


What is Sound?

As the tines move back and forth they exert pressure on the air around them.

(a) The first displacement of the tine compresses the air molecules causing high pressure.

(b) Equal displacement of the tine in the opposite direction forces the molecules to widely disperse themselves and so, causes low pressure.


(c) These rapid variations in pressure over time form a pattern which propagates itself through the air as a wave. Points of high and low pressure are sometimes referred to as 'compression' and 'rarefaction' respectively.



Simple Harmonic Motion -- a Pendulum

- When a pendulum approaches equilibrium it doesn't slow down; it simply travels a smaller distance from the point of rest.
- Any body undergoing simple harmonic motion moves periodically with uniform speed.
- If the tuning fork is moving periodically then the pressure variations it creates will also be periodic.

The time taken to get from position a to b in all three cases is the same




The Unit Circle

These pressure patterns can be represented using as a circle.

Imagine the journey of the pendulum or the tine in four stages:


- from its point of rest to its first point of maximum displacement...
- its first point of maximum displacement back through the point of rest...
- ... to its second point of maximum displacement...
- ... and back from there through its point of rest again

We can map that journey to a circle. This is called the **Unit Circle**. The **sine wave** represents this journey around and around the unit circle *over time*.



Sine Waves


The **sine wave** or **sinusoid** or **sinusoidal signal** is probably the most commonly used graphic representation of sound waves. The diagram below shows one cycle or 'period' of a wave, i.e., the build-up from equilibrium to maximum high pressure, to maximum low pressure, to equilibrium again.



Sine Waves

The specific properties of a sine wave are described as follows.


- Frequency** = the number of cycles per second (this wave has a frequency of 6 hertz)
- Amplitude** = variations in air pressure (measured in decibels)
- Wavelength** = physical length of 1 period of a wave (measured in metres per second)
- = The starting point of a wave along the y-axis (measured in degrees)



Frequency

Frequency refers to the number of cycles of a wave per second. This is measured in **Hertz**. So if a sinusoid has a frequency of 100hz then one period of that wave repeats itself every 1/100th of a second. Humans can hear frequencies between 20hz and 20,000hz (20Khz).

- Frequency is closely related to, *but not the same as!!!*, pitch.
- Frequency does not determine the speed a wave travels at. Sound waves travel at approximately 340metres/second regardless of frequency.
- Frequency is inherent to, and determined by the vibrating body – not the amount of energy used to set that body vibrating. For example, the tuning fork emits the same frequency regardless of how hard we strike it.



Amplitude

Amplitude describes the size of the pressure variations. It is measured along the vertical y-axis. Amplitude is closely related to *but not the same as!!!*, loudness.

(a) Two signals of equal frequency and varying amplitude

(b) Two signals of varying frequency and equal amplitude

Amplitude Envelope

The amplitude of a wave changes or 'decays' over time as it loses energy.

These changes are normally broken down into four stages; **Attack**, **Decay**, **Sustain** and **Release**.

Collectively, the four stages are described as the **amplitude envelope**.

Introduction to Sampling

What is SAMPLING?

- Process by which an analog signal is measured in order to convert the analog signal to digital.


Sampling

- Value of the analog signal is read at evenly spaced time intervals.
- Sample rate (frequency) is measured in kilohertz.
- 1 kHz = 1,000 cps.
- (Cycles per second).

Sampling


Quantization

- The digital signal is defined only at the points at which it is sampled.




Quantization

- The height of each vertical bar can take on only certain values, shown by horizontal dashed lines, which are sometimes higher and sometimes lower than the original signal, indicated by the dashed curve.



Quantization

- If the graphic has 11 quantization levels, how many bits are needed to encode each sample?



Quantization

- 4 bits... why?
- 1 bit would allow up to 2 levels
- 2 bits would allow up to 4 levels
- 3 bits would allow up to 8 levels
- 4 bits would allow up to 16 levels

Quantization

- The difference between a quantized representation and an original analog signal is called the *quantization noise*.
- The more bits for quantization of a signal, the more closely the original signal is reproduced.

Quantization

- Using higher sampling frequency and more bits for quantization will produce better quality digital audio.
- But for the same length of audio, the file size will be much larger than the low quality signal.

