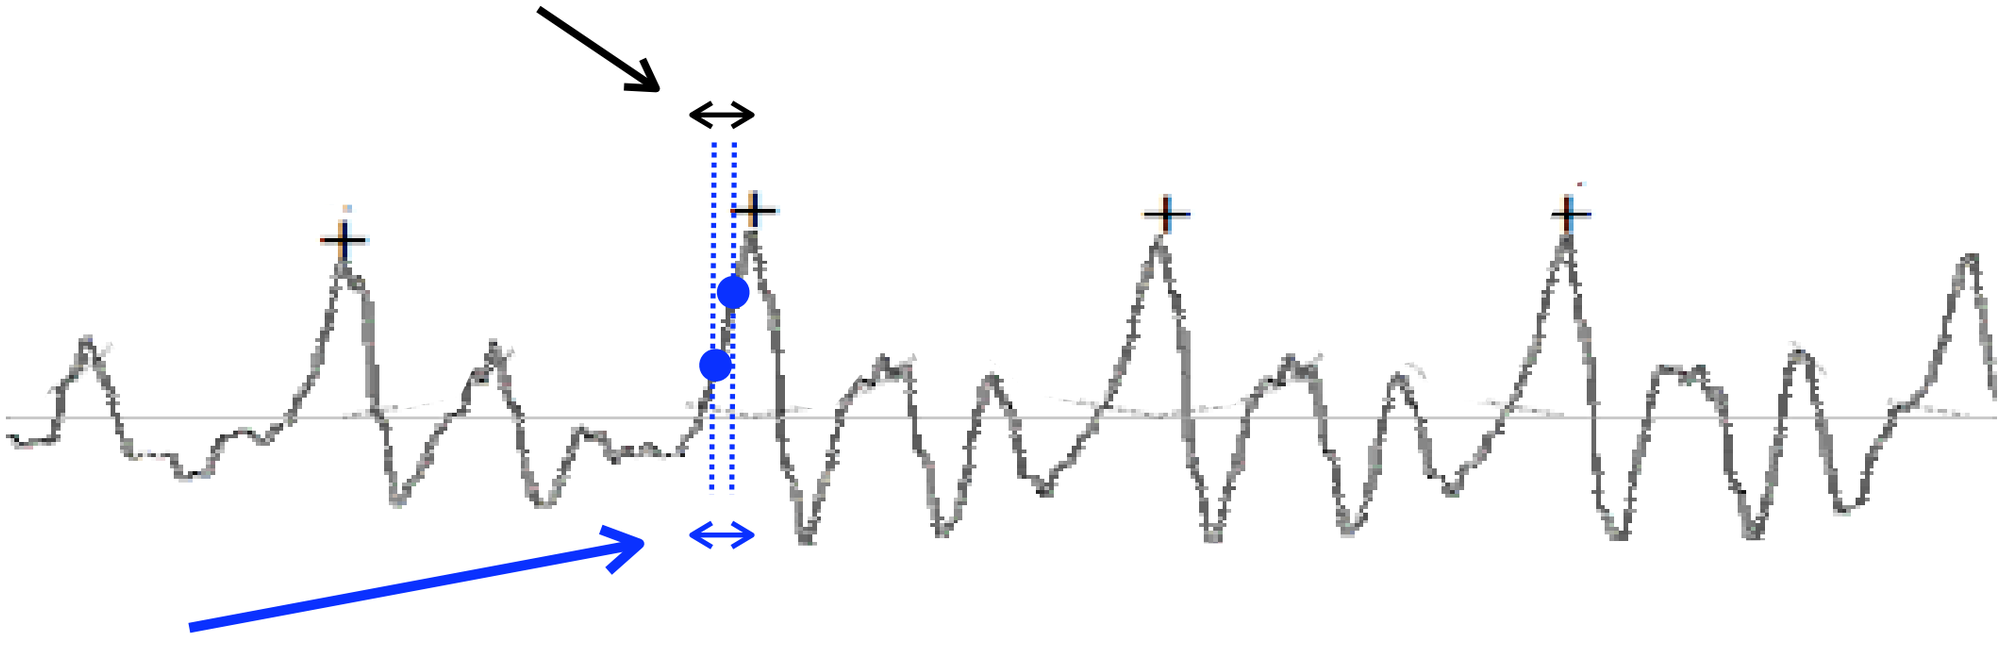

Redefine sample period: $1/(1.33483985 \times 44,100\ \text{Hz})$

Then, interpolate the sampling rate back to 44 100 Hz, being careful to avoid aliasing.

UC Regents Spring 2006 © UCB

Problem #1: Sample is now shorter!

Before:

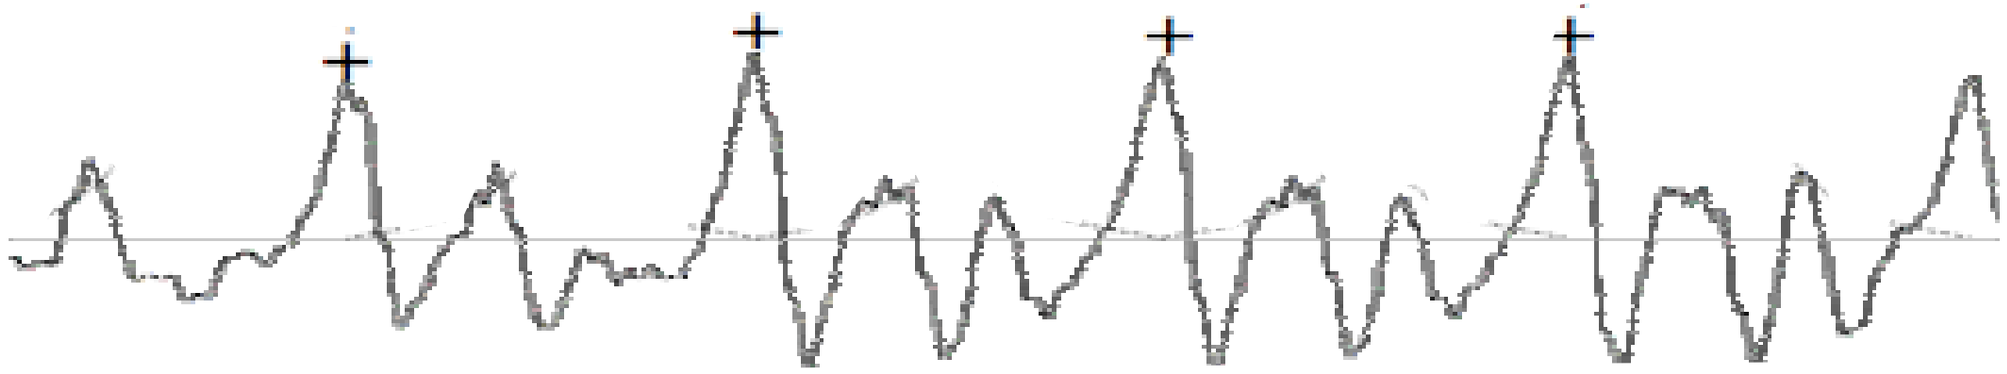

After:

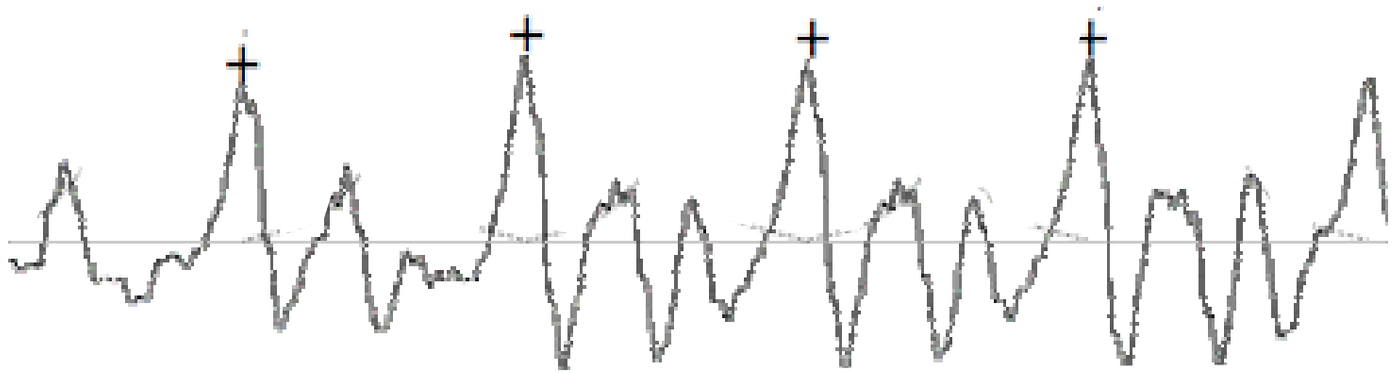

✳ Note and rest durations should be unchanged.

Solution: Time-stretch "after" waveform (last lecture).

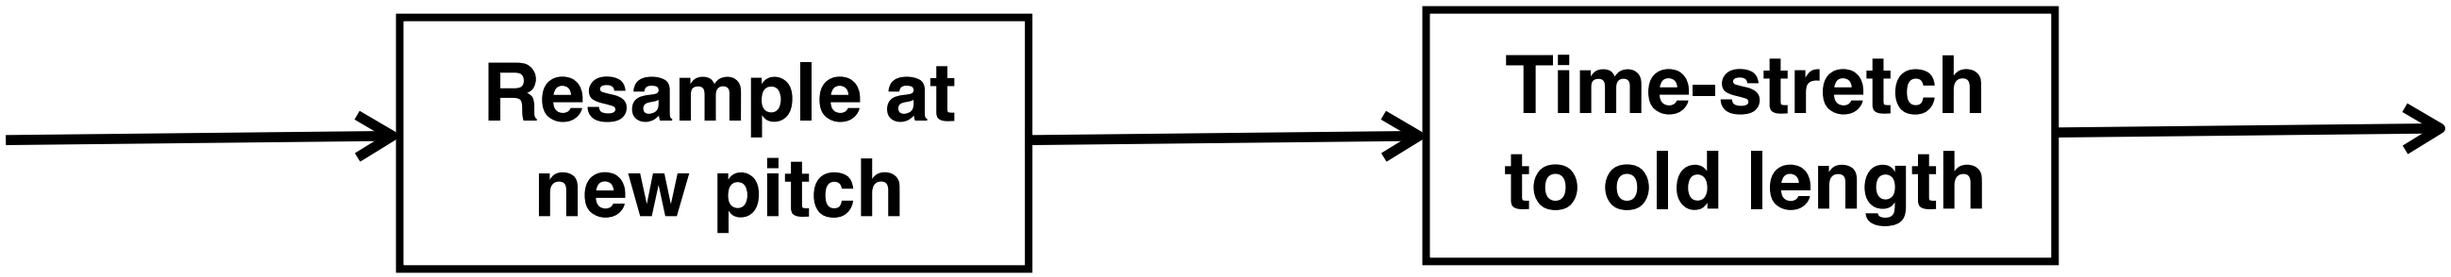

UC Regents Spring 2006 © UCB

Problem #2: Timbre

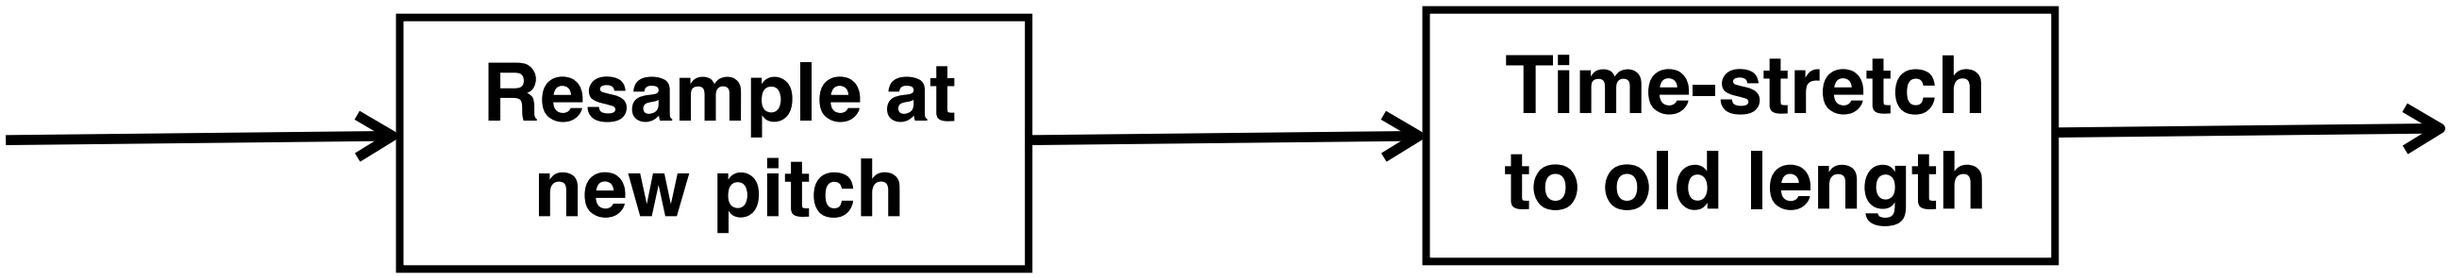

Resampling preserves the spectral magnitude of each partial.

Is this what we want? Depends on the instrument.

