CMPT 365 Multimedia Systems

Media Representations
- Audio

Spring 2017

Outline

- Audio Signals
  - Sampling
  - Quantization
- Audio file format
  - WAV/MIDI
- Human auditory system

What is Sound ?

- Sound is a wave phenomenon, involving molecules of air being compressed and expanded under the action of some physical device.
  - A speaker (or other sound generator) vibrates back and forth and produces a *longitudinal* pressure wave that perceived as sound.
  - Since sound is a pressure wave, it takes on continuous values, as opposed to digitized ones.
    - If we wish to use a digital version of sound waves, we must form digitized representations of audio information.
    - Link to physical description of sound waves

Sound Recording and Reproducing

- Thomas Edison's Phonograph 1877
  - first device to record and reproduce sound
  - Medium: a tinfoil sheet phonograph cylinder.
- Alexander Graham Bell's improvement in 1880s

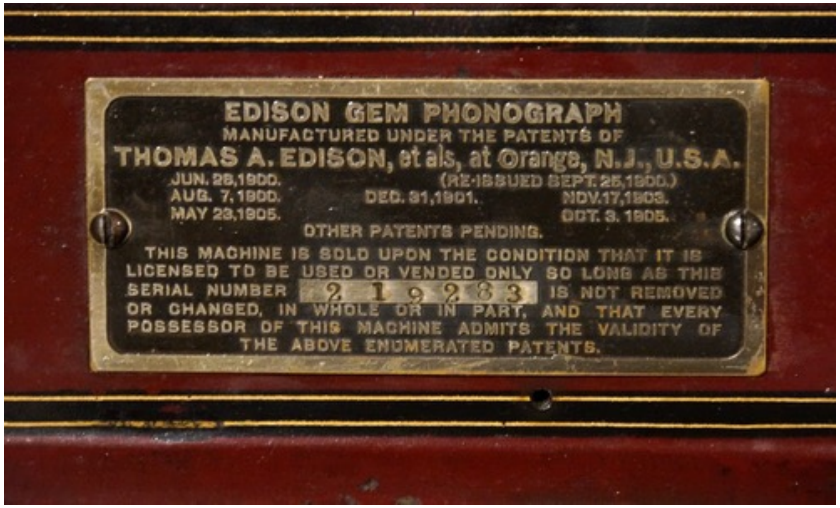

Sound Recording and Reproducing

- Thomas Edison's Phonograph 1877
  - first device to record and reproduce sound
  - Medium: a tinfoil sheet phonograph cylinder.
- Alexander Graham Bell's improvement in 1880s
- Emile Berliner’s gramophone
  - double-sided discs
- Audio tapes, and later Compact Disc (CD)

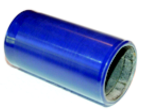

Physical World is often Analog !

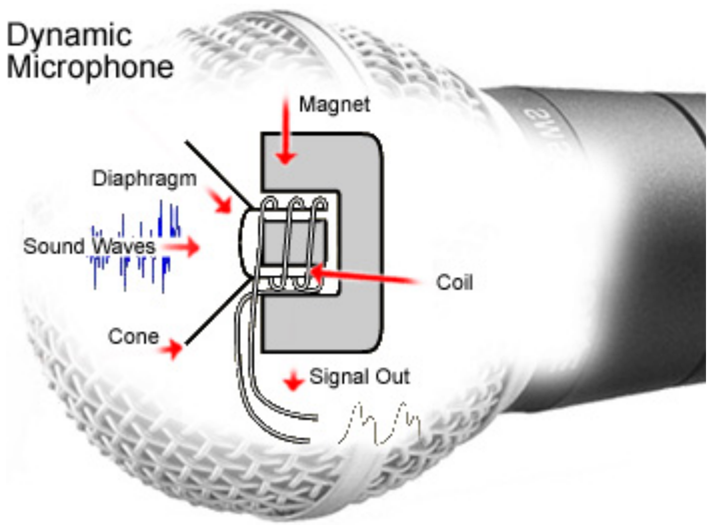

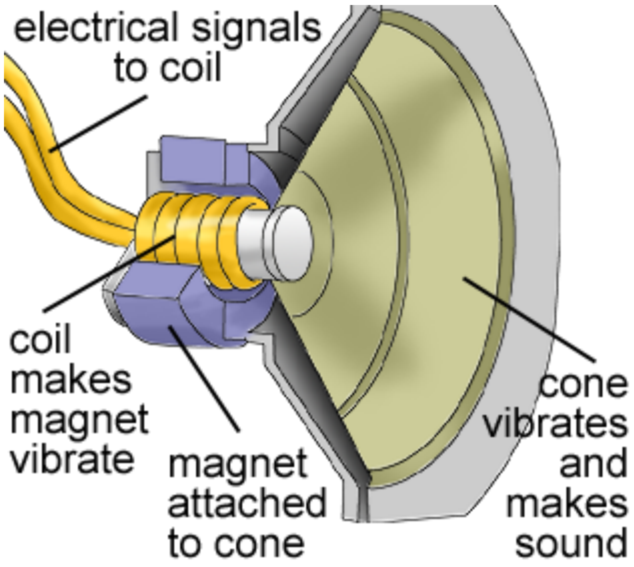

Digitization

- 1-dimensional nature of sound: **amplitude** (sound pressure/level) depend on a 1D variable, the **time**.
  - Input from microphone is analog signal

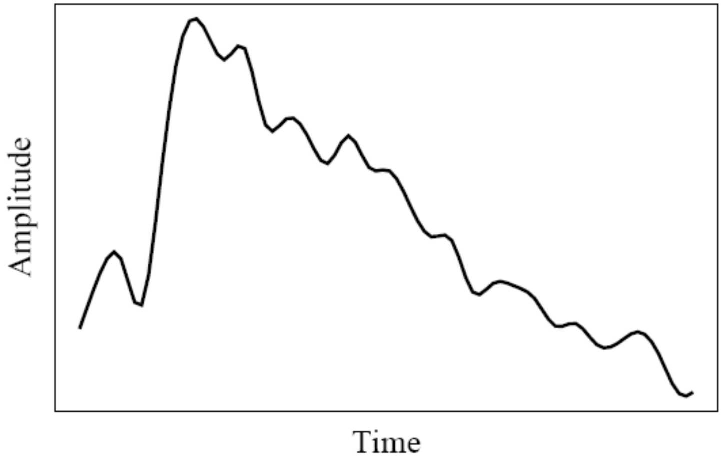

- **Digitization**: conversion to a stream of numbers, and preferably these numbers should be integers for efficiency.

Digitization cont'd

- Digitization must be in both time and amplitude
  - **Sampling**: measuring the quantity we are interested in, usually at evenly-spaced intervals
- First kind of sampling, using measurements only at evenly spaced time intervals, is simply called ***sampling***.
  - The rate is called the *sampling frequency*
  - For audio, typically from 8 kHz (8,000 samples per second) to 48 kHz (determined by Nyquist theorem discussed later).
- Sampling in the amplitude or voltage dimension is called **quantization**

Sampling and Quantization

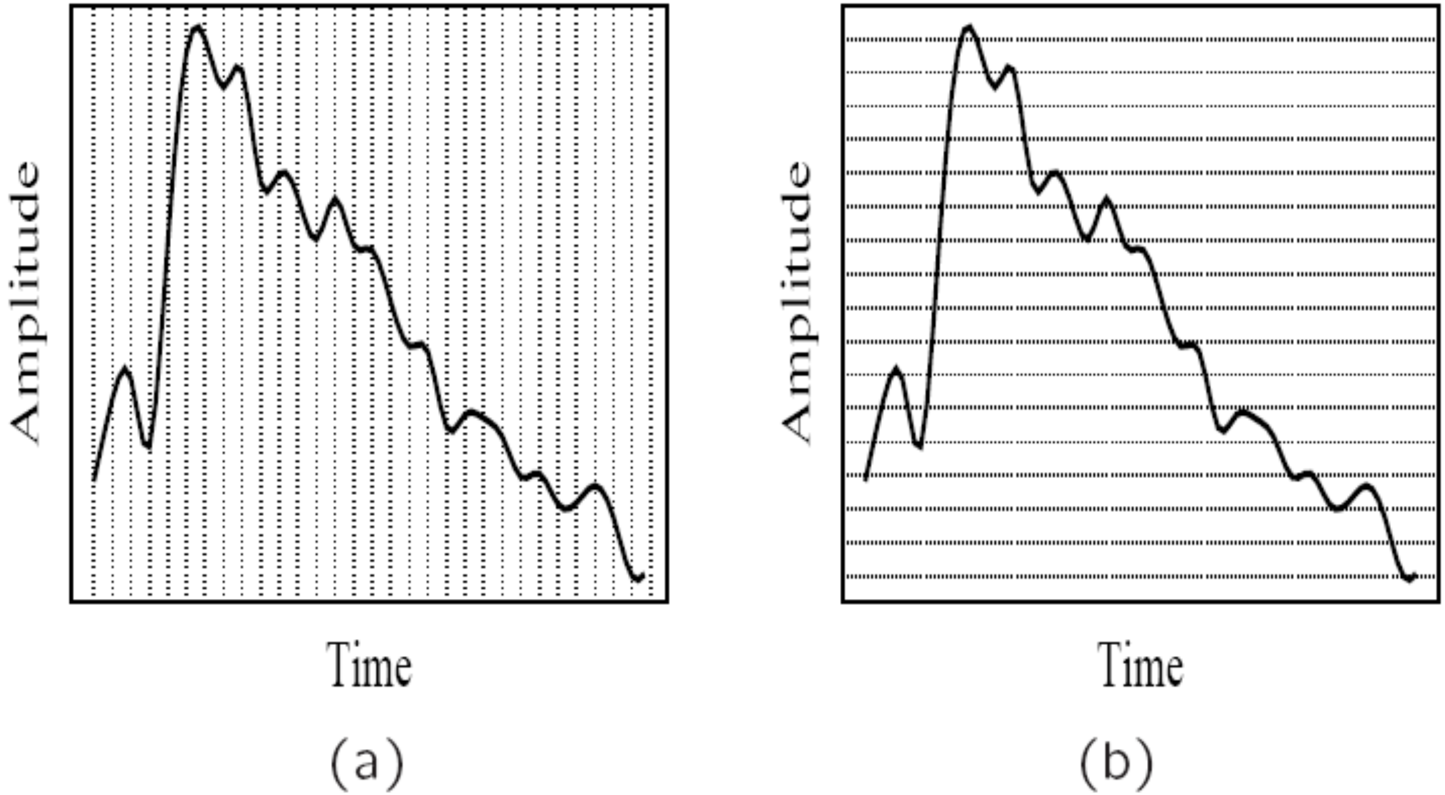

Audio Digitization (PCM)

