



Google  
Developers



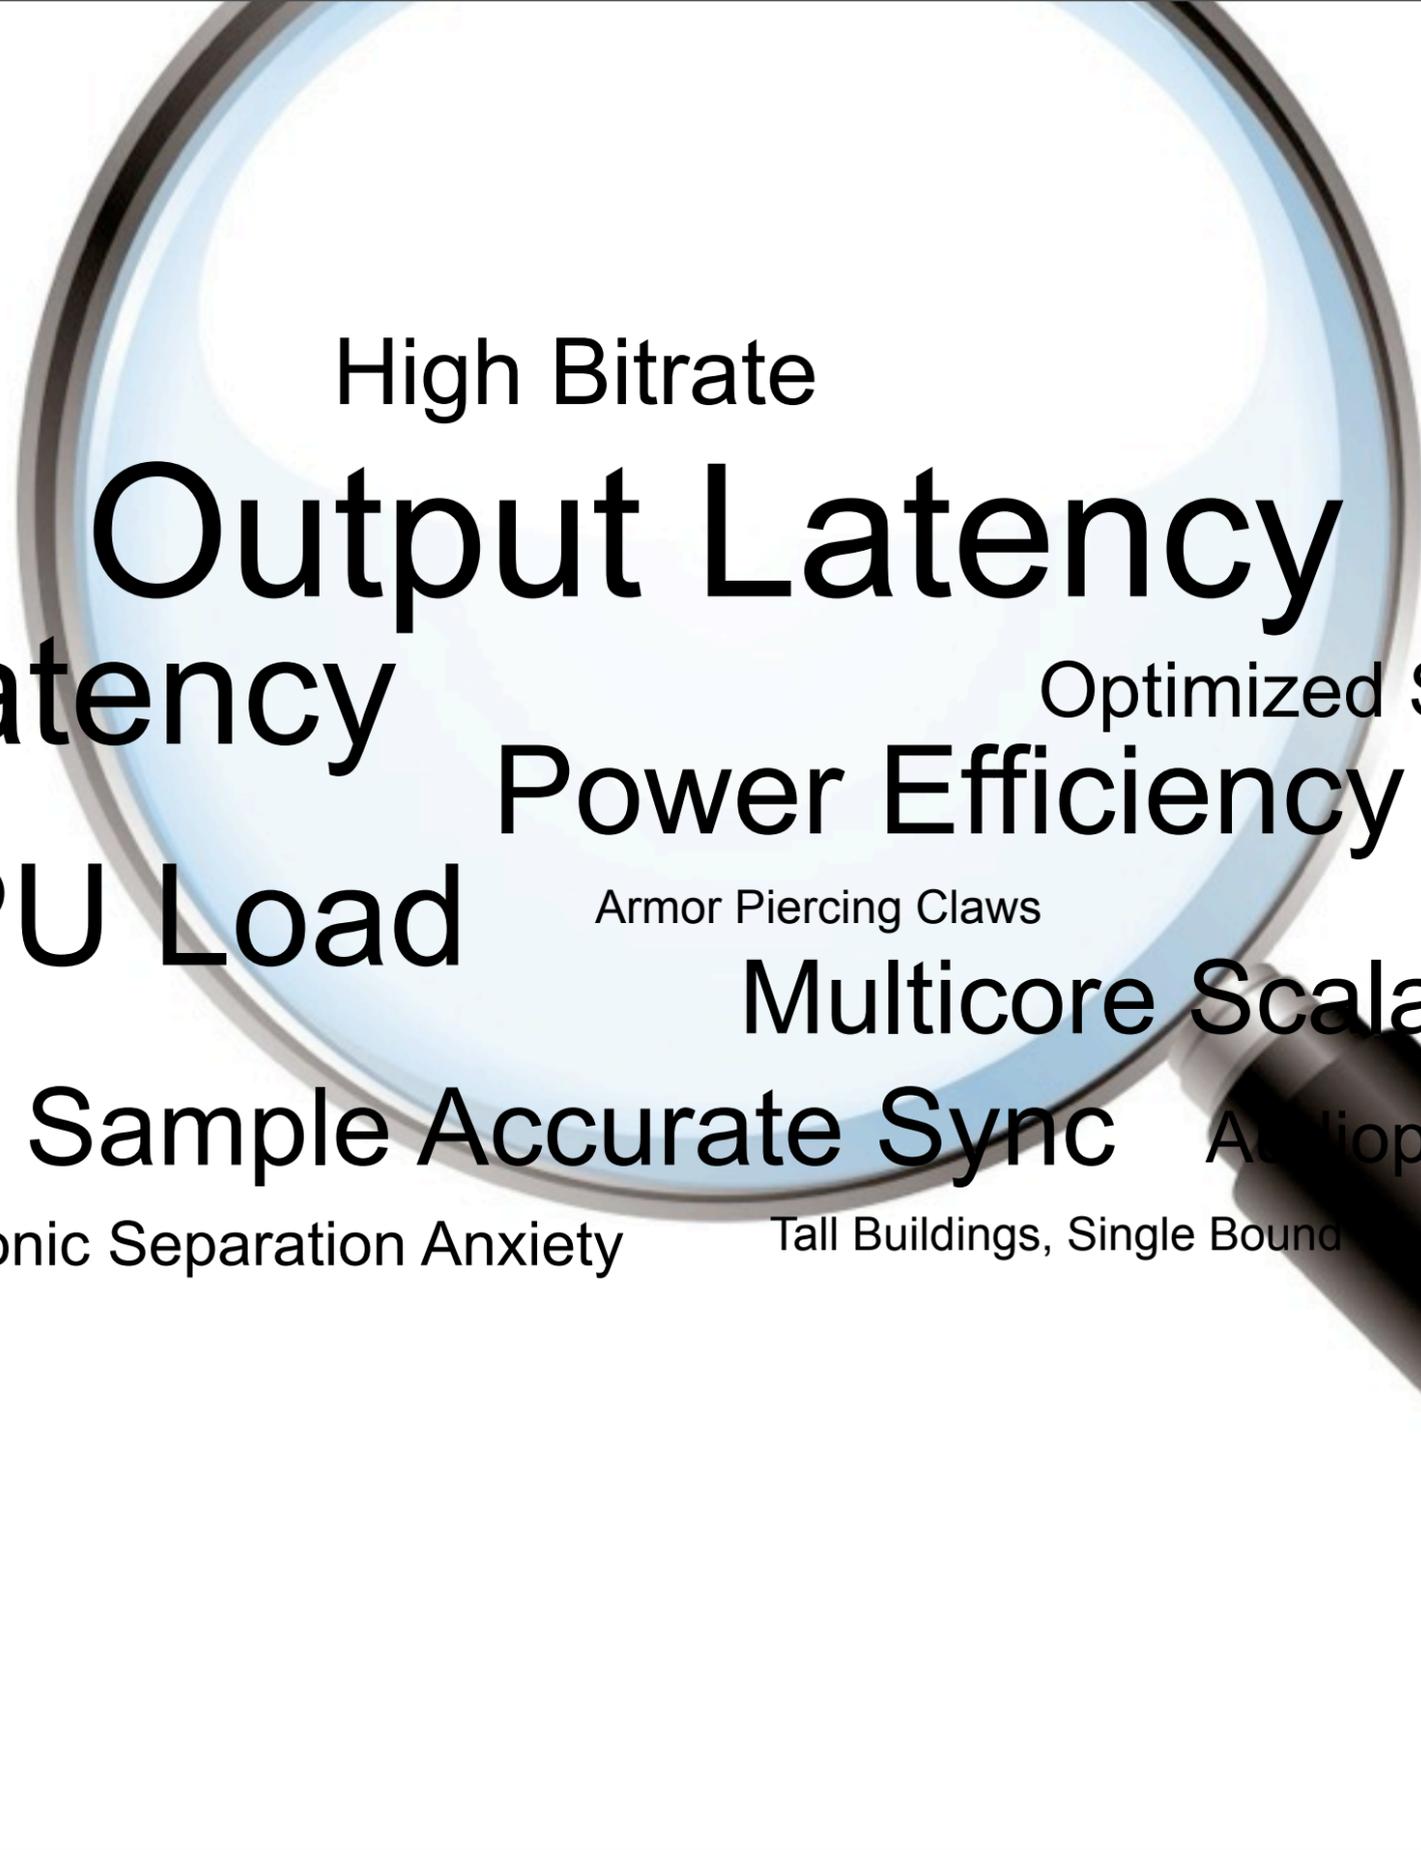




# High Performance Audio on Android

Glenn Kasten, Raph Levien, Ian Ni-Lewis





High Bitrate

**Output Latency**

Almost Negative

**Input Latency**

Optimized Signal Processing

**Power Efficiency**

**Low CPU Load**

Armor Piercing Claws

High Speed Ultrasound

**Multicore Scalability**

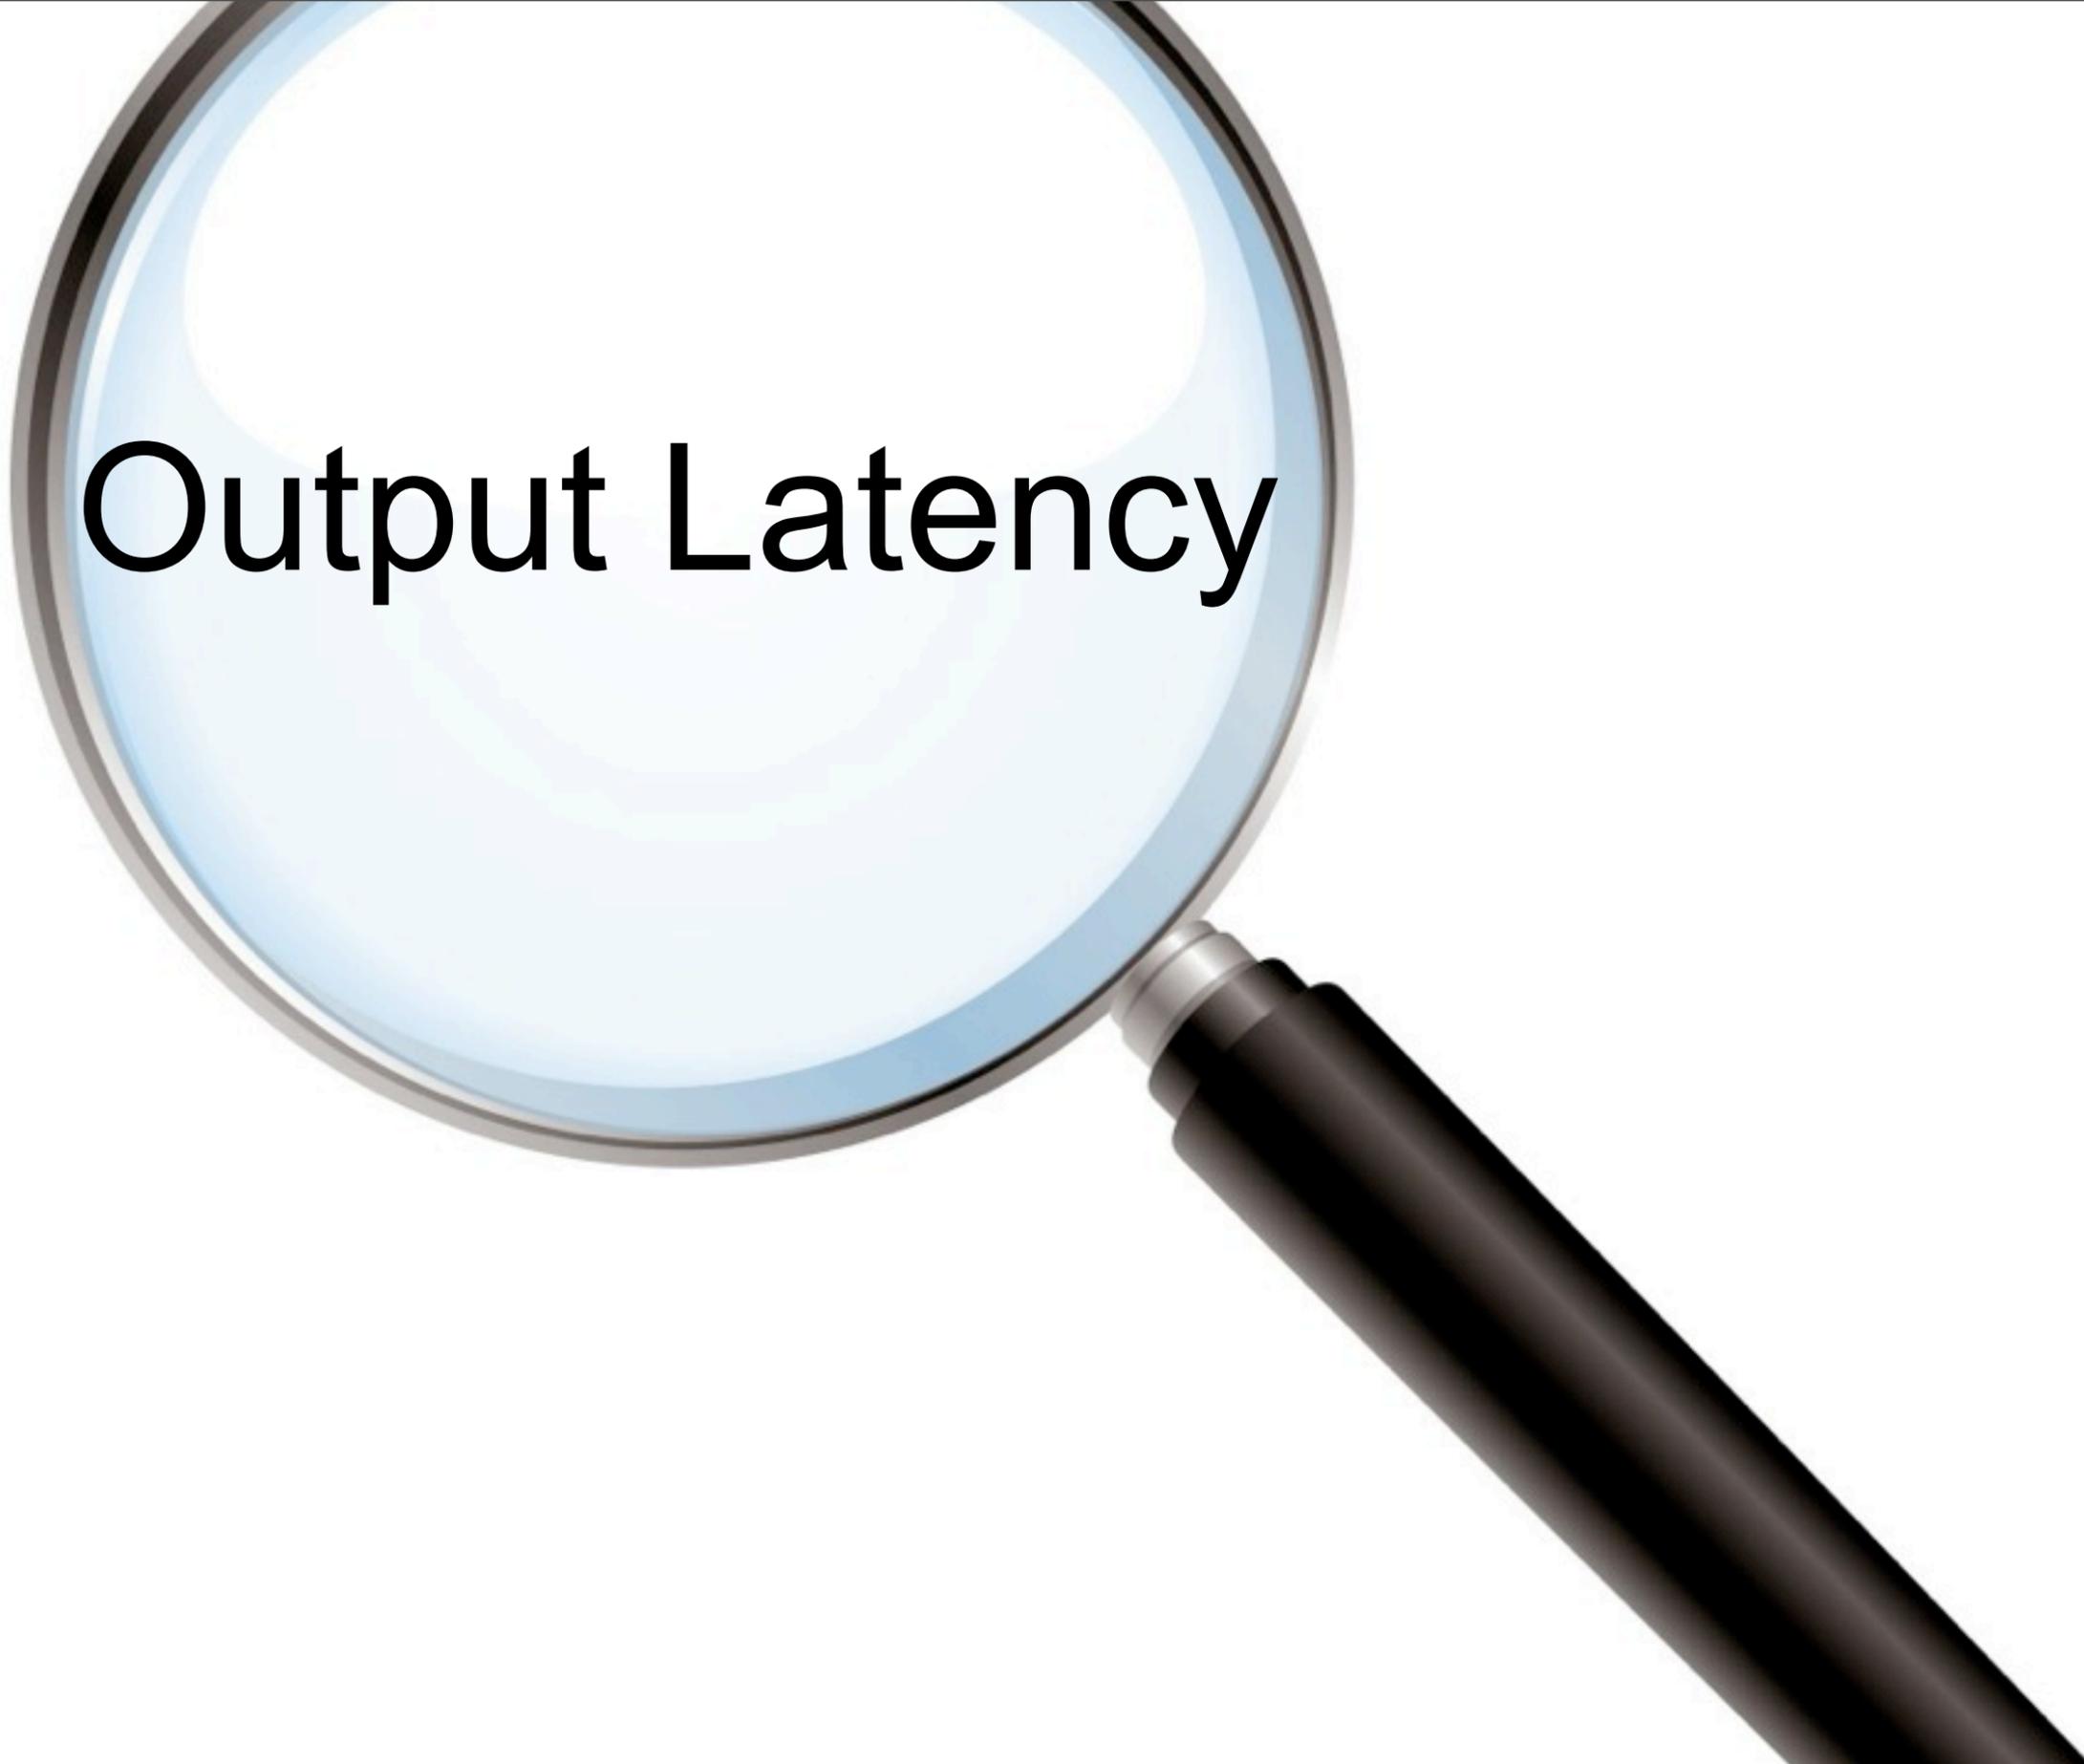
**Sample Accurate Sync**

Audiophile-quality output

True Stereophonic Separation Anxiety

Tall Buildings, Single Bound

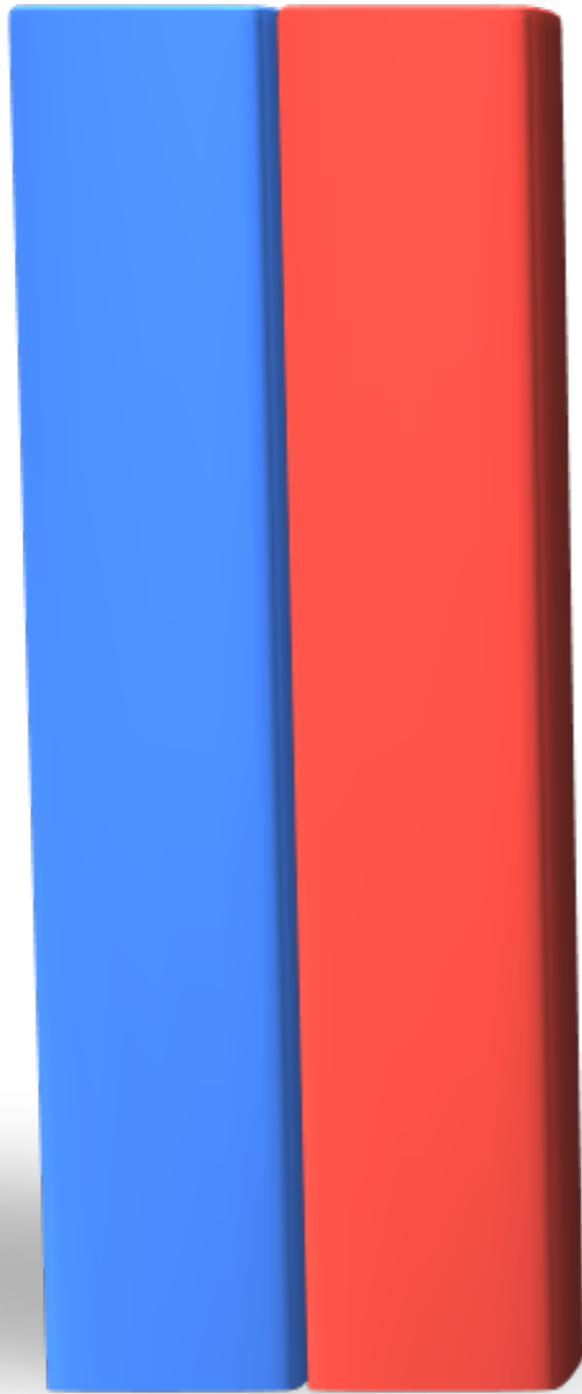




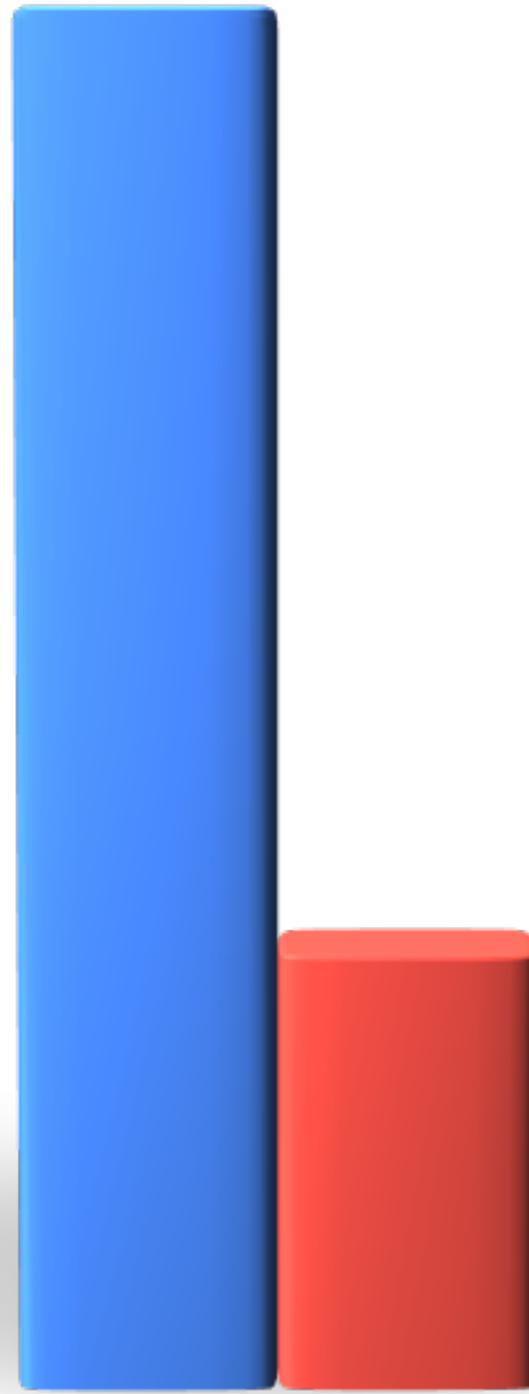
# Output Latency



*Approximate Latency (ms)*