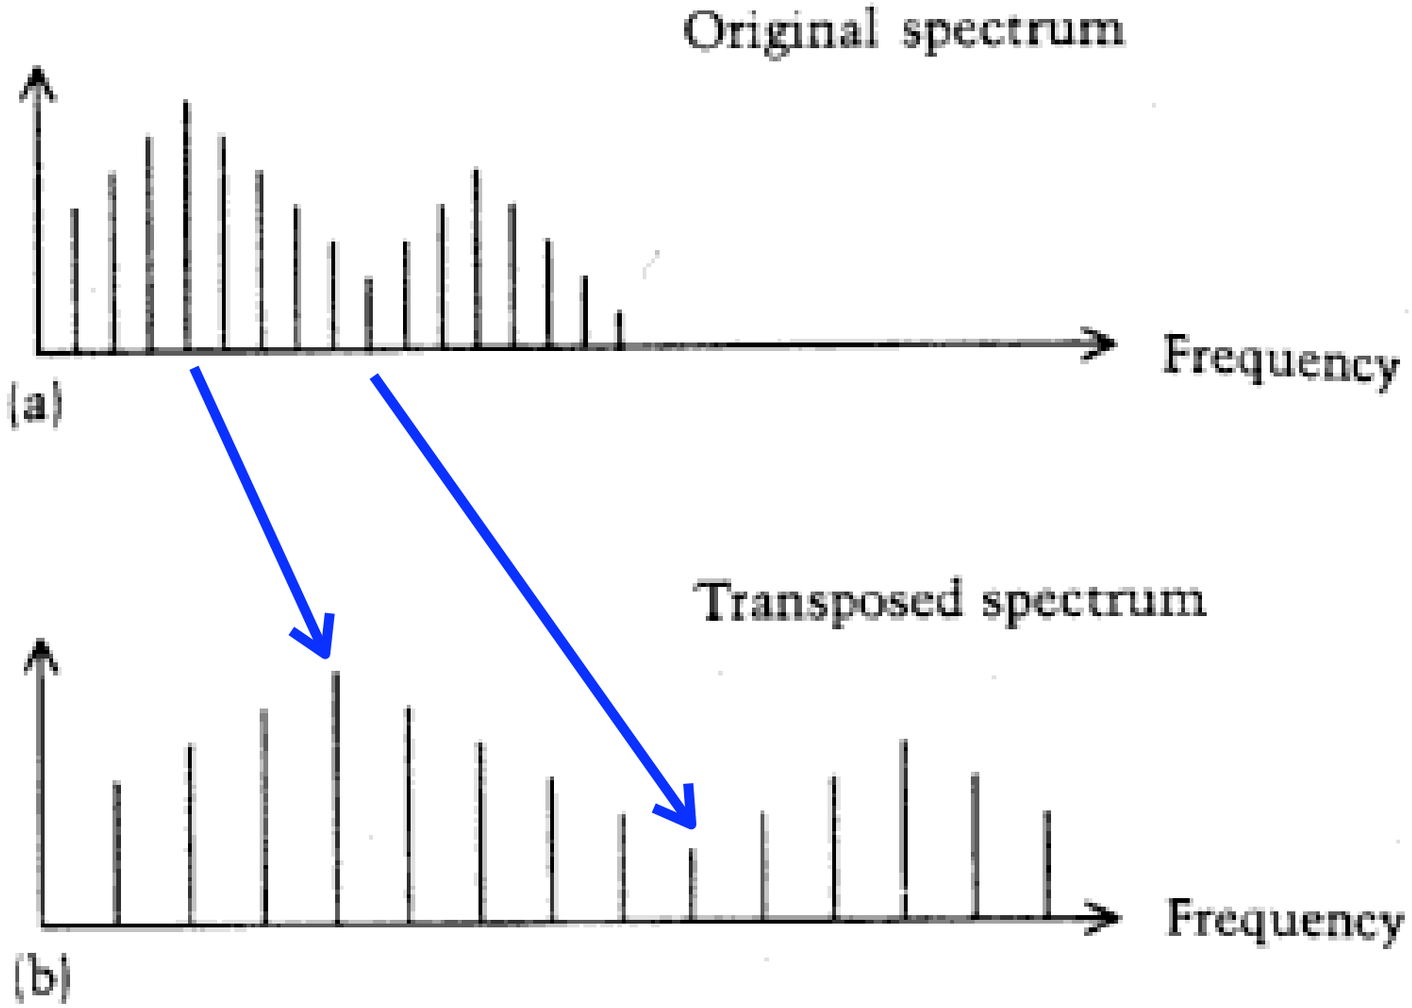

UC Regents Spring 2006 © UCB

Recall: Our pitch shifting definition

✳ **Timbre should sound as it does when playing transposed parts on the instrument.**

For a few instruments, spectral shape scales relative to the pitch, and resampling produces the desired timbre.

Play.

Marimba

UC Regents Spring 2006 © UCB

Voice: Formant control orthogonal to pitch

Cavity sizes affect formants

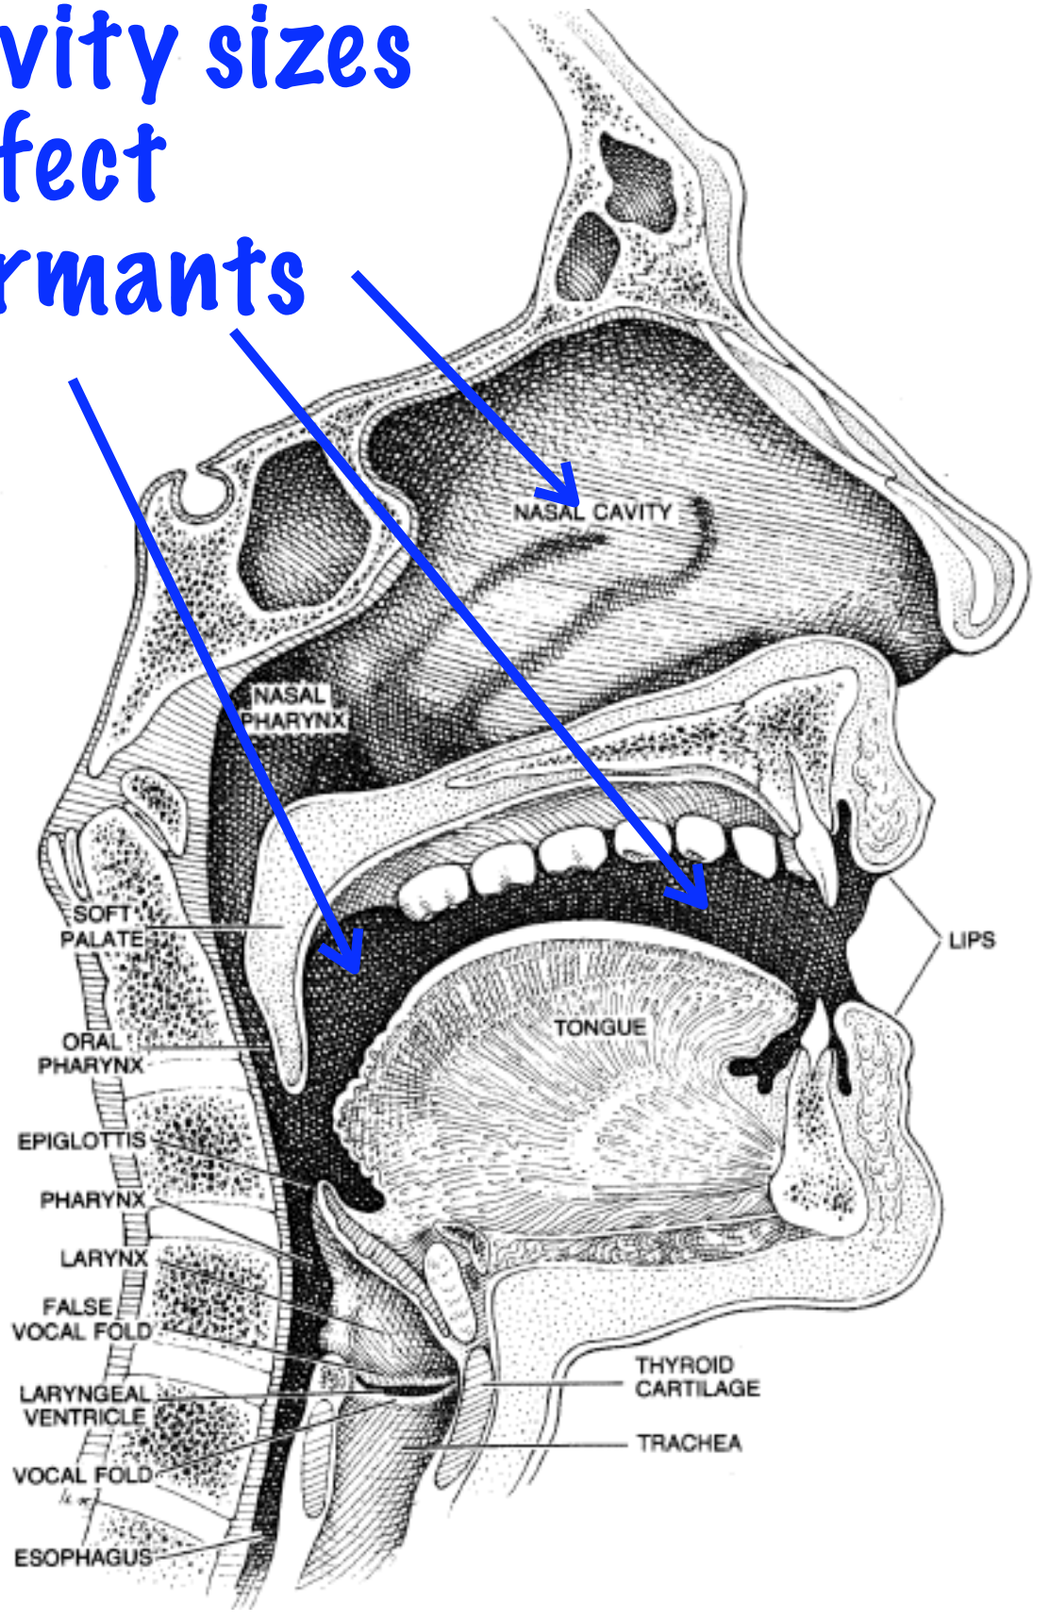

UC Regents Spring 2006 © UCB

Violin Body Resonances

As excited by a bowed bridge.

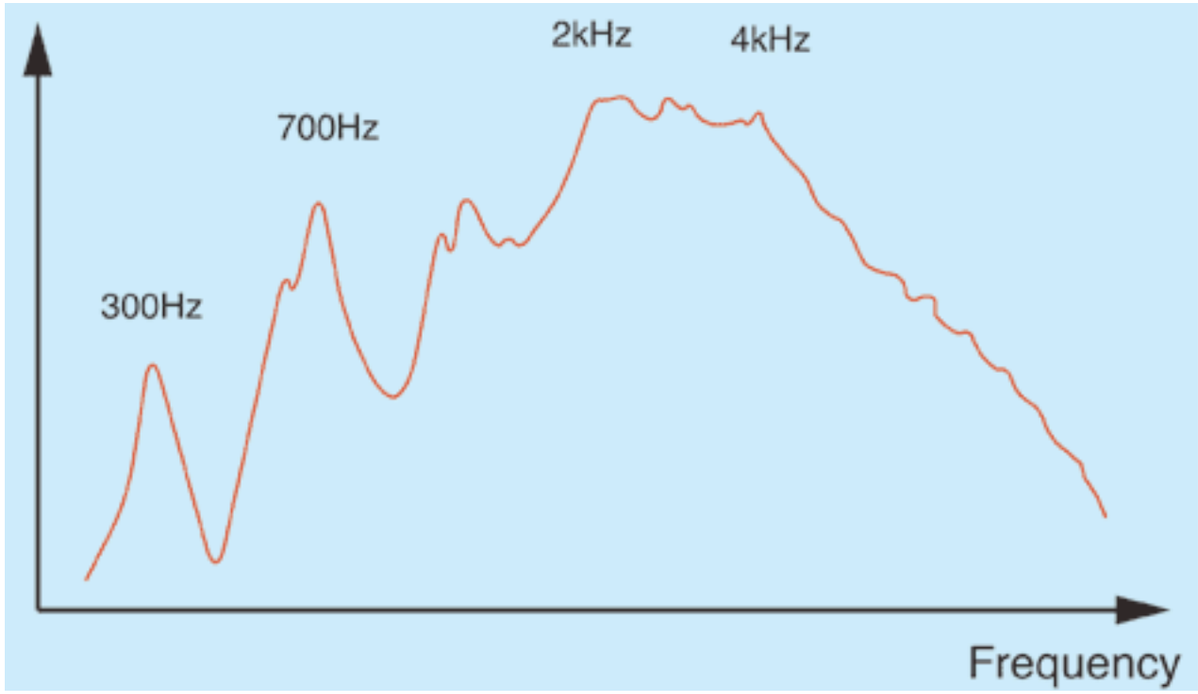

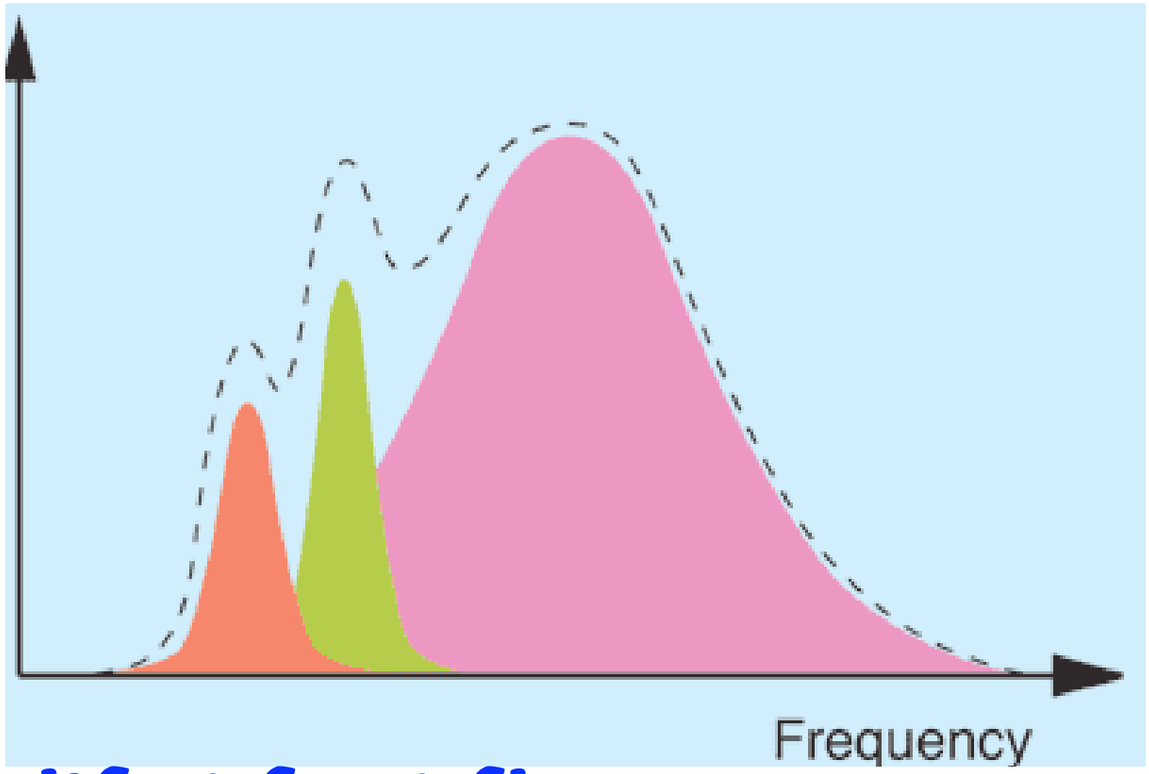

Simplified fixed filters.