Quantization

- The number of bits available to describe sampling values determines the resolution or accuracy of quantization.
- For example, if you have 8-bit analog to digital converters, the varying analog voltage must be quantized to 1 of 256 discrete values;
- a 16-bit converter has 65,536 values.

Nyquist Theorem

- A theorem, developed by Harry Nyquist, which states that an analog signal waveform may be uniquely reconstructed, without error, from samples taken at equal time intervals.

Nyquist Theorem

- The sampling rate must be equal to, or greater than, twice the highest frequency component in the analog signal.

Nyquist Theorem

- Stated differently:
- The highest frequency which can be accurately represented is one-half of the sampling rate.

Error


- Sampling an analog signal can introduce ERROR.
- ERROR is the difference between a computed, estimated, or measured value and the true, specified, or theoretically correct value.

Nyquist Theorem

- By sampling at TWICE the highest frequency:
 - One number can describe the positive transition, and...
 - One number can describe the negative transition of a single cycle.


Nyquist Theorem

The vertical lines are sample intervals, and the white dots are the crossing points - the actual samples taken by the conversion process.




Nyquist Theorem

The sampling rate was below the Nyquist frequency, so the reconstructed waveform does not accurately reproduce the original:



Nyquist Theorem




This under-sampling results in **aliasing** which shows up as noise in digitized sound.

To correct the aliasing, A/D converters use lowpass filters to remove all signals above the Nyquist frequency.

To eliminate aliasing and to get high-fidelity sound, use a high sample rate.

upper => sampling 6 times per cycle ($f_s = 6f$);
middle => sampling 3 times per cycle ($f_s = 3f$);
lower => sampling 6 times in 5 cycles, from [1]



Audio Synthesis Basics

Analog Synthesis
Intro to Digital Oscillators

Analog Synthesis Overview

- Sound is created by controlling electrical current within synthesizer, and amplifying result.
- Basic components:
 - Oscillators
 - Filters
 - Envelope generators
 - Noise generators
- Voltage control



Oscillators

- Creates periodic fluctuations in current, usually with selectable waveform.
- Different waveforms have different harmonic content, or *frequency spectra*.



Filters

- Given an input signal, attenuate or boost a frequency range to produce an output signal
- Basic Types:
 - Low pass
 - High pass
 - Band pass
 - Band reject (notch)



Envelope Generators

- Generate a control function that can be applied to various synthesis parameters, including amplitude, pitch, and filter controls.



Noise Generators

- Generate a random, or semi-random fluctuation in current that produces a signal with all frequencies present.



Digital Synthesis Overview

- Sound is created by manipulating numbers, converting those numbers to an electrical current, and amplifying result.
- Numerical manipulations are the same whether they are done with software or hardware.
- Same capabilities (components) as analog synthesis, plus significant new abilities



Digital Oscillators

- Everything is a Table
 - A table is an indexed list of elements (or values)
 - The index is the address used to find a value

Generate a Sine Tone Digitally (1)

- Compute the sine in real time, every time it is needed.
 - equation:

$$signal(t) = r \sin(\omega t)$$

- t = a point in time; r = the radius, or amplitude of the signal; w (ω) = $2\pi \times f$ the frequency
- Advantages: It's the perfect sine tone. Every value that you need will be the exact value from the unit circle.
- Disadvantages: must generate every sample of every oscillator present in a synthesis patch from an algorithm. This is very expensive computationally, and most of the calculation is redundant.

Generate a Sine Tone Digitally (2)

- Compute the sine tone once, store it in a table, and have all oscillators look in the table for needed values.
 - Advantages: Much more efficient, hence faster, for the computer. You are not, literally, re-inventing the wheel every time.
 - Disadvantages: Table values are discrete points in time. Most times you will need a value that falls somewhere in between two already computed values.

Table Lookup Synthesis

- Sound waves are very repetitive.
- For an oscillator, compute and store one cycle (period) of a waveform.
- Read through the wavetable repeatedly to generate a periodic sound.

Changing Frequency

- The Sample Rate doesn't change within a synthesis algorithm.
- You can change the speed that the table is scanned by skipping samples.
- skip size is the increment, better known as the phase increment.
phase increment is a very important concept

Algorithm for a Digital Oscillator

- Basic, two-step program:
 - $phase_index = \text{mod}_i(\text{previous_phase} + \text{increment})$
 - $output = \text{amplitude} \times \text{wavetable}[phase_index]$
- $increment = \frac{\text{TableLength} \times \text{DesiredFrequency}}{\text{SampleRate}}$

If You're Wrong, it's Noise


- What happens when the phase increment doesn't land exactly at an index location in the table?
 - It simply looks at the last index location passed for a value. In other words, the phase increment is truncated to the integer.
- Quantization
- Noise
- The greater the error, the more the noise.