Where we were



Where we are



Where we're going



# The Story



# The Story





National Cancer Institute/Wikimedia Commons

**The latency**

**is too darn high**

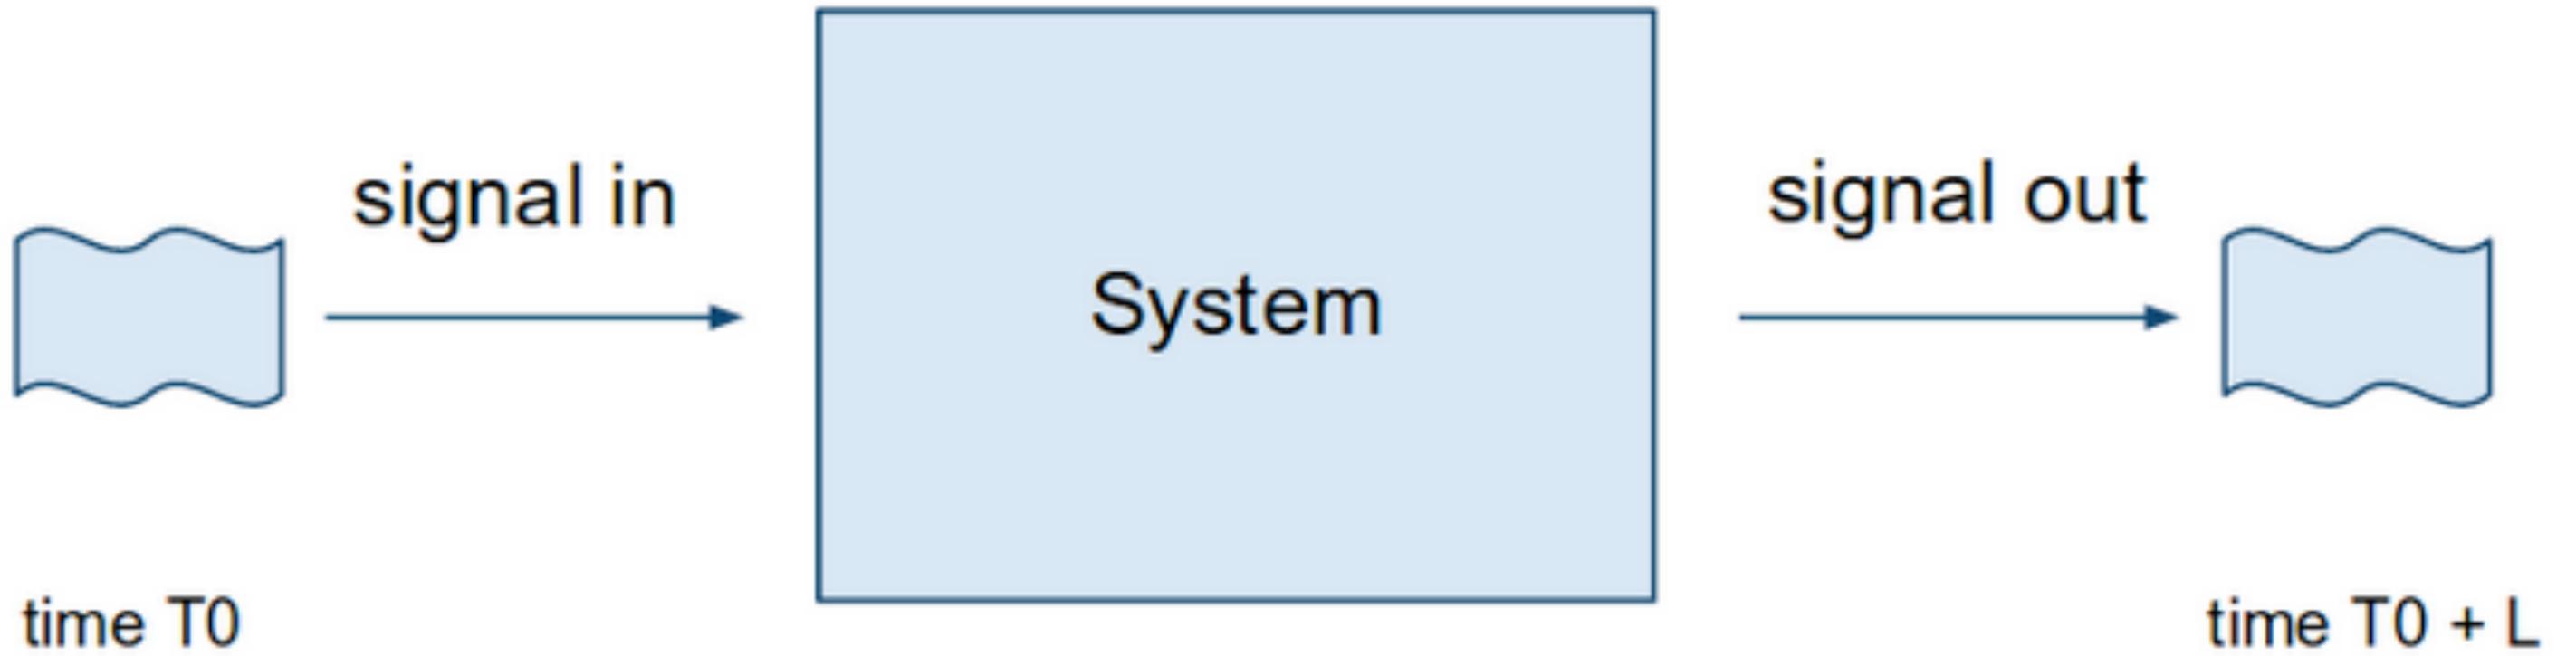


**The latency**

**is too darn high**

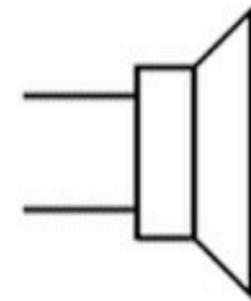
**( ISSUE 3434 )**





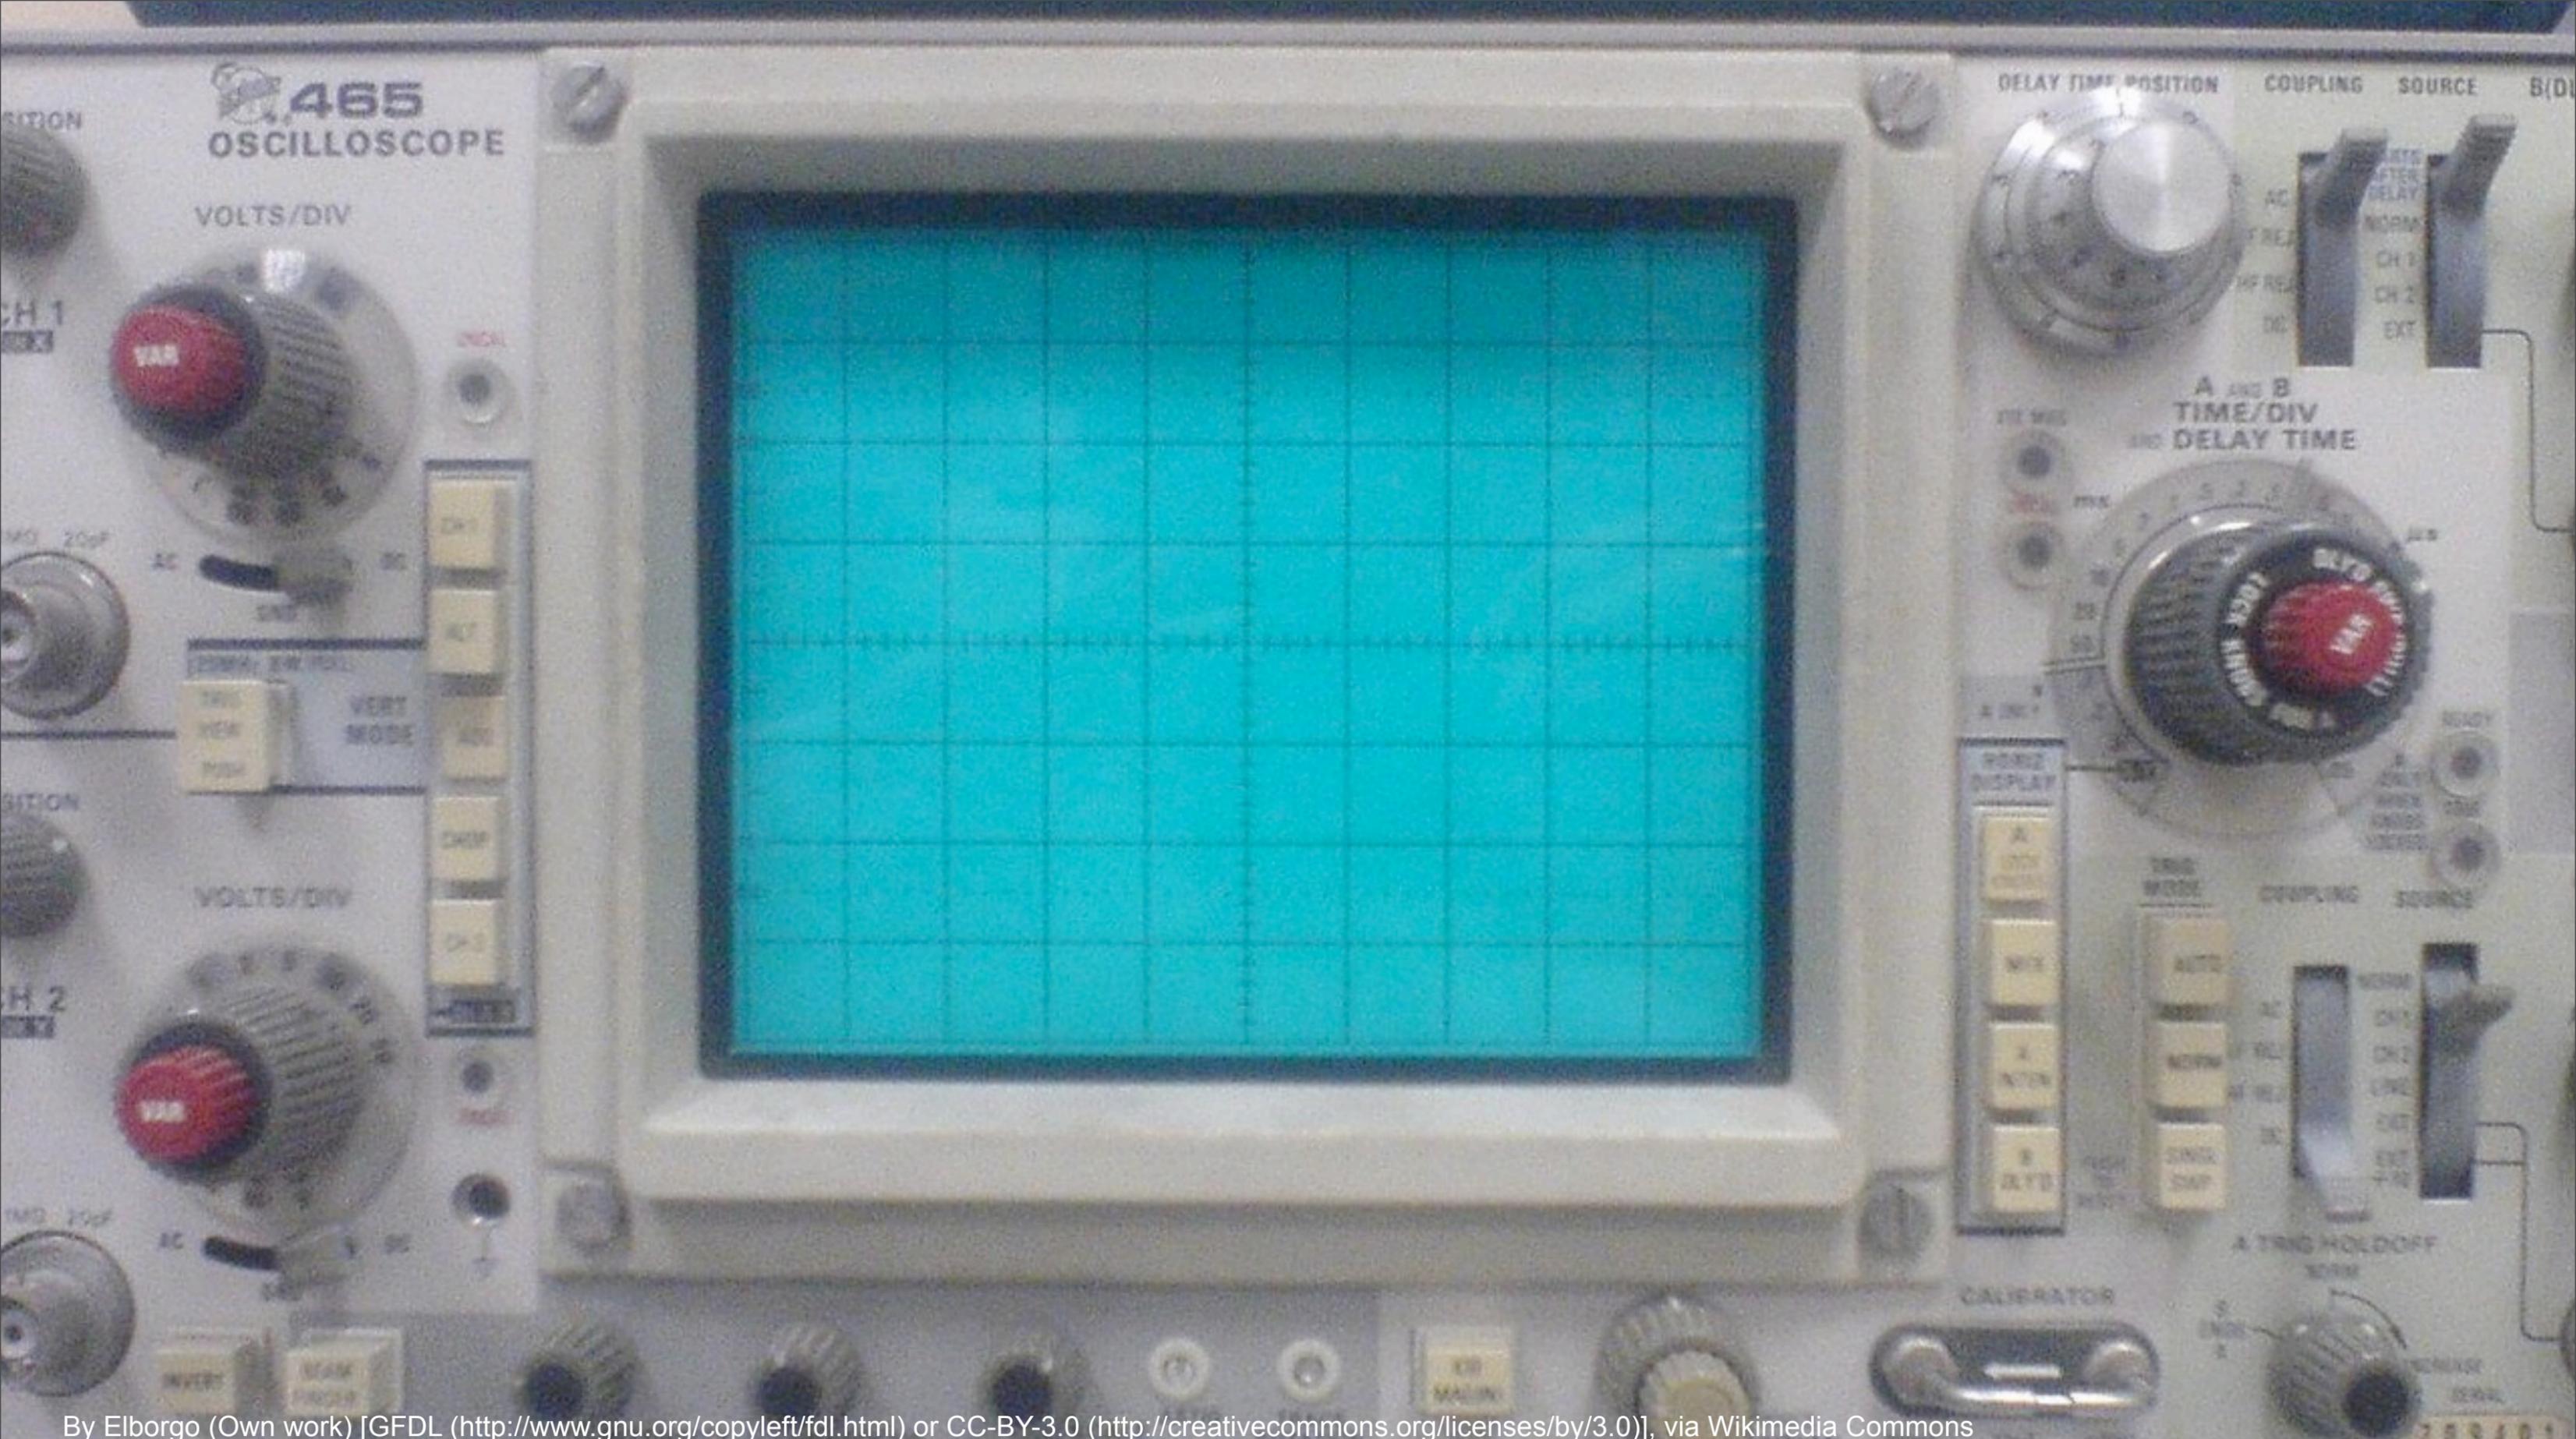
# App

```
AudioTrack at = ...;  
short[] data = ...;  
for (int i = 0; i < data.length; ++i)  
    short[i] = (i * 1000) % 10000 - 5000;  
at.write(data, 0, data.length);
```

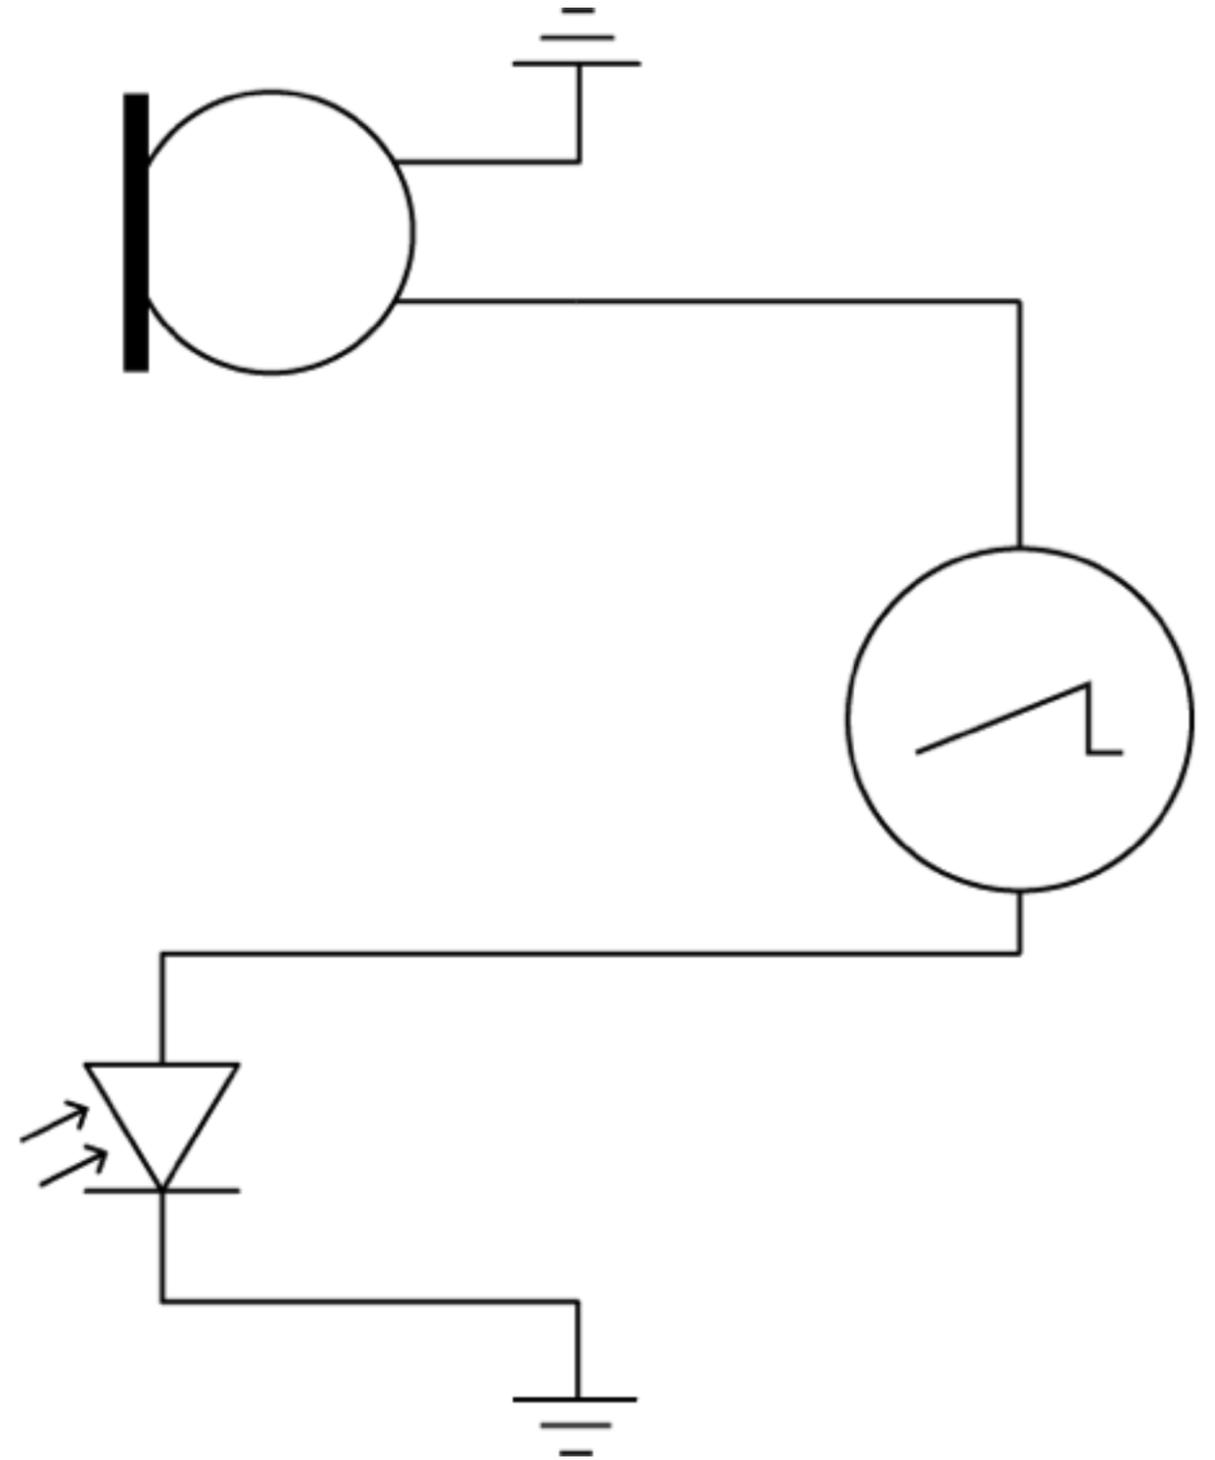


speaker

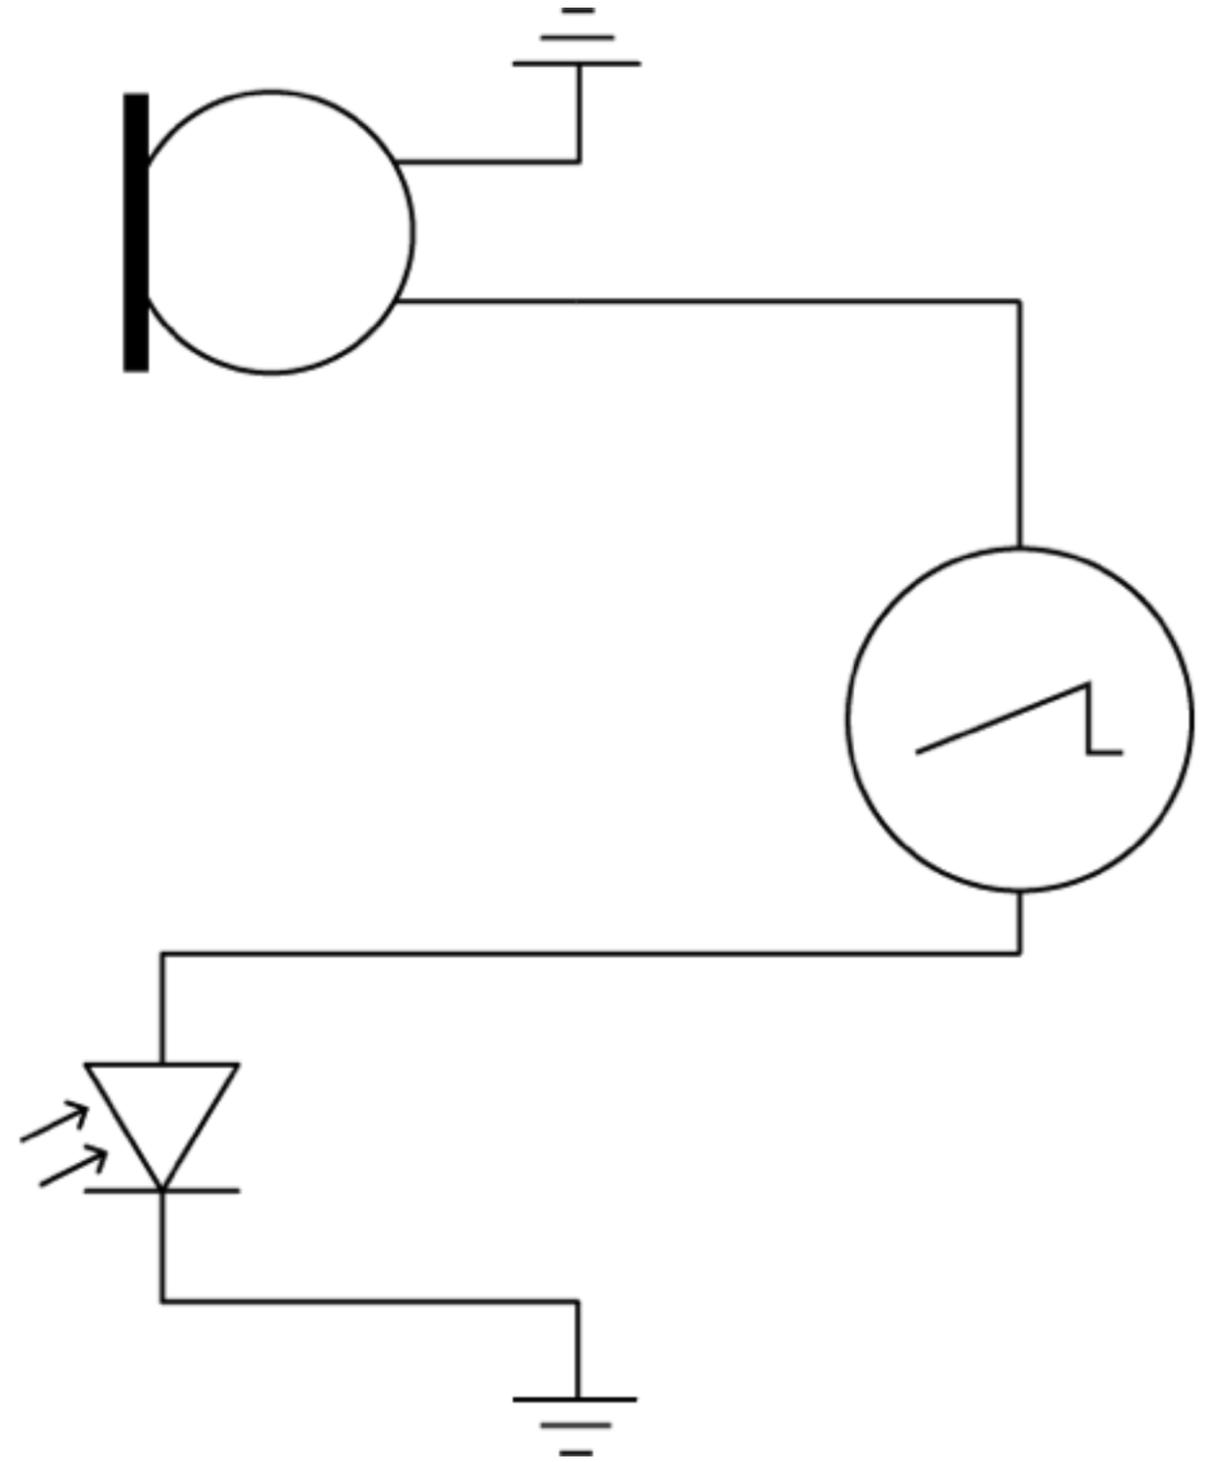
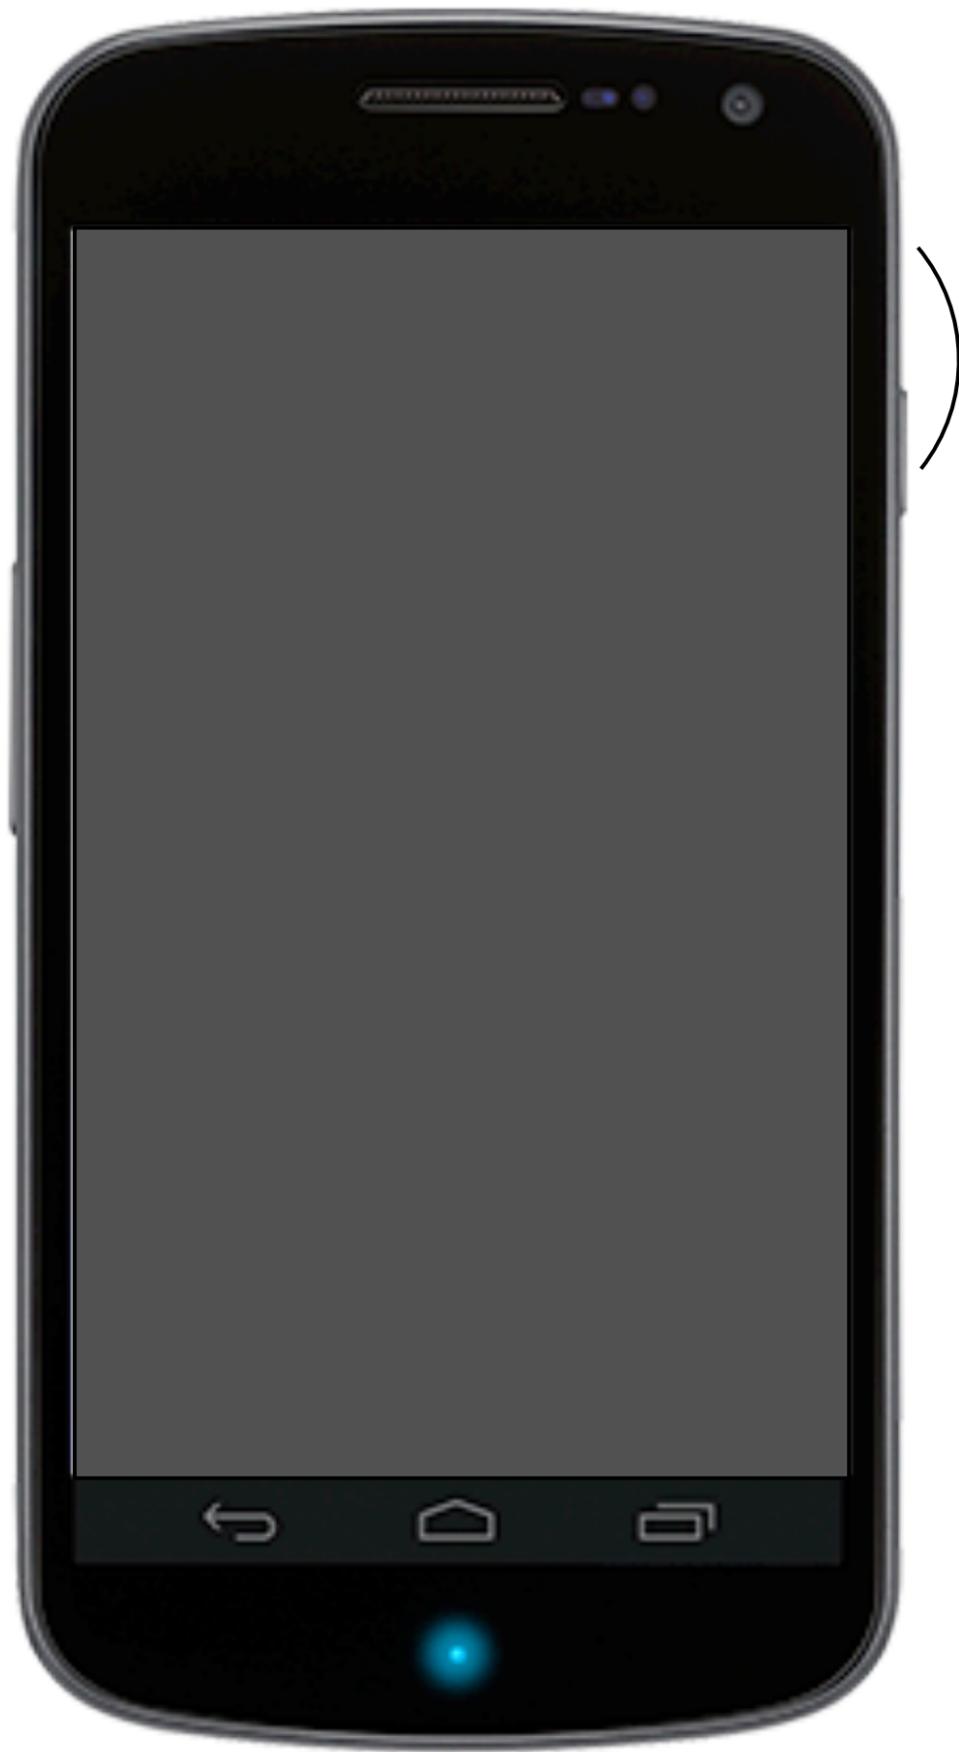