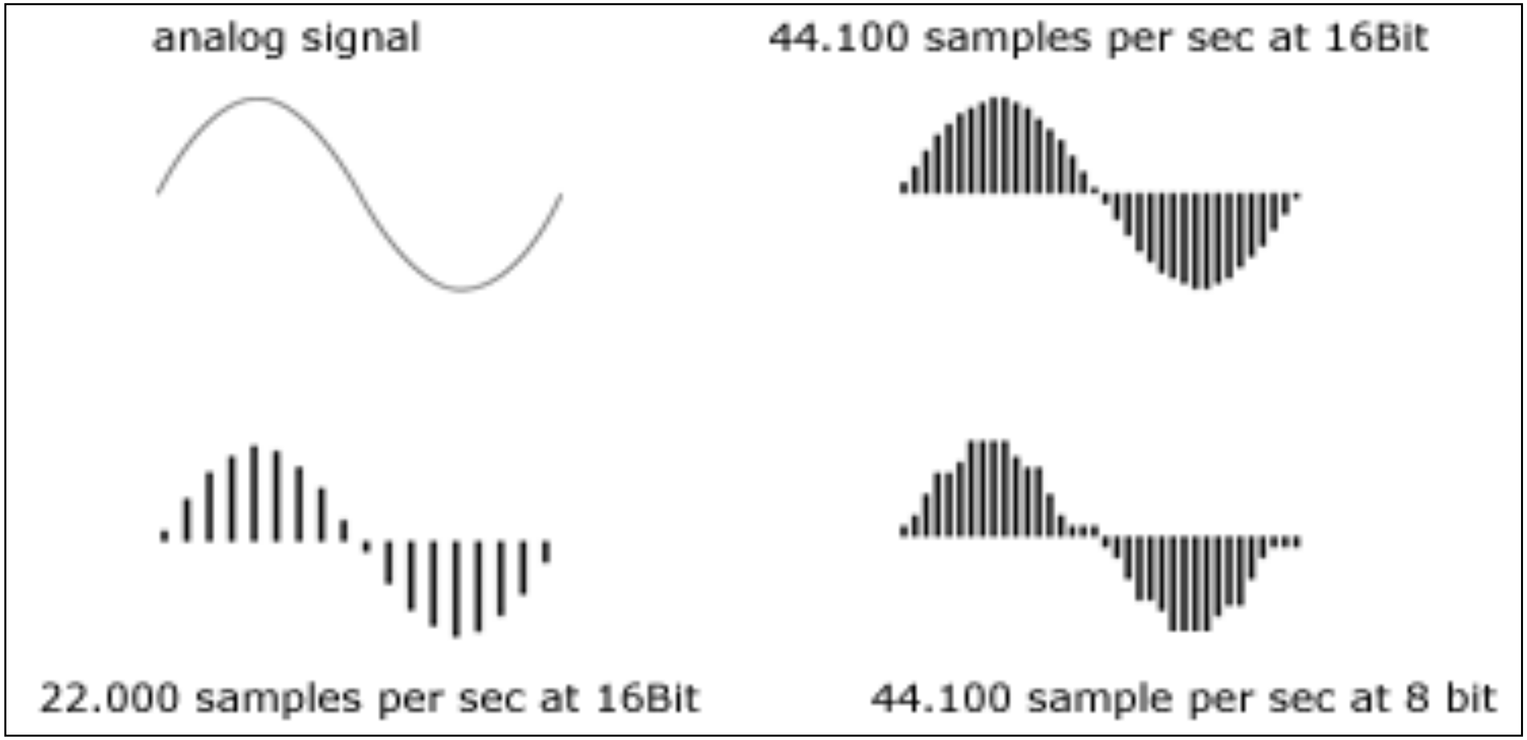

PCM: Pulse coded modulation

Parameters in Digitizing

- To decide how to digitize audio data we need to answer the following questions:

    1. What is the sampling rate?

    2. How finely is the data to be quantized, and is quantization uniform?

    3. How is audio data formatted? (file format)

Outline

- Audio Signals
  - Sampling
  - Quantization
- Audio file format
  - WAV/MIDI
- Human auditory system

Sampling Rate

- Signals can be decomposed into a sum of sinusoids.
  -- weighted sinusoids can build up quite a complex signals
  (recall Calculus and linear algebra)

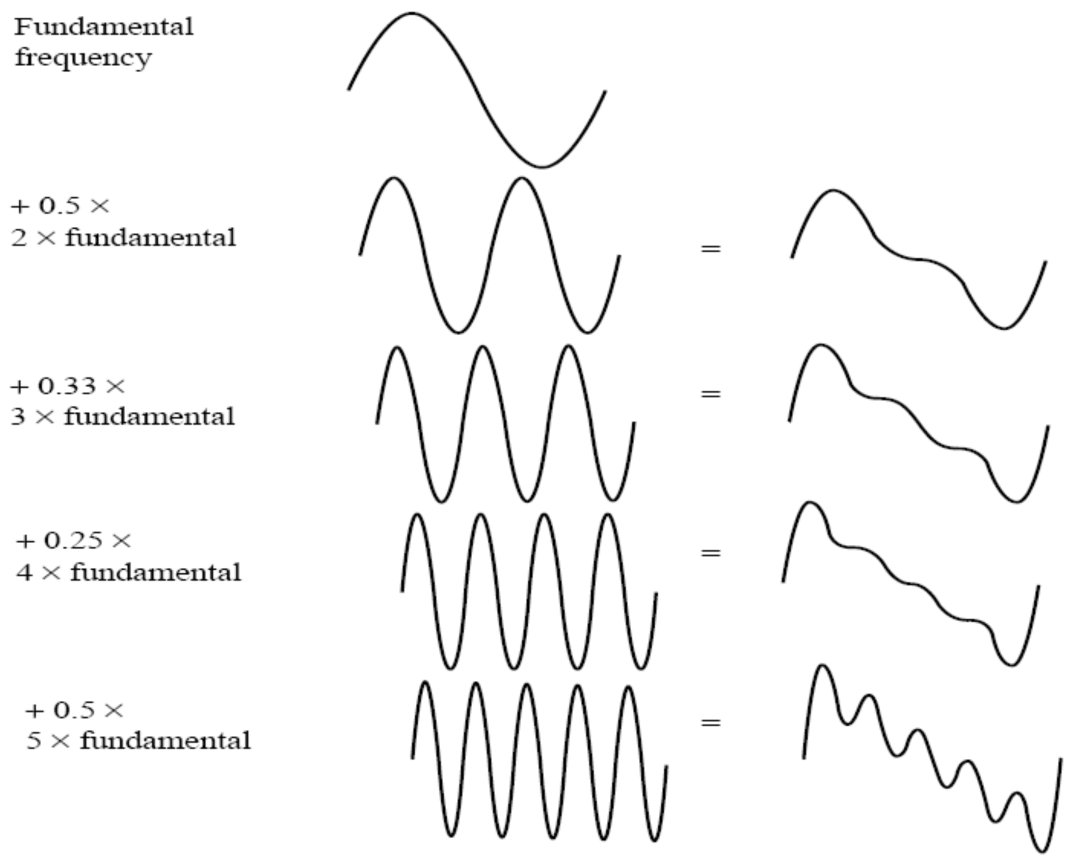

Sampling Rate cont'd

- If sampling rate just equals the actual frequency
  - a false signal (constant) is detected
- If sample at 1.5 times the actual frequency
  - an incorrect (**alias**) frequency that is lower than the correct one
    - it is half the correct one -- the wavelength, from peak to peak, is double that of the actual signal

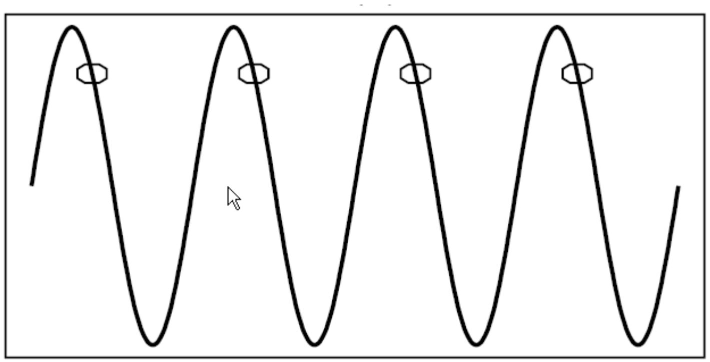

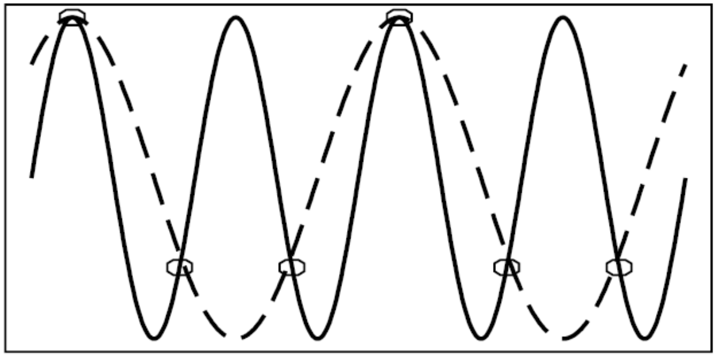

Sampling Rate cont'd

- For correct sampling we must use a sampling rate equal to at least *twice the maximum frequency* content in the signal. This rate is called the **Nyquist rate**.
- The relationship among the Sampling Frequency, True Frequency, and the Alias Frequency is as follows:
  - $f_{alias} = f_{sampling} - f_{true}$, for $f_{true} < f_{sampling} < 2 \times f_{true}$

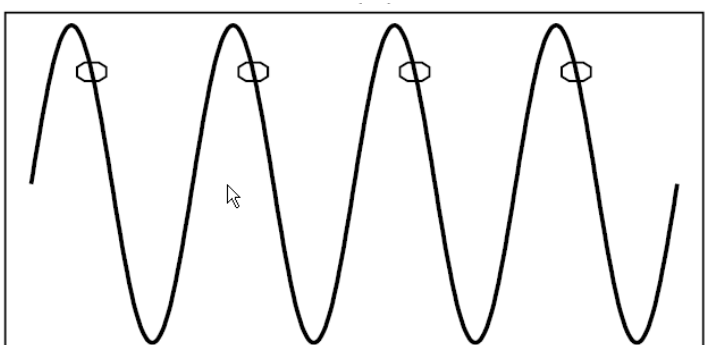

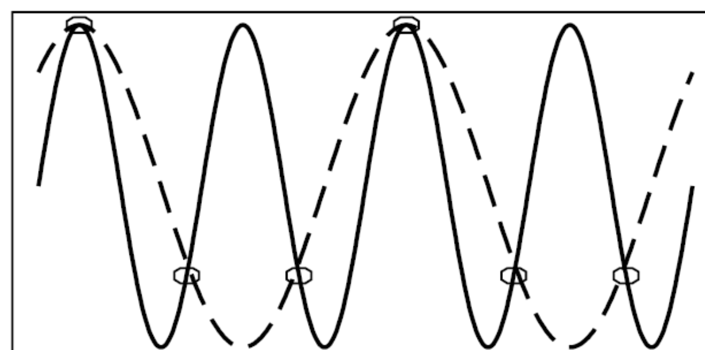

Sampling Rate cont'd

- Fig. 6.5 shows the relationship of the apparent frequency to the input frequency.