UC Regents Spring 2006 © UCB

For fixed formants, resampling a bad match

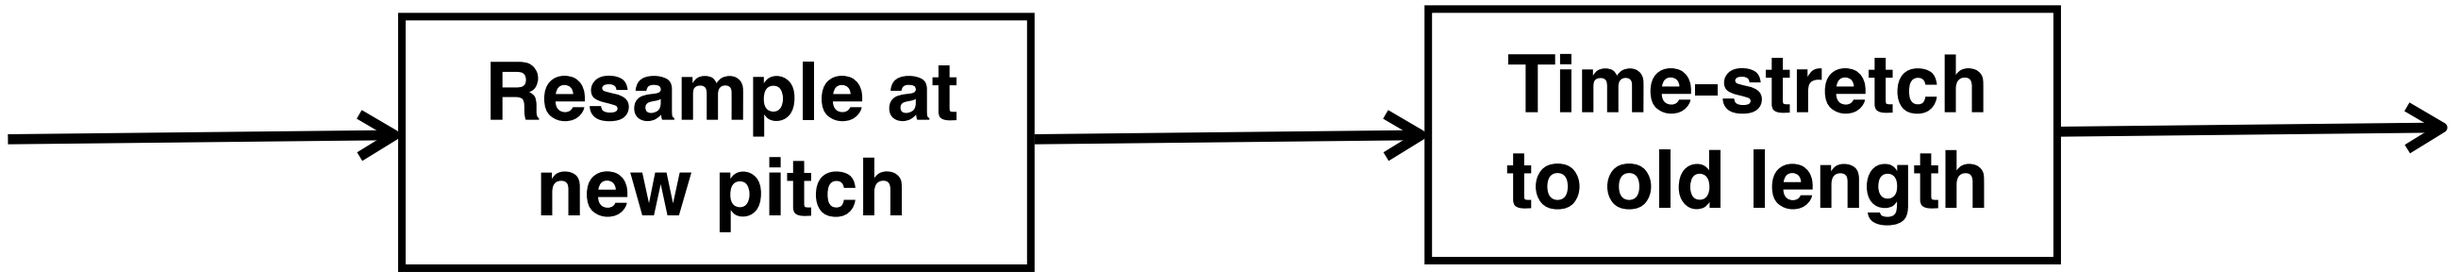

Resampling preserves the magnitude of each partial.

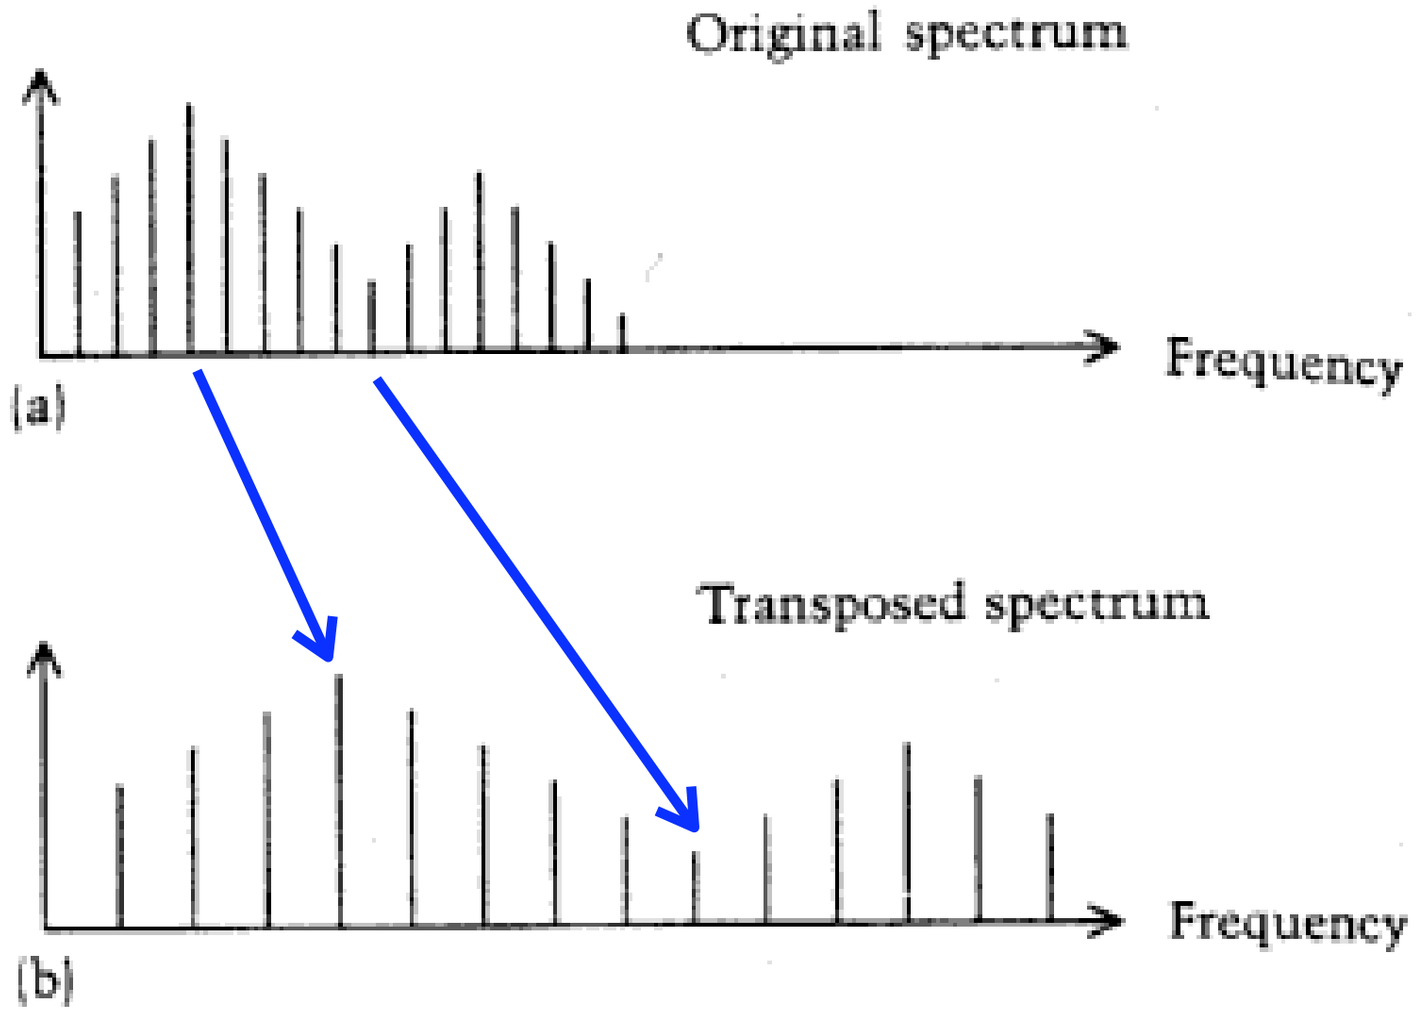

Instead, we want formant frequencies to stay fixed ...

UC Regents Spring 2006 © UCB

Desired behavior for these instruments

Spectral peak and valley positions do not change ... fixes Chipmunk effect on vocals ...

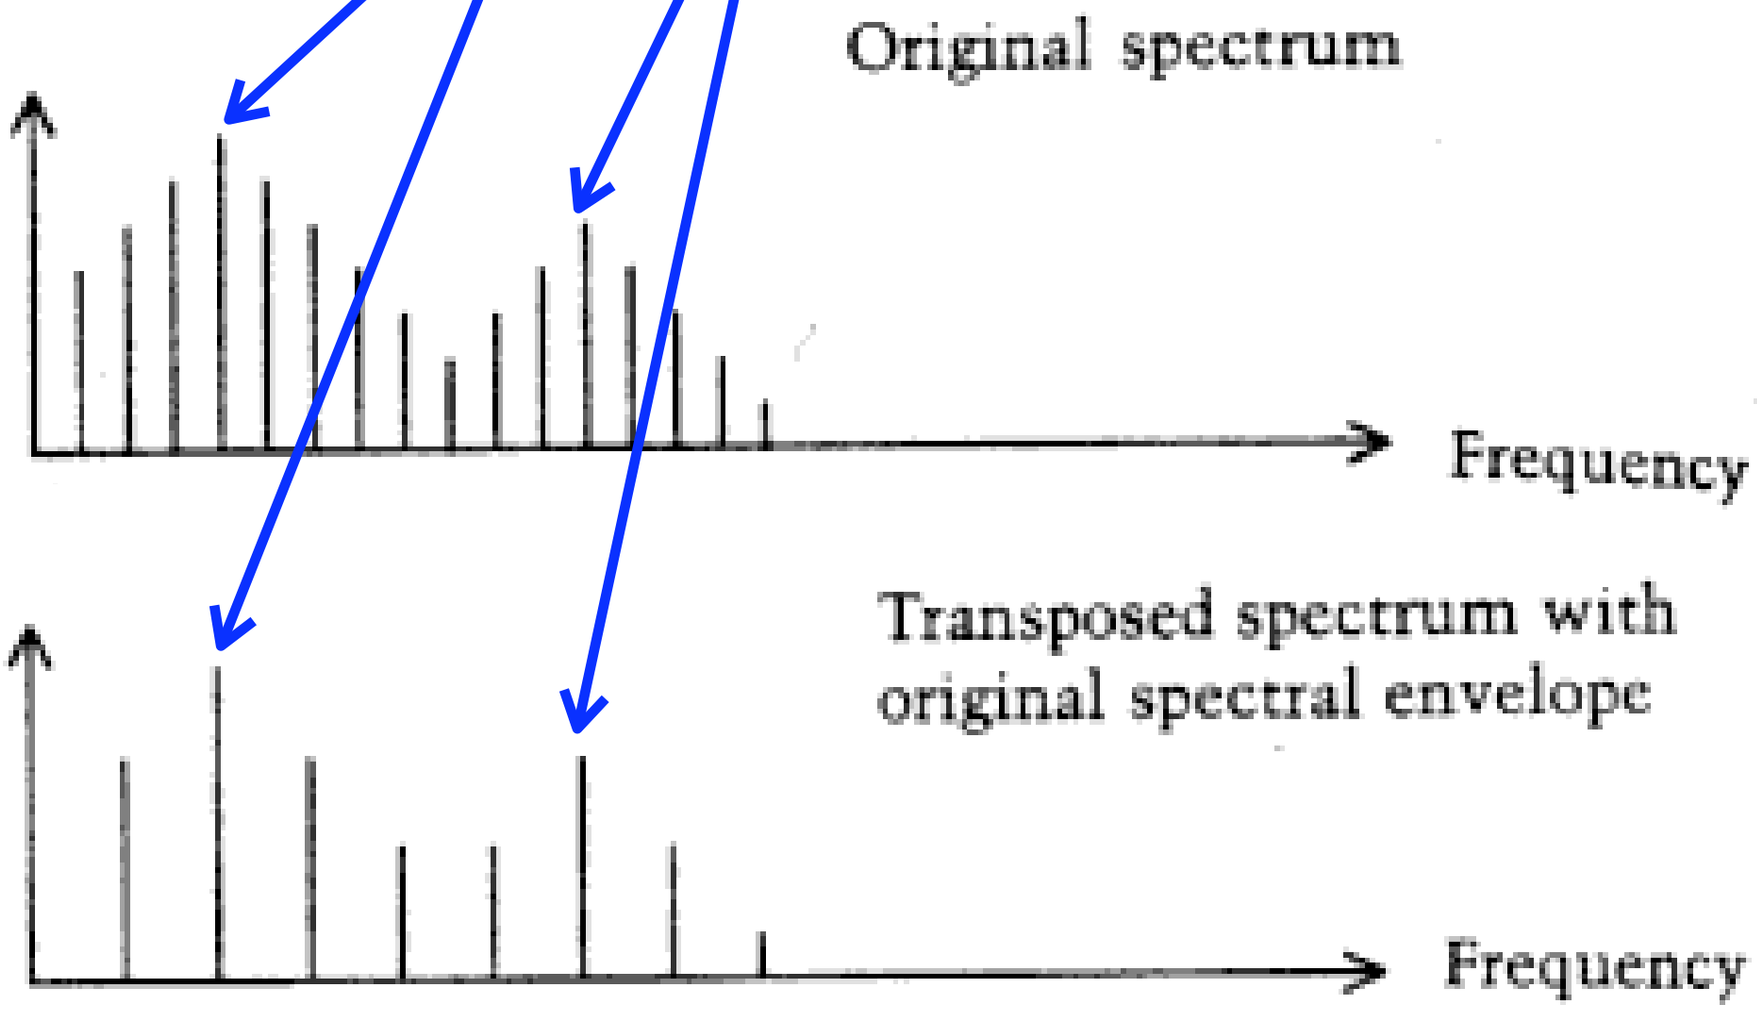

Play.

Still a simplification: Some aspects of timbre scale with pitch and others do not. But it is a big improvement ....

UC Regents Spring 2006 © UCB

Basic idea: Beat-slice pitch periods

To preserve spectral shape, do not change the timebase or waveshape of the pitch period slices.