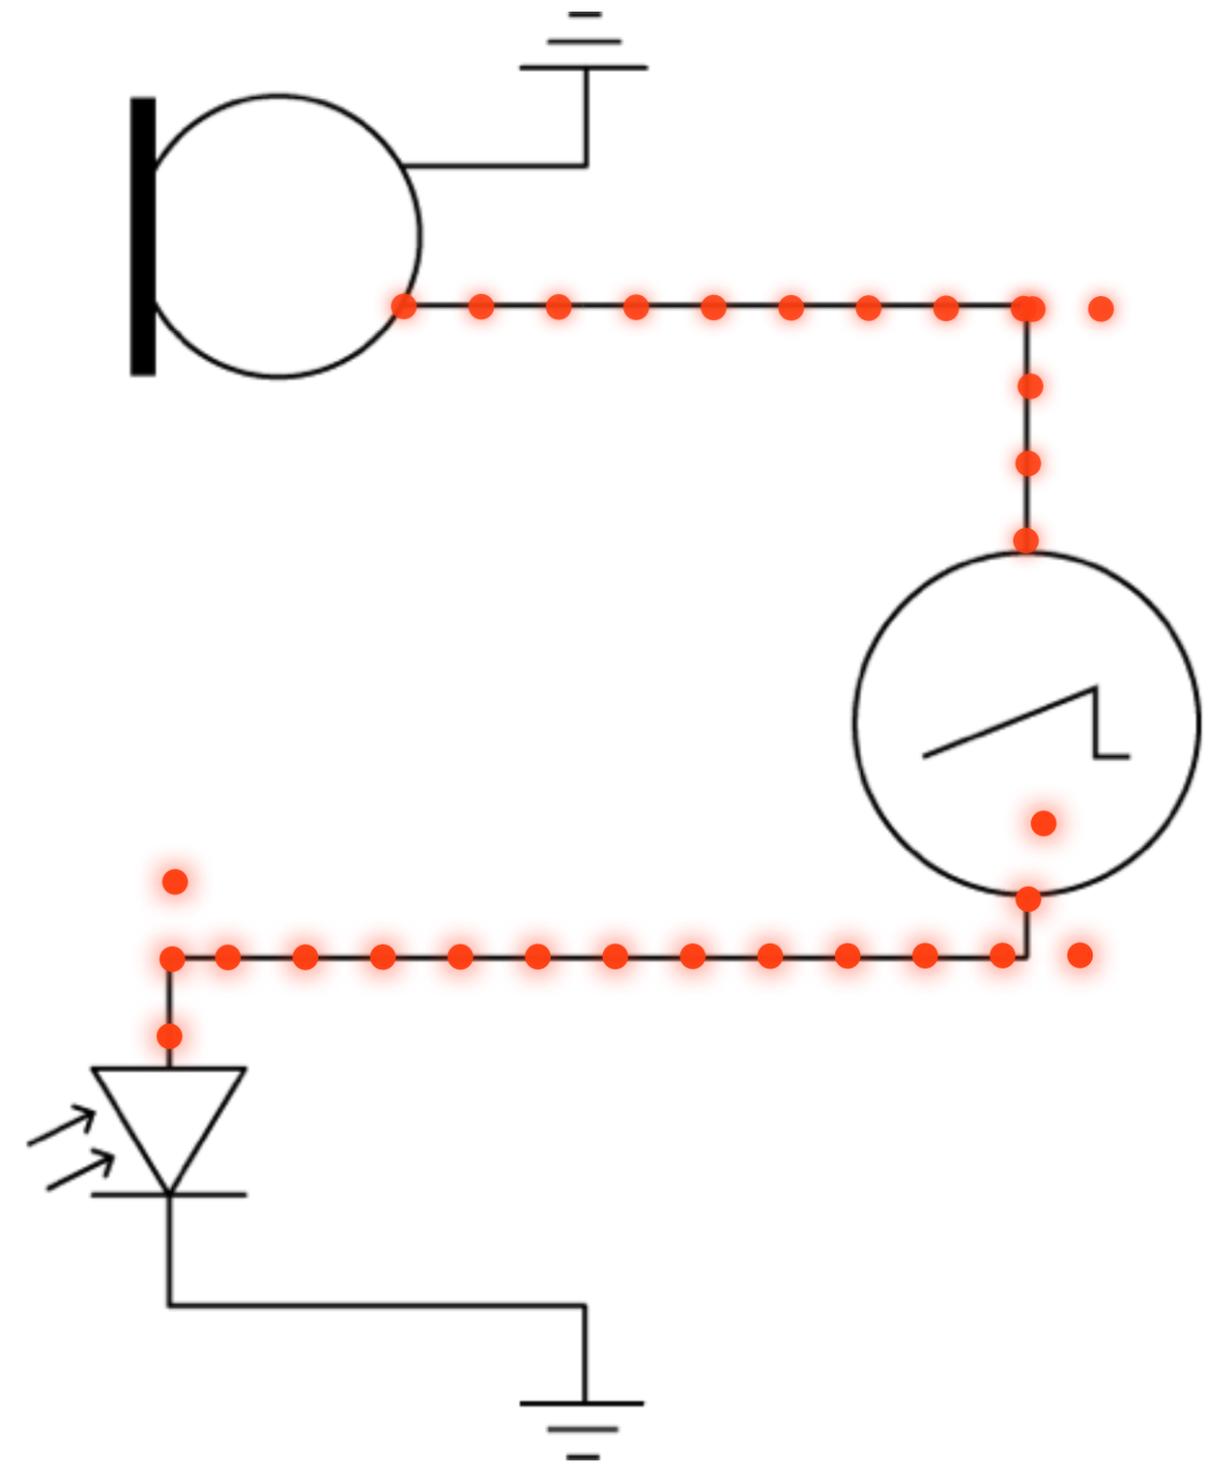
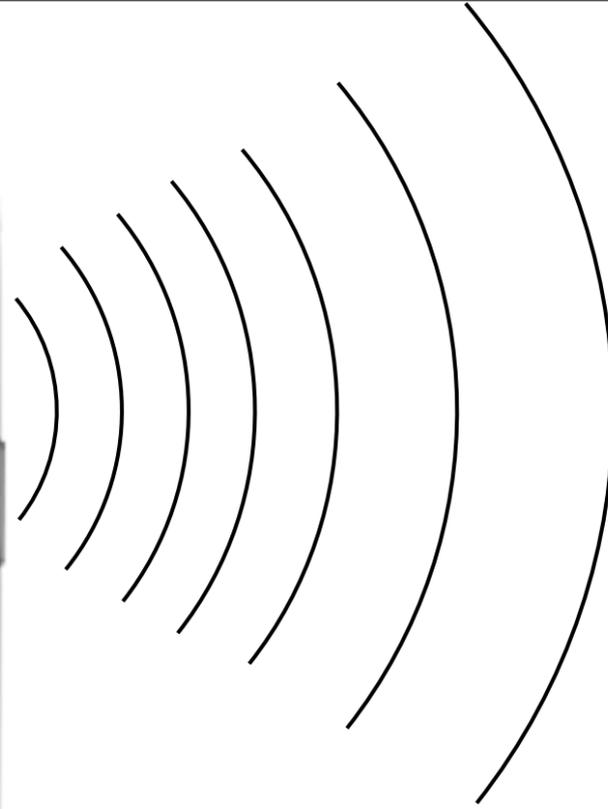
By Elborgo (Own work) [GFDL (<http://www.gnu.org/copyleft/fdl.html>) or CC-BY-3.0 (<http://creativecommons.org/licenses/by/3.0/>)], via Wikimedia Commons



photodiode: Omegatron/Wikimedia Commons  
Microphone: J.P. Lon / Wikimedia Commons  
Oscilloscope: MovGPO / Wikimedia Commons (de)  
Galaxy Nexus: R. Nurik, I. Ni-Lewis / Google

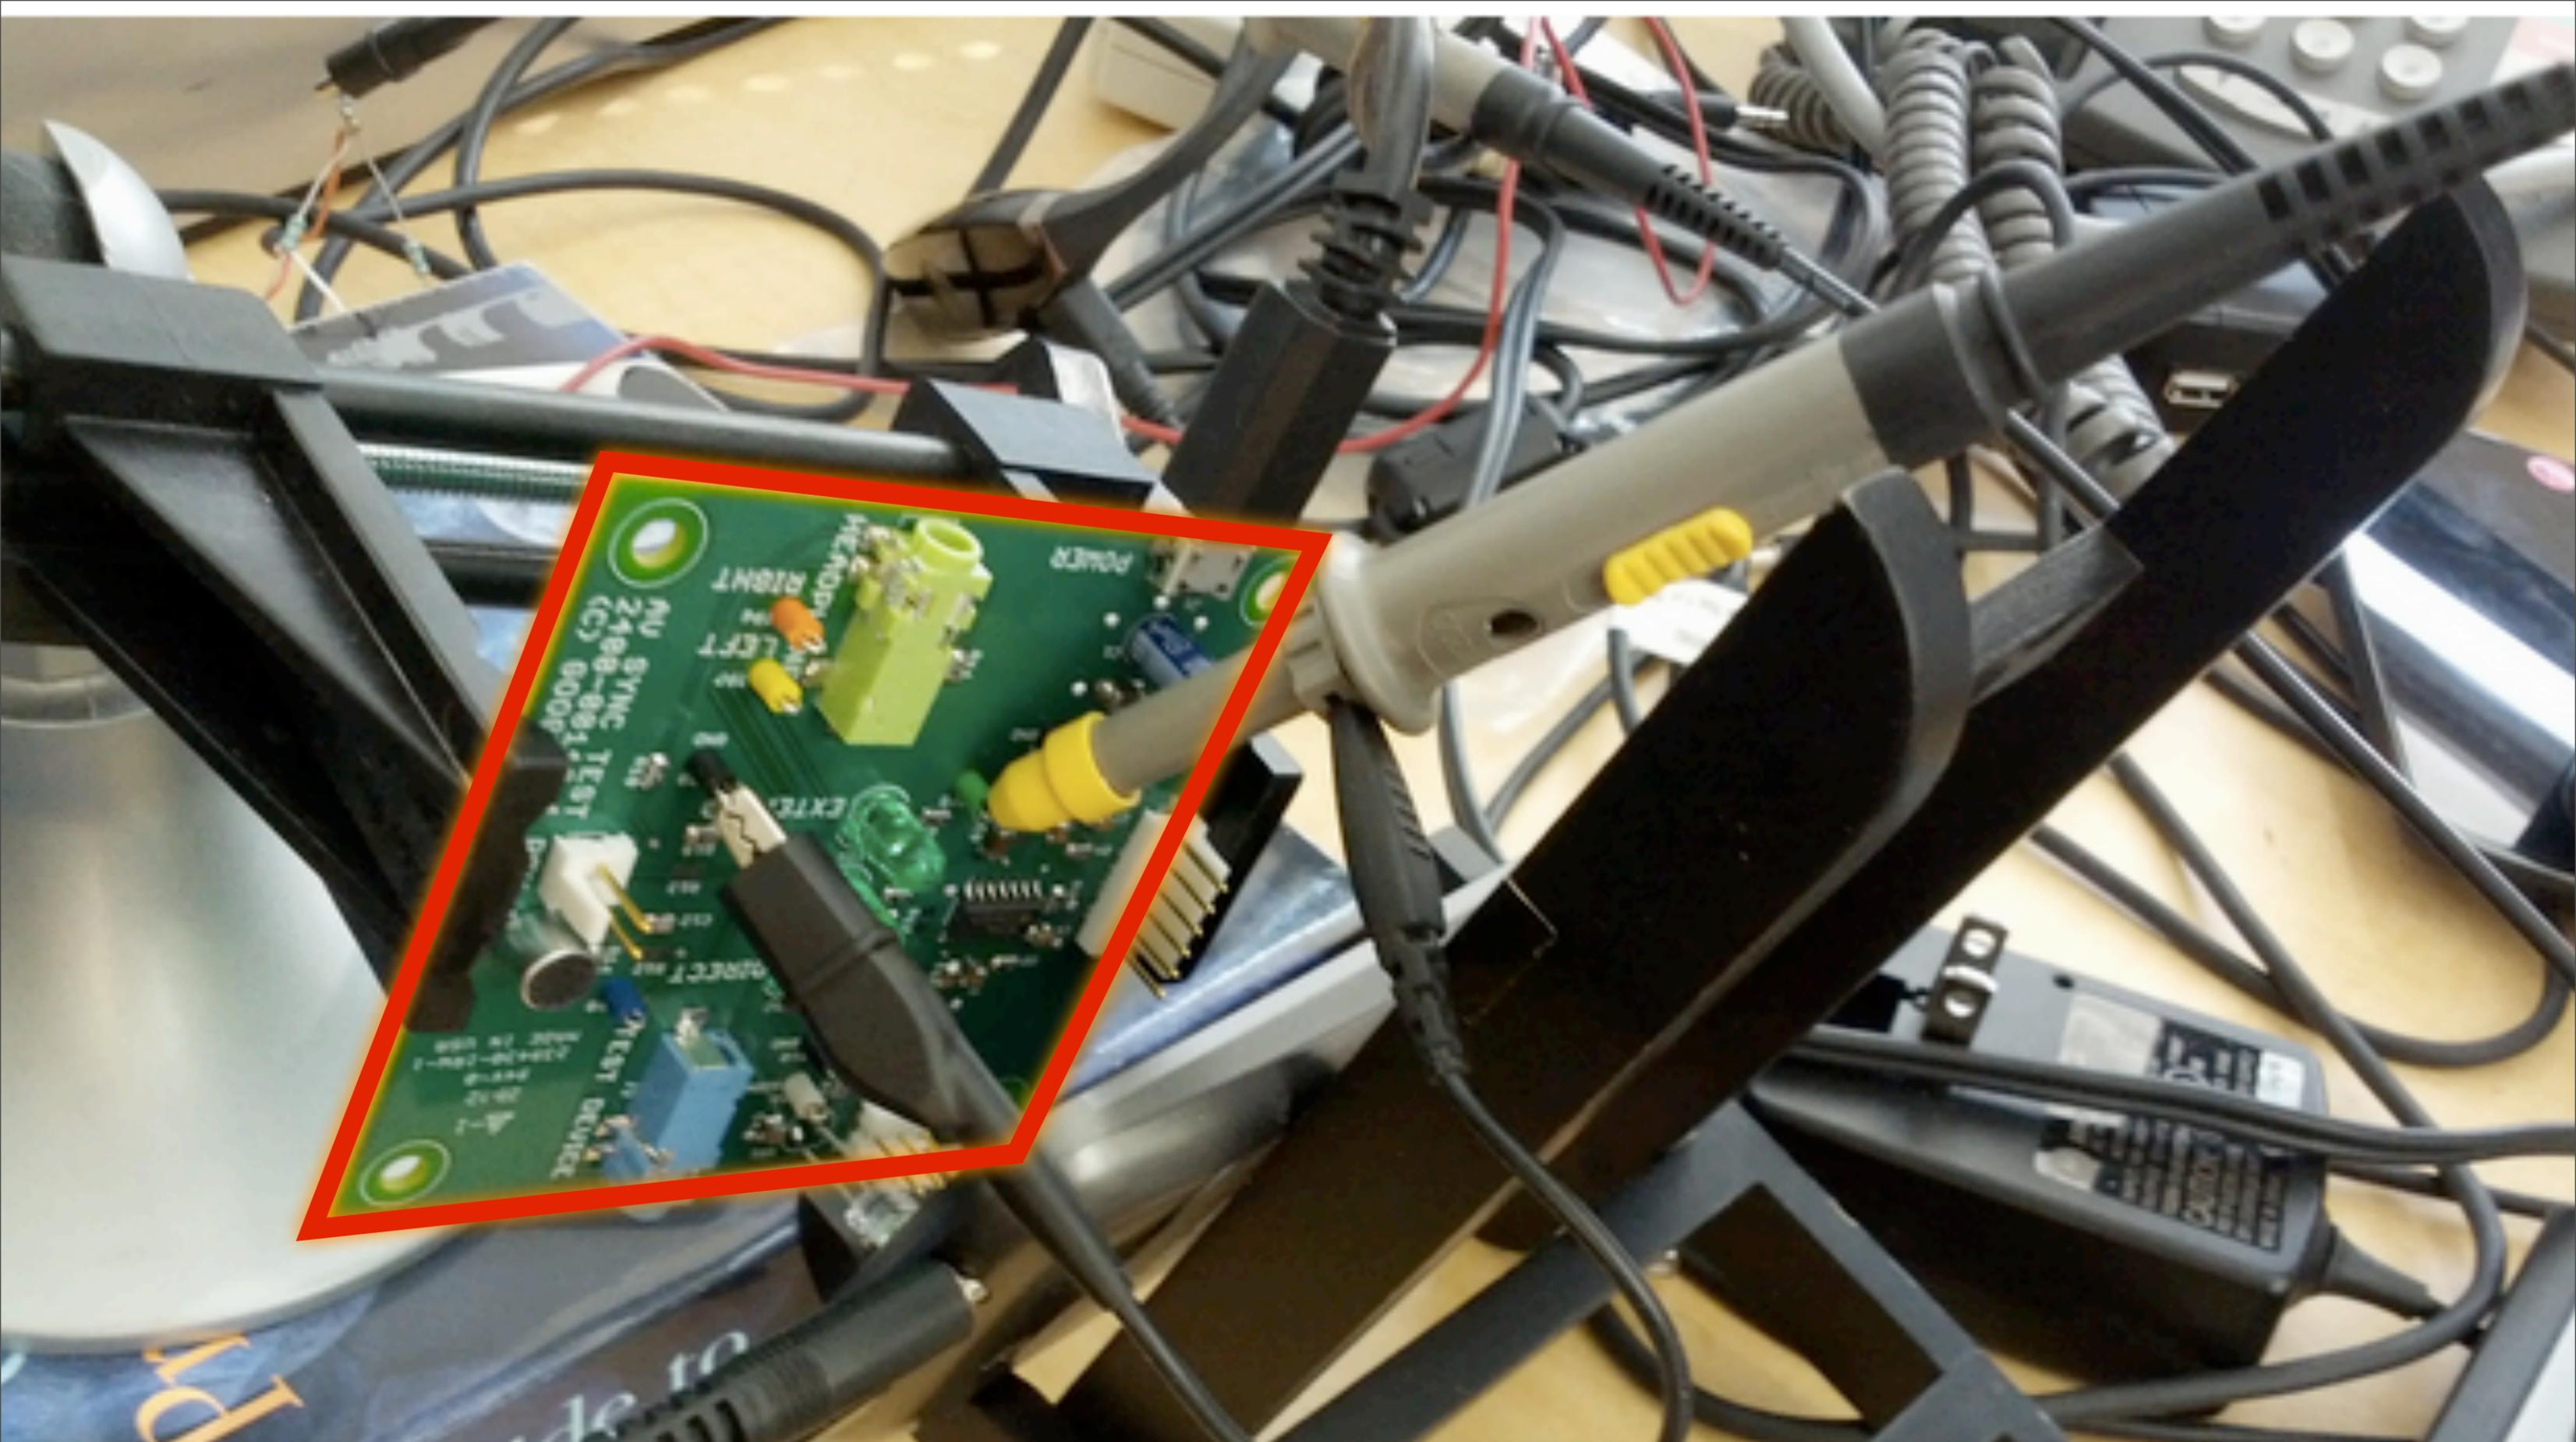


photodiode: Omegatron/Wikimedia Commons  
Microphone: J.P. Lon / Wikimedia Commons  
Oscilloscope: MovGPO / Wikimedia Commons (de)  
Galaxy Nexus: R. Nurik, I. Ni-Lewis / Google



photodiode: Omegatron/Wikimedia Commons  
Microphone: J.P. Lon / Wikimedia Commons  
Oscilloscope: MovGPO / Wikimedia Commons (de)  
Galaxy Nexus: R. Nurik, I. Ni-Lewis / Google