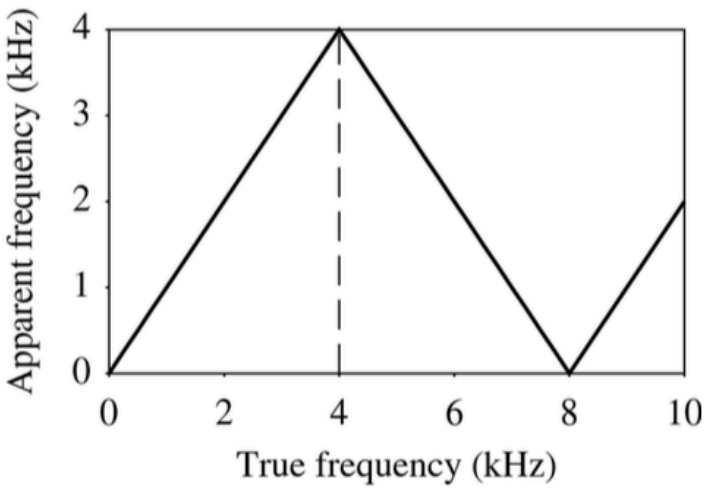

Fig. 6.5: Folding of sinusoid frequency which is sampled at 8,000 Hz. The folding frequency, shown dashed, is 4,000 Hz.

Sampling Rate cont'd

- **Nyquist frequency**: half of the Sampling rate
  - Since it would be impossible to recover frequencies higher than Nyquist frequency in any event, most systems have an **antialiasing filter** that restricts the frequency content in the input to the sampler to a range at or below Nyquist frequency.

Nyquist Theorem

- Sampling theory – Nyquist theorem
  - If a signal is **band-limited**, i.e., there is a lower limit *f*1 and an upper limit *f*2 of frequency components in the signal, then the sampling rate should be at least 2(*f*2 – *f*1).

Proof and more math:
https://en.wikipedia.org/wiki/Nyquist-Shannon_sampling_theorem
https://en.wikipedia.org/wiki/Undersampling

Outline

- Audio Signals
  - Sampling
  - Quantization
- Audio file format
  - WAV/MIDI
- Human auditory system

Quantization (Pulse Code Modulation)

- At every time interval the sound is converted to a digital equivalent
- Using 2 bits the following sound can be digitized
  - Tel: 8 bits
  - CD: 16 bits

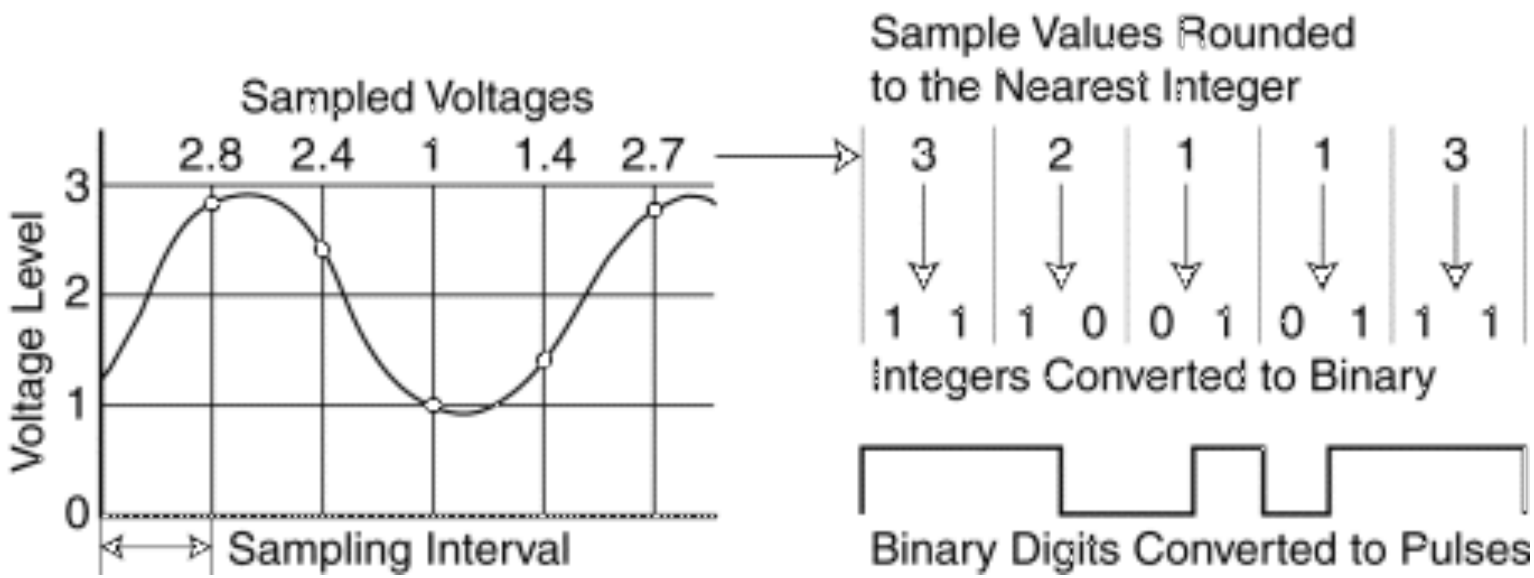

Digitize audio

- Each sample quantized, i.e., rounded
  - e.g., 2^8=256 possible quantized values
- Each quantized value represented by bits
  - 8 bits for 256 values
- Example: 8,000 samples/sec, 256 quantized values --> 64,000 bps
- Receiver converts it back to analog signal:
  - some quality reduction

Example rates

- CD: 1.411 Mbps
- MP3: 96, 128, 160 kbps (with compression)
- Internet telephony: 5.3 - 13 kbps (with compression)

Audio Quality vs. Data Rate

Quality	Sample Rate (KHz)	Bits per Sample	Mono/ Stereo	Data Rate (uncompressed) (kB/sec)	Frequency Band (KHz)
Telephone	8	8	Mono	8	0.200-3.4
AM Radio	11.025	8	Mono	11.0	0.1-5.5
FM Radio	22.05	16	Stereo	88.2	0.02-11
CD	44.1	16	Stereo	176.4	0.005-20
DAT	48	16	Stereo	192.0	0.005-20
DVD Audio	192 (max)	24 (max)	6 channels	1,200.0 (max)	0-96 (max)

More on Quantization

- Quantization is lossy !
- Roundoff errors => **quantization noise/error**

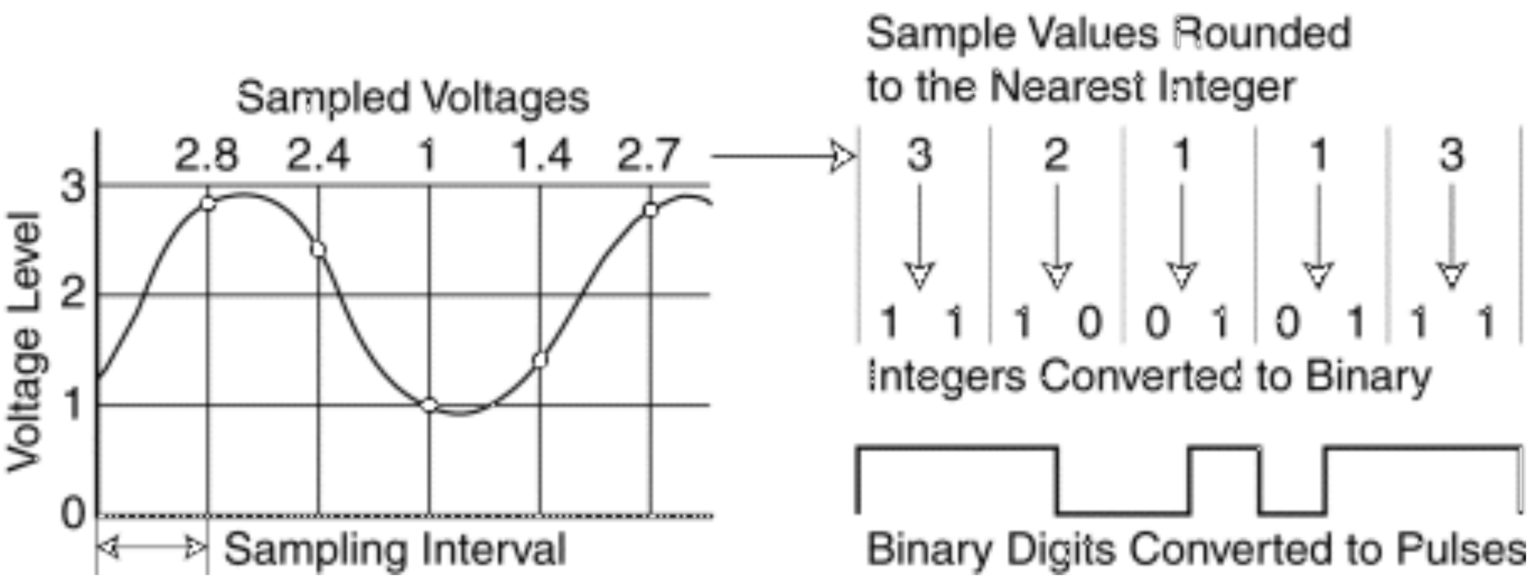

Quantization Noise

- **Quantization noise**: the difference between the actual value of the analog signal, for the particular sampling time, and the nearest quantization interval value.
  - At most, this error can be as much as half of the interval.
- The quality of the quantization is characterized by the Signal to Quantization Noise Ratio (**SQNR**).
  - A special case of SNR (Signal to Noise Ratio)

Signal to Noise Ratio (SNR)

- *Signal to Noise Ratio* **(SNR)**: the ratio of the power of the correct signal and the noise
  - A common measure of the quality of the signal
  - The ratio can be huge and often non-linear
- So practically, SNR is usually measured in log-scale: decibels **(dB)**, where 1 dB is 1/10 **Bel**. The SNR value, in units of dB, is defined in terms of base-10 logarithms of squared voltages, as follows:

$$SNR = 10 \log_{10} \frac{V_{signal}^2}{V_{noise}^2} = 20 \log_{10} \frac{V_{signal}}{V_{noise}}$$

Signal to Noise Ratio (SNR) cont'd

$$SNR = 10 \log_{10} \frac{V_{signal}^2}{V_{noise}^2} = 20 \log_{10} \frac{V_{signal}}{V_{noise}}$$

- The actual power in a signal is proportional to the square of the voltage. For example, if the signal voltage V_{signal} is 10 times the noise, then the SNR is 20 $\log_{10}(10)$=20dB.
  - if the power from ten violins is ten times that from one violin playing, then the ratio of power is 10dB, or 1B.