Interpolation

- Rather than truncate the phase location...
 - look at the values stored before and after the calculated phase location
 - calculate what the value would have been at the calculated phase location if it had been generated and stored.
- Interpolate
- More calculations, but a much cleaner signal.

Linear interpolation


- Interpolate between two audio samples
 - ```
double inbetween = fmod(sample, 1);
return (1. - inbetween) * wave[int(sample)] +
 inbetween * wave[int(sample) + 1];
```
- More accurate, yet still efficient



## Envelopes


- We commonly will make samples with fixed amplitudes, then make a synthetic *envelope* for the sound event.

## Attack and Release




## ADSR

- ADSR: Attack, decay, sustain, release



## PWM: Pulse Width Modulation



- Signal is compared to a sawtooth wave producing a pulse width proportional to amplitude



## What Can Be Done With PWM?

- Question: What happens if voltages like the ones above are connected to an LED?
- Answer: The longer the duty cycle, the longer the LED is on and the brighter the light.

## What Can Be Done With PWM?

- Average power can be controlled
- Average flows can also be controlled by fully opening and closing a valve with some duty cycle

## PULSE WIDTH MODULATION

- Pulse Width Modulation (PWM) involves the generation of a series of pulses at a fixed period and frequency.
- The duty cycle defines the width of each pulse which is varied to generate waveforms.

## PULSE WIDTH MODULATION

- A simple low pass filter is then used to generate an output voltage directly proportional to the average time spent in the HIGH state.
- (i.e., 50% duty cycle is equal to 2.5 volts when  $V_{DD} = 5.0V$ ).

## RC low pass filter (integrator)

Choosing the -3 dB point at 4 kHz, and using the Relation:

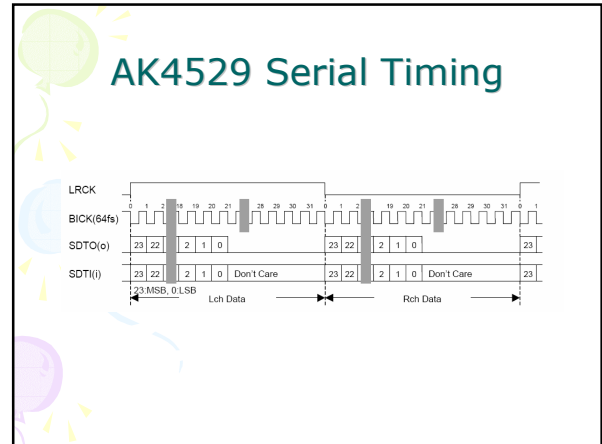
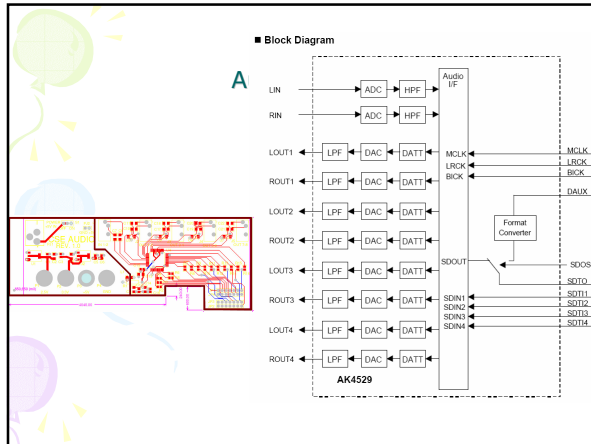
$$RC = 1/(2 \cdot \pi \cdot f)$$

we get  $R = 4\text{ k}$ , if  $C$  is chosen as  $0.01\text{ mF}$ :

$R = 4.0\text{ k}$   
 $C = 0.01\text{ mF}$

## OPB Audio Controller

- The opb\_audio\_controller unit is used to communicate with the audio boards mounted on the AFX BG560 boards. The chip requires that all communications with it be in a serial format.
- This controller allows you to have memory mapped I/O to the component so that you can build software to communicate with the audio chip.