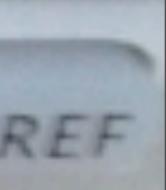




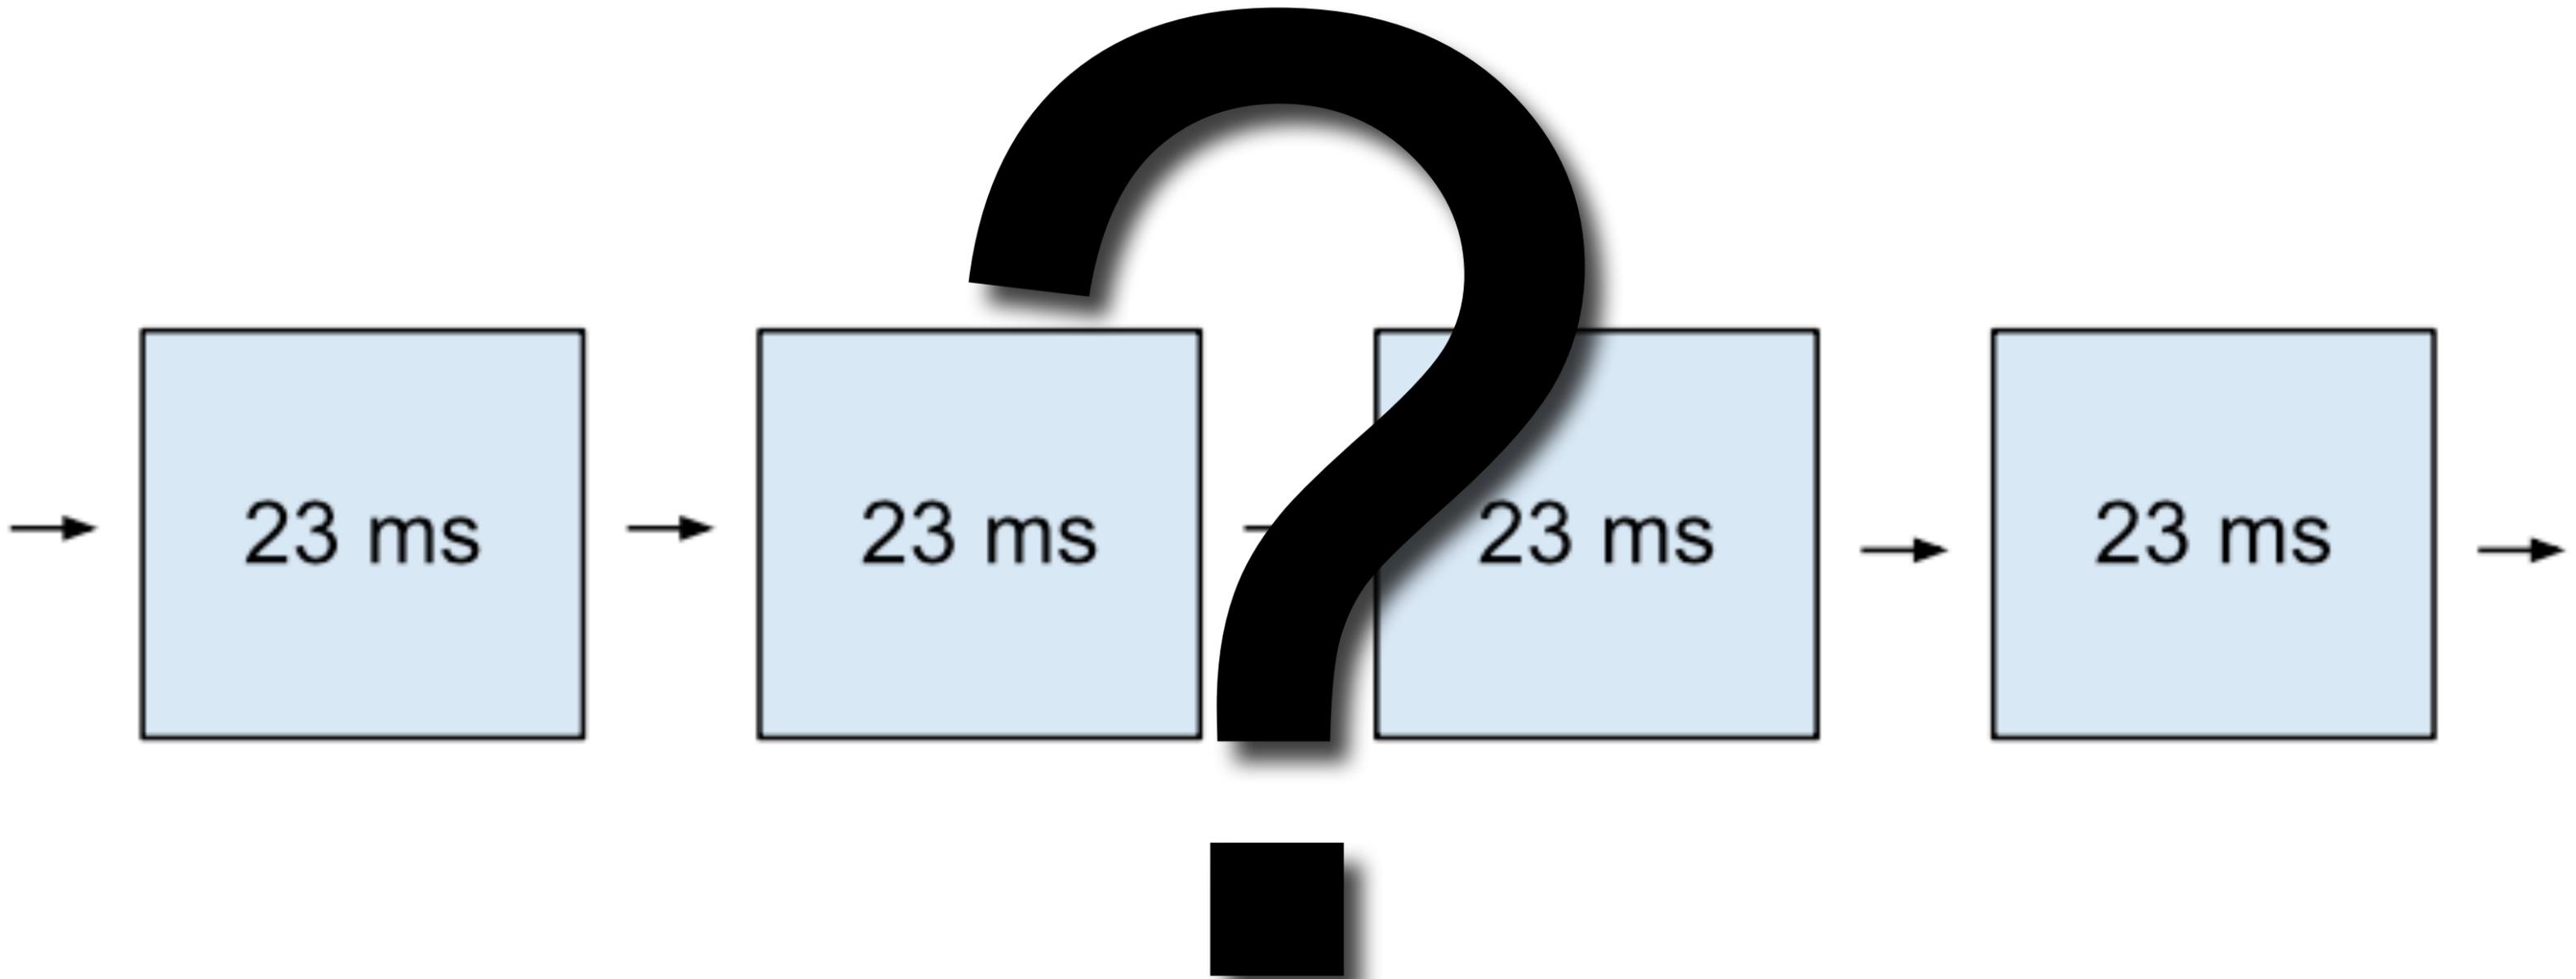
DS1102E  
DIGITAL OSCILLOSCOPE

2 Channel  
100MHz 1GSa/s

MENU  
ON/OFF

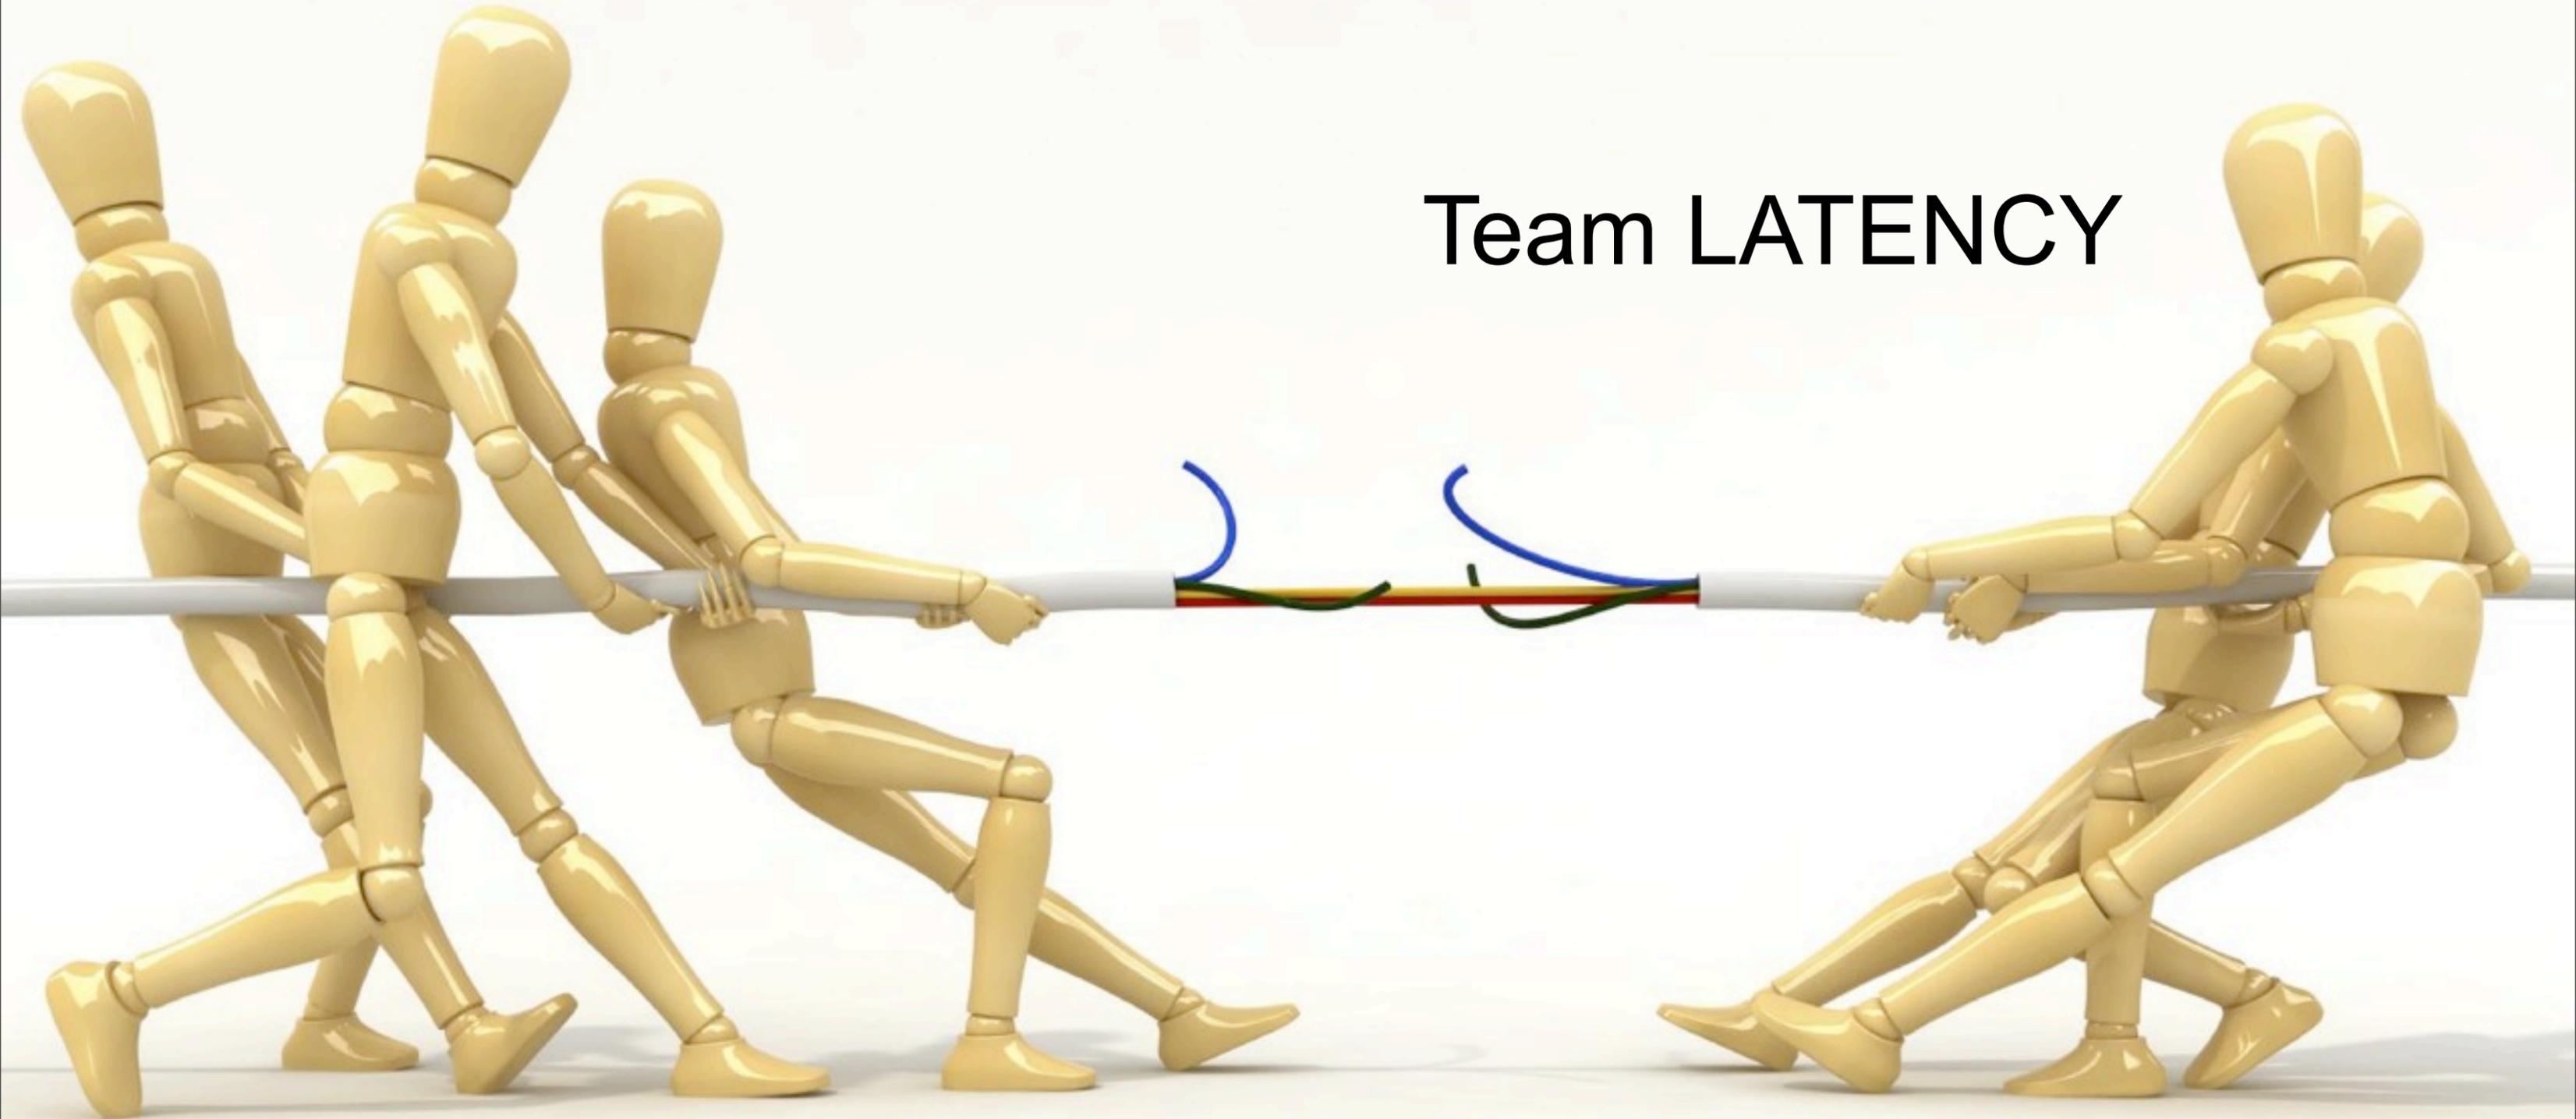


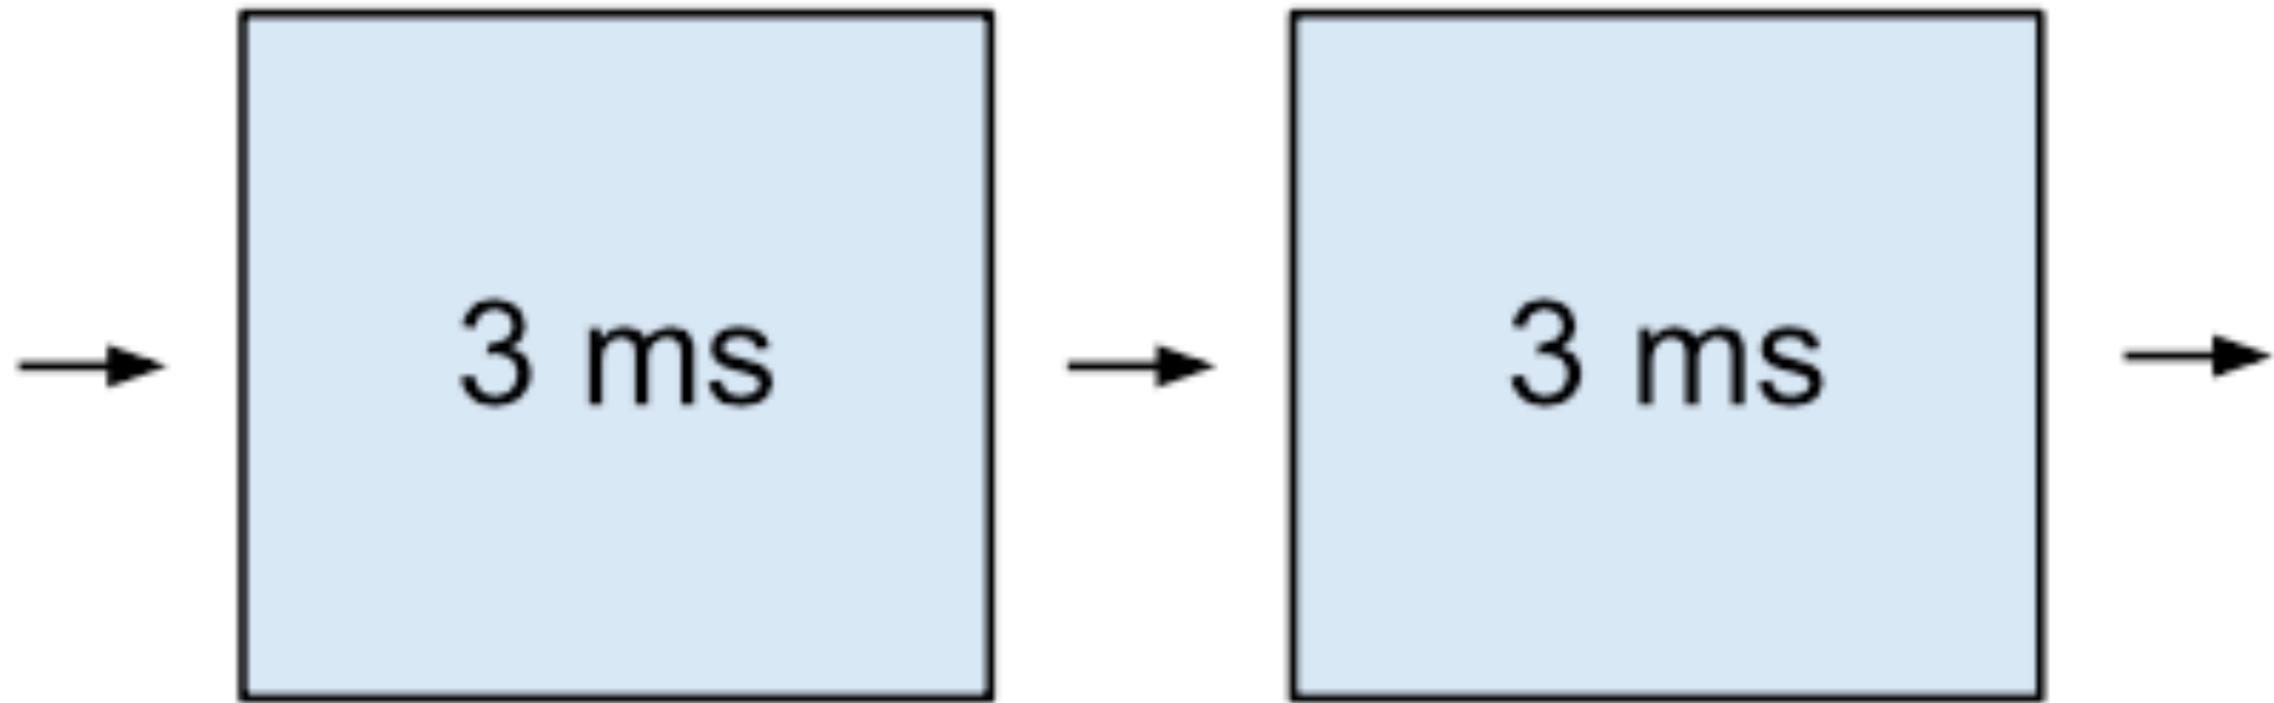




# Team POWER

# Team LATENCY







background\_with\_evil\_pixie\_dust

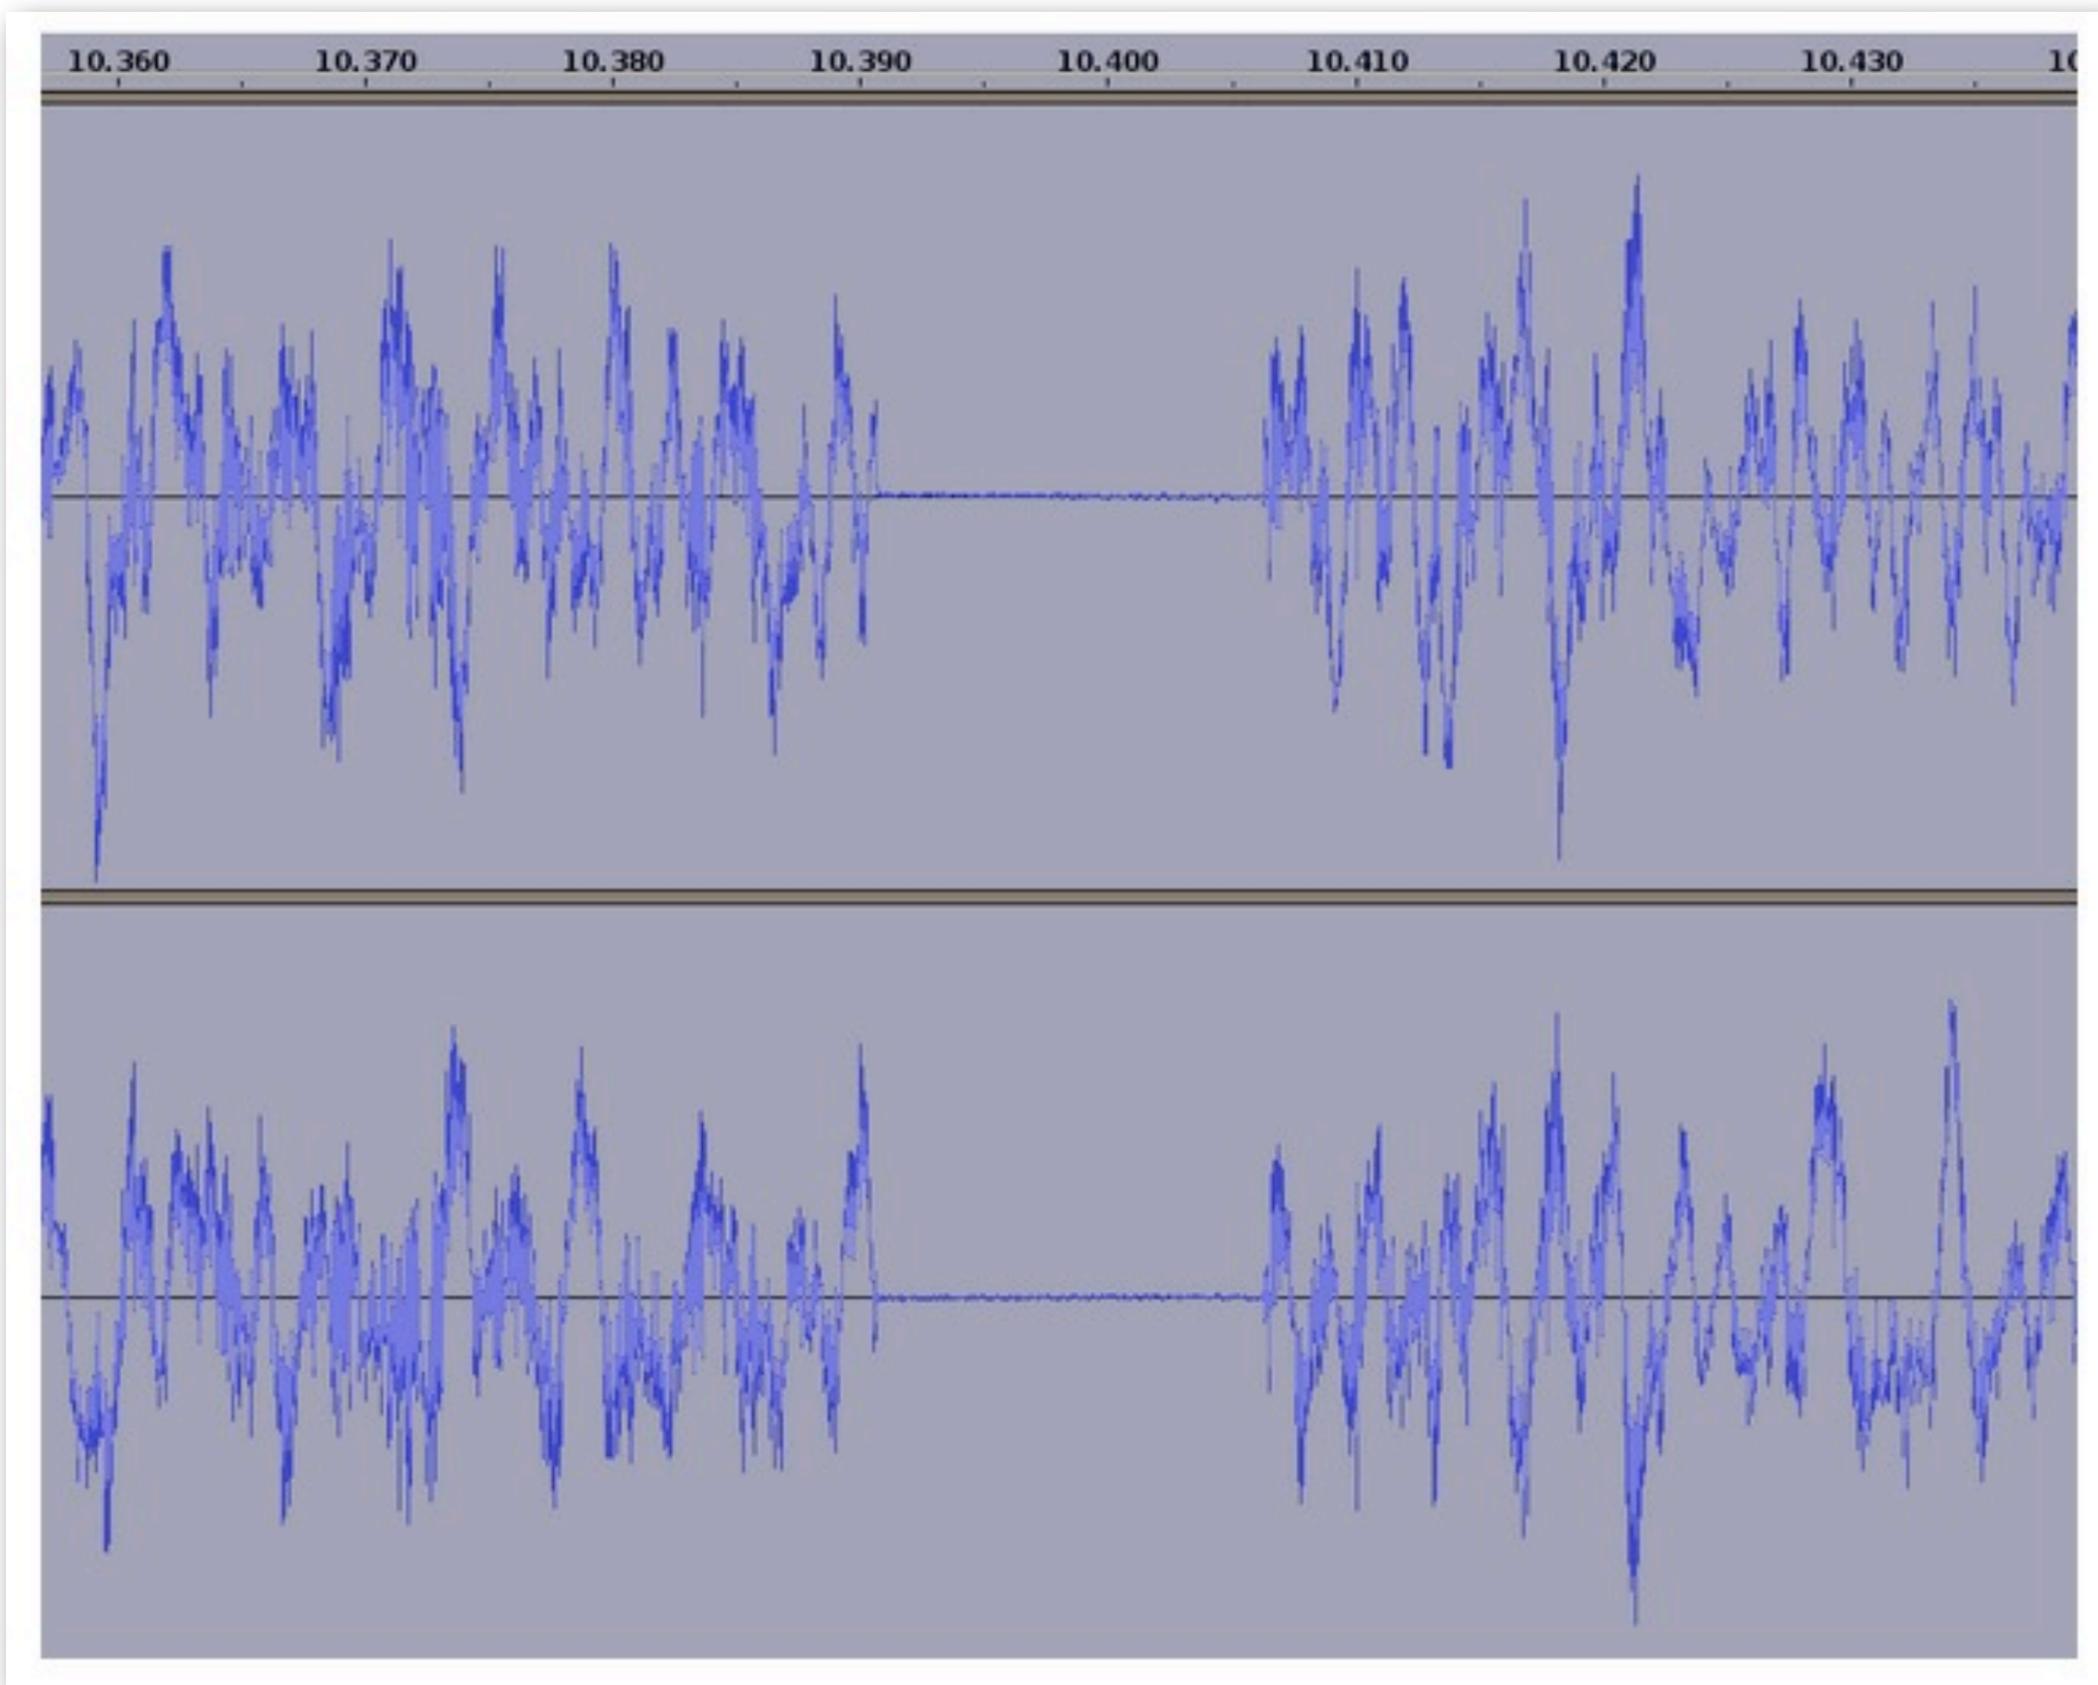
Core Au... Soundflower (2ch) Built-in Microph 1 (Mono)...