Common sound levels

Table 6.1: Magnitude levels of common sounds, in decibels

Threshold of hearing	0
Rustle of leaves	10
Very quiet room	20
Average room	40
Conversation	60
Busy street	70
Loud radio	80
Train through station	90
Riveter	100
Threshold of discomfort	120
Threshold of pain	140
Damage to ear drum	160

Quantization Noise Ratio

- Aside from any noise that may have been present in the original analog signal, there is also an additional error that results from quantization.

  - (a) If voltages are actually in 0 to 1 but we have only 8 bits in which to store values, then effectively we force all continuous values of voltage into only 256 different values.

  - (b) This introduces a roundoff error. It is not really "noise". Nevertheless it is called **quantization noise** (or quantization error).

Signal-to-Quantization Noise Ratio (SQNR)

- The quality of the quantization is characterized by the Signal to Quantization Noise Ratio (**SQNR**).

  (a) **Quantization noise**: the difference between the actual value of the analog signal, for the particular sampling time, and the nearest quantization interval value.

  - (b) At most, this error can be as much as half of the interval.

Signal-to-Quantization Noise Ratio (SQNR) cont'd

- For a quantization accuracy of N bits per sample, the peak SQNR can be simply expressed:

$$SQNR = 20\log_{10}\frac{V_{signal}}{V_{quan\_noise}} = 20\log_{10}\frac{2^{N-1}}{\frac{1}{2}}$$

$$= 20 \times N \times \log 2 = 6.02\,N(\text{dB}) \qquad (6.3)$$

- $6.02N$ is the worst case.

Note: We map the maximum signal to $2^{N-1} - 1$ ($\simeq 2^{N-1}$) and the most negative signal to -2^{N-1}.

Dynamic range : the ratio of maximum to minimum absolute values of the signal: V_{max}/V_{min}. The max abs. value V_{max} gets mapped to $2^{N-1} - 1$; the min abs. value V_{min} gets mapped to 1. V_{min} is the smallest positive voltage that is not masked by noise. The most negative signal, $-V_{max}$, is mapped to -2^{N-1}.

Linear and Non-linear Quantization

- **Linear format**: samples are typically stored as uniformly quantized values.
- **Non-uniform quantization**: set up more finely-spaced levels where humans hear with the most acuity.
  - Weber's Law stated formally says that equally perceived differences have values proportional to absolute levels:

$$\Delta\text{Response} \propto \Delta\text{Stimulus}/\text{Stimulus} \qquad (6.5)$$

  - Inserting a constant of proportionality k, we have a differential equation that states:

$$dr = k\,(1/s)\,ds \qquad (6.6)$$

with response r and stimulus s.

– Integrating, we arrive at a solution

$$r = k \ln(s) + C \tag{6.7}$$

with constant of integration C.

Stated differently, the solution is

$$r = k \ln(s/s_0) \tag{6.8}$$

s_0 = the lowest level of stimulus that causes a response (r = 0 when s = s_0).

- Nonlinear quantization works by first transforming an analog signal from the raw s space into the theoretical r space, and then uniformly quantizing the resulting values.
- Such a law for audio is called μ-**law** encoding, (or **u-law**). A very similar rule, called A-**law**, is used in telephony in Europe.
- The equations for these very similar encodings are as follows:

- μ-law:

$$r = \frac{\operatorname{sgn}(s)}{\ln(1+\mu)} \ln\left\{1 + \mu\left|\frac{s}{s_p}\right|\right\}, \qquad \left|\frac{s}{s_p}\right| \le 1 \tag{6.9}$$

- A-law:

$$r = \begin{cases} \dfrac{A}{1+\ln A}\left(\dfrac{s}{s_p}\right), & \left|\dfrac{s}{s_p}\right| \le \dfrac{1}{A} \\[2ex] \dfrac{\operatorname{sgn}(s)}{1+\ln A}\left[1 + \ln A\left|\dfrac{s}{s_p}\right|\right], & \dfrac{1}{A} \le \left|\dfrac{s}{s_p}\right| \le 1 \end{cases} \tag{6.10}$$

$$\text{where } \operatorname{sgn}(s) = \begin{cases} 1 & \text{if } s > 0, \\ -1 & \text{otherwise} \end{cases}$$

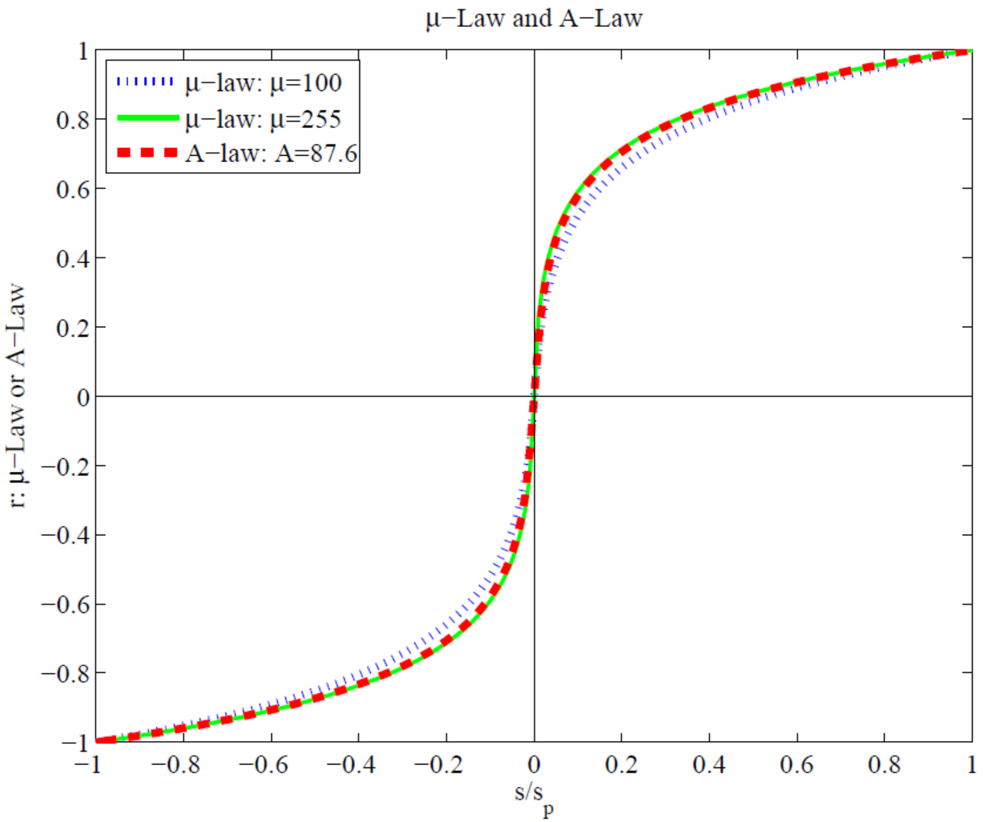

- **Fig. 6.6**: Nonlinear transform for audio signals.

The parameter μ is set to $\mu = 100$ or $\mu = 255$; the parameter A for the A-law encoder is usually set to $A = 87.6$.

- The μ-law in audio is used to develop a nonuniform quantization rule for sound: uniform quantization of r gives finer resolution in s at the quiet end.

Bit allocation

- In μ-law, we would like to put the available bits where the most perceptual acuity (sensitivity to small changes) is.
    1. Savings in bits can be gained by transmitting a smaller bit-depth for the signal.
    2. μ-law often starts with a bit-depth of 16 bits, but transmits using 8 bits.
    3. And then expands back to 16 bits at the receiver.
- Let μ=255.
- Now, we want s in $[-1, 1]$. The input is in -2^{15} to $(+2^{15}-1)$, we divide by 2^{15} to normalize.
- Then the μ-law is applied to turn s into r.
- Now go down to 8-bit samples, using $\hat{r} = \text{sign}(s) * \text{floor}(128 * r)$.
- Now the 8-bit signal $\hat{r}$ is transmitted.
- At the receiver side, we normalize $\hat{r}$ by dividing by 2^7, and then apply the inverse μ-law function:

- at the receiver side, we normalize $\hat{r}$ by dividing by 2^7, and then apply the inverse μ-law function:

$$\hat{s} = \mathrm{sign}(s)\left(\frac{(\mu+1)^{|\hat{r}|}-1}{\mu}\right)$$

- Finally, we expand back up to 16 bits:

$$\tilde{s} = ceil\left(2^{15} \times \hat{s}\right)$$

Companding puts the most accuracy at the quiet end near zero.

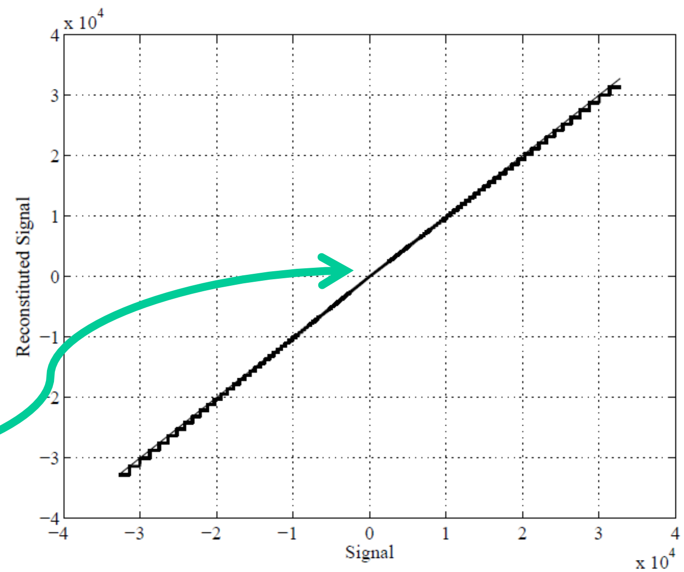

Fig. 6.7: Nonlinear quantization by companding.