### Audio Controller Register Map

| BASEADDR + 0x00 | Control register      | R/W |
|-----------------|-----------------------|-----|
| BASEADDR + 0x04 | Input Left channel    | R   |
| BASEADDR + 0x08 | Input Right channel   | R   |
| BASEADDR + 0x0C | Output1 Left channel  | R/W |
| BASEADDR + 0x10 | Output1 Right channel | R/W |

- All registers are 32 bits wide, however, only the lower order 24 bits are used in the input and output registers. They are signed 24-bit values, however they are sign extended to 32-bits when you read them.

### Audio Controller Register Map

| BASEADDR + 0x00 | Control register      | R/W |
|-----------------|-----------------------|-----|
| BASEADDR + 0x04 | Input Left channel    | R   |
| BASEADDR + 0x08 | Input Right channel   | R   |
| BASEADDR + 0x0C | Output1 Left channel  | R/W |
| BASEADDR + 0x10 | Output1 Right channel | R/W |

- There are only 2 flags in the control register. The lowest order bit (0x00000001) is the enable bit and enables the codec. It is tied to the PDN pin. The next bit (0x00000002) is the interrupt enable. When this is high, an interrupt is generated every time the codec is ready for a new sample. The interrupt is cleared by writing or reading from any register in the codec.

### OPB Audio Controller

- The codec requires several different clocks so this component contains a DLL to generate these clock signals.
- All other component's clock inputs to use an internal net that is connected to the SYS\_CLK output of the opb\_audio\_controller.
- Connect the XTAL\_CLK input of the opb\_audio\_controller to the off-chip crystal clock source (probably an external net that gets connected to AL17).

### Lab- Interrupt routine

```

void audio_interrupt_handler(void *InstancePtr)
{
 /*
 * TODO: calculate the next sample and give it to the audio controller.
 * Note that this interrupt will happen every 23 microseconds, or
 * at 43.4kHz (the sample rate of the codec)
 */

 /* this currently makes a triangle wave */
 if(curr_value == HIGH_VALUE) climbing = 0;
 if(curr_value == LOW_VALUE) climbing = 1;
 if(climbing)
 {
 curr_value++;
 }
 else
 {
 curr_value--;
 }

 XAudio_mWriteOutput(XPAR_AUDIO_BASEADDR, LEFT, 0, curr_value << SHIFT_AMOUNT);
}

```

## Lab- Main

```
int main()
{
 /* TODO: initialization code should go here */

 curr_value = 0;
 climbing = 1;

 /* register for the interrupts */
 XIntc_InterruptVectorTable[0].Handler = audio_interrupt_handler;
 XIntc_InterruptVectorTable[0].CallBackRef = NULL;
 XIntc_mEnableIntr(XPAR_INTC_SINGLE_BASEADDR, XPAR_AUDIO_INTERRUPT_MASK);
 XIntc_mMasterEnable(XPAR_INTC_SINGLE_BASEADDR);

 /* globally enable the interrupts on the microblaze */
 microblaze_enable_interrupts();

 /* enable the audio codec and make it interrupt me every time it wants new data */
 XAudio_mSetControlReg(XPAR_AUDIO_BASEADDR, AUDIO_CR_INT_ENABLE_MASK
 AUDIO_CR_ENABLE_MASK);

 for(;;)
 {
 /* do nothing, just let the interrupts handle the rest of the work */
 }
 return 0; /* never reached */
}
```

## Tidbits

- The sample rate of the codec is 43.4kHz. Use a table length of 256 entries.
- The Microblaze has no floating point operations. If you use floating point, the compiler will emulate it, however, it will be extremely slow, so you should not use floating point.
- The Microblaze does not have a hardware multiplier, so multiplies are done in software. As a result, they are very slow. You probably have about enough time between samples to do about 2 multiplies, maybe 3.
- Don't even think about doing a divide by a number other than a power of 2 (bit shift).
- The audio codec takes 24-bit signed numbers, centered at 0. This means that:
  - The highest value it can receive is  $2^{23} - 1$ , or 8388607, or 0x007FFFFF.
  - The lowest value it can receive is  $-(2^{23})$  or -8388608, or 0xFF800000.