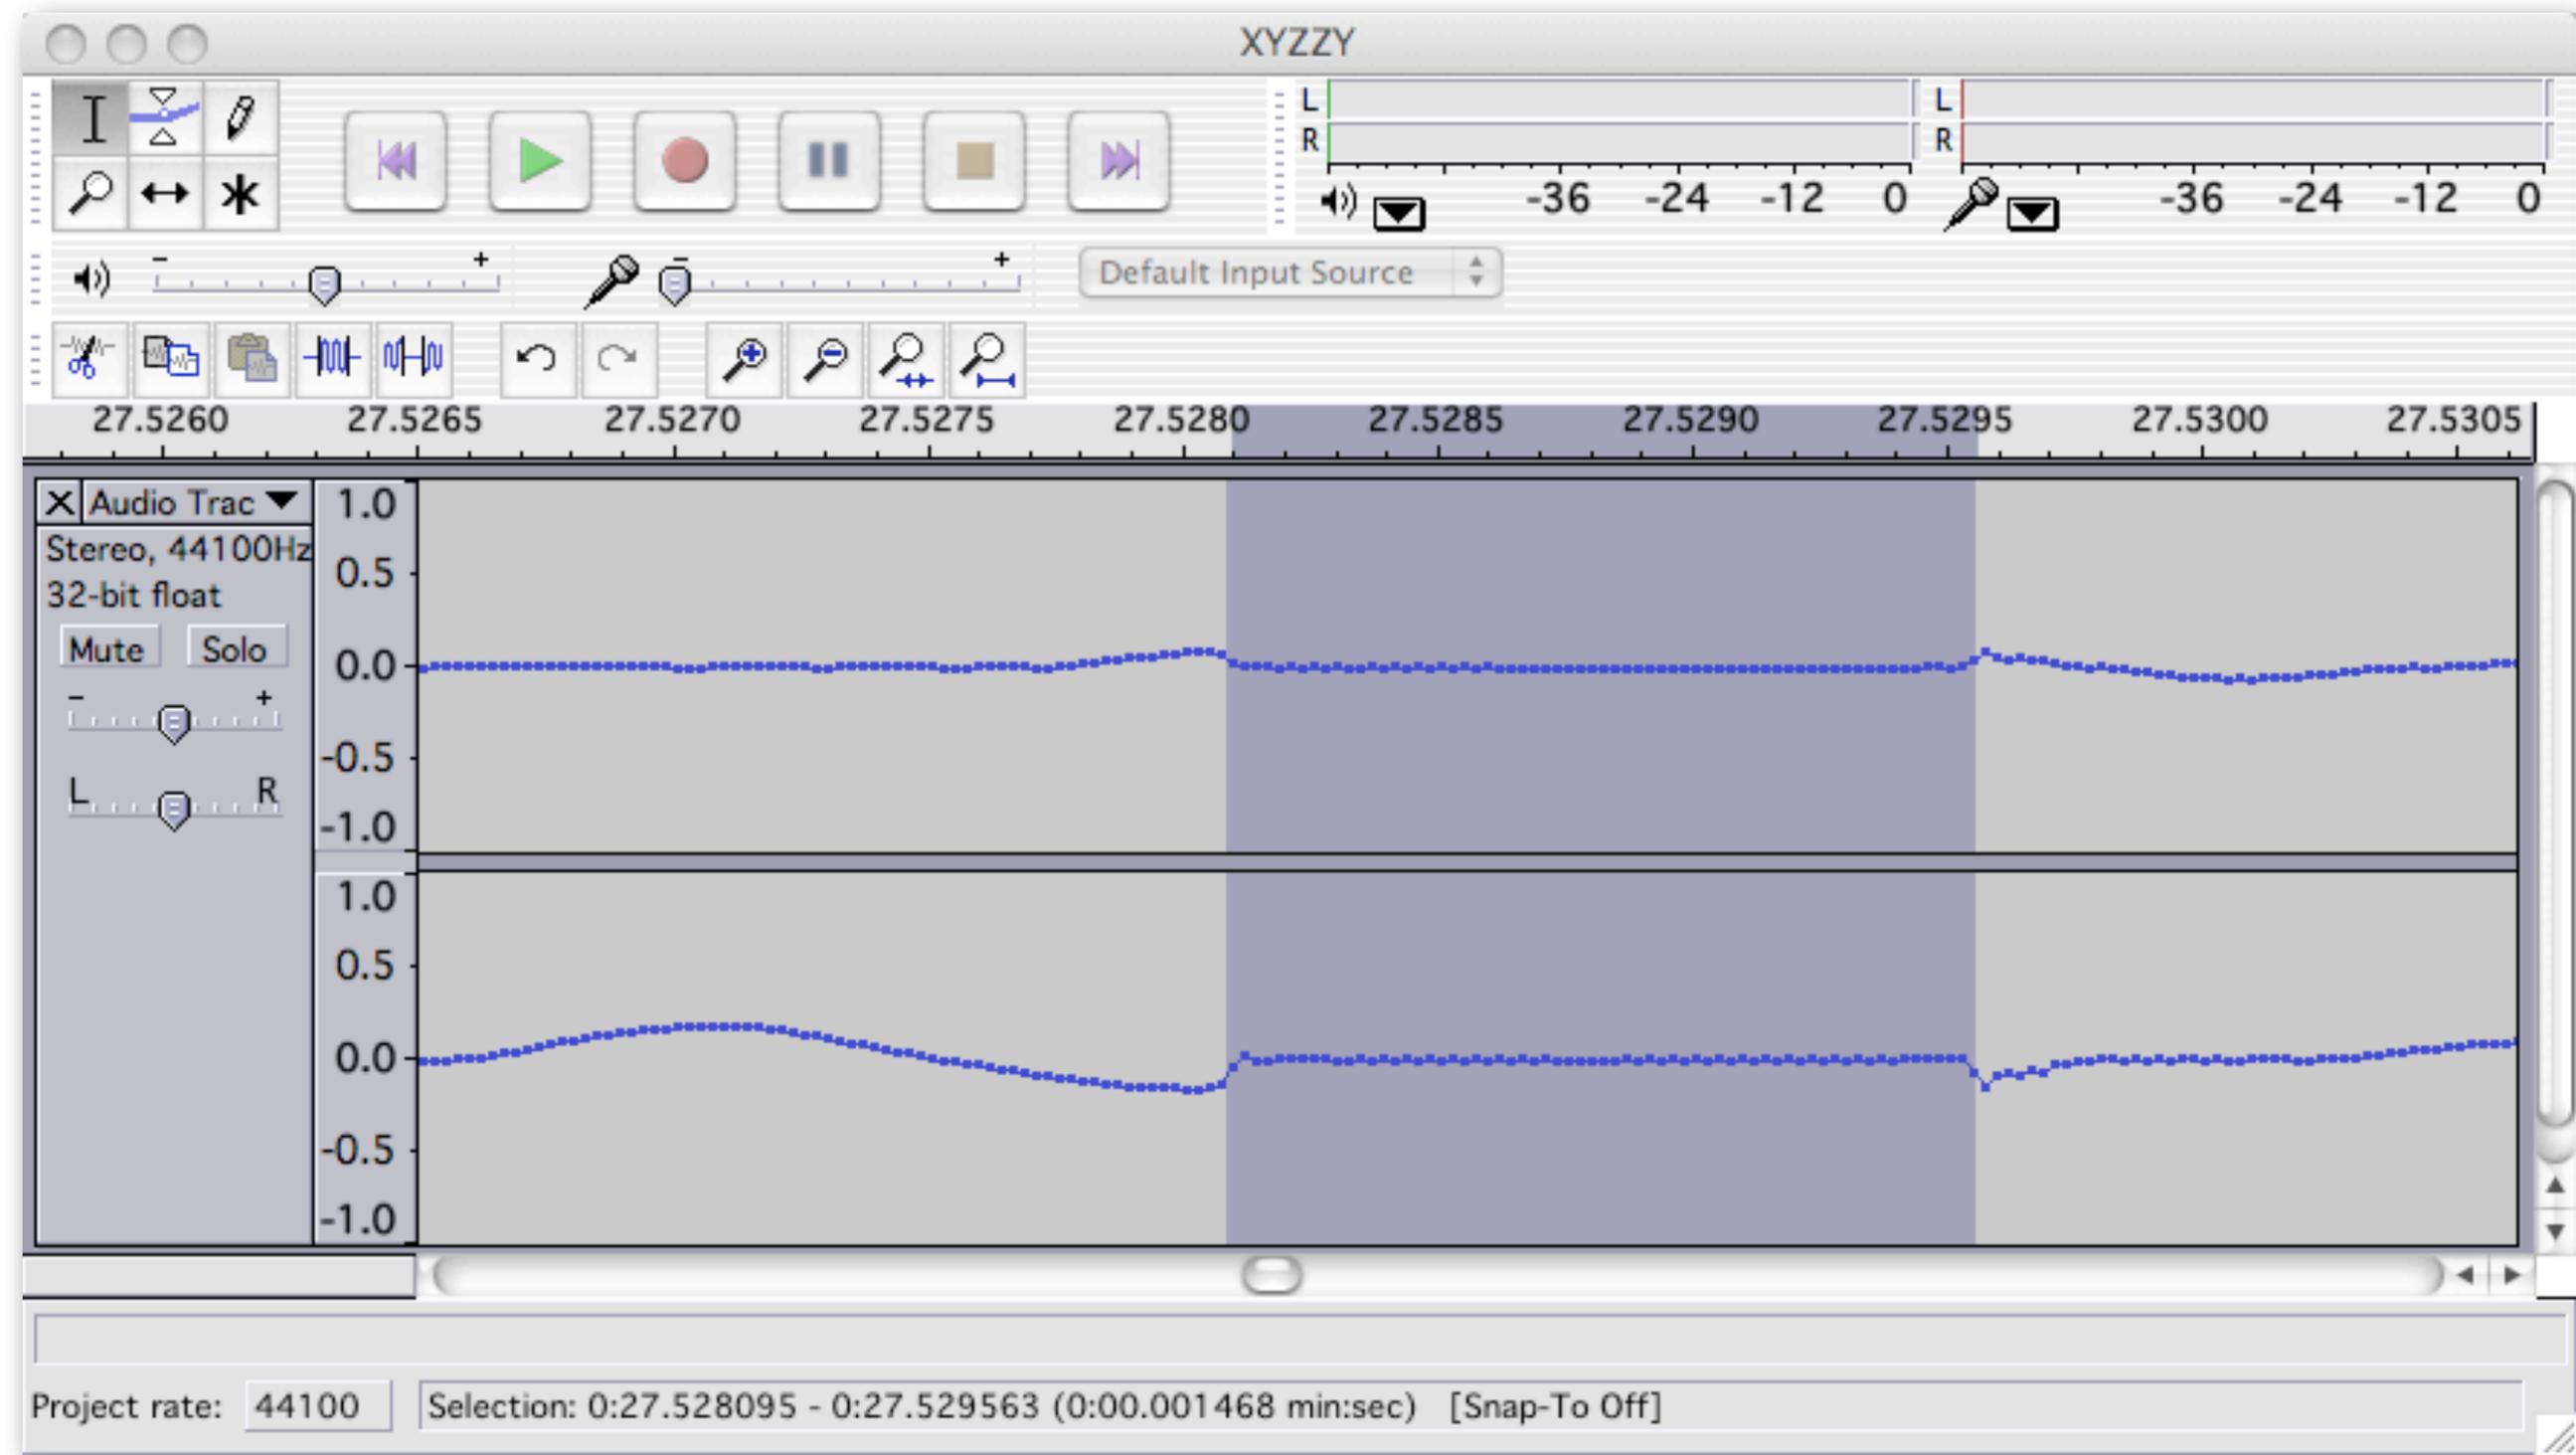
15 30 45 1:00 1:15 1:30 1:45

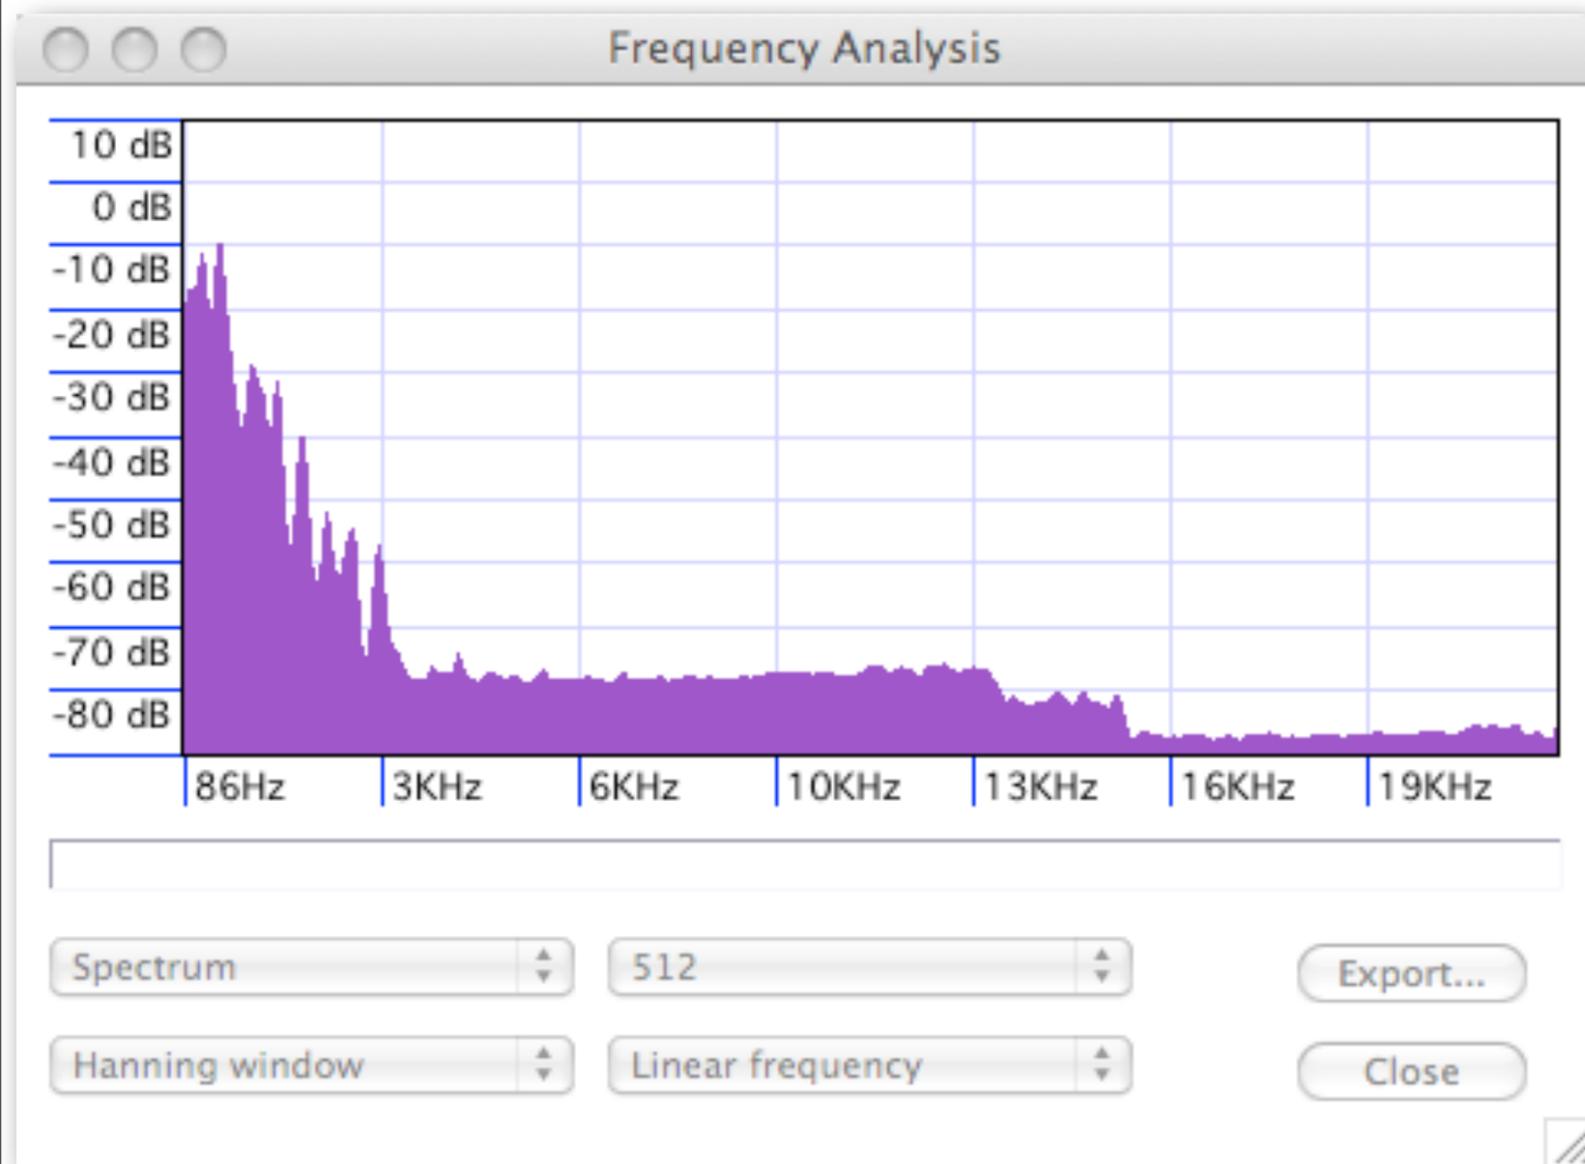
background  
Stereo, 44100Hz  
32-bit float  
Mute Solo  
L R

The screenshot shows an audio editing application window titled "background\_with\_evil\_pixie\_dust". The interface includes a top toolbar with playback controls (play, stop, previous, next), editing tools (insert, delete, copy, paste), and volume meters for left and right channels. Below the toolbar is a transport bar with a play button, a volume slider, and input/output device selection (Core Audio, Soundflower (2ch), Built-in Microph, 1 (Mono)). A timeline at the top shows time markers from -15 to 1:45. The main area contains two tracks of audio waveforms, both showing a similar pattern of noise with a slight dip around the 1:00 mark. The left track has a volume of 1.0. The bottom of the window features a control panel with "Project Rate (Hz): 44100", "Selection Start" and "End" fields set to "00 h 00 m 00.000 s", and "Audio Position" set to "00 h 00 m 33.599 s". A "Play (Shift for Loop Play)" button and "Actual Rate: 44100" are also visible.

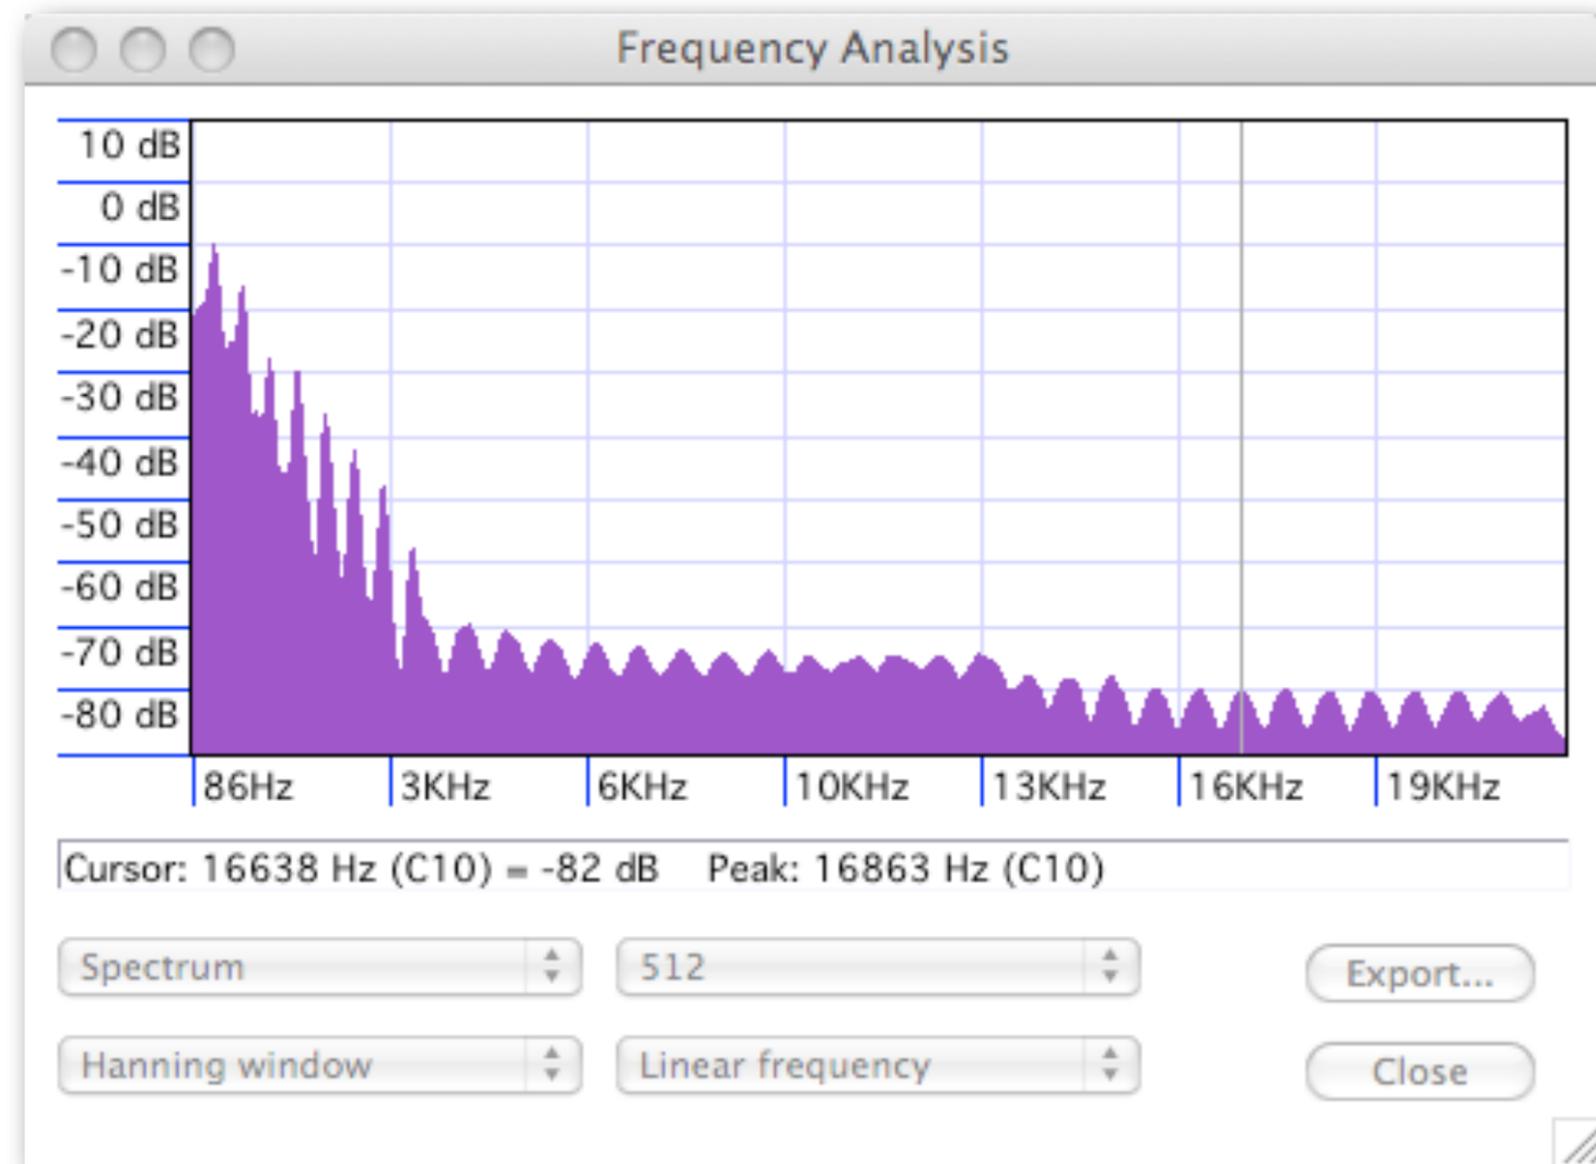
Project Rate (Hz): 44100  
Selection Start: 00 h 00 m 00.000 s  
End: 00 h 00 m 00.000 s  
Audio Position: 00 h 00 m 33.599 s  
Play (Shift for Loop Play) Actual Rate: 44100