Outline

- Audio Signals
  - Sampling
  - Quantization
- Audio file format
  - WAV/MIDI
- Human auditory system

MIDI: Musical Instrument Digital Interface

- Use the sound card's defaults for sounds: ⇒ use a simple scripting language and hardware setup called **MIDI**.

- **MIDI Overview**

  a) MIDI is a scripting language — it codes "events" that stand for the production of sounds. E.g., a MIDI event might include values for the pitch of a single note, its duration, and its volume.

  b) MIDI is a standard adopted by the electronic music industry for controlling devices, such as synthesizers and sound cards, that produce music.

Midi: https://www.youtube.com/watch?v=SUUxmJ84dnI
Example: https://onlinesequencer.net/#263068

a) The MIDI standard is supported by most synthesizers, so sounds created on one synthesizer can be played and manipulated on another synthesizer and sound reasonably close.

b) Computers must have a special MIDI interface, but this is incorporated into most sound cards. The sound card must also have both D/A and A/D converters.

MIDI Concepts

- MIDI **channels** are used to separate messages.
  - (a) There are 16 channels numbered from 0 to 15. The channel forms the last 4 bits (the least significant bits) of the message.
  - (b) Usually a channel is associated with a particular instrument: e.g., channel 1 is the piano, channel 10 is the drums, etc.
  - (c) Nevertheless, one can switch instruments midstream, if desired, and associate another instrument with any channel.

- **System messages**
  - (a) Several other types of messages, e.g. a general message for all instruments indicating a change in tuning or timing.
  - (b) If the first 4 bits are all 1s, then the message is interpreted as a **system common** message.

- The way a synthetic musical instrument responds to a MIDI message is usually by simply ignoring any **play sound** message that is not for its channel.
  - If several messages are for its channel, then the instrument responds, provided it is **multi-voice**, i.e., can play more than a single note at once.

MIDI Terminology

- Synthesizer:
  - was, and still can be, a stand-alone sound generator that can vary pitch, loudness, and tone color.
  - Units that generate sound are referred to as *tone modules* or sound modules.
- Sequencer:
  - started off as a special hardware device for storing and editing a *sequence* of musical events, in the form of MIDI data.
  - Now it is more often a software *music editor* on the computer.
- MIDI Keyboard:
  - produces no sound, instead generating sequences of MIDI in- structions, called *MIDI messages*
  - MIDI messages are rather like assembler code and usually consist of just a few bytes

MIDI Terminology

- Timbre vs Vioce
  - Timbre is MIDI terminology for just what instrument that is trying to be emulated, e.g. a piano as opposed to a violin: it is the quality of the sound.
  - Vioce is used in MIDI to mean every different timbre and pitch that the tone module can produce at the same time. Synthesizers can have many (typically 16, 32, 64, 256, etc.) voices. Each voice works independently and simultaneously to produce sounds of different timbre and pitch.
- *Polyphony*
  - Refers to the number of voices that can be produced at the same time

MIDI Specifics

Question:
How different timbres are produced digitally ?

- Different timbres are produced digitally by using a **patch** — the set of control settings that define a particular timbre. Patches are often organized into databases, called **banks**.

- **General MIDI**: A standard mapping specifying what instruments (what patches) will be associated with what channels.
    a) In General MIDI, channel 10 is reserved for percussion instruments, and there are 128 patches associated with standard instruments.
    b) For most instruments, a typical message might be a Note On message (meaning, e.g., a keypress and release), consisting of what channel, what pitch, and what "velocity" (i.e., volume).
    c) For percussion instruments, however, the pitch data means which kind of drum.
    d) A Note On message consists of "status" byte — which channel, what pitch — followed by two data bytes. It is followed by a Note Off message, which also has a pitch (which note to turn off) and a velocity (often set to zero).

- → Link to General MIDI Instrument Patch Map
- → Link to General MIDI Percussion Key Map

- The data in a MIDI status byte is between 128 and 255; each of the *data bytes* is between 0 and 127. Actual MIDI bytes are 10-bit, including a 0 start and 0 stop bit.

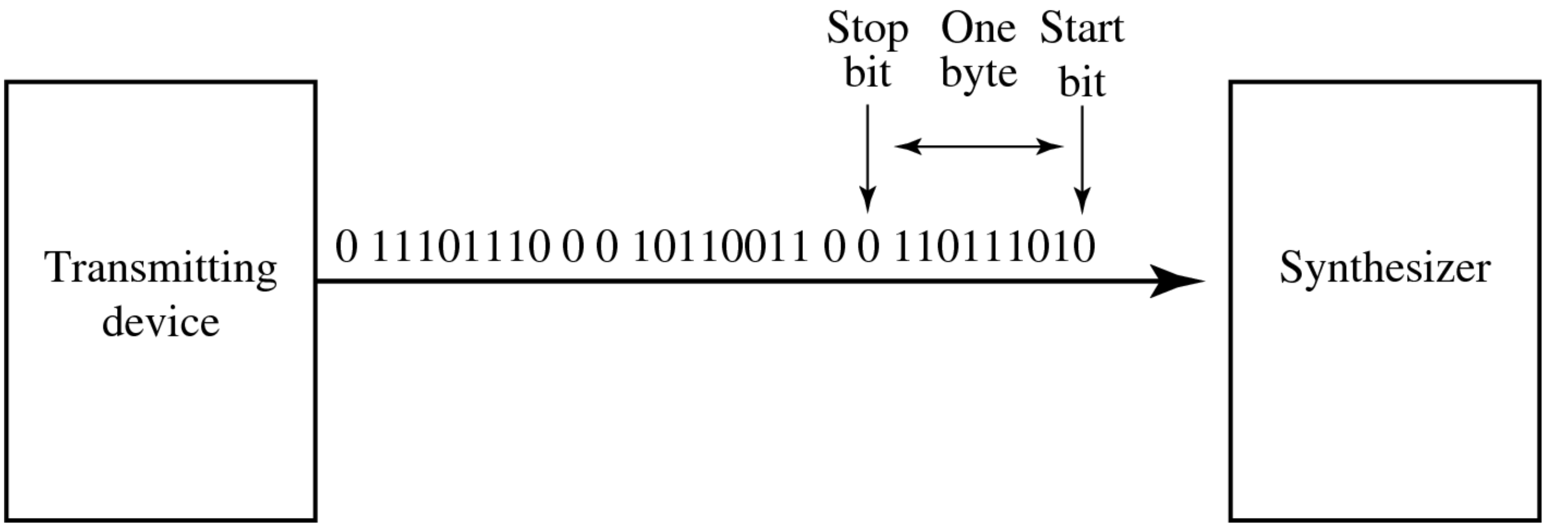

Fig. 6.9: Stream of 10-bit bytes; for typical MIDI messages, these consist of {Status byte, Data Byte, Data Byte} = {Note On, Note Number, Note Velocity}

- A MIDI device often is capable of **programmability**, and also can change the **envelope** describing how the amplitude of a sound changes over time.
- Fig. 6.10 shows a model of the response of a digital instrument to a Note On message:

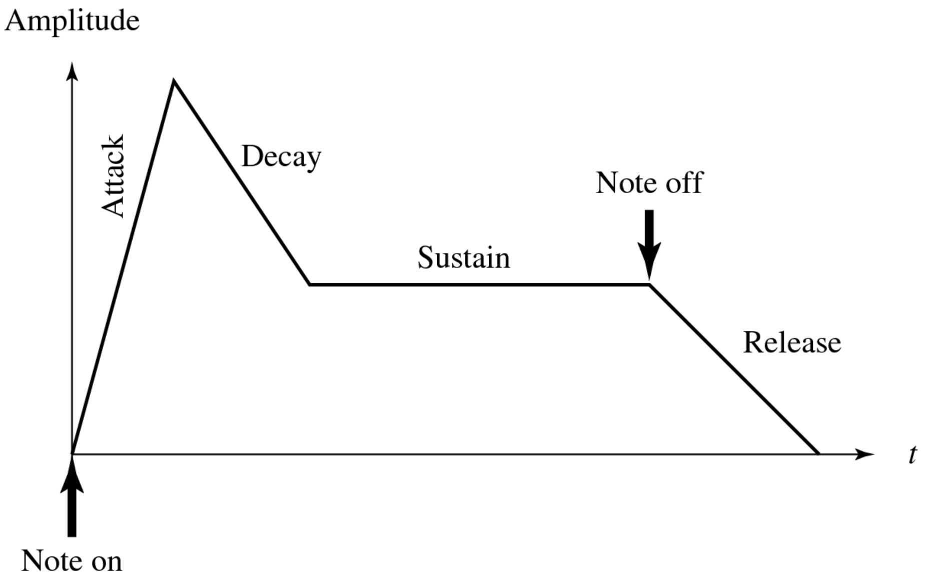

❒ **Fig. 6.10:** Stages of amplitude versus time for a music note

6.2.2 Hardware Aspects of MIDI

- The MIDI hardware setup consists of a 31.25 kbps serial connection. Usually, MIDI-capable units are either Input devices or Output devices, not both.
- A traditional synthesizer is shown in Fig. 6.11:

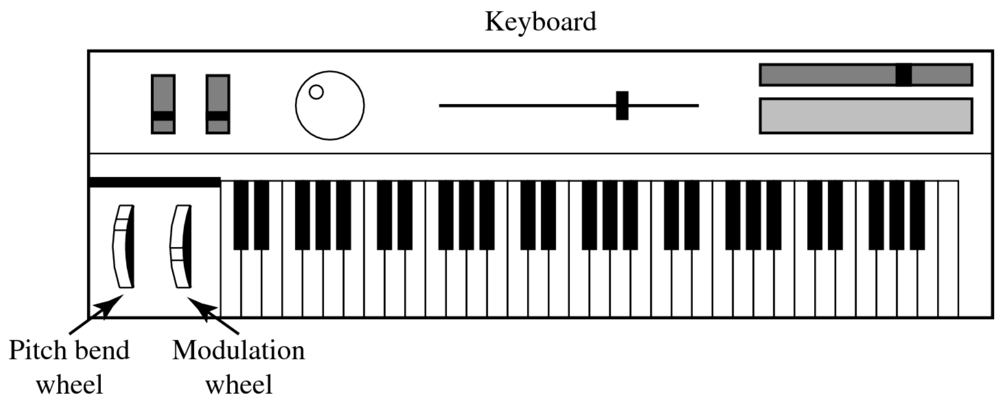

Fig. 6.11: A MIDI synthesizer

- The physical MIDI ports consist of 5-pin connectors for IN and OUT, as well as a third connector called THRU.

  a) MIDI communication is half-duplex.

  b) MIDI IN is the connector via which the device receives all MIDI data.

  c) MIDI OUT is the connector through which the device transmits all the MIDI data it generates itself.

  d) MIDI THRU is the connector by which the device echoes the data it receives from MIDI IN. Note that it is only the MIDI IN data that is echoed by MIDI THRU — all the data generated by the device itself is sent via MIDI OUT.