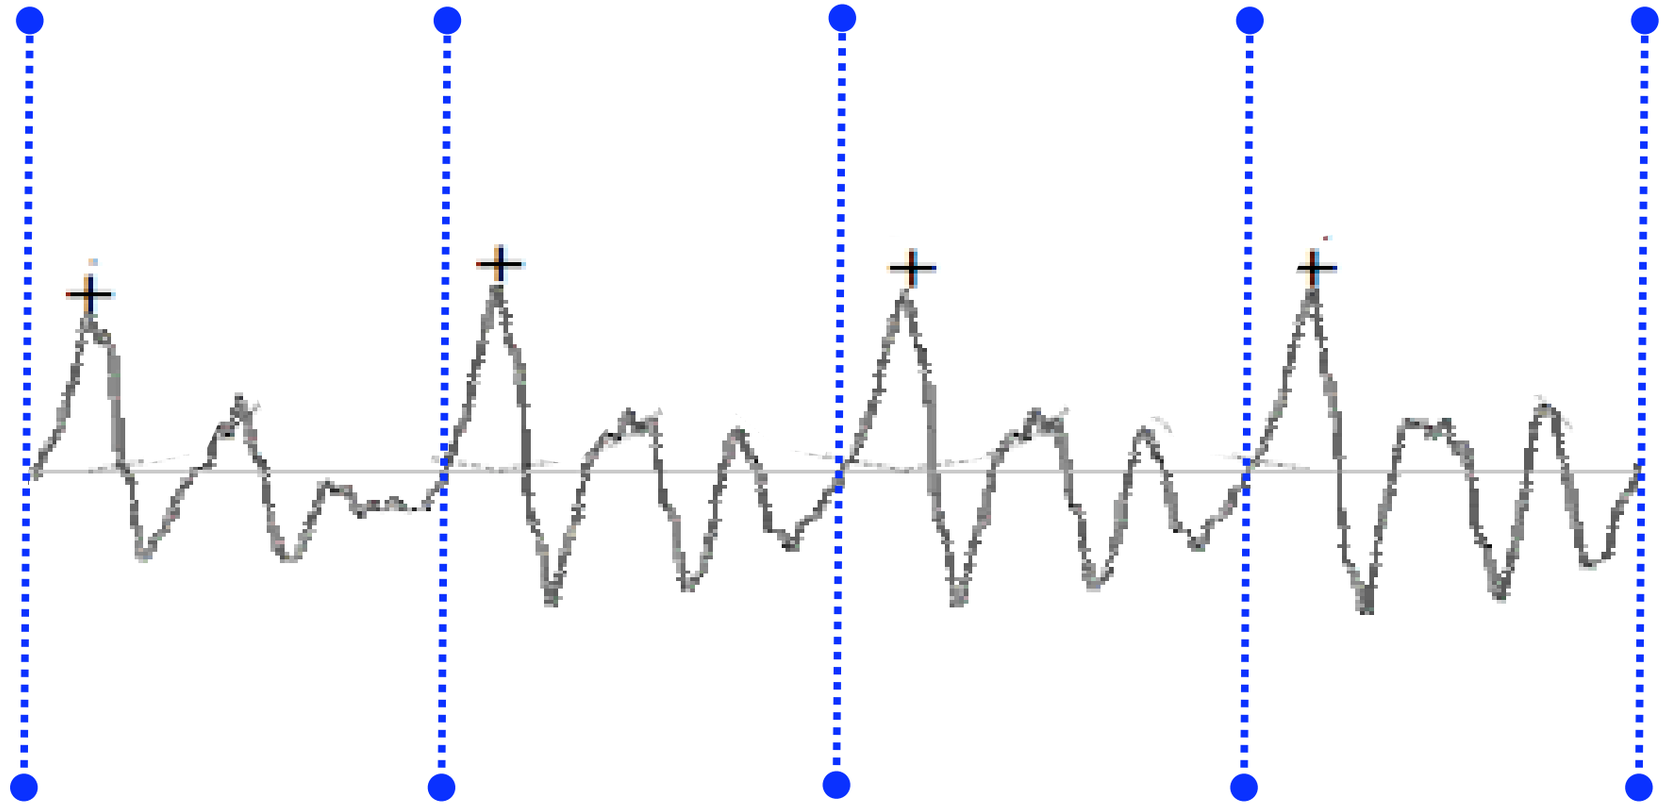

But we need to shift the pitch!

What transformations are permitted?

UC Regents Spring 2006 © UCB

Zero-pad each pitch period ...

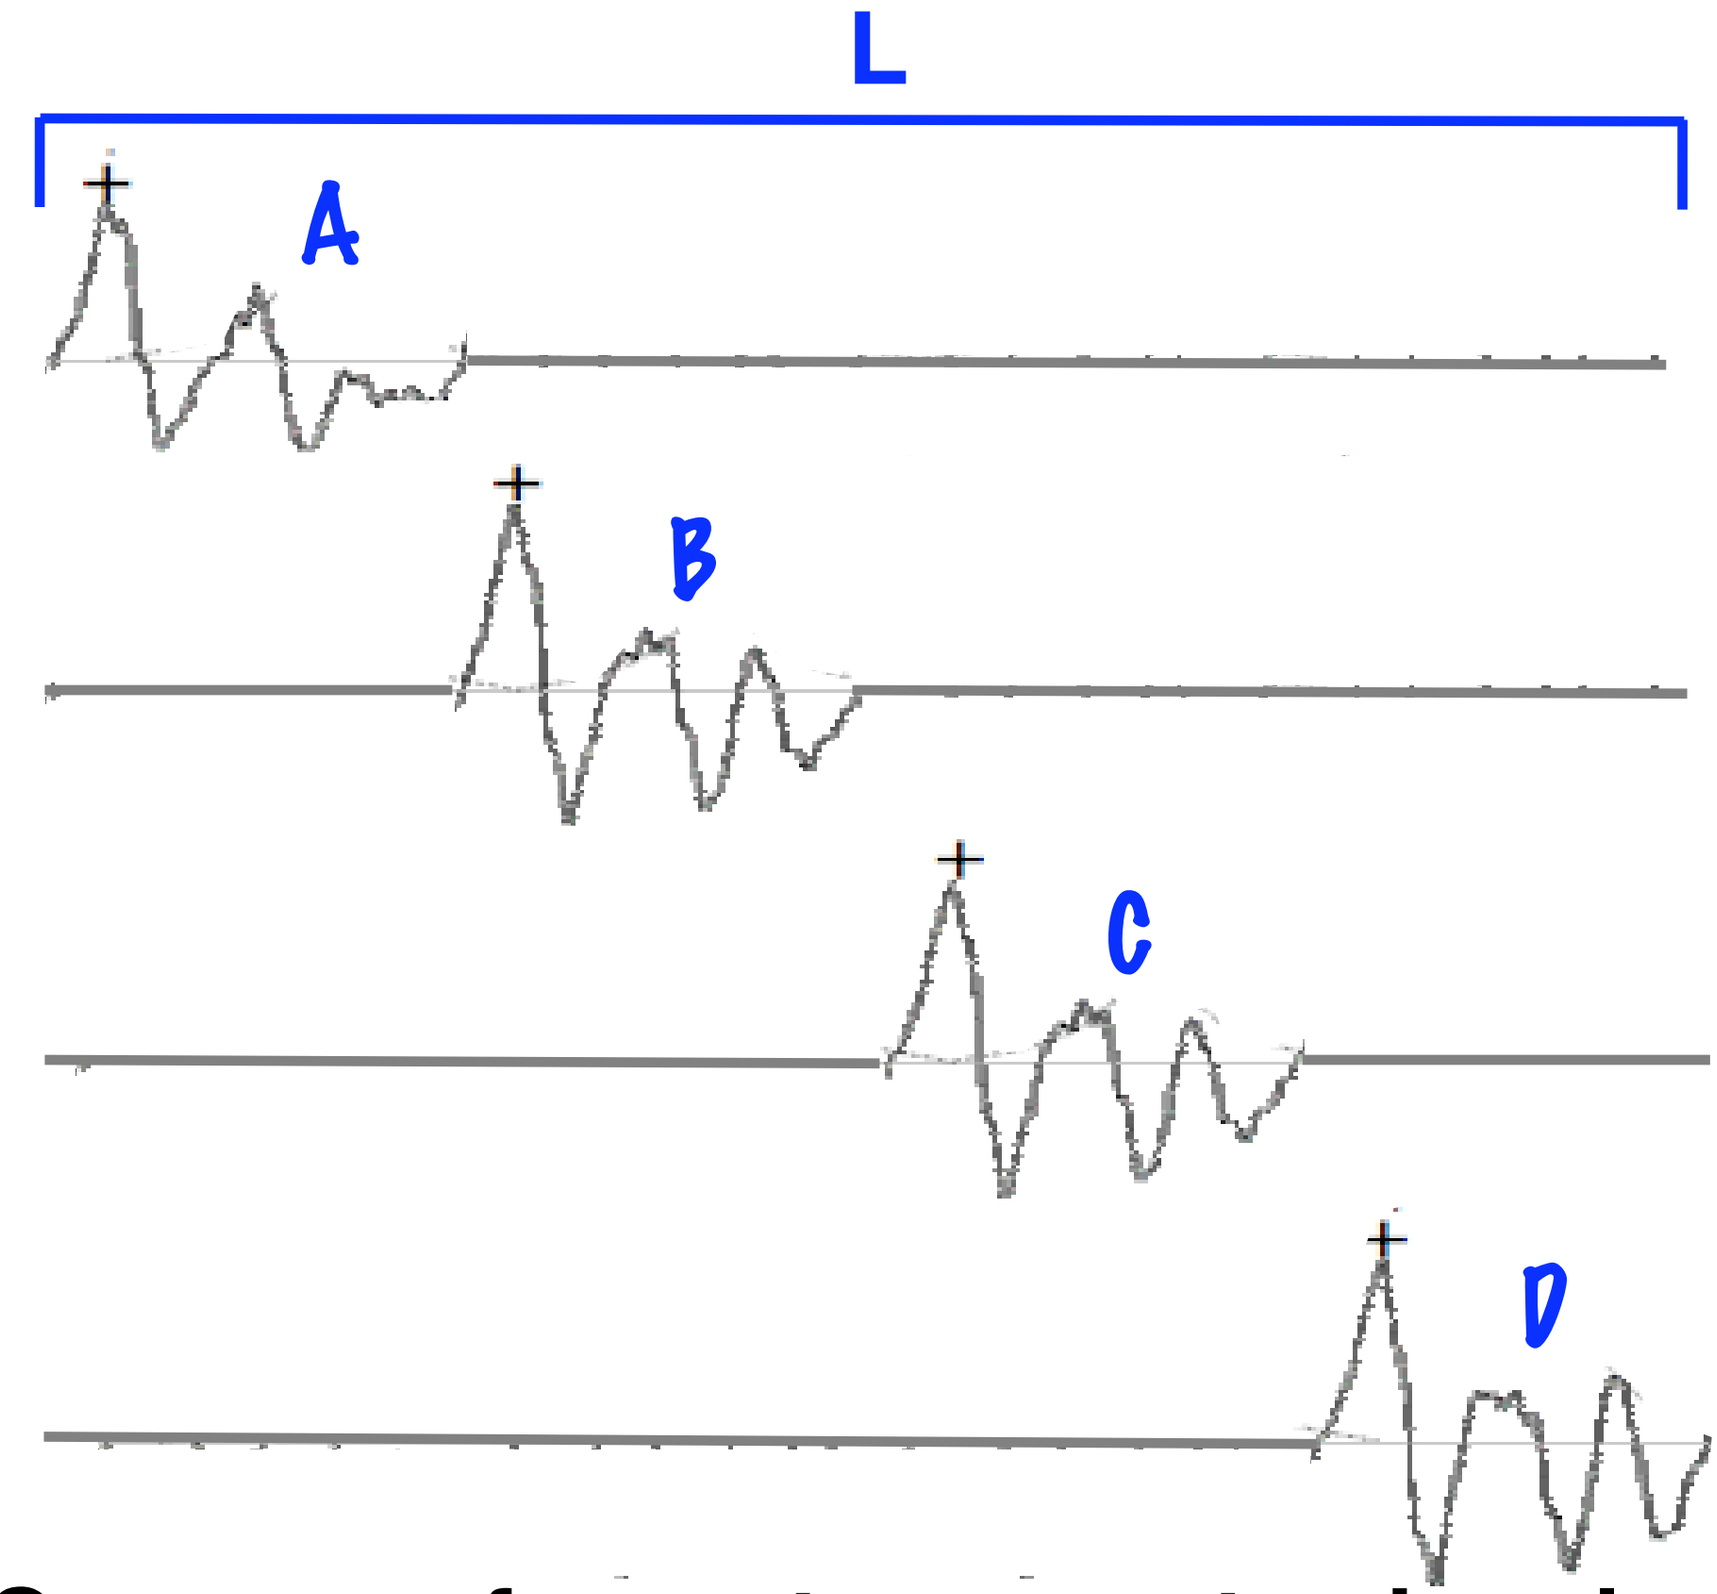

Original local pitch: Four pitch periods per unit "L"

Sum waveforms to recreate signal.