Good



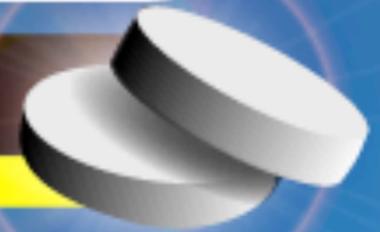
Bad



*Extra Long*

**BUFFERING**

**BATTERY EXTENDER**



*Extra Long*

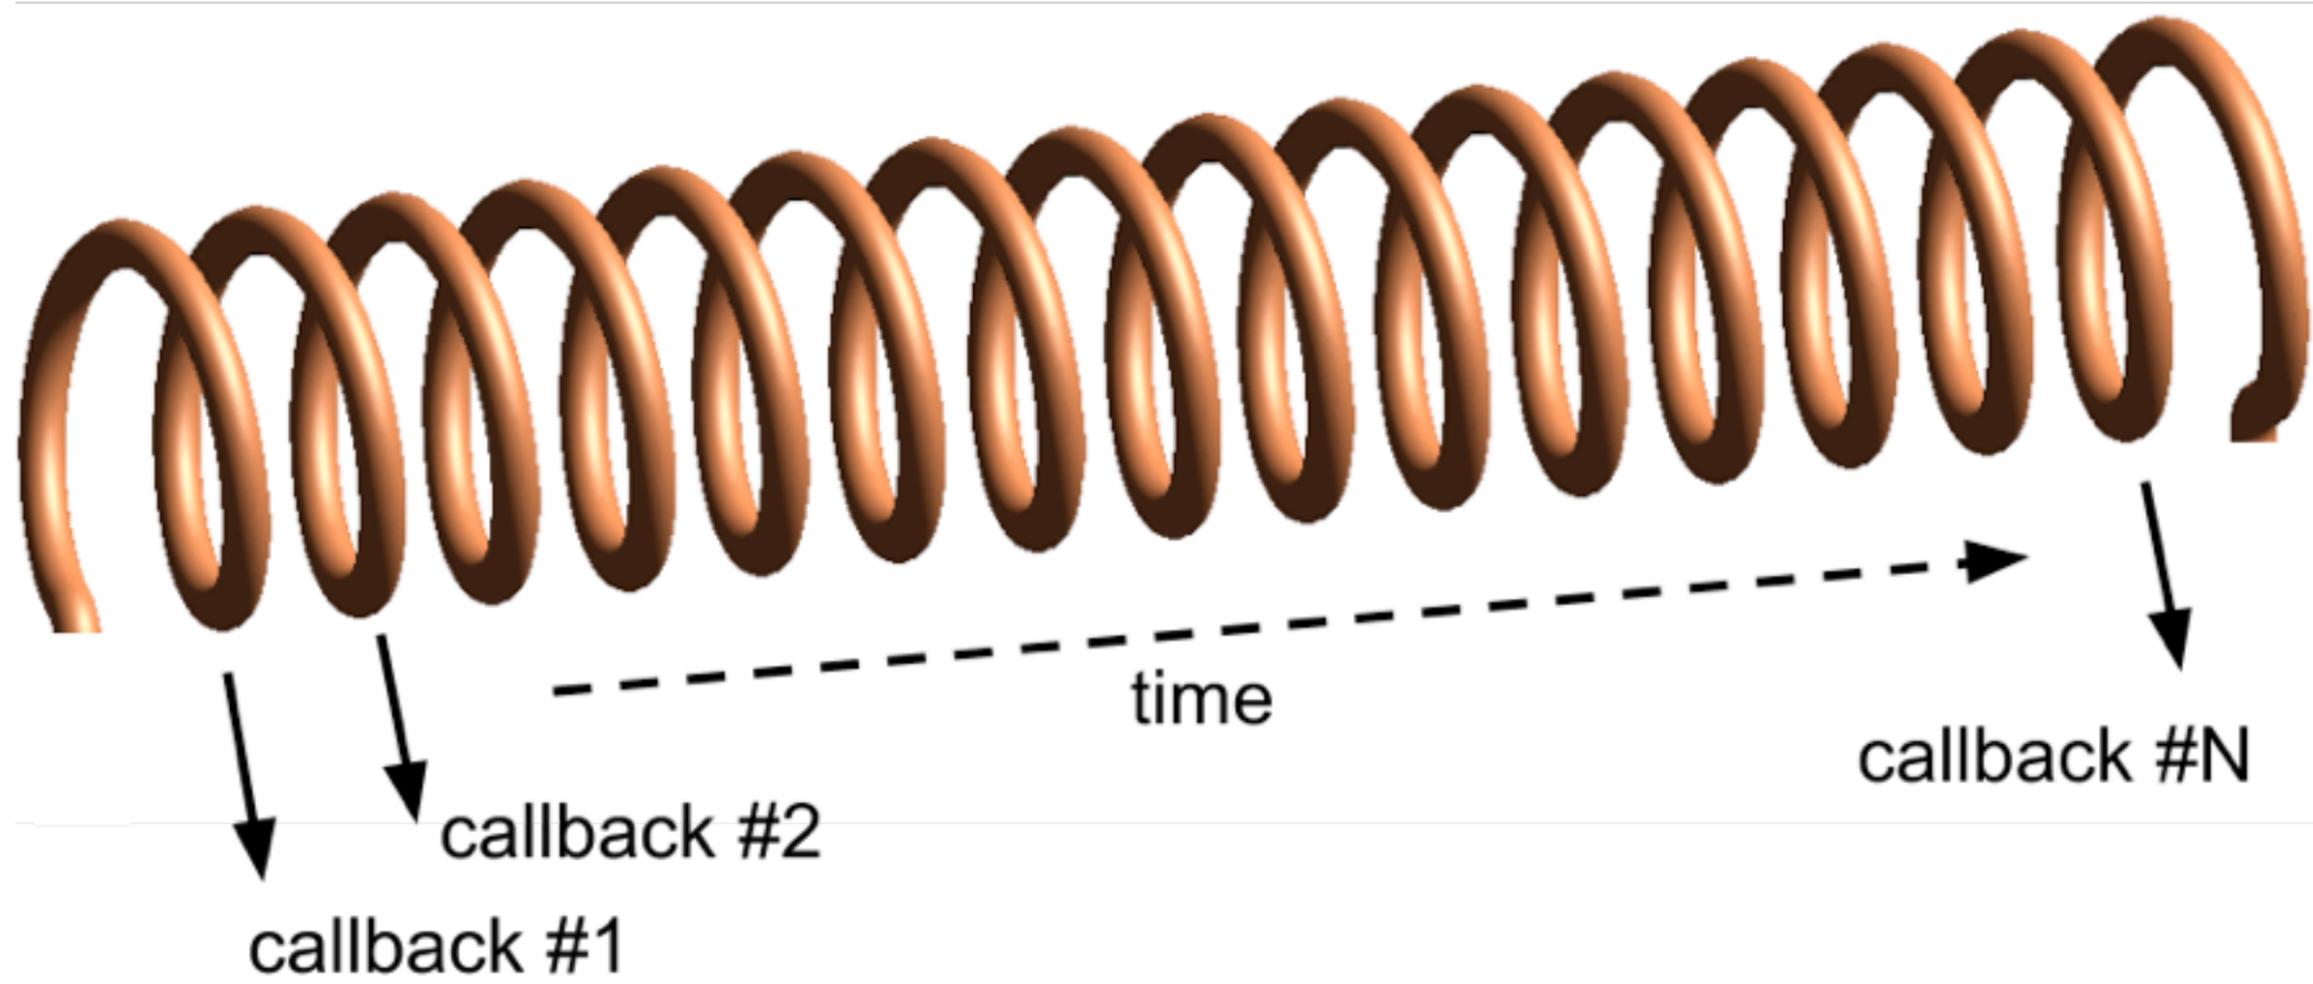
**BUFFERING**

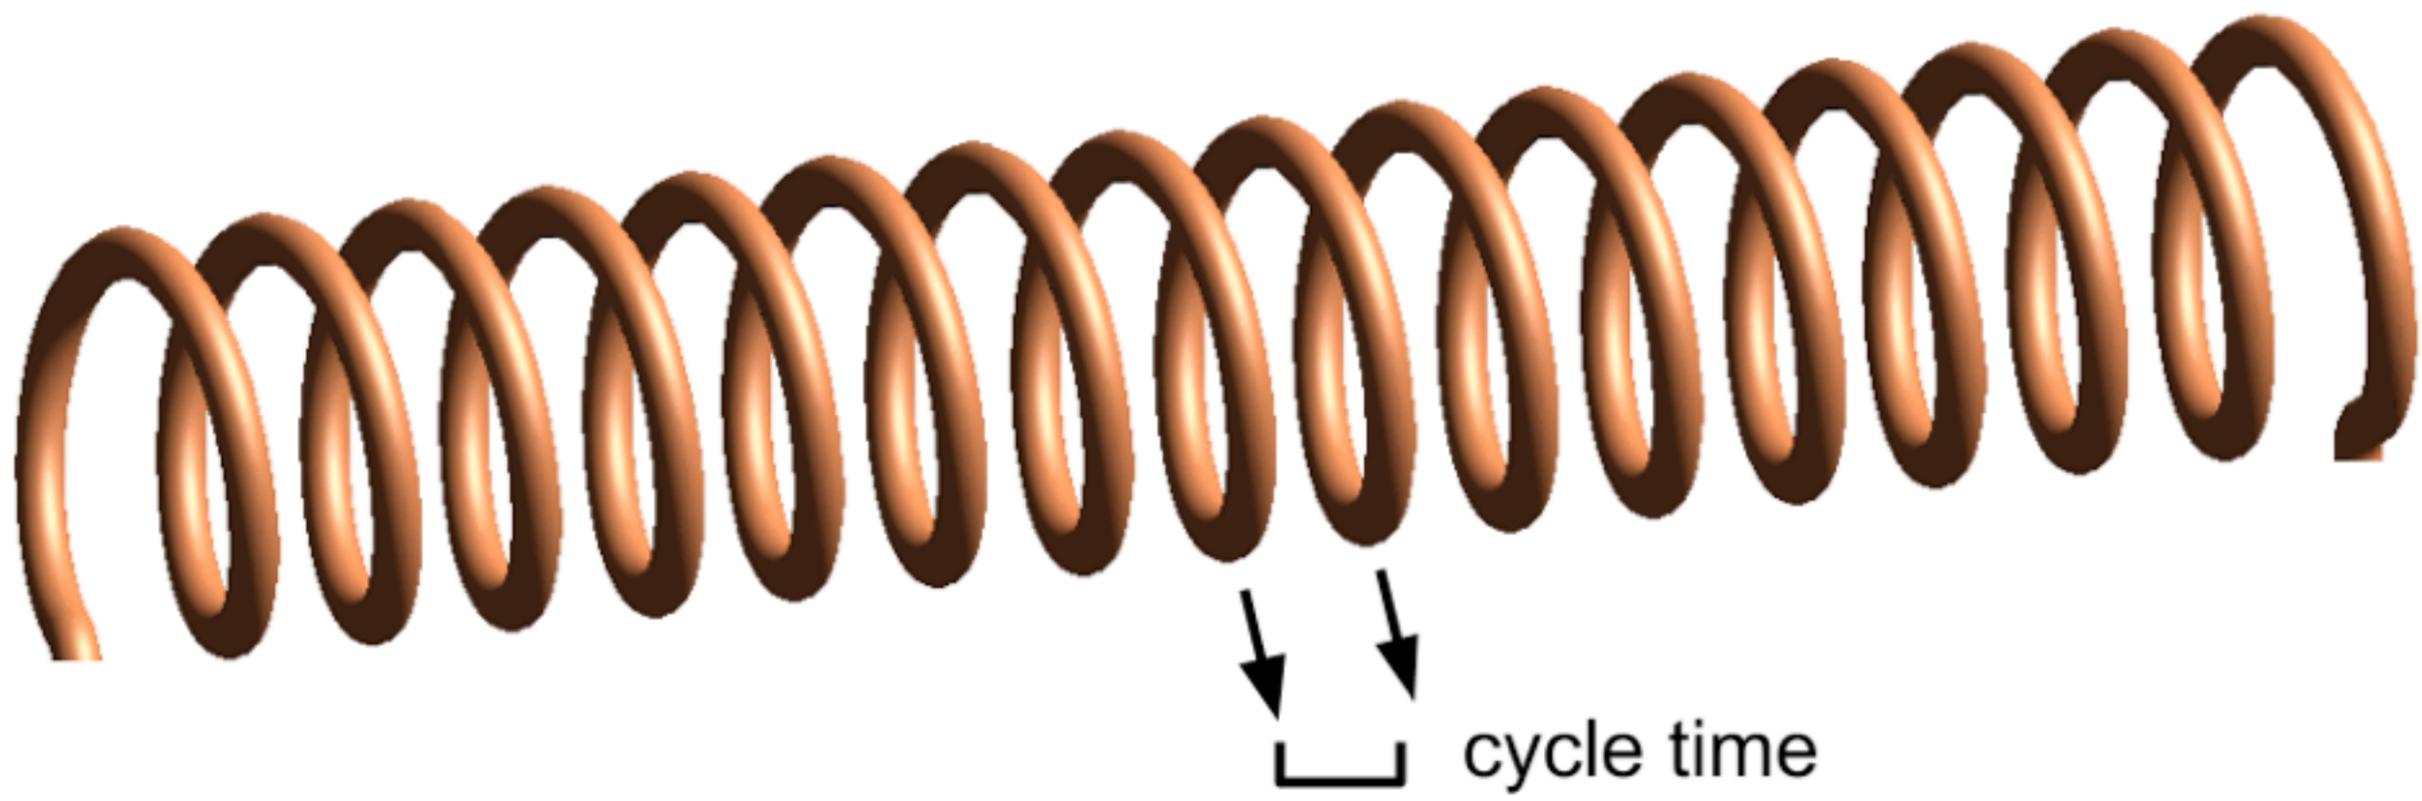
BATTERY EXTENDER \* GLITCH REDUCER

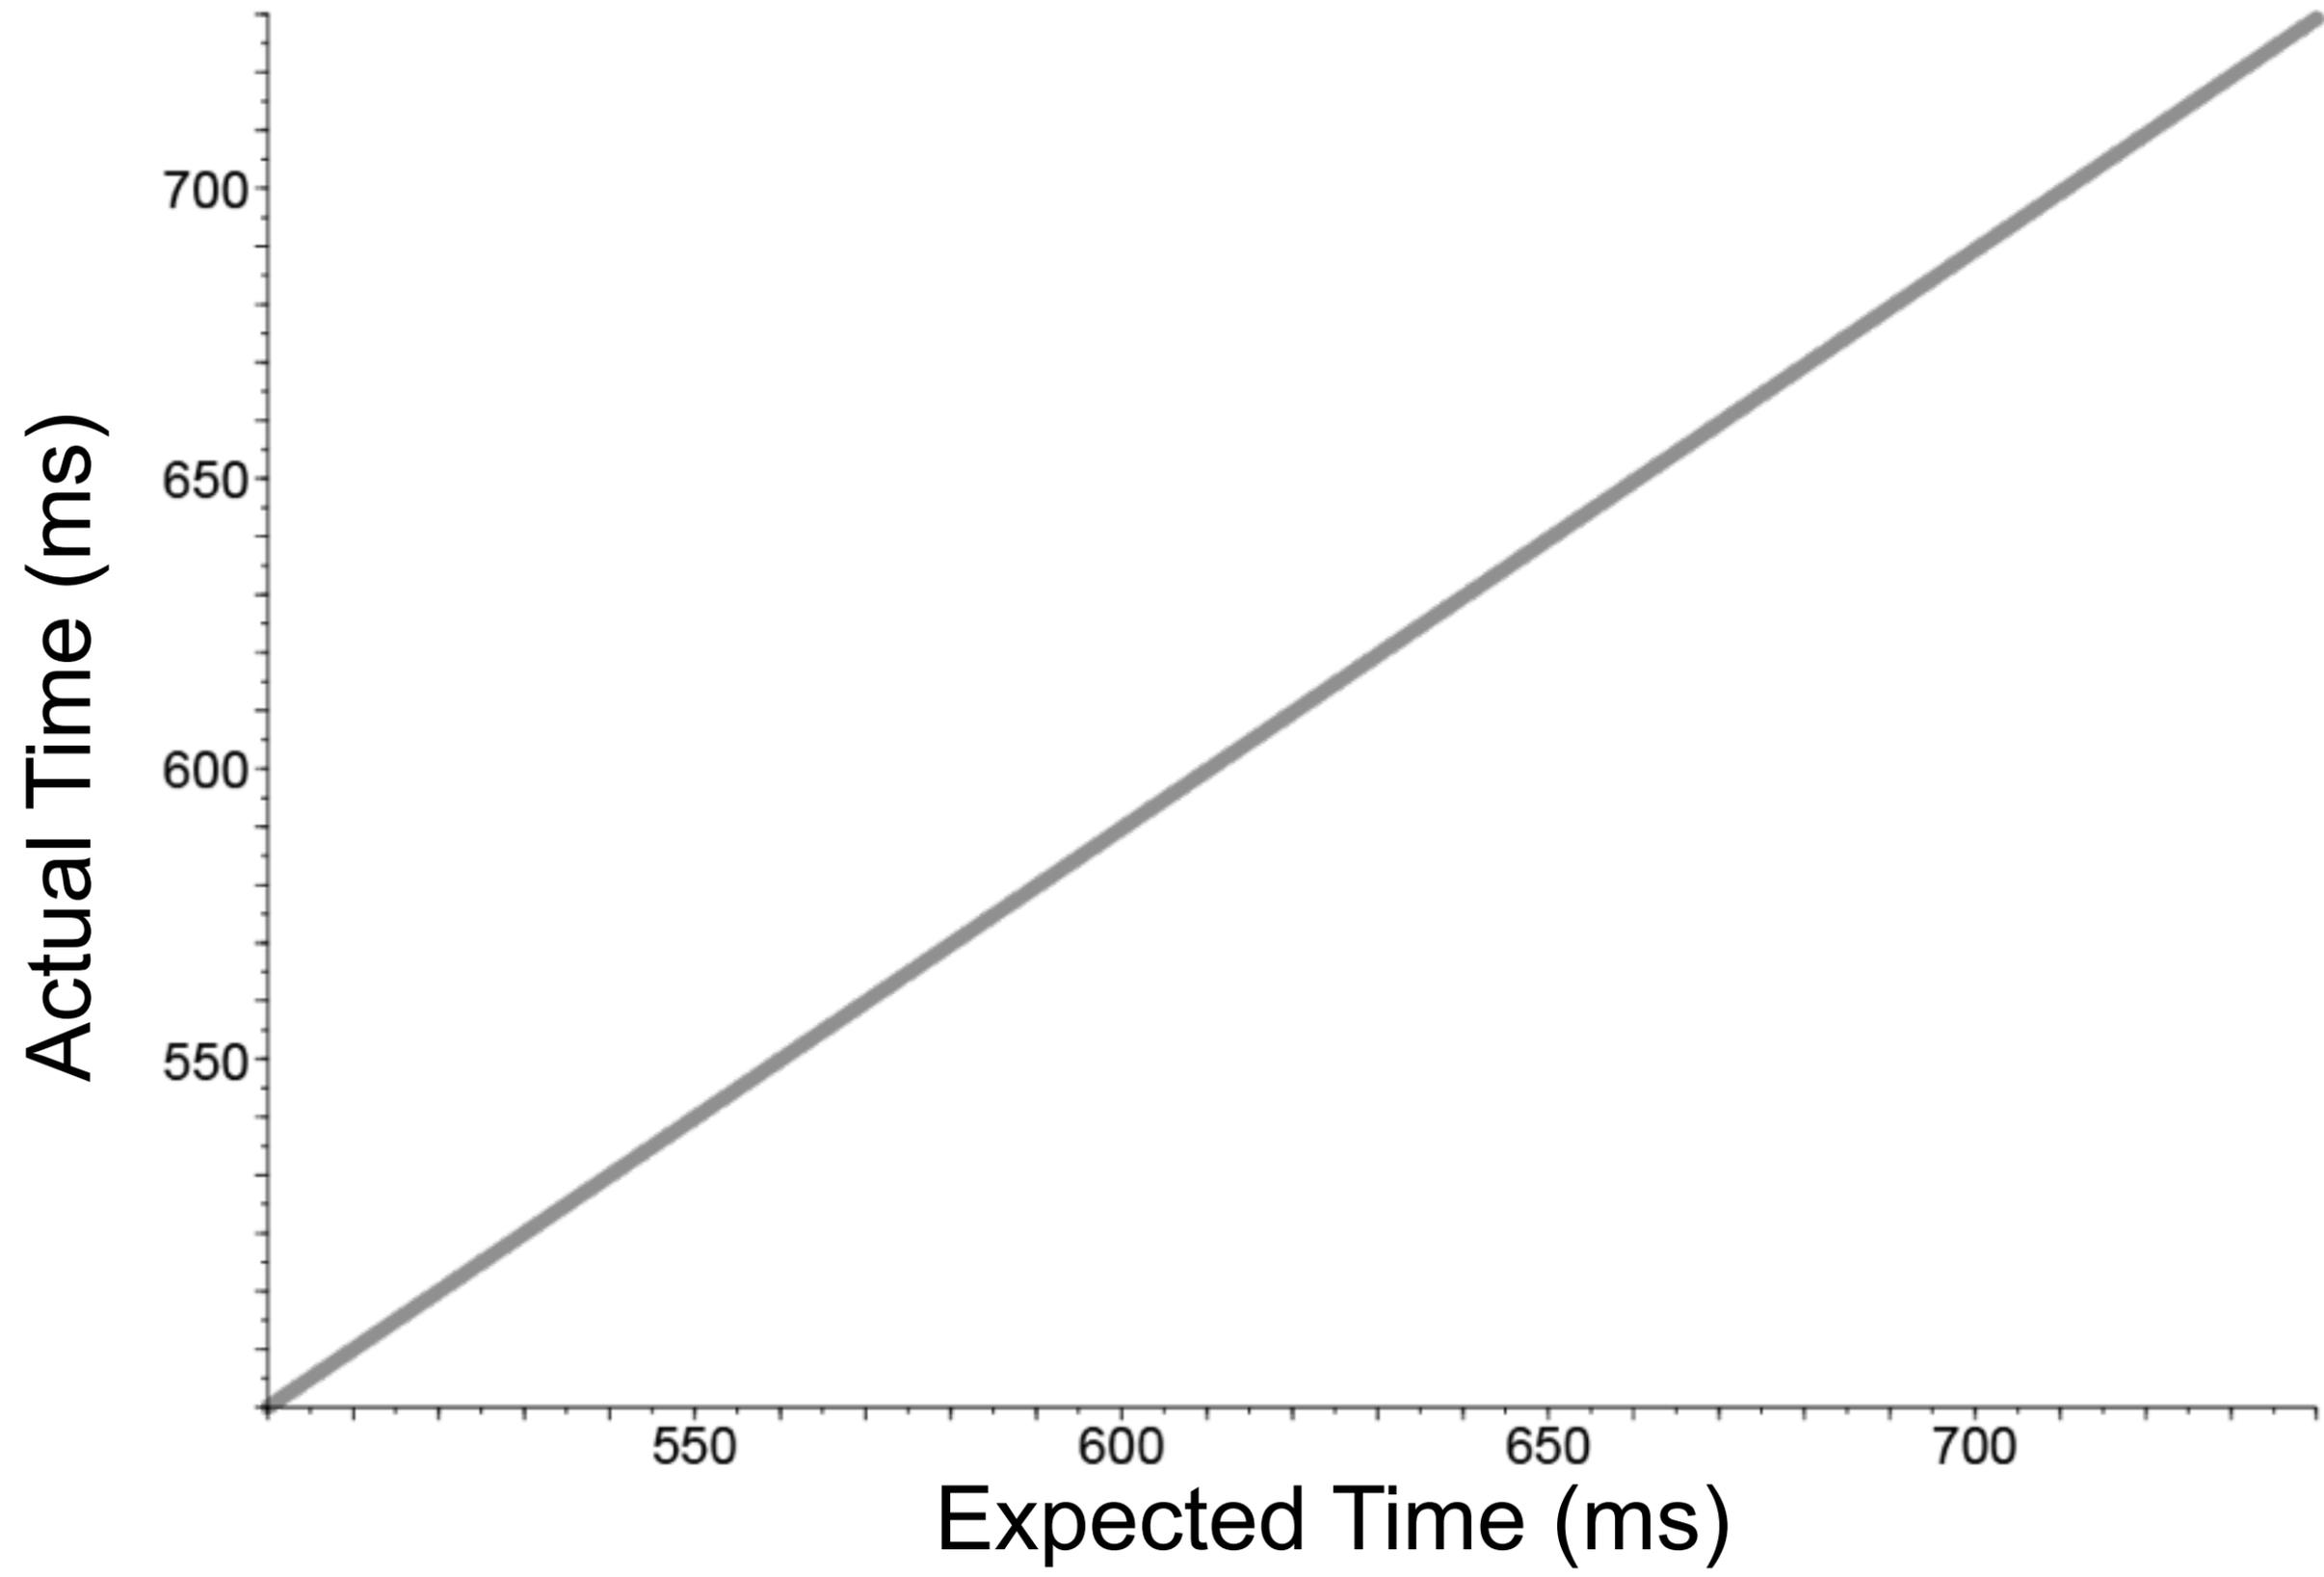


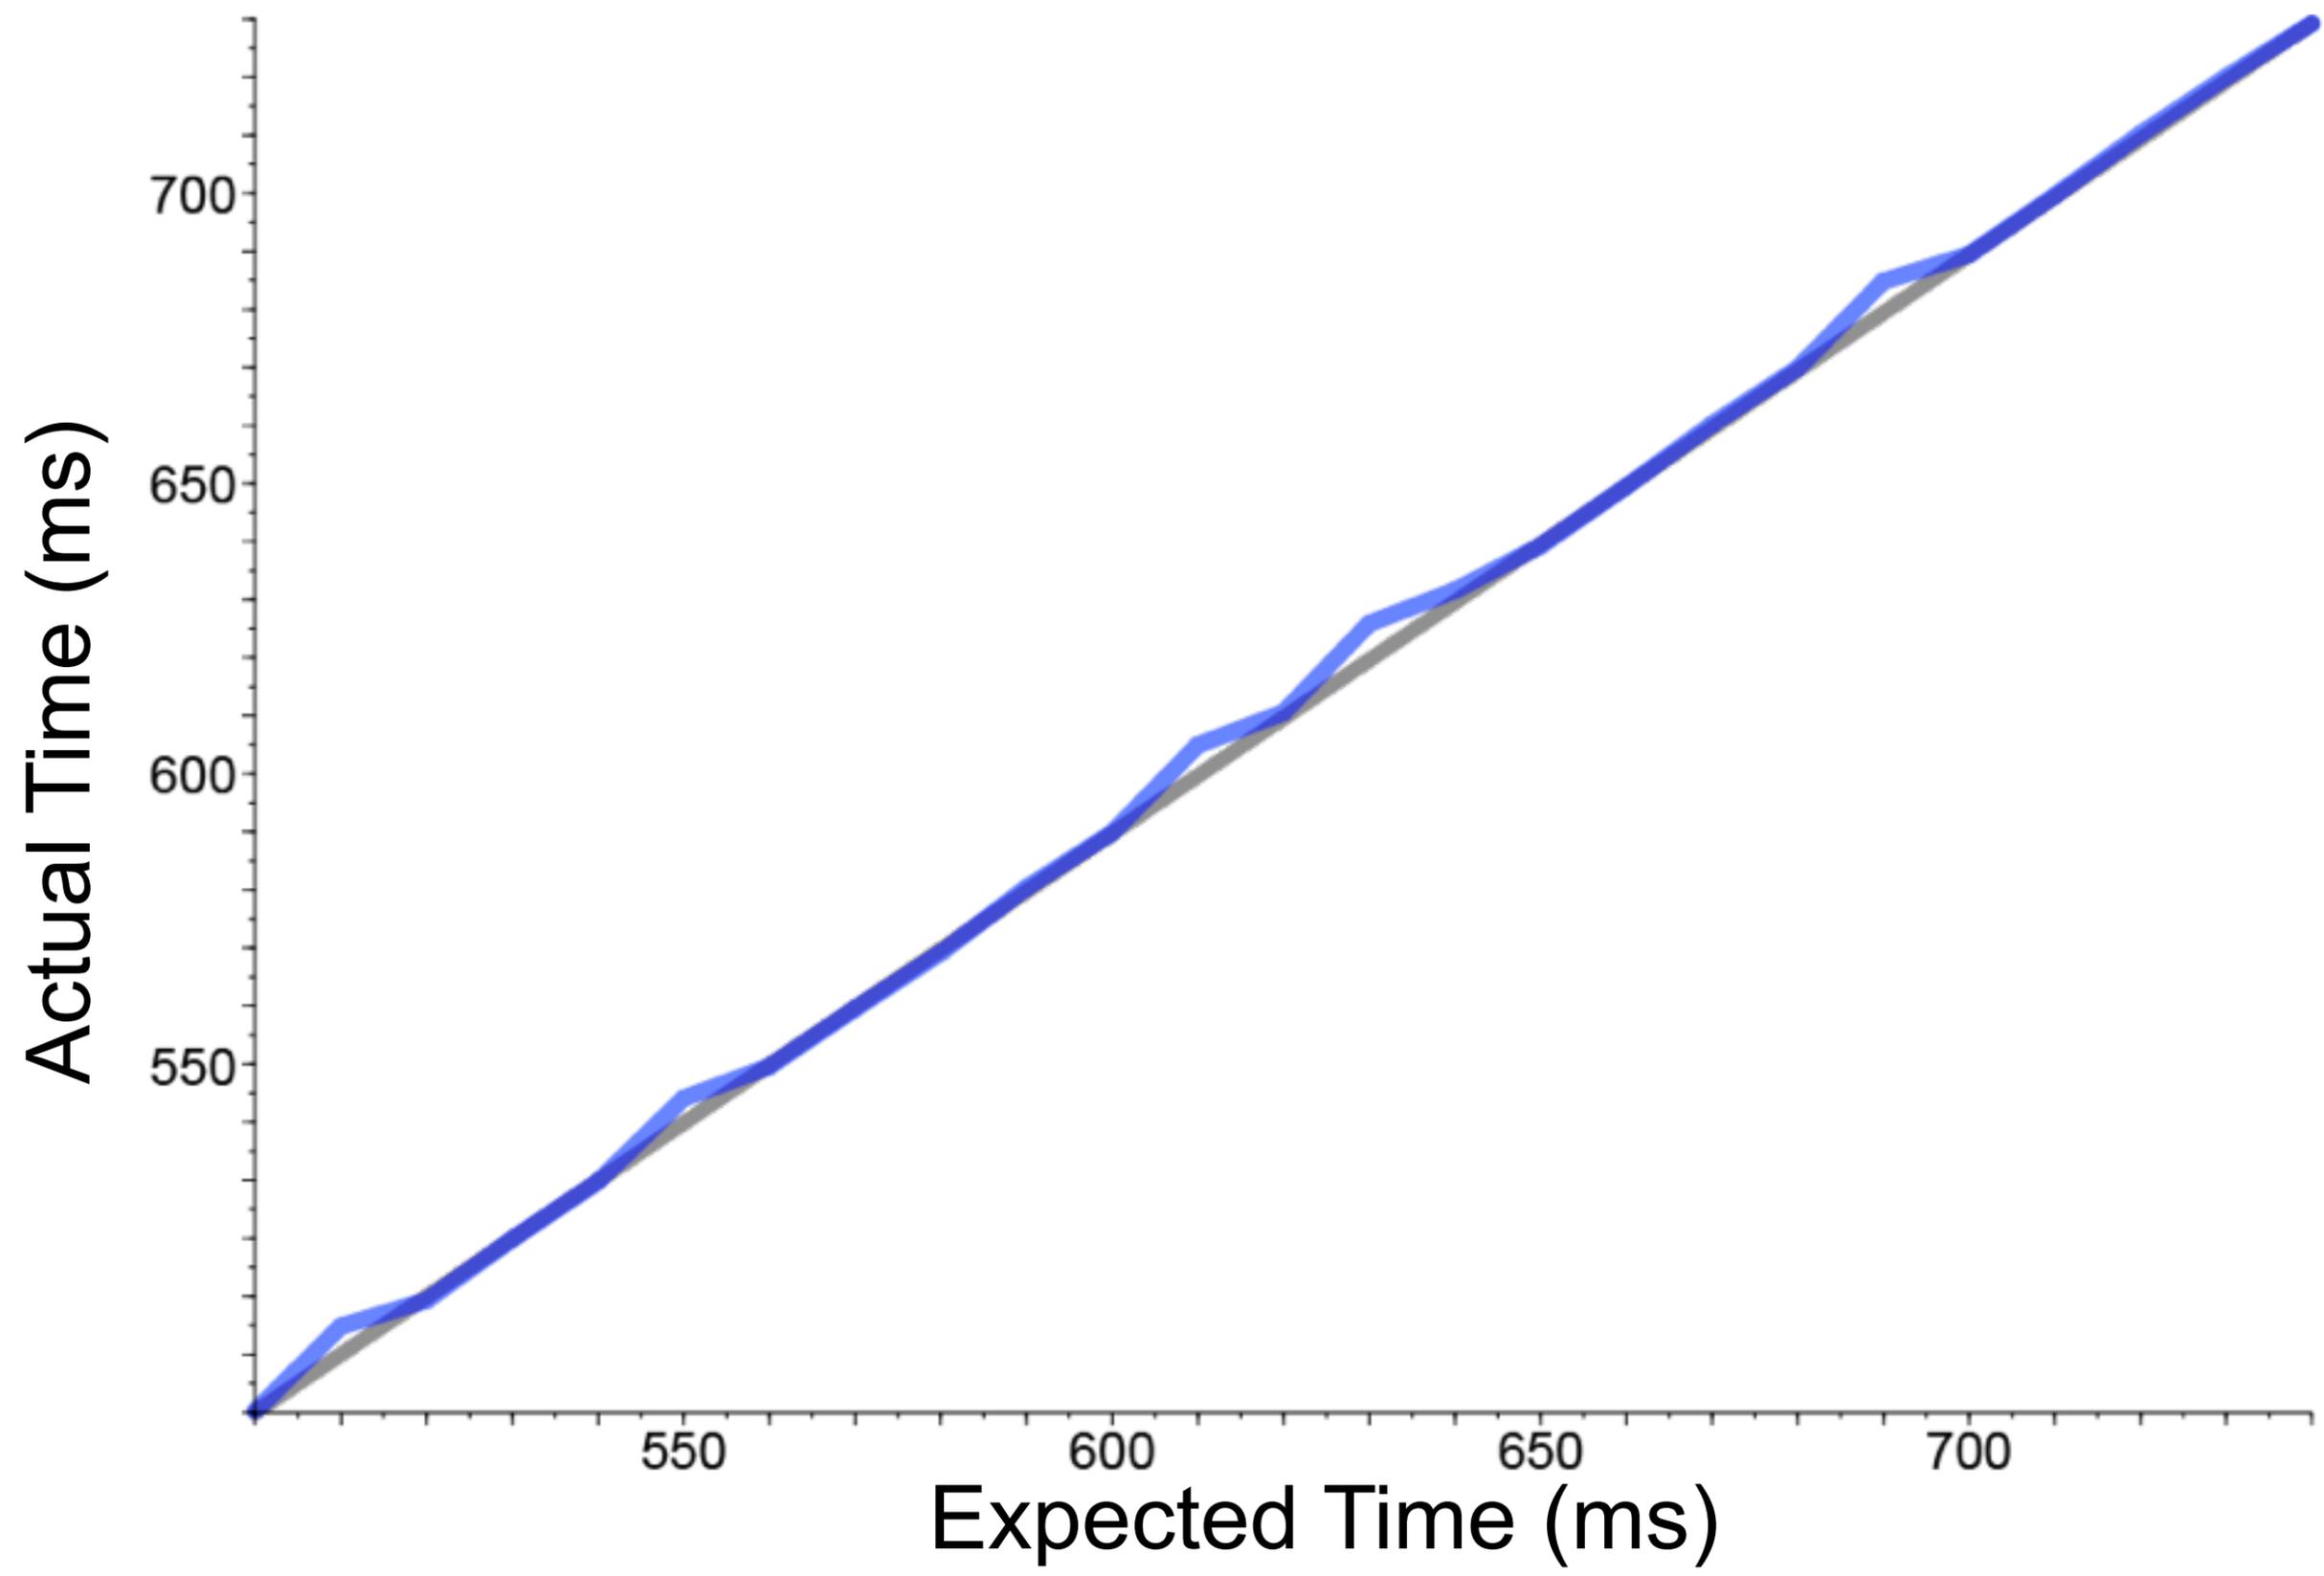
```
OpenSL_callback() {  
    start_time = clock_gettime(CLOCK_MONOTONIC);  
    calculate one buffer full of audio  
    end_time = clock_gettime(CLOCK_MONOTONIC);  
    buffer_queue->Enqueue();  
    log start_time and end_time  
}
```