- A typical MIDI sequencer setup is shown in Fig. 6.12:

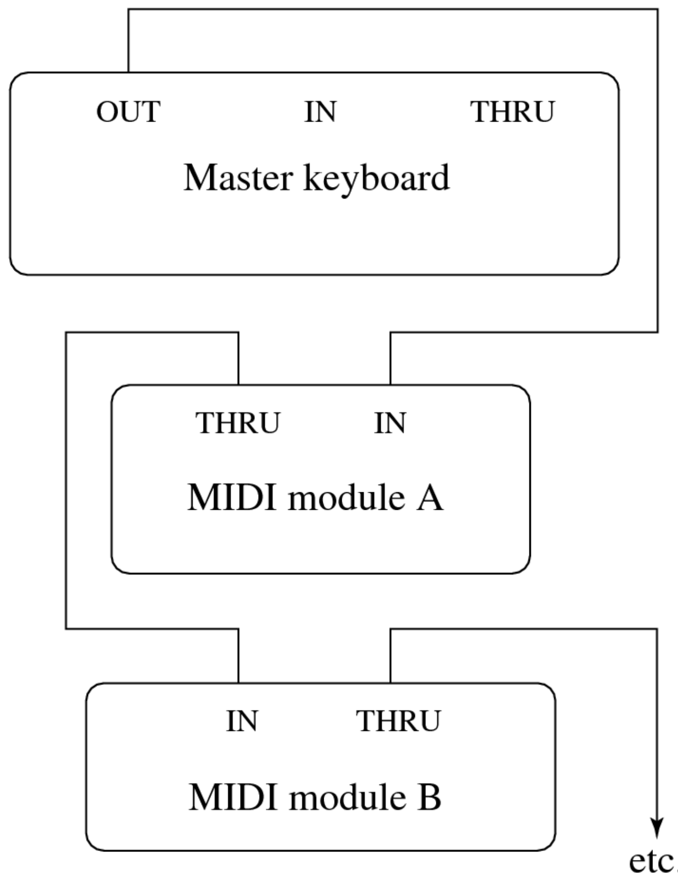

Fig. 6.12: A typical MIDI setup

6.2.3 Structure of MIDI Messages

- MIDI messages can be classified into two types: channel messages and system messages, as in Fig. 6.13:

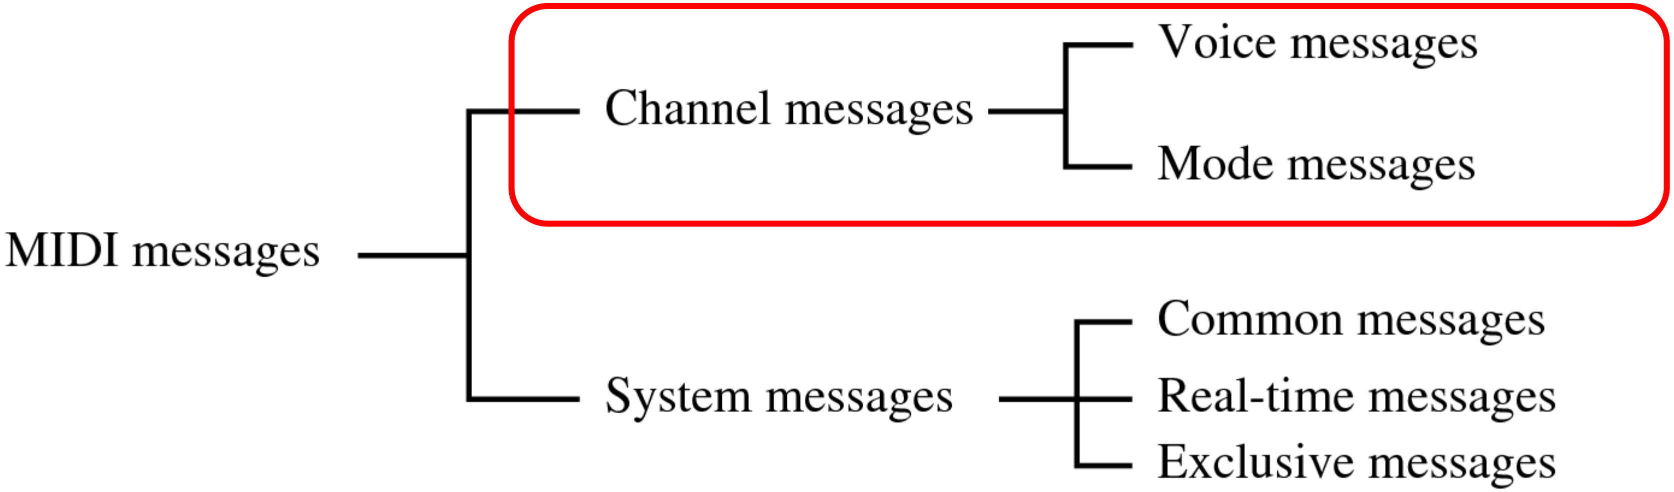

- **Fig. 6.13:** MIDI message taxonomy

- **A. Channel messages**: can have up to 3 bytes:
  a) The first byte is the status byte (the opcode, as it were); has its most significant bit set to **1**.
  b) The 4 low-order bits identify which channel this message belongs to (for 16 possible channels).
  c) The 3 remaining bits hold the message. For a data byte, the most significant bit is set to **0**.
- **A.1. Voice messages**:
  a) This type of channel message controls a voice, i.e., sends information specifying which note to play or to turn off, and encodes key pressure.
  b) Voice messages are also used to specify controller effects such as sustain, vibrato, tremolo, and the pitch wheel.

❍ Table 6.3 lists these operations.

- **Table 6.3: MIDI voice messages**

Voice Message	Status Byte	Data Byte1	Data Byte2
Note Off	&H8n	Key number	Note Off velocity
Note On	&H9n	Key number	Note On velocity
Poly. Key Pressure	&HAn	Key number	Amount
Control Change	&HBn	Controller num.	Controller value
Program Change	&HCn	Program number	None
Channel Pressure	&HDn	Pressure value	None
Pitch Bend	&HEn	MSB	LSB

- (** &H indicates hexadecimal, and 'n' in the status byte hex value stands for a channel number. All values are in 0..127 except Controller number, which is in 0..120)

❒ • A.2. Channel mode messages:

a) Channel mode messages: special case of the Control Change message → opcode B (the message is &HBn, or 1011nnnn).

b) However, a Channel Mode message has its first data byte in 121 through 127 (&H79–7F).

c) Channel mode messages determine how an instrument processes MIDI voice messages: respond to all messages, respond just to the correct channel, don't respond at all, or go over to local control of the instrument.

d) The data bytes have meanings as shown in Table 6.4.

Table 6.4: MIDI mode messages

1st Data Byte	Description	Meaning of 2nd Data Byte
&H79	Reset all controllers	None; set to 0
&H7A	Local control	0 = off; 127 = on
&H7B	All notes off	None; set to 0
&H7C	Omni mode off	None; set to 0
&H7D	Omni mode on	None; set to 0
&H7E	Mono mode on (Poly mode off)	Controller number
&H7F	Poly mode on (Mono mode off)	None; set to 0

6.2.3 Structure of MIDI Messages

- MIDI messages can be classified into two types: channel messages and system messages, as in Fig. 6.13:

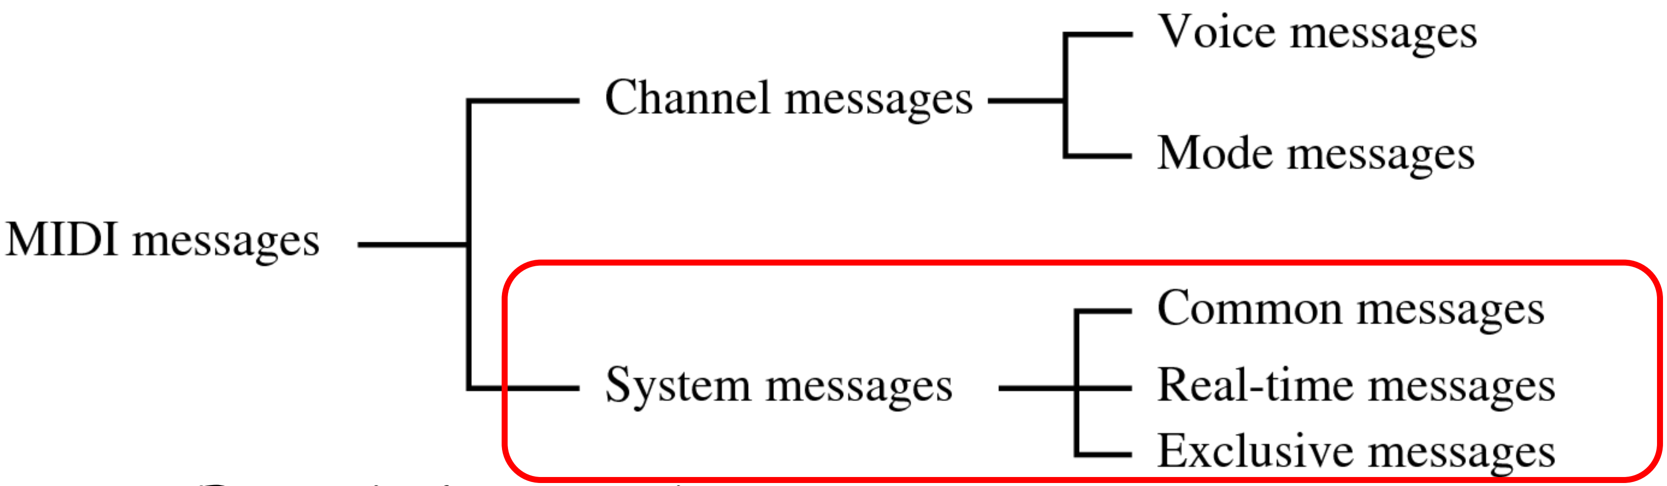

- **Fig. 6.13: MIDI message taxonomy**

- B. **System Messages:**
  a) System messages have no channel number — commands that are not channel specific, such as timing signals for synchronization, positioning information in pre-recorded MIDI sequences, and detailed setup information for the destination device.
  b) Opcodes for all system messages start with &HF.
  c) System messages are divided into three classifications, according to their use:

System Common Message

System Common Message	Status Byte	Number of Data Bytes
MIDI Timing Code	&HF1	1
Song Position Pointer	&HF2	2
Song Select	&HF3	1
Tune Request	&HF6	None
EOX (terminator)	&HF7	None

- **System real-time messages**: related to synchronization.