UC Regents Spring 2006 © UCB

Lower Pitch: Delete C and re-space A/B/D

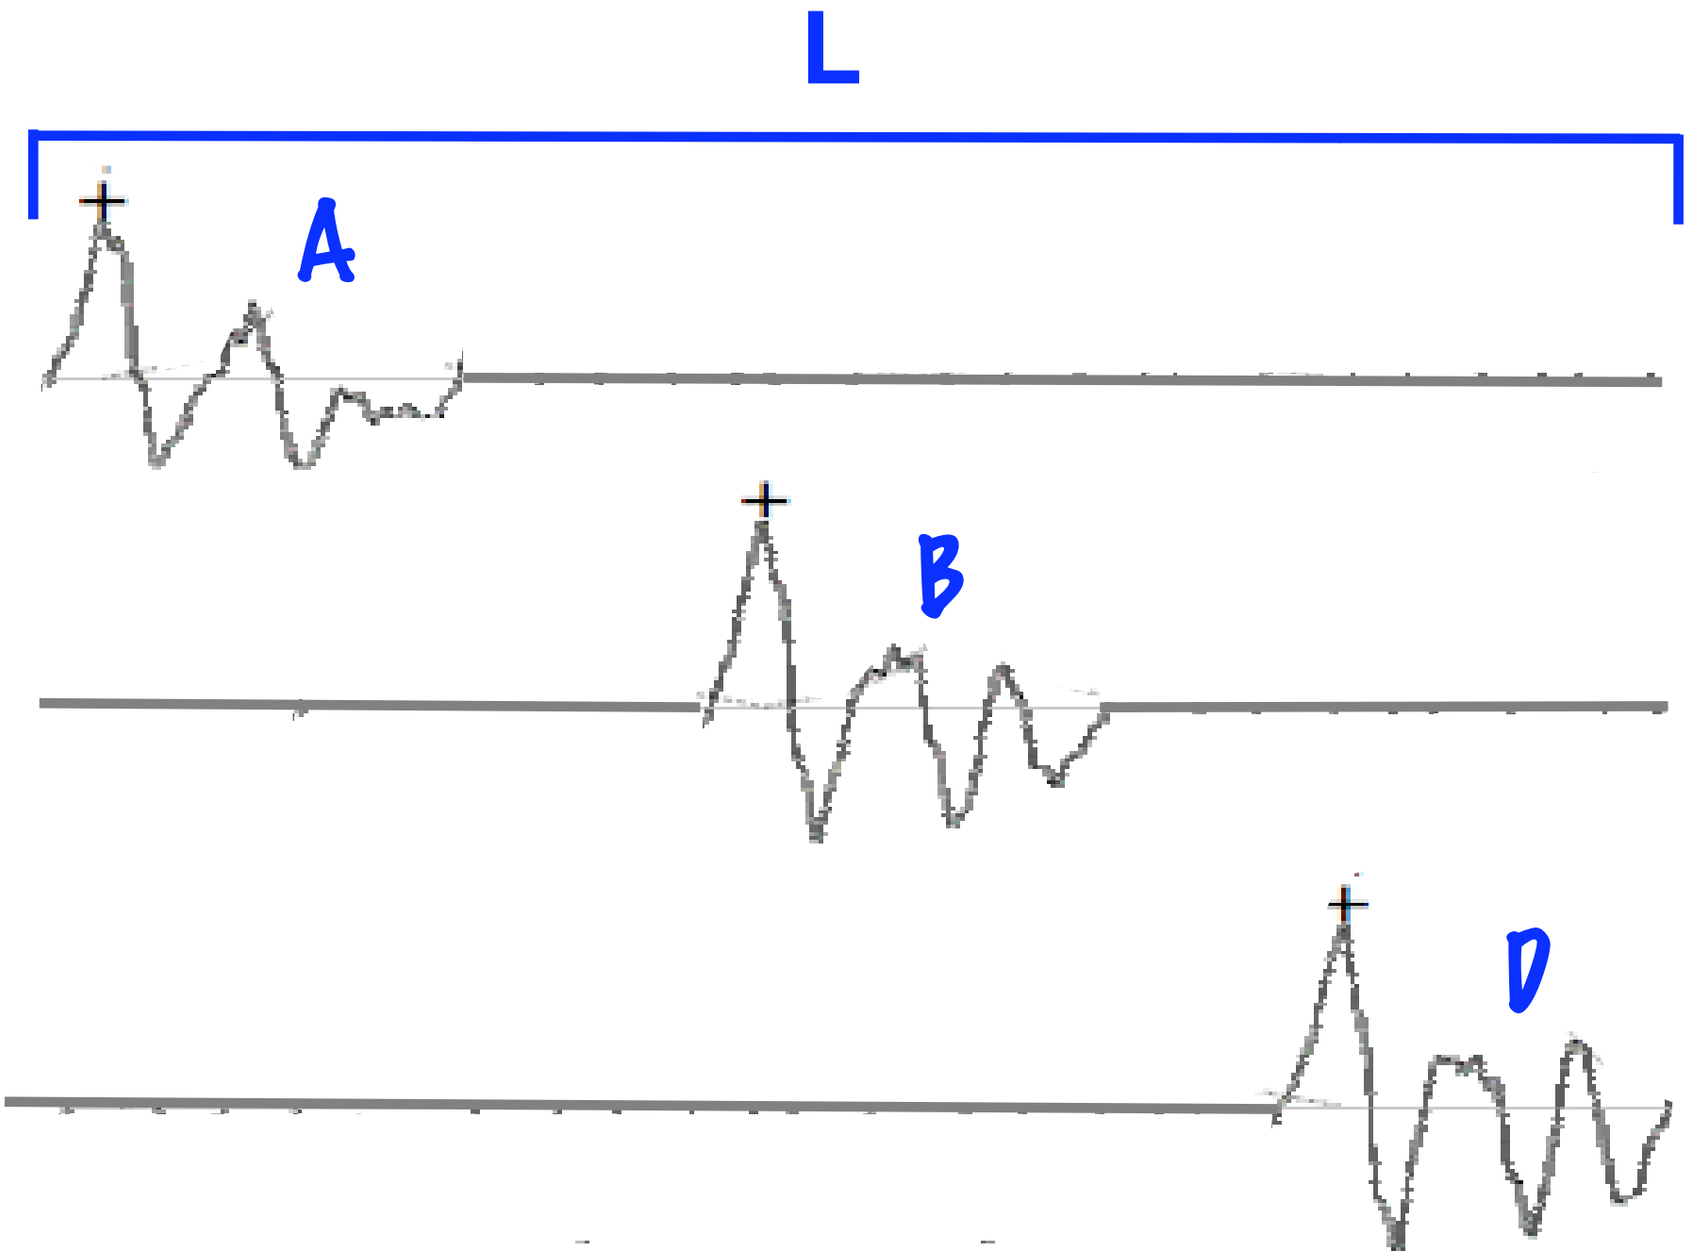

New local pitch: Three pitch periods per unit "L"

Apart from "edge artifacts", spectral shape is not changed by this operation ...

UC Regents Spring 2006 © UCB

Raise pitch: duplicate C and re-space ...

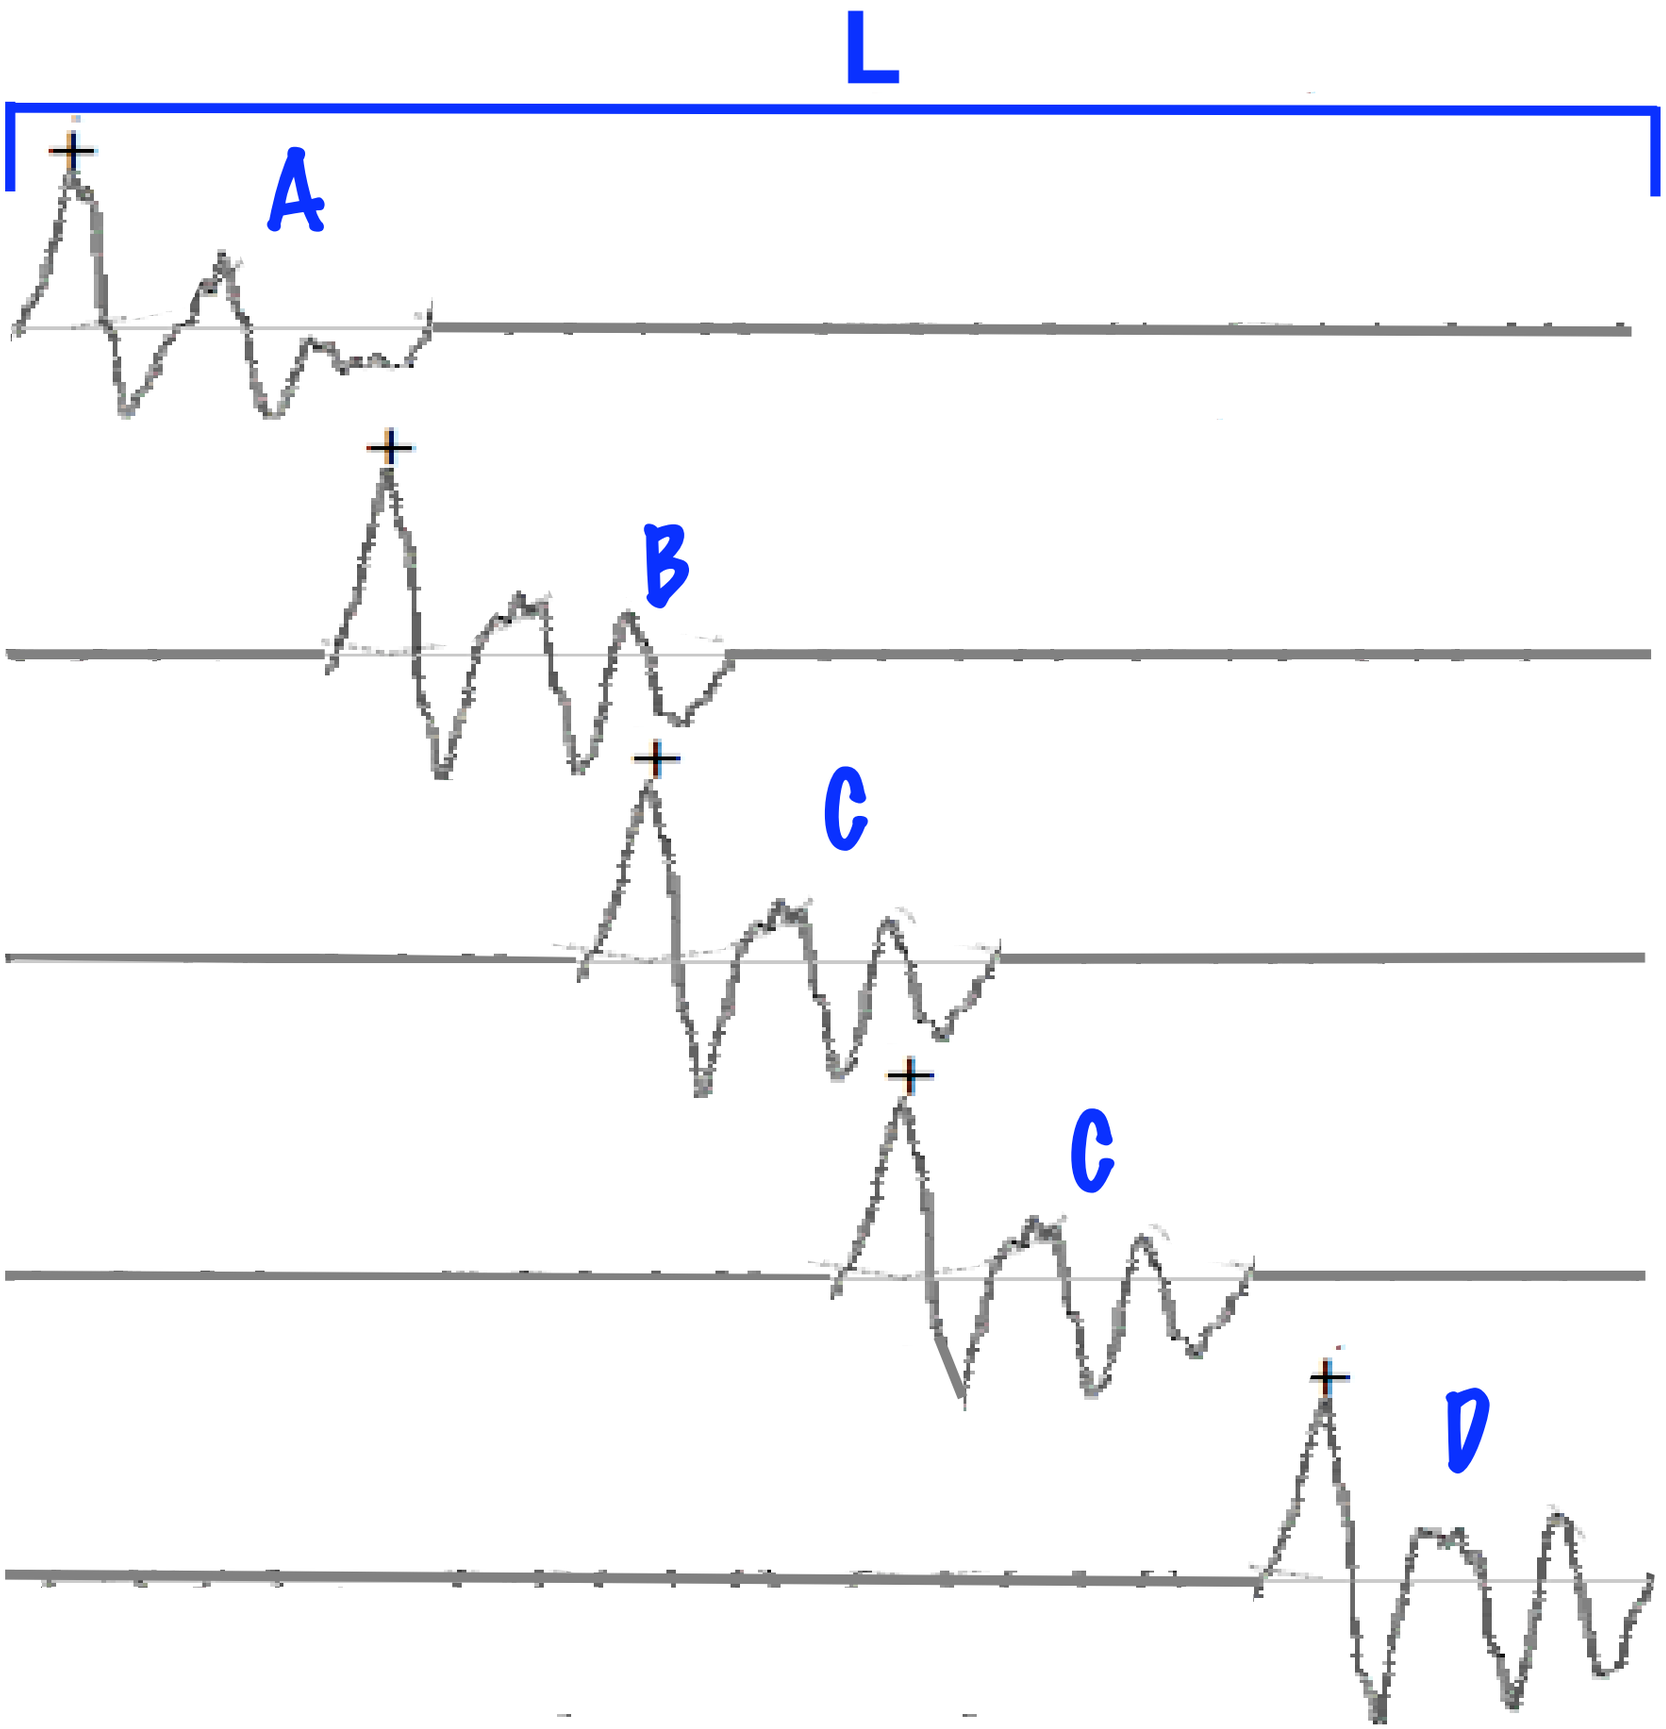

New local pitch: Five pitch periods per unit "L"

UC Regents Spring 2006 © UCB

Why it works ...