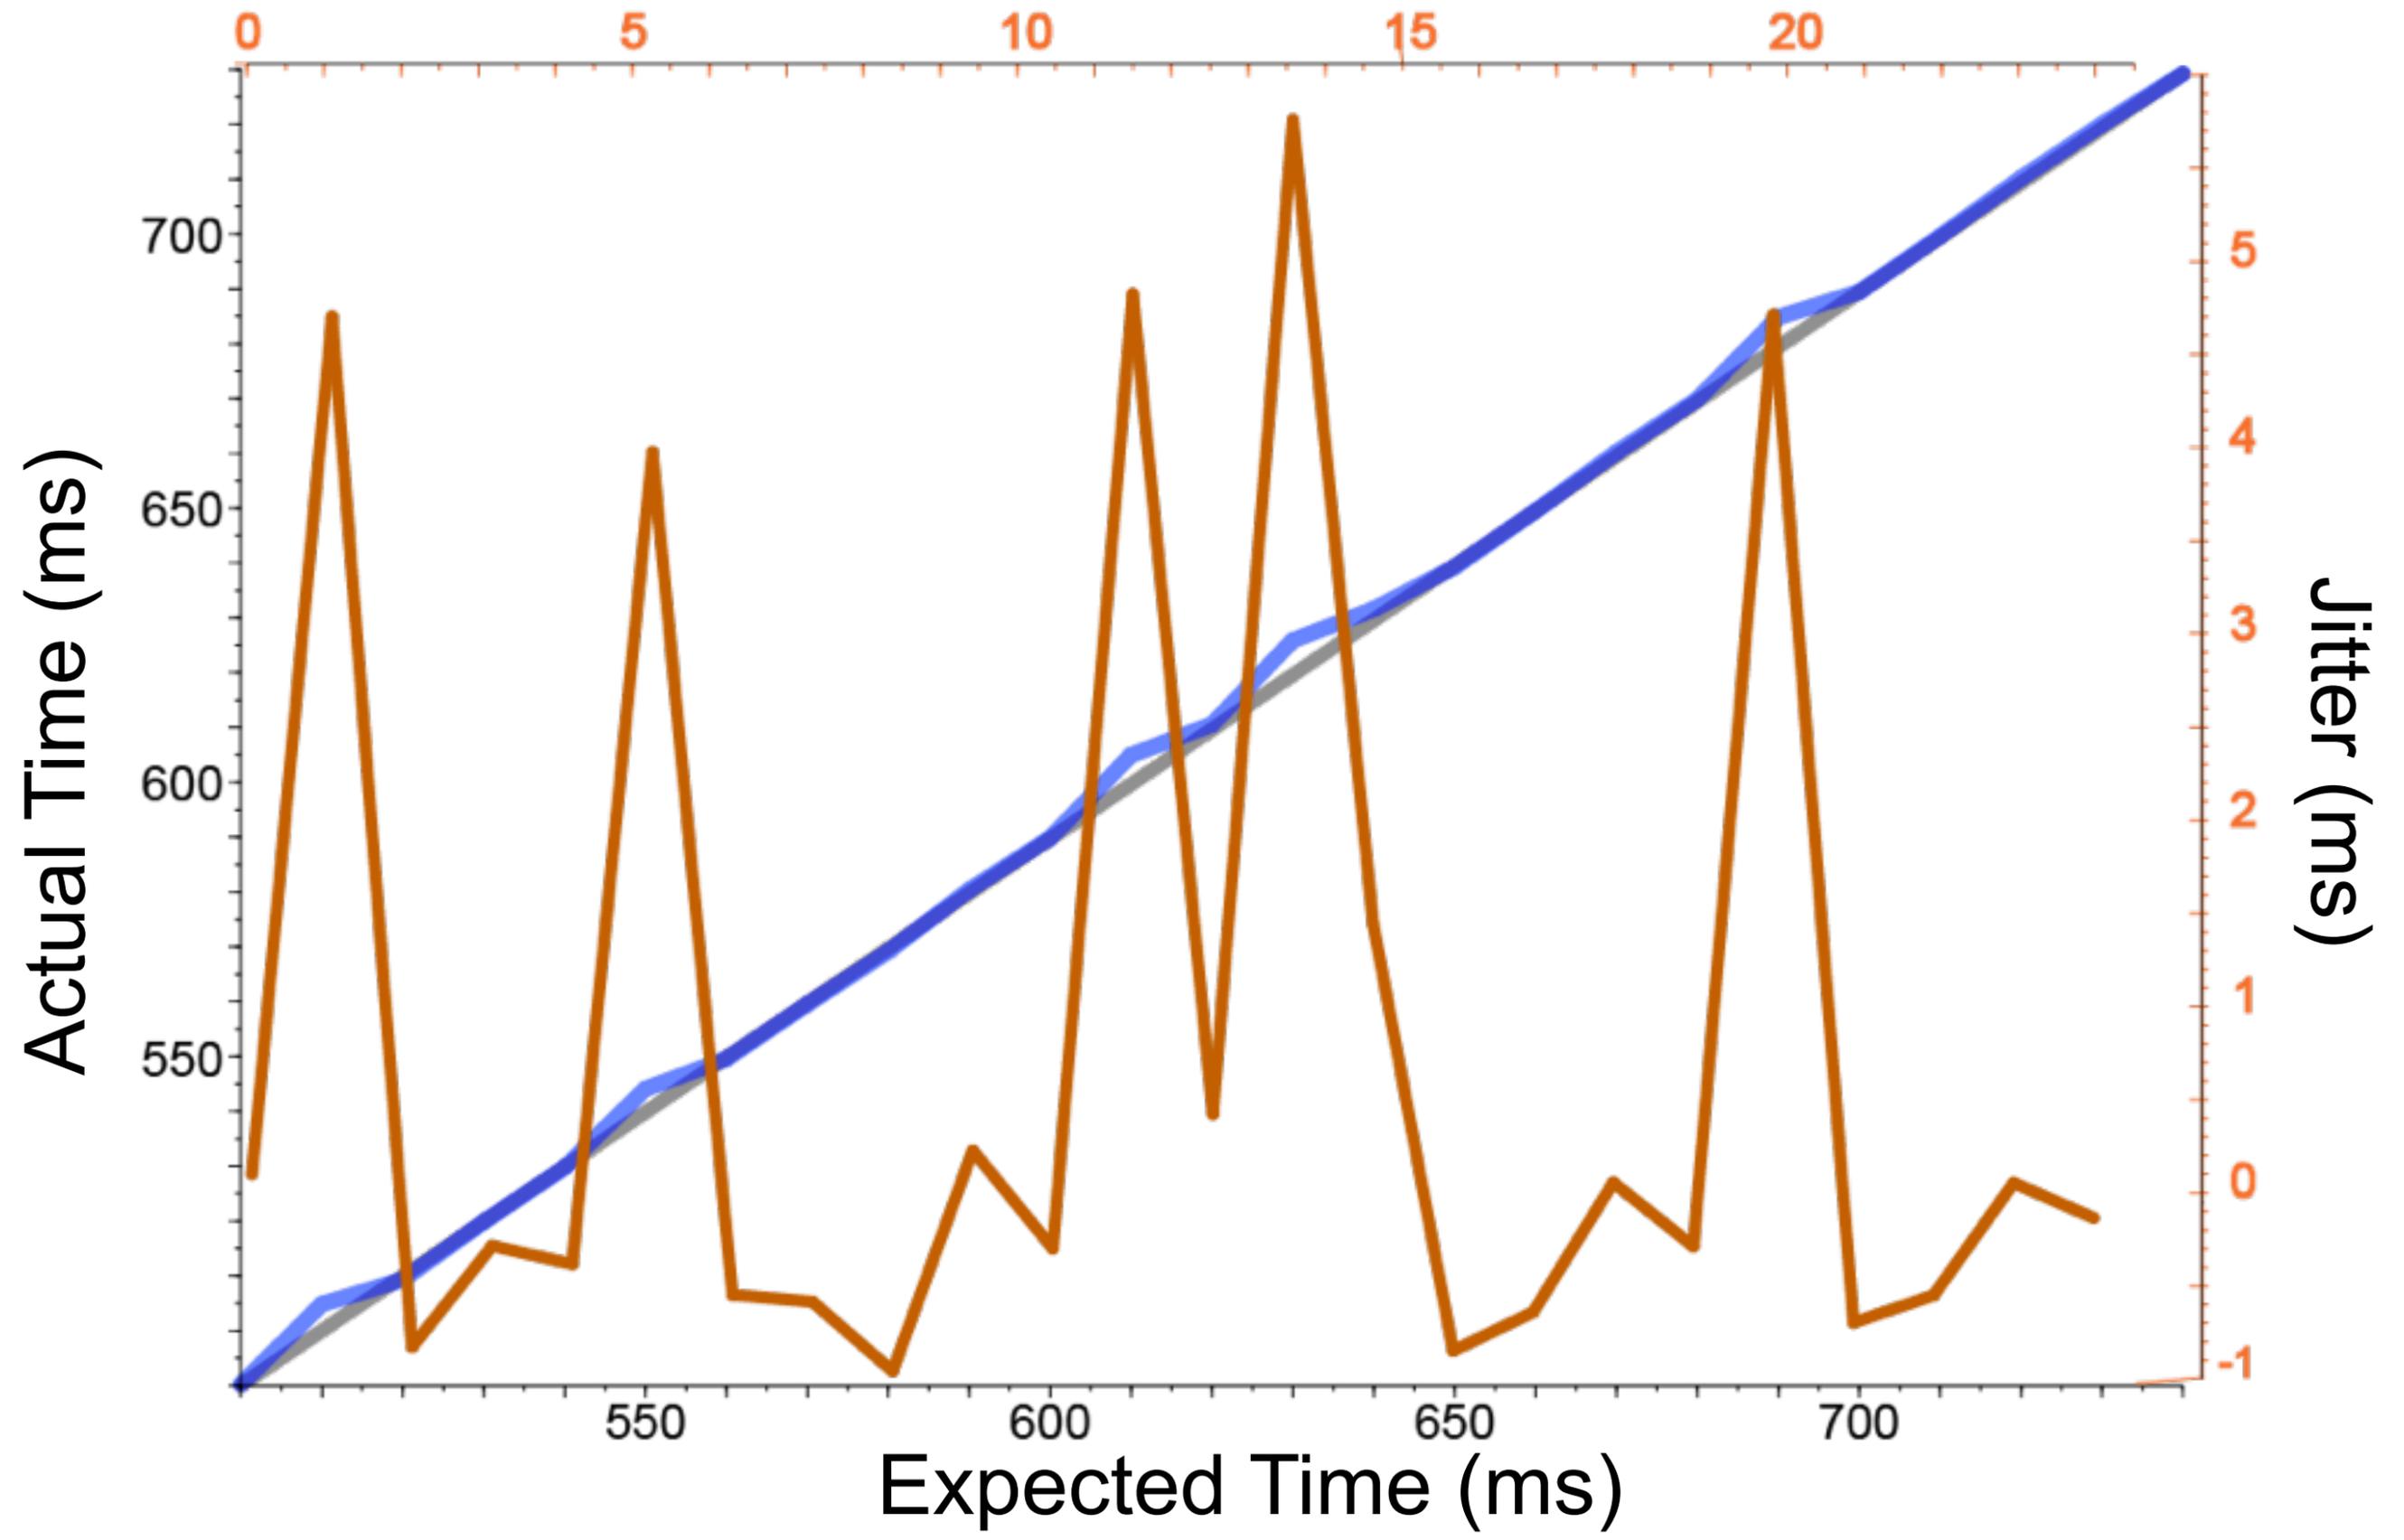




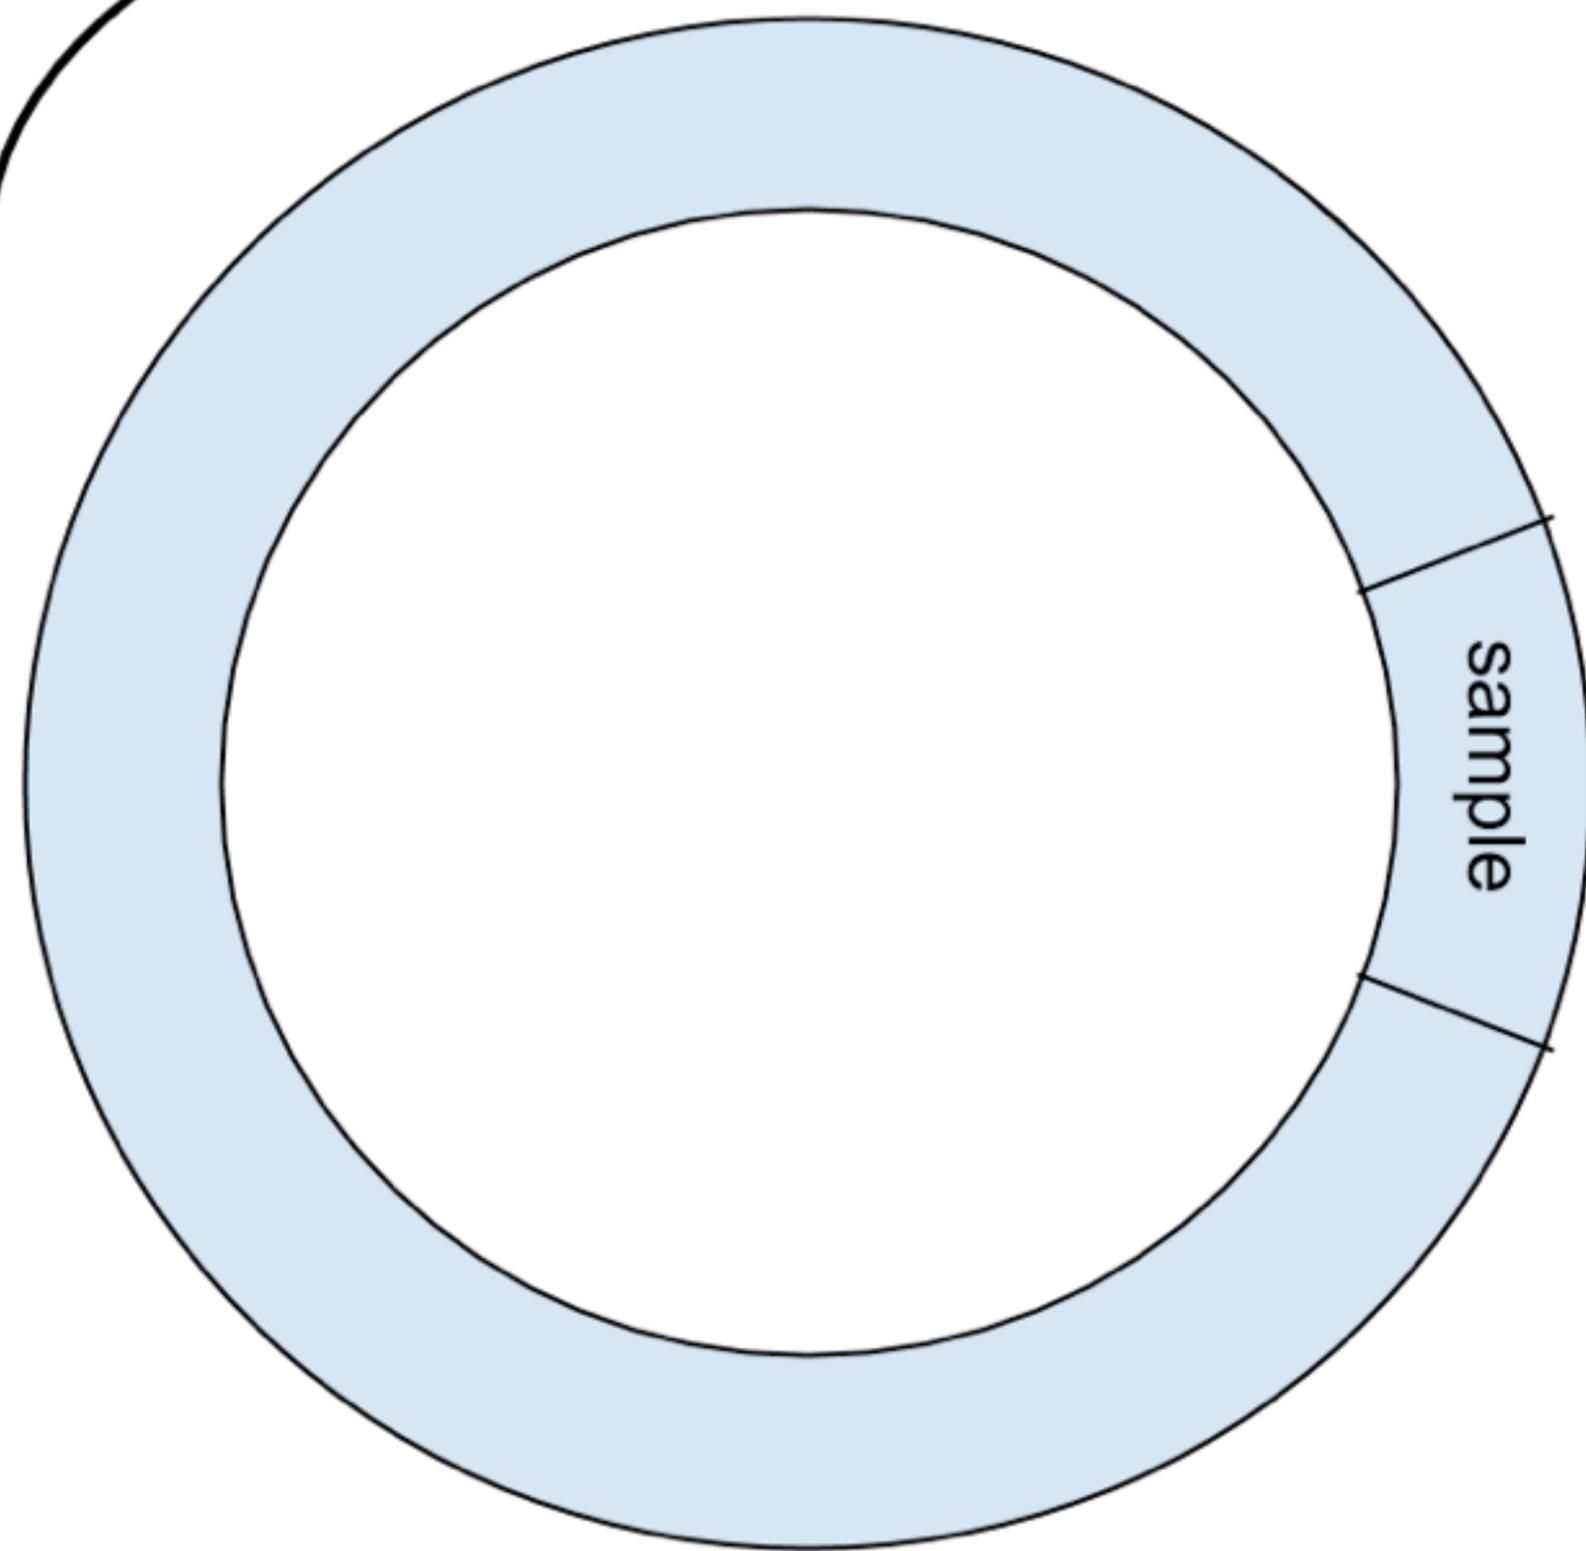
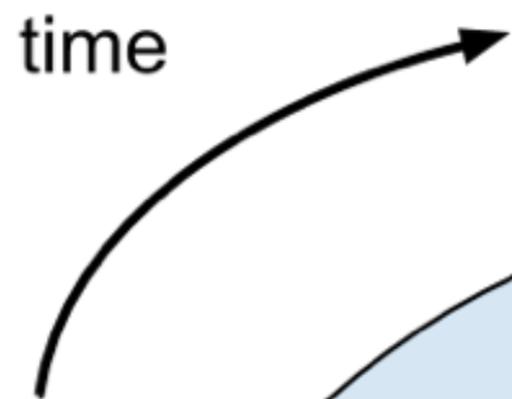




# Callback Count



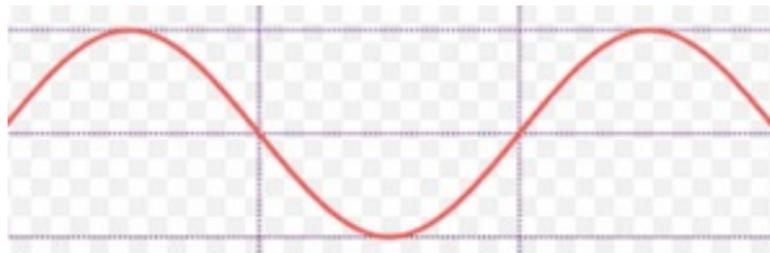
time



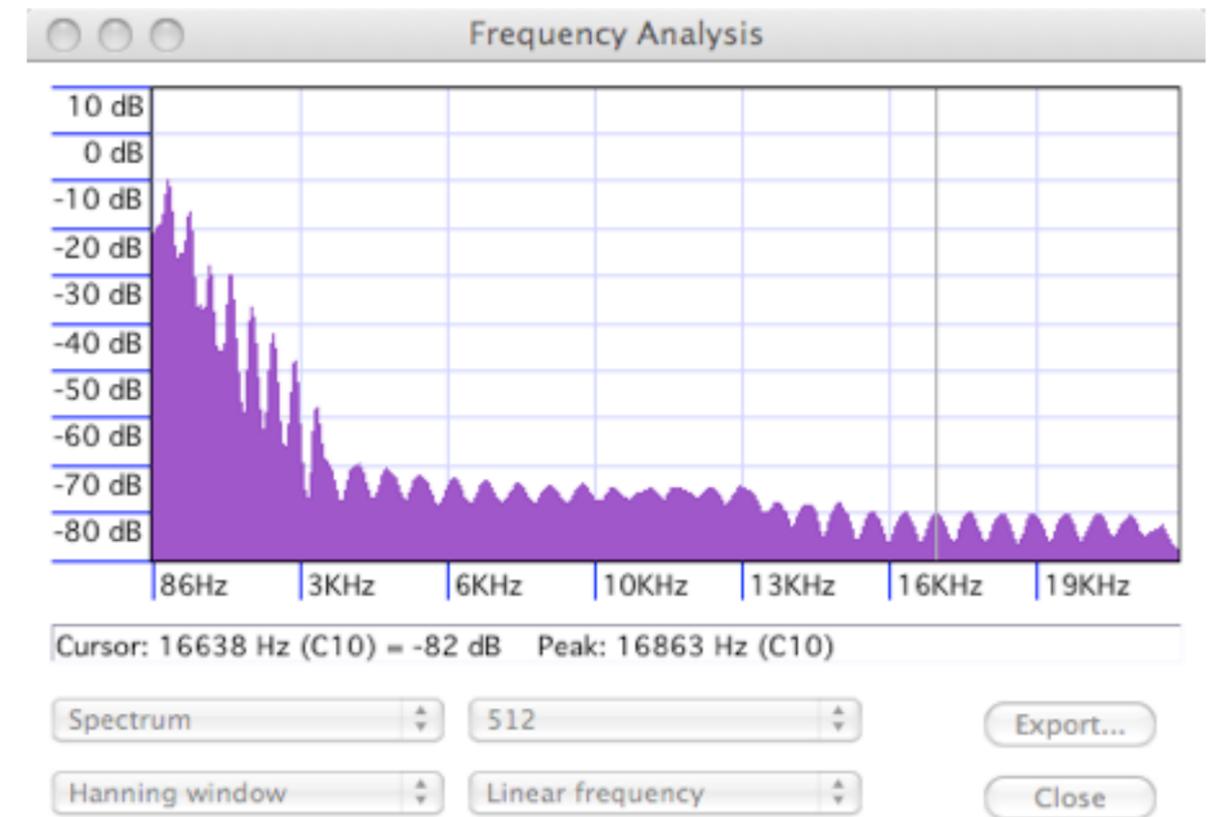
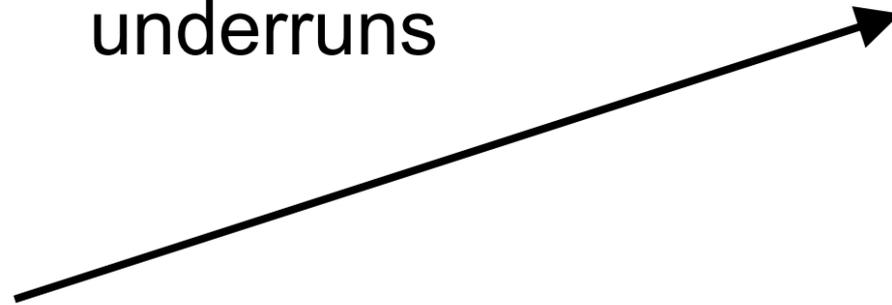
each context  
switch:  
current time  
+  
task ID