  - **Table 6.6:** MIDI System Real-Time messages.

System Real-Time Message	Status Byte
Timing Clock	&HF8
Start Sequence	&HFA
Continue Sequence	&HFB
Stop Sequence	&HFC
Active Sensing	&HFE
System Reset	&HFF

- B.3. **System exclusive message**: included so that the MIDI standard can be extended by manufacturers.
  a) After the initial code, a stream of any specific messages can be inserted that apply to their own product.
  b) A System Exclusive message is supposed to be terminated by a terminator byte &HF7, as specified in Table 6.
  c) The terminator is optional and the data stream may simply be ended by sending the status byte of the next message.

6.2.4 General MIDI

- General MIDI is a scheme for standardizing the assignment of instruments to patch numbers.
  a) Patch 1 should always be a piano
  b) A standard percussion map specifies 47 percussion sounds. Where a "note" appears on a musical score determines what percussion instrument is being struck: a bongo drum, a cymbal.
  c) Other requirements for General MIDI compatibility: MIDI device must support all 16 channels; a device must be multitimbral (i.e., each channel can play a different instrument/program); a device must be polyphonic (i.e., each channel is able to play many voices); and there must be a minimum of 24 dynamically allocated voices.
- **General MIDI Level2**: An extended general MIDI was defined in 1999 and updated in 2003, with a standard `.smf` "Standard MIDI File" format defined — inclusion of extra character information, such as karaoke lyrics.

6.2.5 MIDI to WAV Conversion

- Some programs, such as early versions of Premiere, cannot include `.mid` files — instead, they insist on `.wav` format files.
  a) Various shareware programs exist for approximating a reasonable conversion between MIDI and WAV formats.
  b) These programs essentially consist of large lookup files that try to substitute pre-defined or shifted WAV output for MIDI messages, with inconsistent success.

Outline

- Audio Signals
  - Sampling
  - Quantization
- Audio file format
  - WAV/MIDI
- Human auditory system

Computer vs. Ear

- Multimedia signals are interpreted by humans!
  - Need to understand human perception
- Almost all original multimedia signals are analog signals:
  - A/D conversion is needed for computer processing

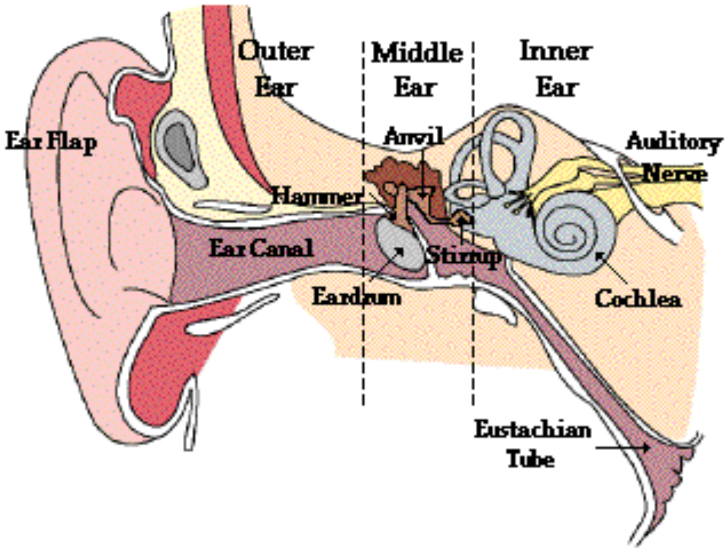

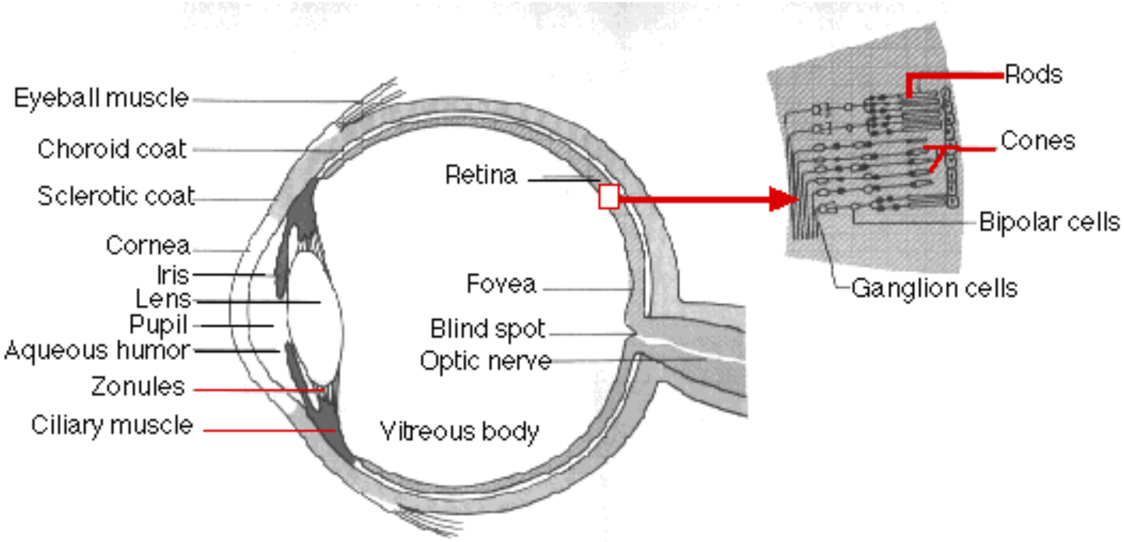

Properties of HAS: Human Auditory System

- Range of human' hearing: 20Hz - 20kHz
  - ➔ Minimal sampling rate for music: 40 kHz (Nyquist rate)
  - CD Audio:
    - 44.1 kHz sampling rate
    - each sample is represented by a 16-bit signed integer
    - 2 channels are used to create stereo system

    ➔44100 * 16 * 2 = 1,411,200 bits / second (bps)
  - Speech signal: 300 Hz – 4 KHz
    - → Minimum sampling rate is 8 KHz (as in telephone system)
    - The extremes of the human voice
      - http://www.noiseaddicts.com/2009/04/extremes-of-human-voice/

Properties of Human Auditory System

- Hearing threshold varies dramatically at different frequencies
- Most sensitive around 2KHz

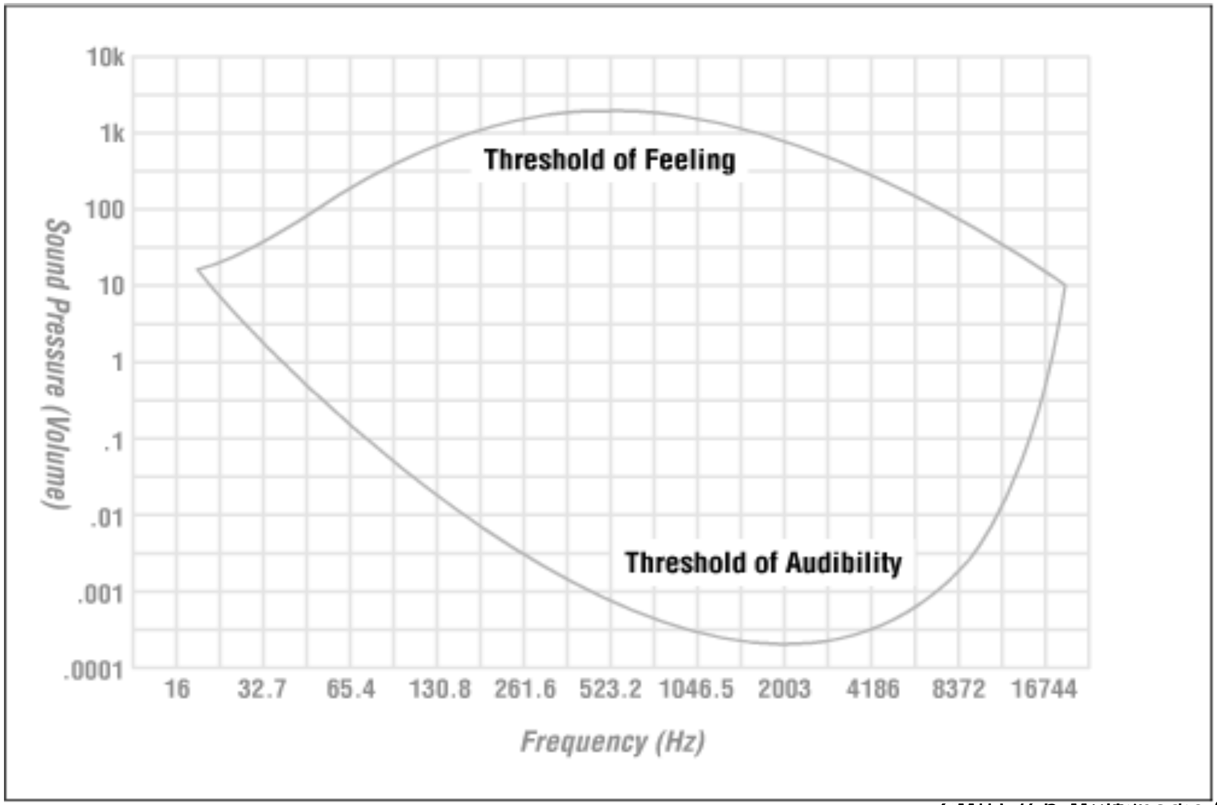

Properties of Human Auditory System

- Hearing Loss Test
  - http://www.noiseaddicts.com/2010/10/hearing-loss-test/
  - http://www.freemosquitoringtones.org/hearing_test/

- Can you hear like an audio engineer ?
  - http://www.noiseaddicts.com/2010/03/can-you-hear-like-an-audio-engineer/
- Can you hear which is louder ?
  - www.noiseaddicts.com/2010/03/sound-challenge-can-you-hear-which-is-louder/

Properties of Human Auditory System

- Can I hear ultrasonic ringtones ?
  - http://www.ultrasonic-ringtones.com/
- **Mosquito Ringtones** (>17Khz, not auditable by 30+ age)
  - http://www.noiseaddicts.com/2011/06/mosquito-ringtones/
  - **http://www.freemosquitoringtones.org/**

Properties of Human Auditory System

Ultrasonic noise thwarts loitering teens

Photo 1 of 3 | View Enlarged Photo BUY PHOTO

The Mosquito speaker at the Hot Chocolate Sparrow emits a high-frequency, piercing sound that only young people can hear.