To preserve spectral shape, do not change the timebase or waveshape of the pitch period slices.

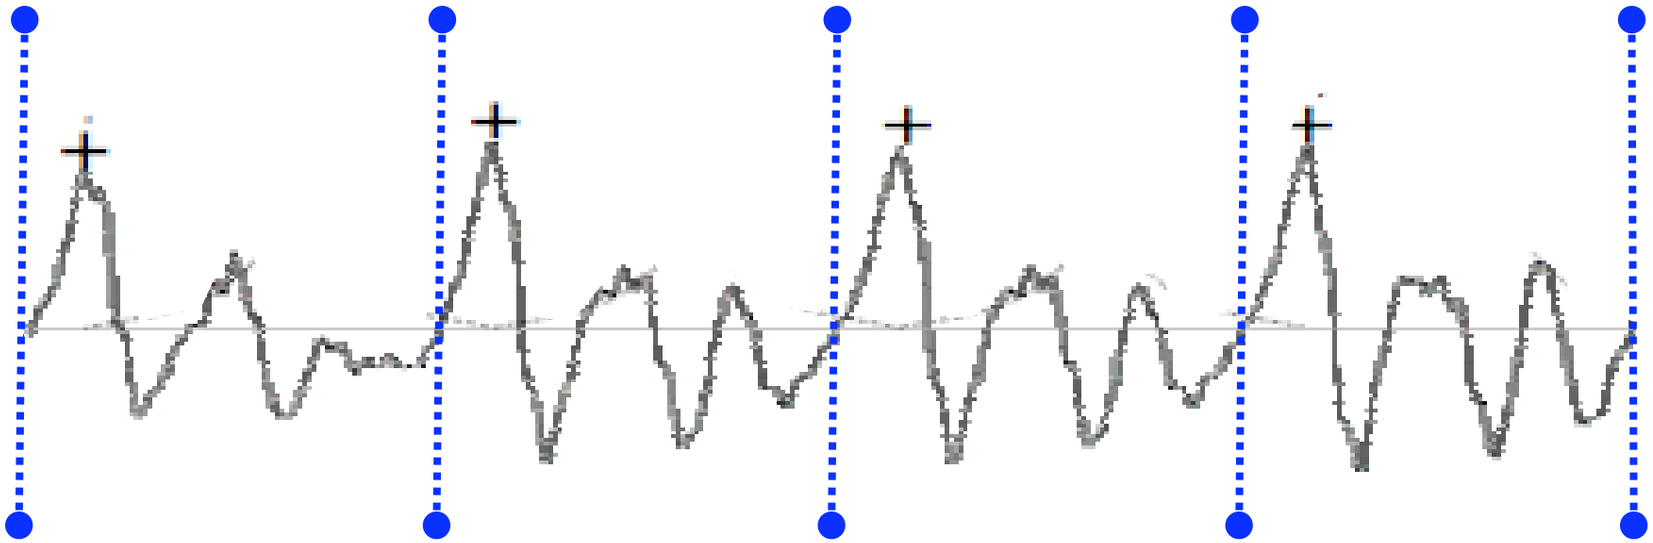

Pitch relies on a repeating pattern: retiling pitch period waveforms keeps repetition.

- **Relies on reliable pitch period detection.**
- **Windowed pitch-period pairs reduce artifacts.**

References (on website): Lent, Bristow-Johnson.

UC Regents Spring 2006 © UCB

Can resampling preserve formants?

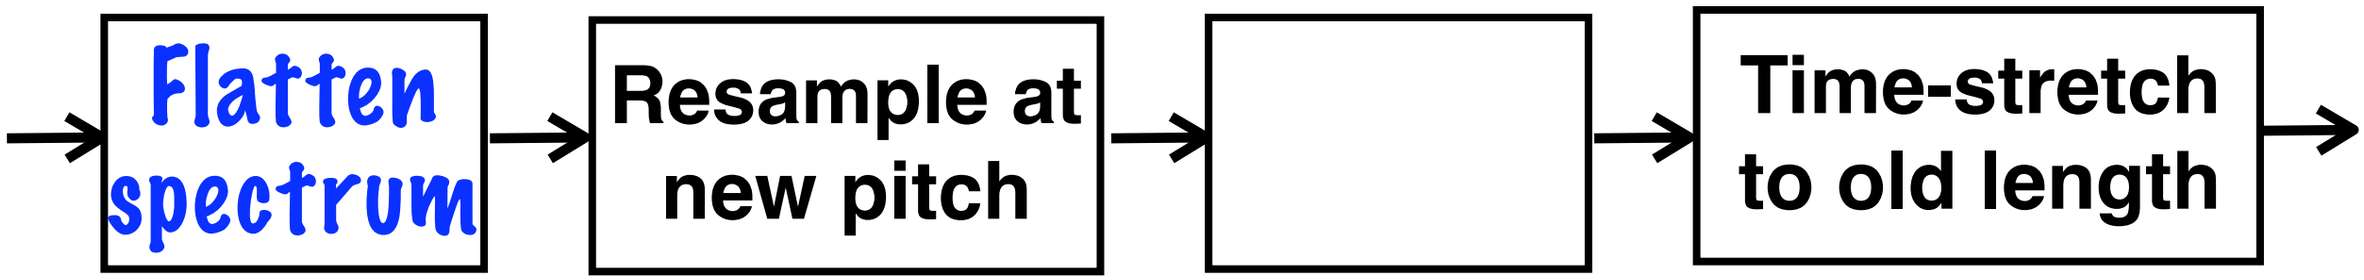

Yes.

Begin by flattening the spectrum for a short waveform block.

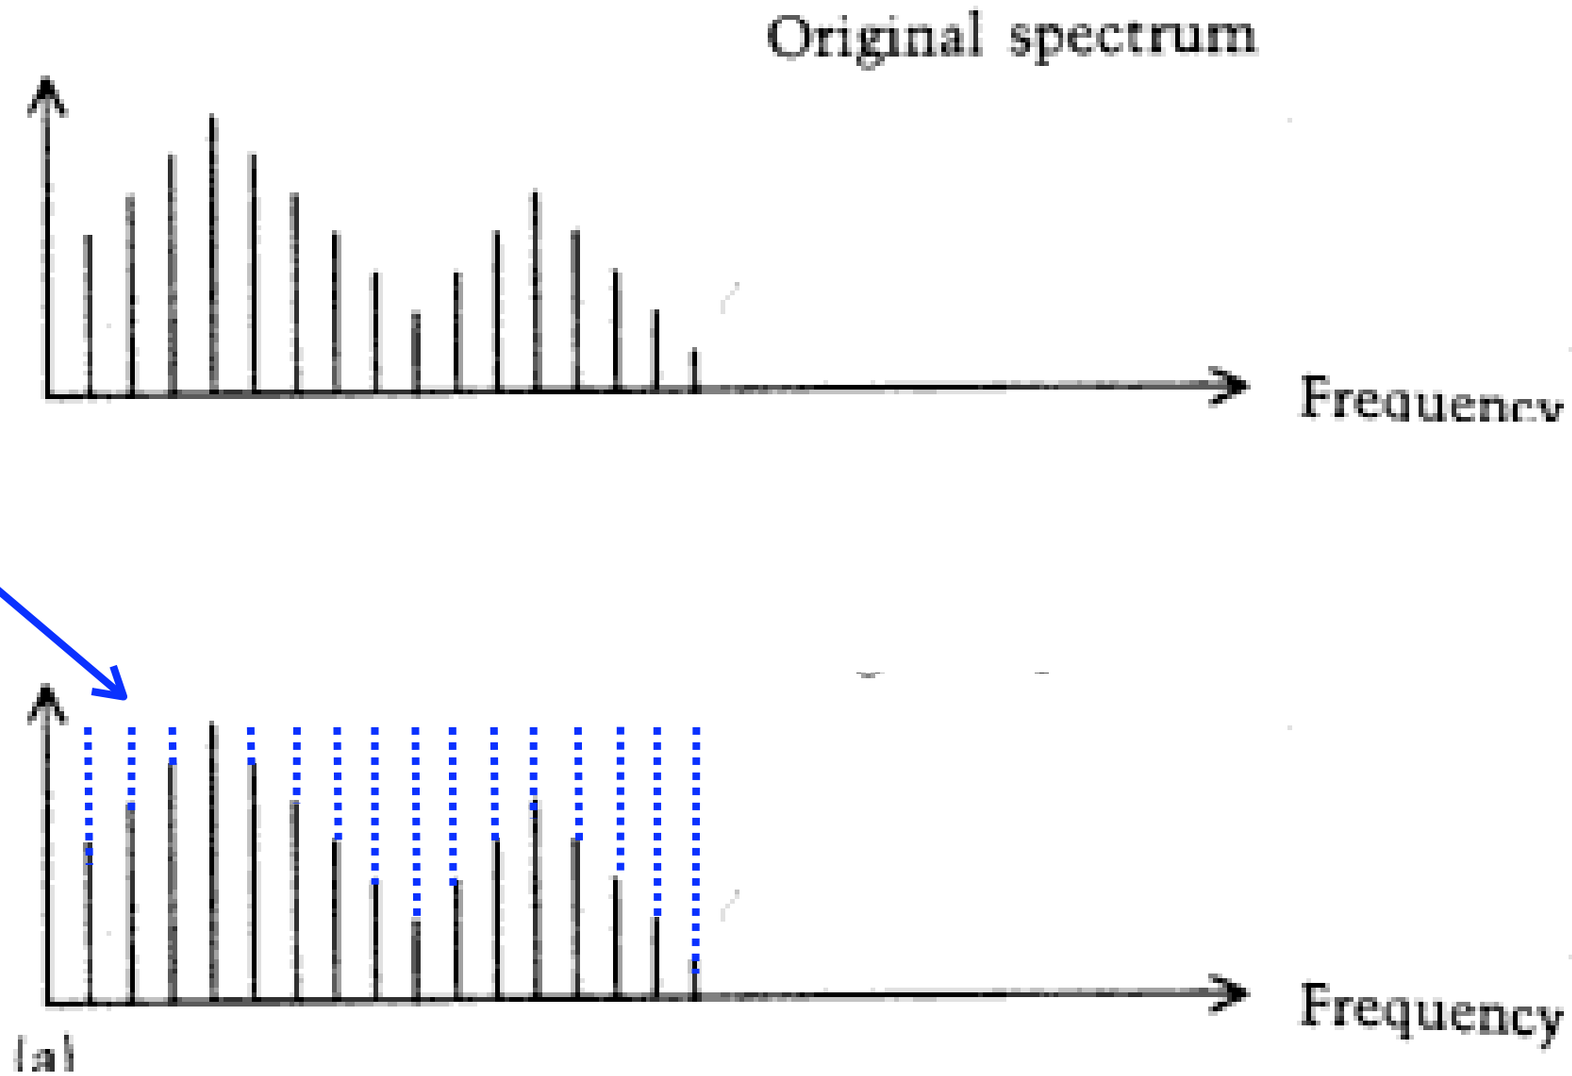

UC Regents Spring 2006 © UCB

Then ...

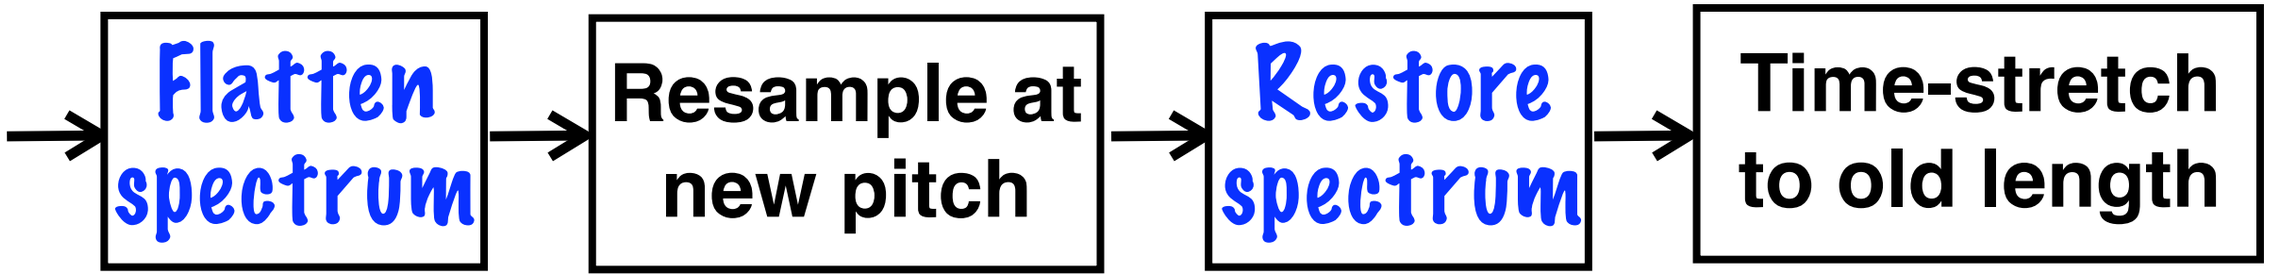

Spectral modification ??? Several options for algorithms (LPC, Fourier, etc).

Restore original spectral envelope.