Cape Cod Times/Susan Milton

Text Size: A | A | A

Print this Article Email this Article

Share This

By Susan Milton
STAFF WRITER
April 29, 2008

ORLEANS — Among teens, the buzz in Orleans is really an irritating screech.

The noise comes from a small, black box mounted near a second-story window at the Hot Chocolate Sparrow. It's an early version of the Mosquito, the "ultrasonic teen repellent," distributed in the United States by a company called Kids Be Gone. The company has sold nearly 1,000 Mosquitoes for $1,500 each.

The outside speaker broadcasts sound at a high frequency for three to 10 minutes. It annoys and drives away people 12 to 25 years old who are within 30 to 50 feet.

Related Links

- Listen to, or try to hear, the teen

Properties of Human Auditory System

Critical Bands:

- Our brains perceive the sounds through 25 distinct *critical bands*. The bandwidth grows with frequency (above 500Hz).
- At 100Hz, the bandwidth is about 160Hz;
- At 10kHz it is about 2.5kHz in width.

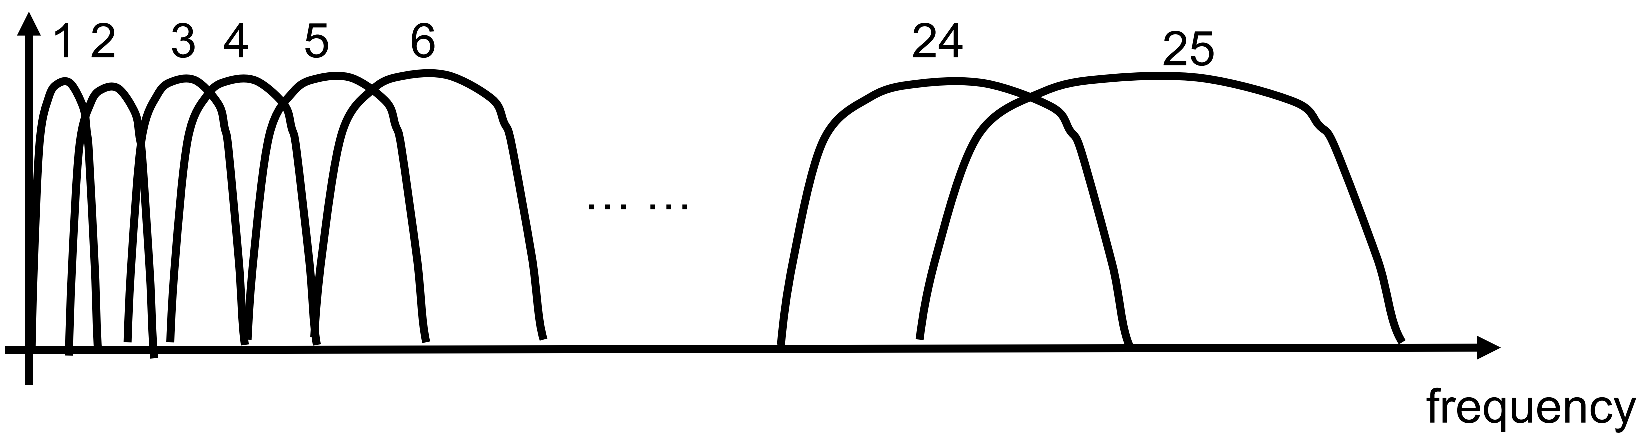

Properties of Human Auditory System

- Masking effect:
  - what we hear depends on what audio environment we are in
  - One strong signal can overwhelm/ hide another

The masking effects in the frequency domain:

A masker inhibits perception of coexisting signals below the masking threshold.

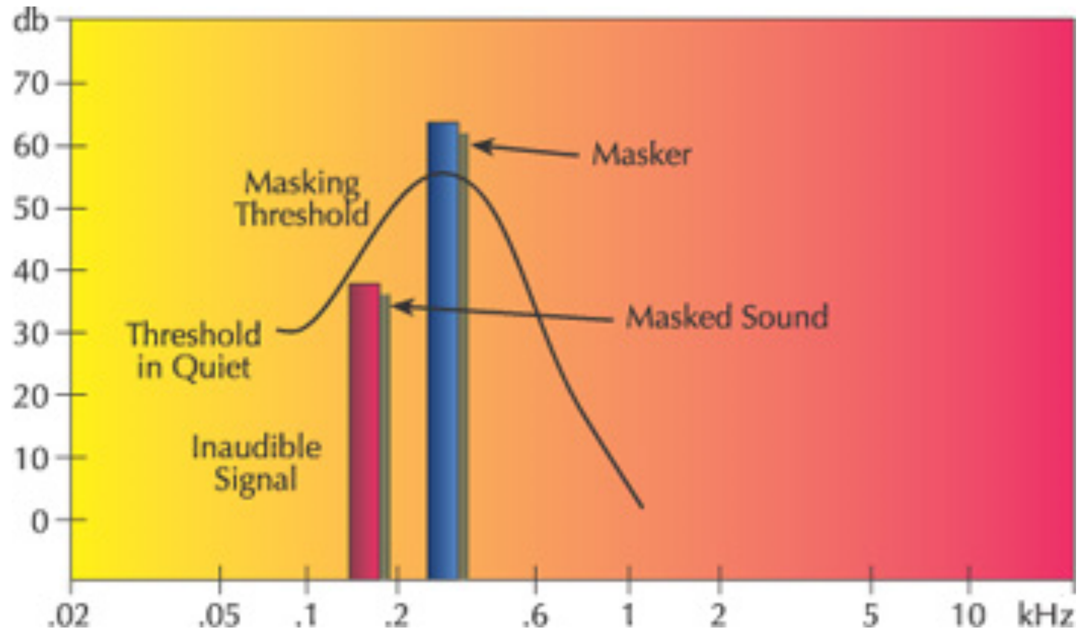

Properties of Human Auditory System

- Masking thresholds in the time domain:

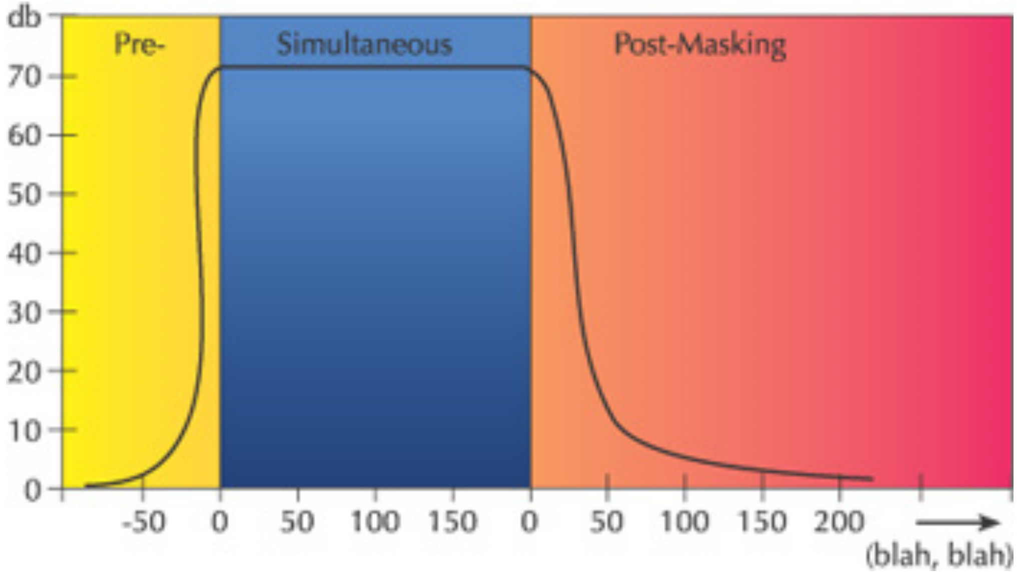

Simultaneous masking: Two sounds occur simultaneously and one is masked by the other.

Backward masking (Pre): A softer sound that occurs prior to a loud one will be masked by the louder sound.

Forward masking (Post): softer sounds that occur as much as 200 milliseconds after the loud sound will also be masked.

HAS: Audio Filtering

- Prior to sampling and AD (Analog-to-Digital) conversion, the audio signal is also usually *filtered* to remove unwanted frequencies.
  - For speech, typically from 50Hz to 10kHz is retained, and other frequencies are blocked by the use of a **band-pass filter** that screens out lower and higher frequencies
  - An audio music signal will typically contain from about 20Hz up to 20kHz
  - At the DA converter end, high frequencies may reappear in the output (Why ?)
    - because of sampling and then quantization, smooth input signal is replaced by a series of step functions containing all possible frequencies
  - So at the decoder side, a **lowpass** filter is used after the DA circuit

HAS: Perceptual audio coding

- The HAS properties can be exploited in audio coding:
  - Different quantizations for different critical bands
    - Subband coding
  - If you can't hear the sound, don't encode it
  - Discard weaker signal if a stronger one exists in the same band (frequency-domain masking)
  - Discard soft sound after a loud sound (time-domain masking)
  - Stereo redundancy: At low frequencies, we can't detect where the sound is coming from. Encode it mono.
- More on later (MP3, APE…)

Review

- Audio Signals
  - Sampling
    - Sampling Rate, Nyquist Rate, Nyquist Frequency, Nyquist Theorem
  - Quantization
    - Uniform Quantization, SNR, SQNR, Non-uniform Quantization
- Audio file format
  - WAV/MIDI
    - Midi Channel, Midi Messages Format, Midi Terminology, Midi Hardware setup, General Midi
- Human auditory system
    - Critical Band, Masking, Coding

Readings

- Chapter 6.1.1-6.1.6
- Chapter 6.2