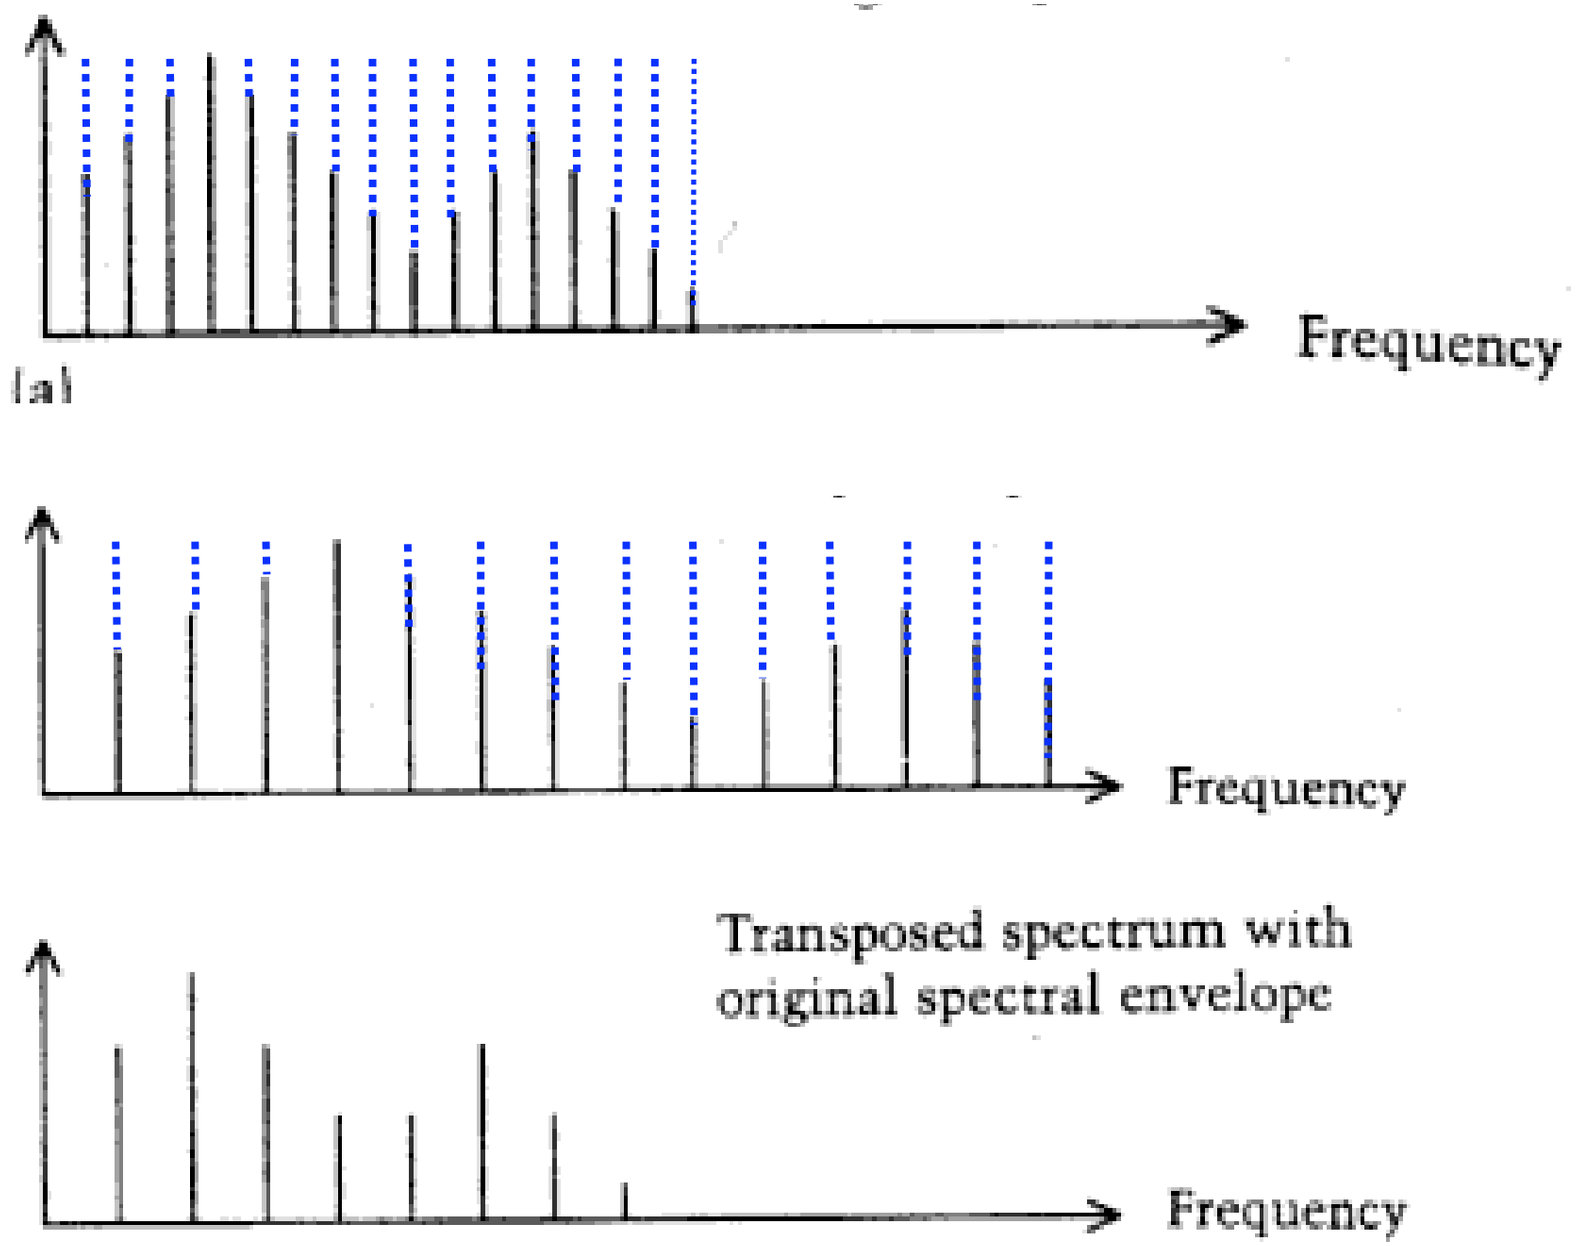

UC Regents Spring 2006 © UCB

Projects Ideas

UC Regents Spring 2006 © UCB

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

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).

- Drum-related Projects
  - Creating Electronic Drum Samples from Acoustic Drum Samples
  - Tools for Automating Drum Track Arrangements
  - Timbre-Space Browsers for Drum Loops and Individual Hits
  - Realistic Retuning of Drum Sounds
  - Real-Time Performance by Retiming Drum Loops
  - Fusing Multiple Drum Hits into a Single Percept
- Wind Instrument Projects
  - Playing Horns from a Keyboard with Improved Articulation
  - Automatic Horn Phase Selection to Match a Track
  - Real-time Timbre Selection with a Wind Controller
- Computer Systems Projects
  - CoreSample: Kernel Database Services for Concatenative Synthesis
- Vocal Projects
  - Synthesis, Analysis, and Algorithmic Composition of Glossolalia Vocals
  - Lyric Design for Phrase-Based Vocal Synthesis

Voice Project Idea #1

Glossolalia Vocals

UC Regents Spring 2006 © UCB

Recall: Speech Singing Synthesis

- **Beauty is easier than intelligibility: foreign language singing sounds more pleasant than an understood language.**

- **Using it with a real-time controller has big obstacles: language-oriented algorithms require lookahead to work well.**

UC Regents Spring 2006 © UCB

Glossolalia Singing Synthesis ...

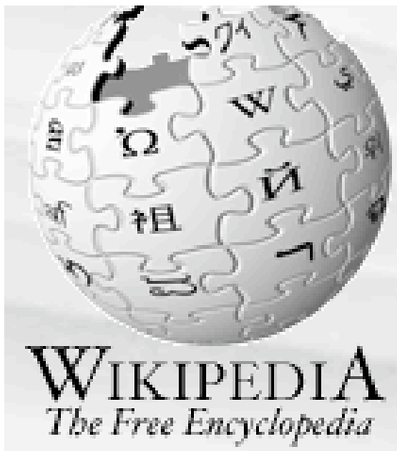

navigation

- Main Page
- Community Portal
- Current events
- Recent changes
- Random article
- Help
- Contact us
- Donations

search

Go Search

toolbox

- What links here
- Related changes
- Upload file
- Special pages
- Printable version
- Permanent link

Sign in / create account

article | discussion | edit this page | history

Your continued donations help Wikipedia grow!

Glossolalia

From Wikipedia, the free encyclopedia

Glossolalia (from the Greek, "γλώσσα" (glossa), tongue and "λαλώ" (lalô), to speak) comprises the utterance of what appears (to the casual listener) either as an unknown foreign language (**xenoglossia**), simply nonsense syllables, or utterance of an unknown mystical language; the utterances sometimes occur as part of religious worship (**religious glossolalia**).

Skeptics dismiss these cases as simply being in a state of trance, self-hypnotism or religious ecstasy. It is notable that in Charismatic/Pentecostal Churches there is often a state of heightened emotionalism which may, in the view of skeptics, itself give rise to instances of glossolalia (what Christians in those churches often refer to as **feeling the spirit**).

Contents [hide]

Play

UC Regents Spring 2006 © UCB

A good match to concatenation ...

- **We can design the language with phonemic transitions that sound good.**
- **There are no native listeners, so no one will hear marginal transitions as synthetic.**
- **If we let lyrics be generated algorithmically, playing the voice from a MIDI controller becomes possible.**

UC Regents Spring 2006 © UCB

Two ways artists approach glossolalia

- **Scientifically.** (example: Elizabeth Frazier, of the Cocteau Twins). A linguist, she designs syntax and semantics for a novel language, then writes lyrics in it.

  Project idea: computer tools to help the design process, perhaps with the goal of making concatenative singing synthesis sound good. in the language (Adrian Freed's idea).

- **Improvisationally.** (example: Lisa Gerrard, of Dead Can Dance).

  Project idea: Sample her a cappella Glossolalia singing, and use it in a concatenative system.

UC Regents Spring 2006 © UCB

Recall: Construct database of complete musical phrases that are browsed via GUI (example: Liquid Saxophone).

Phrase-Based Synthesis

Main Problem: Choosing lyrics that would be useful ....

UC Regents Spring 2006 © UCB

Children's choir: $375. Sold out first run quickly.

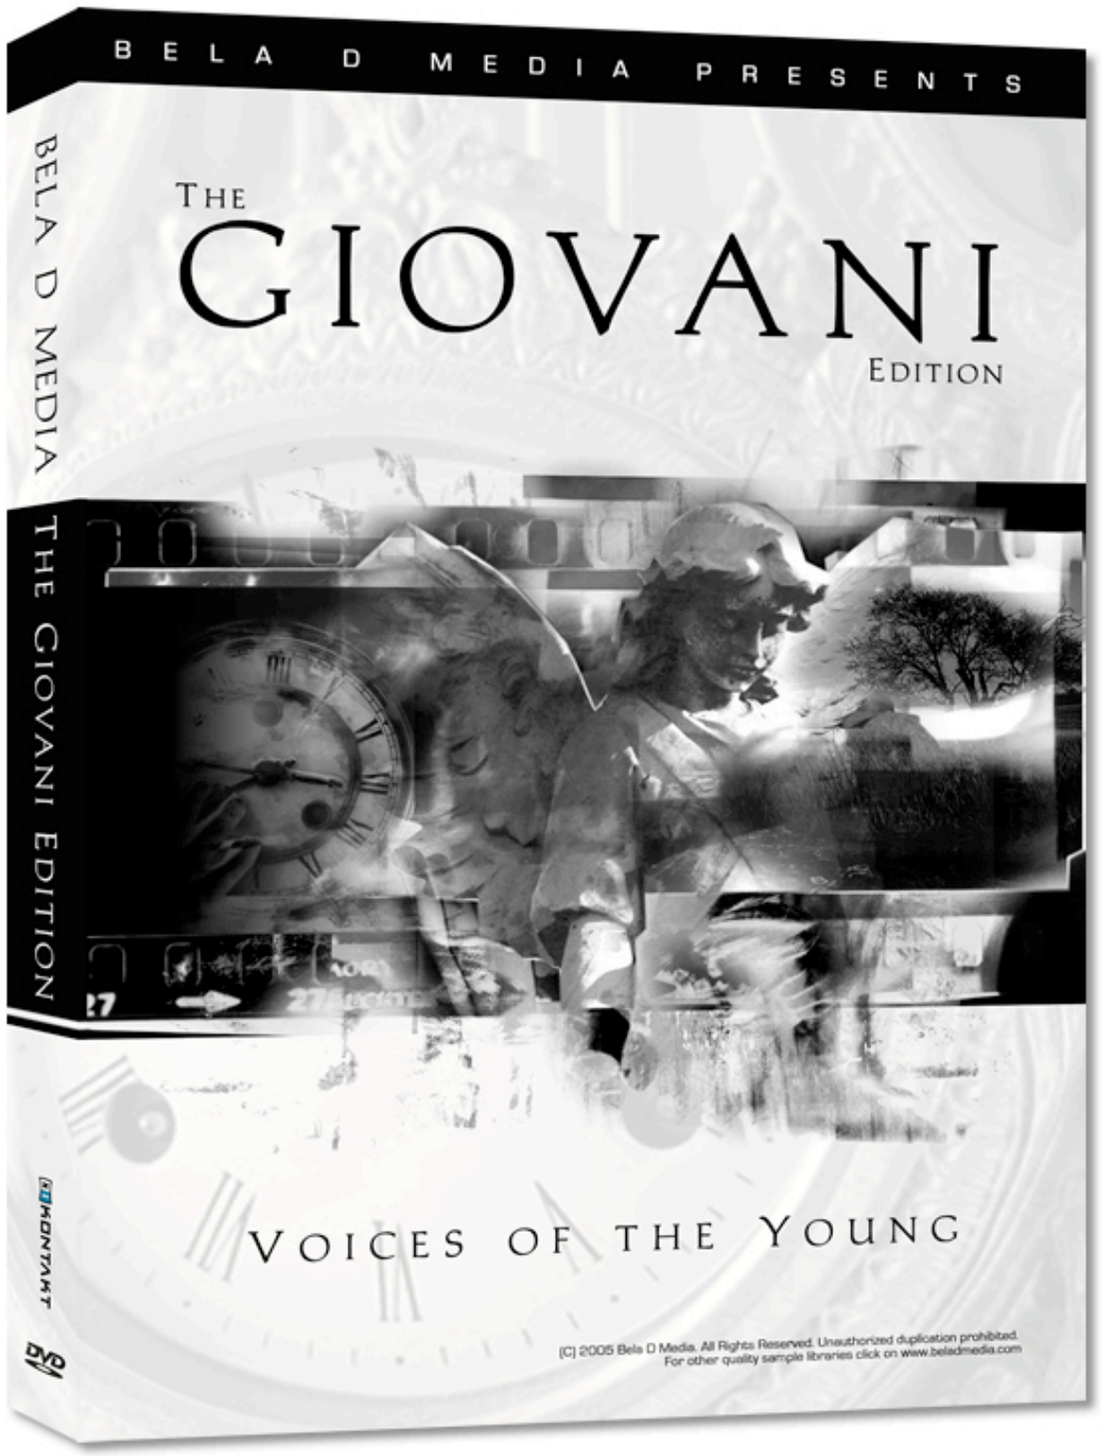

Sampled Latin

- Agnus Dei
- Benedictus
- Dies Irae
- Veritas Domini
- Morte Aeterna
- Peccata Mundi
- Requiem Aeternam

Play

UC Regents Spring 2006 © UCB

Rudimentary phrase concatenation ...

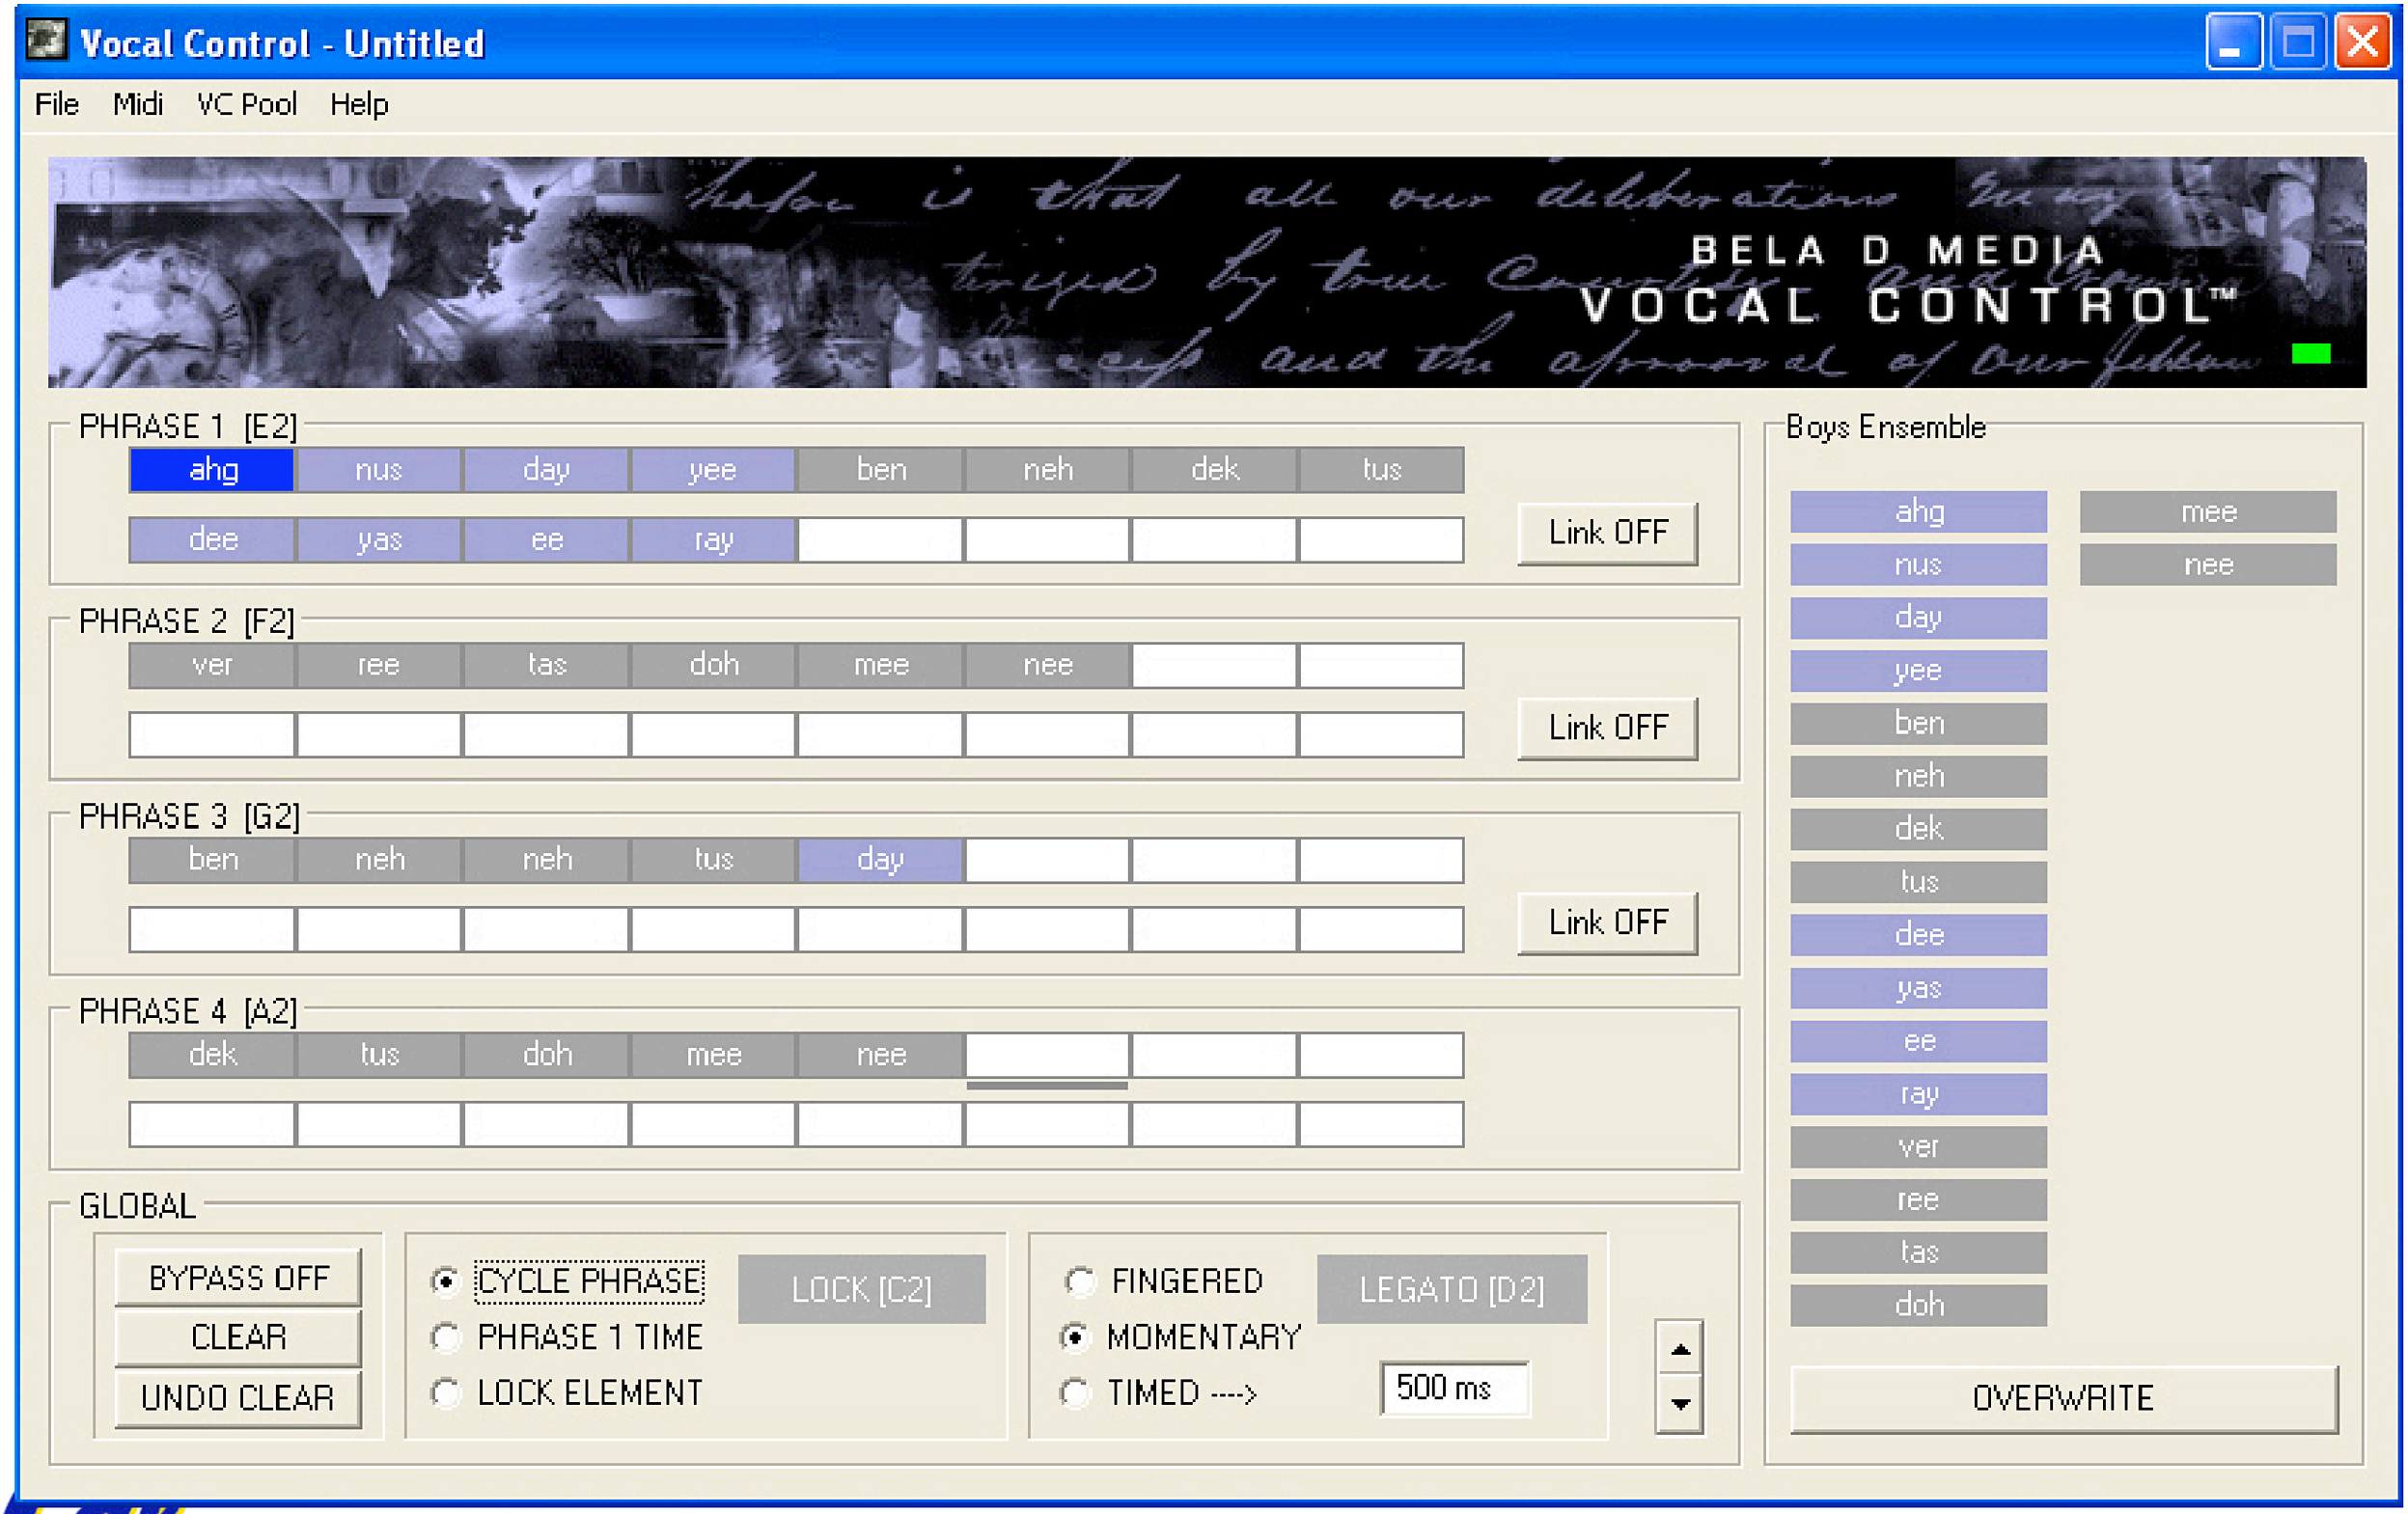

UC Regents Spring 2006 © UCB

Voice Project Idea #2

UC Regents Spring 2006 © UCB

Phrase approach tricky for pop music ...

The Voice Vol. 1 features 300 verbal vocal phrases between 2 and 8 bars focused mainly on pop, dance and RnB productions. All vocal phrases can be combined with each other.

The verbal phrases include: "listen 2 the groove", "keep me movin on", "liftin me higher", "party everybody", "ready 4 my luv", "u make me wanna dance", "universal love", "feel so high", "sexy dancer", "when will u stop playing" and many more.

UC Regents Spring 2006 © UCB

There has to be a better way ...

The verbal phrases include: "listen 2 the groove", "keep me movin on", "liftin me higher", "party everybody", "ready 4 my luv", "u make me wanna dance", "universal love", "feel so high", "sexy dancer", "when will u stop playing" and many more.

Project idea: Come up with a principled idea for creating a useful phrase library (words and melody + signal processing) that is data driven from lyric and MIDI databases on the web.

UC Regents Spring 2006 © UCB