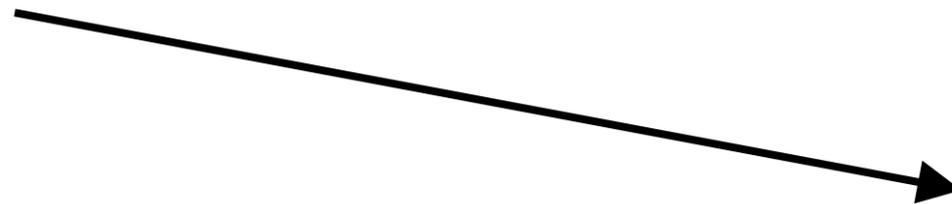
play audio



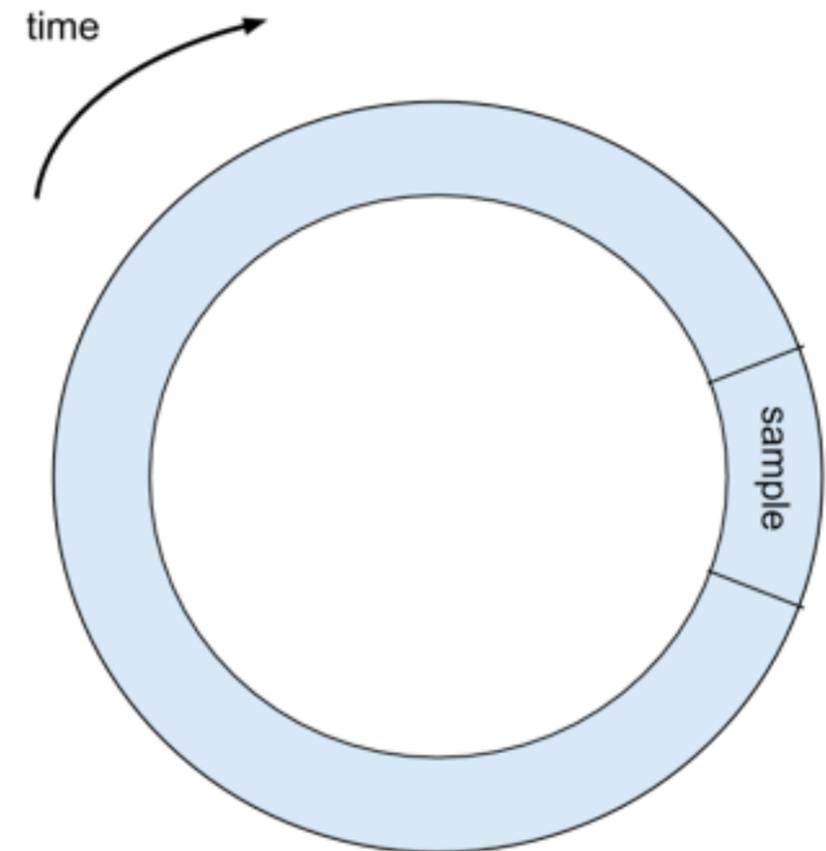
listen for  
underruns



record  
context  
switches



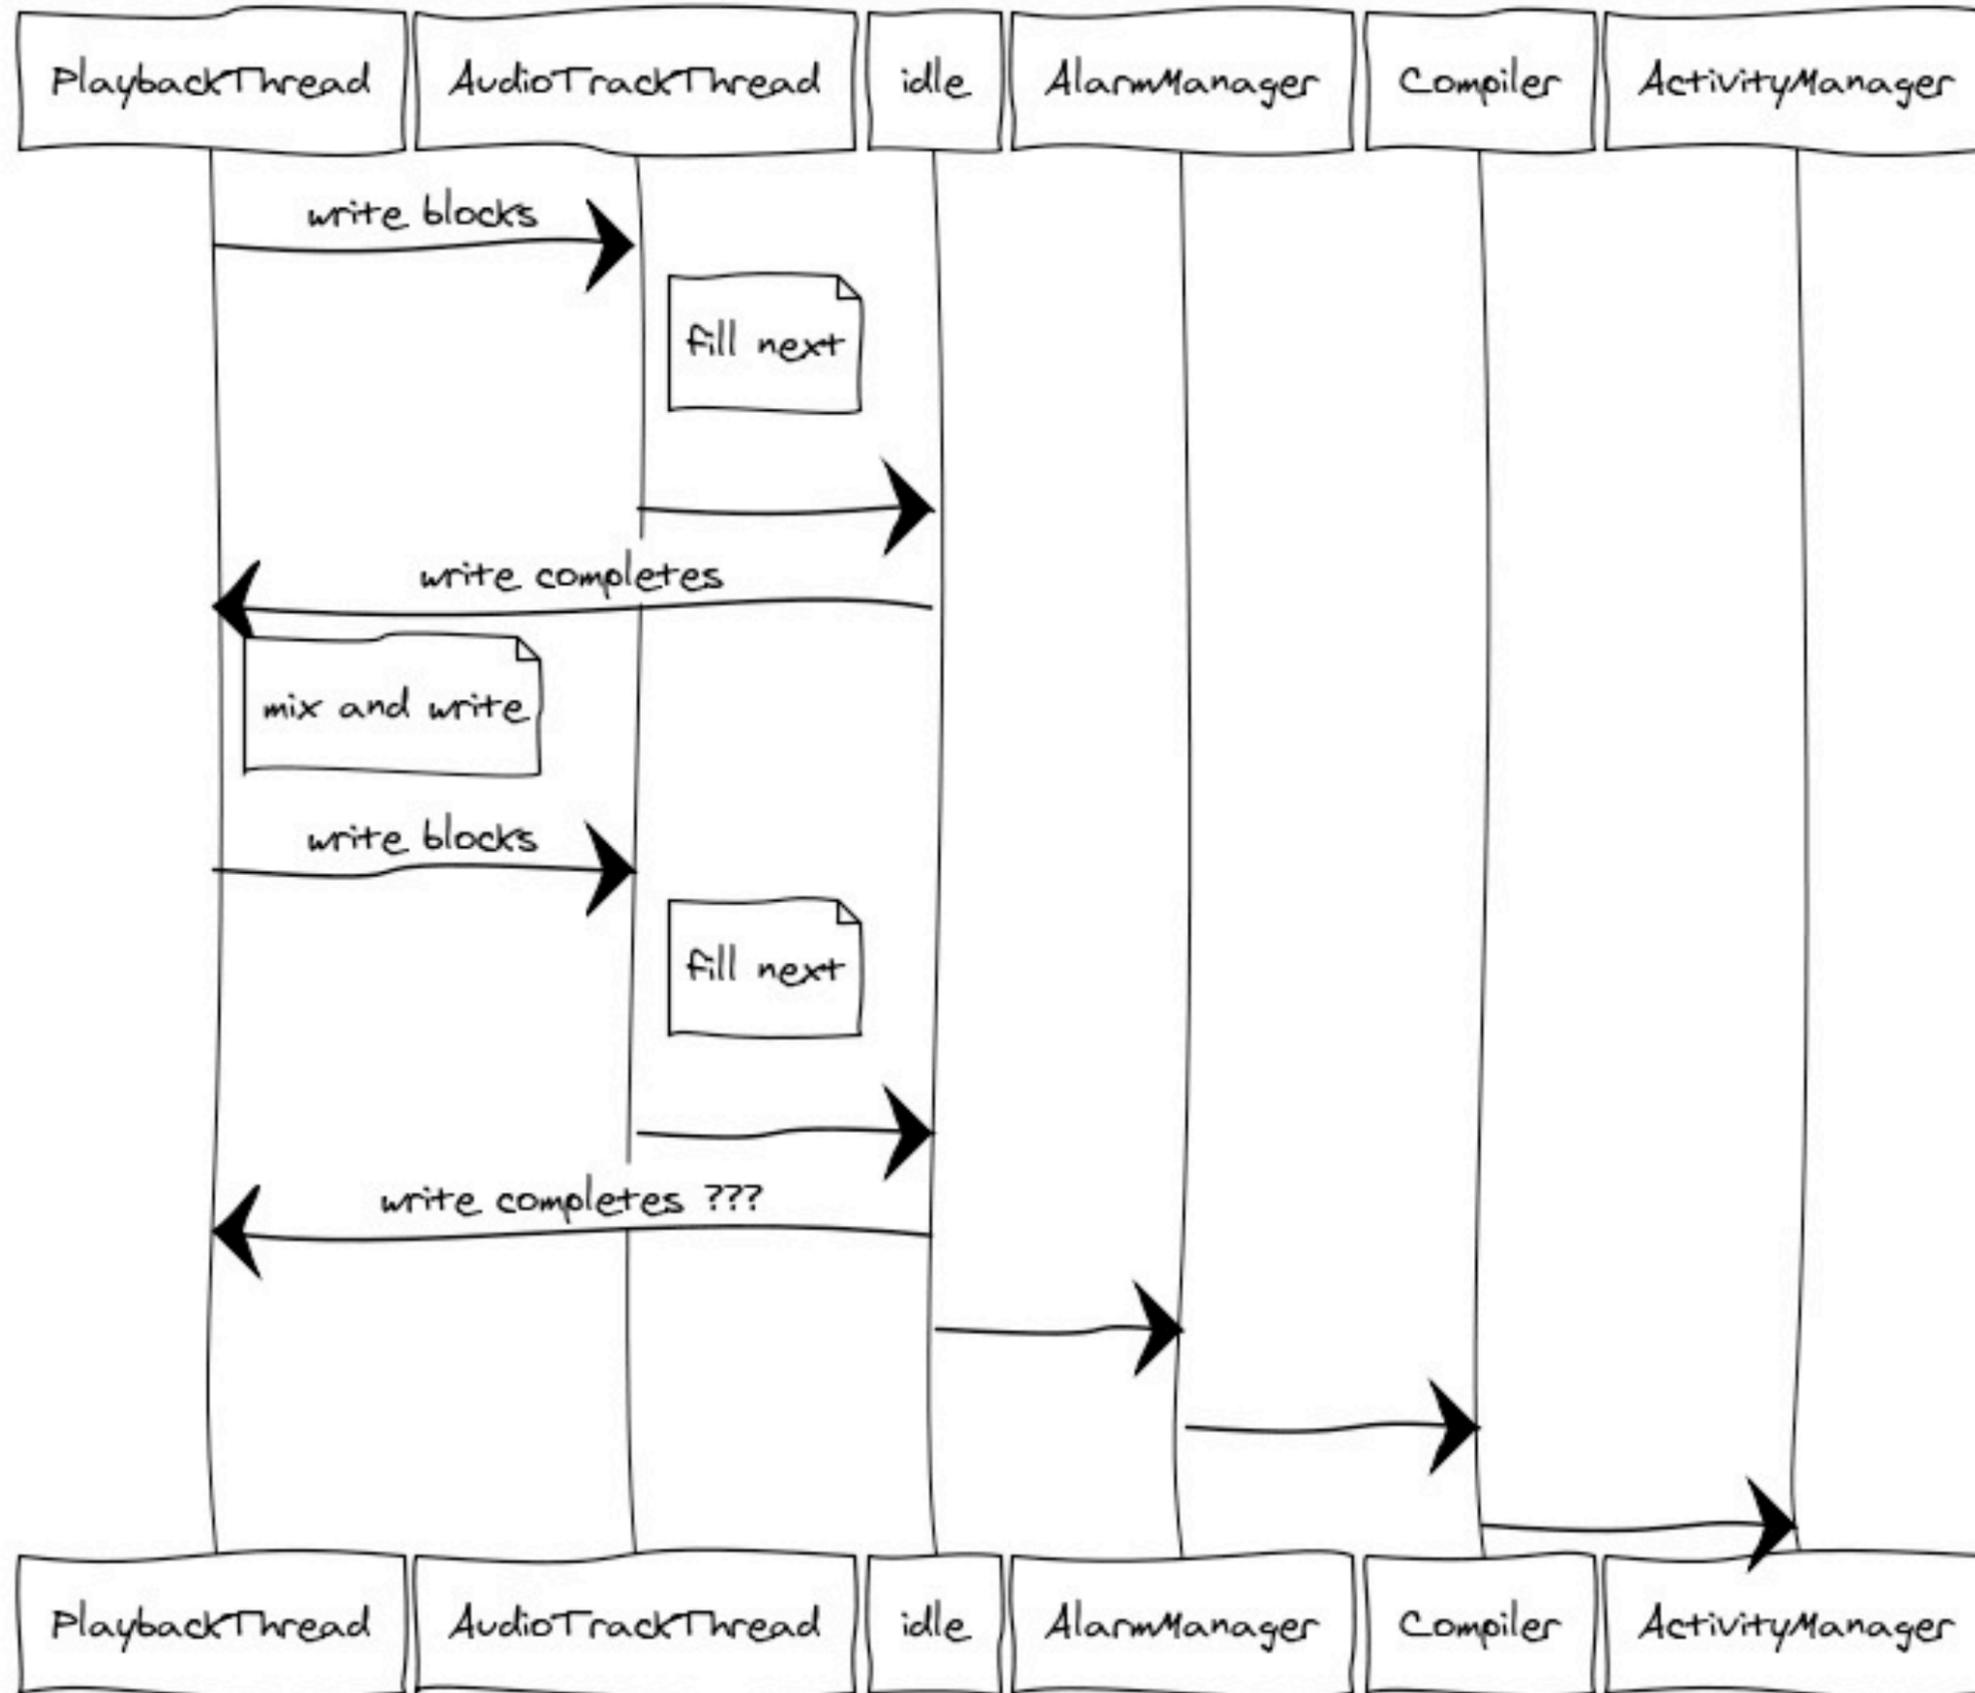
time

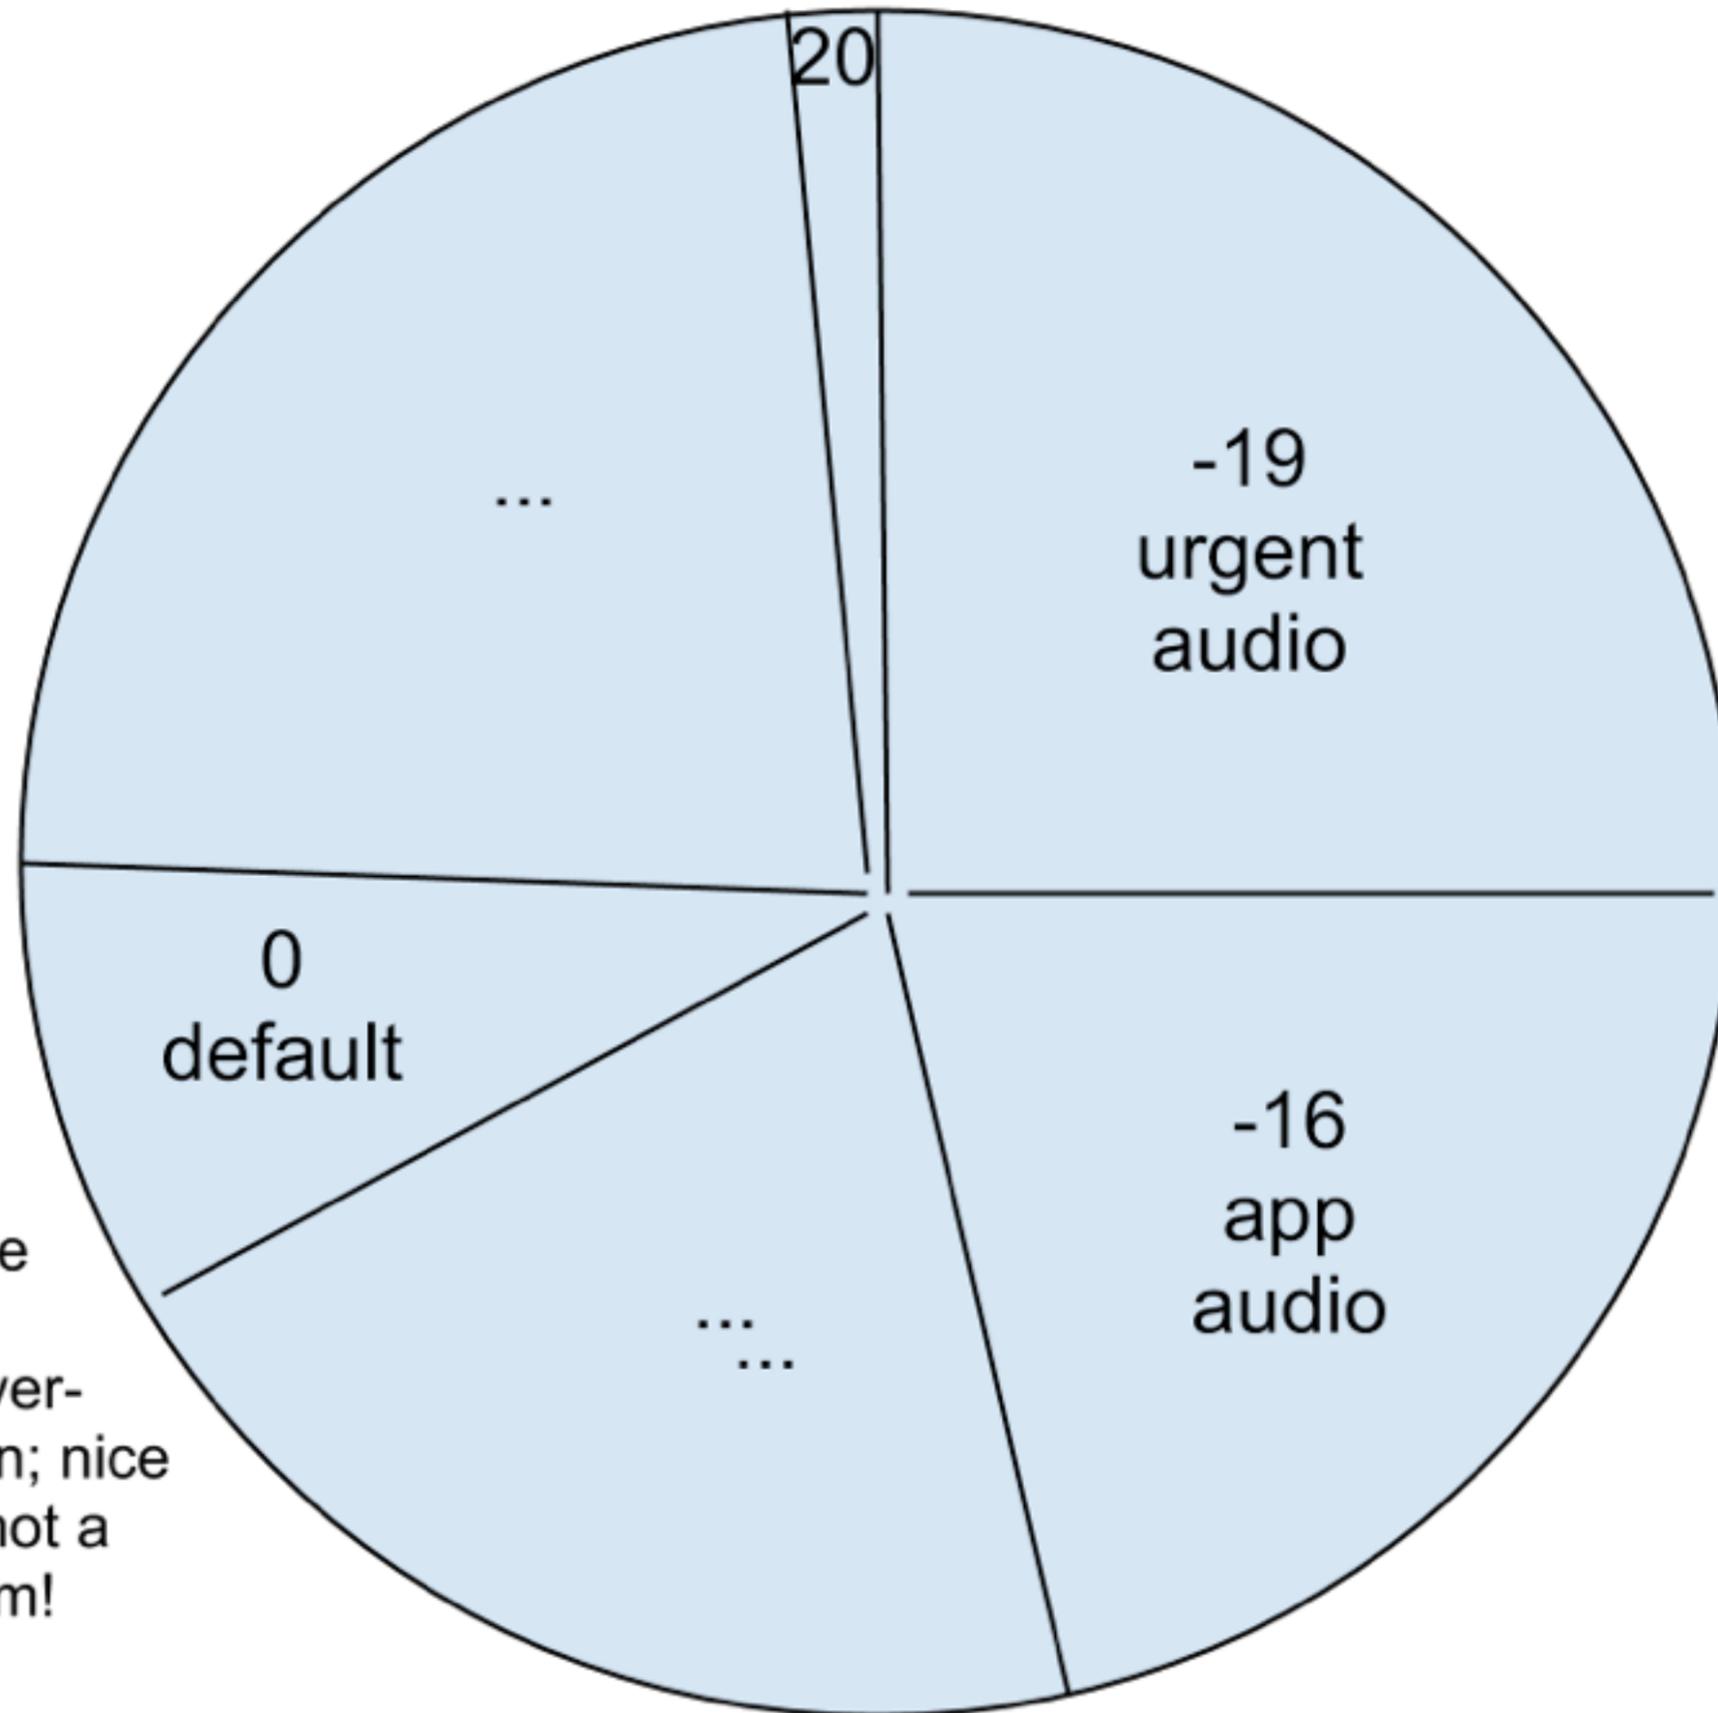


Timestamp	Tid	Name	Delta time
2024372840	146	Playback Thread	
2024416840	900	AudioTrackThrea	44000
2024439840	0	idle	23000
2027275840	146	Playback Thread	2836000
2027319840	900	AudioTrackThrea	44000
2027339840	0	idle	20000
2030177840	146	Playback Thread	2838000
2030221840	900	AudioTrackThrea	44000
2030243840	0	idle	22000
2033081840	146	Playback Thread	2838000
2033126840	900	AudioTrackThrea	45000
2033145840	0	idle	19000
2035997840	192	AlarmManager	2852000
2036728840	160	Compiler	731000
2037092840	168	ActivityManager	364000
2037770840	167	er.ServerThread	678000



# The smoking gun





\* not to scale

\*\* a *huge* over-simplification; nice values are not a quota system!





CAT VIDEOS

**AUDIO**

CAT VIDEOS

**AUDIO**

EVIL PLOTS

**AUDIO**

SLACKING  
OFF

**AUDIO**

CAT VIDEOS

**AUDIO**

CAT VIDEOS

**AUDIO**

THUMB  
TWIDDLING

**AUDIO**

NEWS  
UPDATES

**AUDIO**

CAT VIDEOS

**AUDIO**

THUMB  
TWIDDLING

**AUDIO**

THUMB  
TWIDDLING

**AUDIO**

CAT VIDEOS

**AUDIO**

CAT VIDEOS

**AUDIO**

SLACKING  
OFF

**AUDIO**

CAT VIDEOS

**AUDIO**

NEWS  
UPDATES

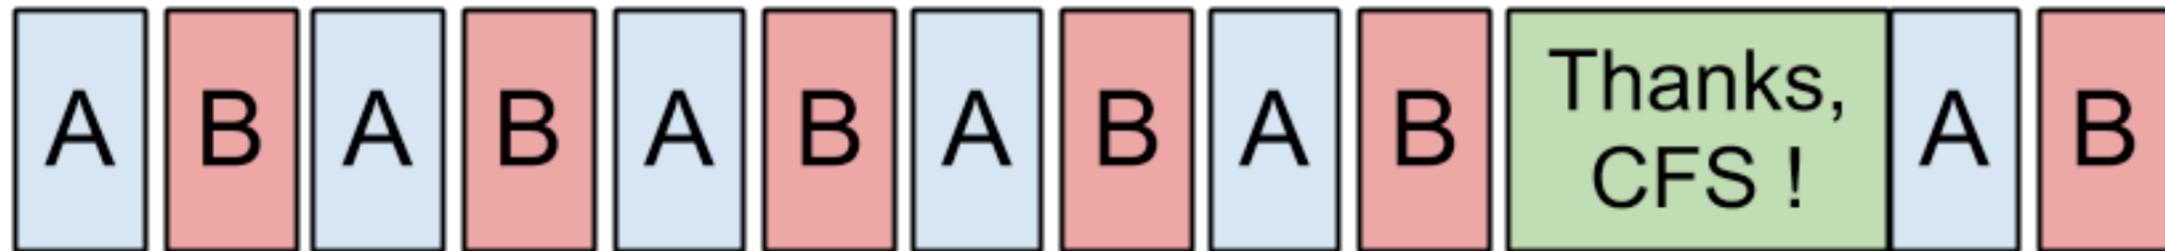
**AUDIO**

EVIL PLOTS



Hey buddy,  
can you fit  
me in ?

Audio Task





Resistor



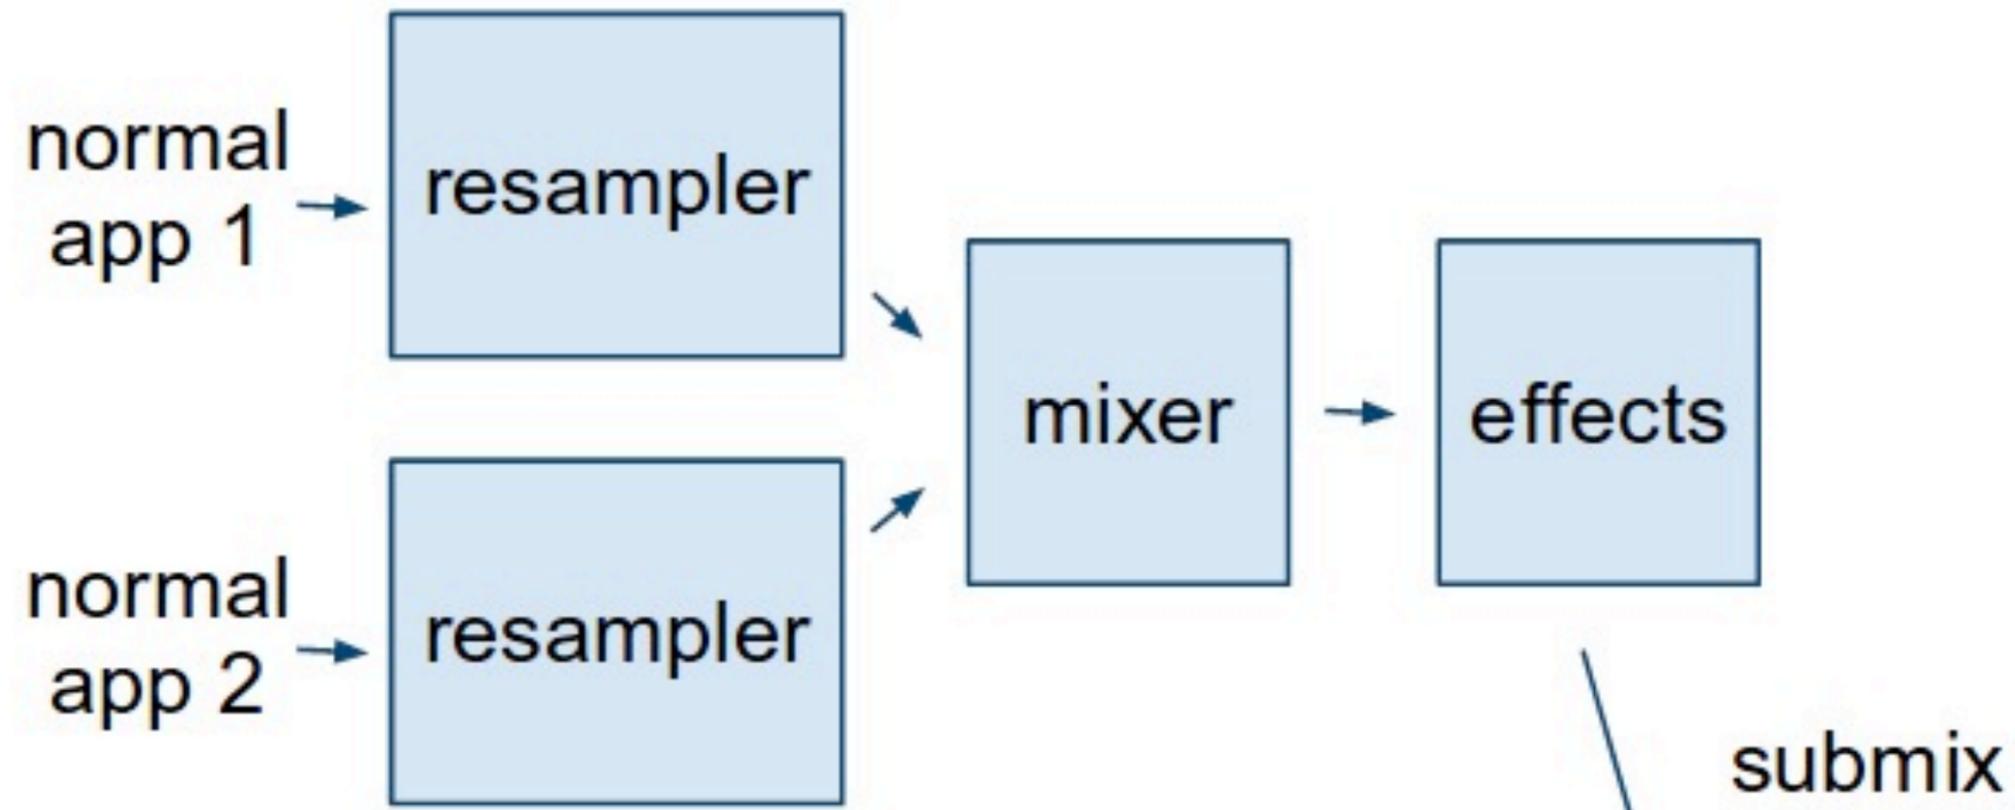
Mixer



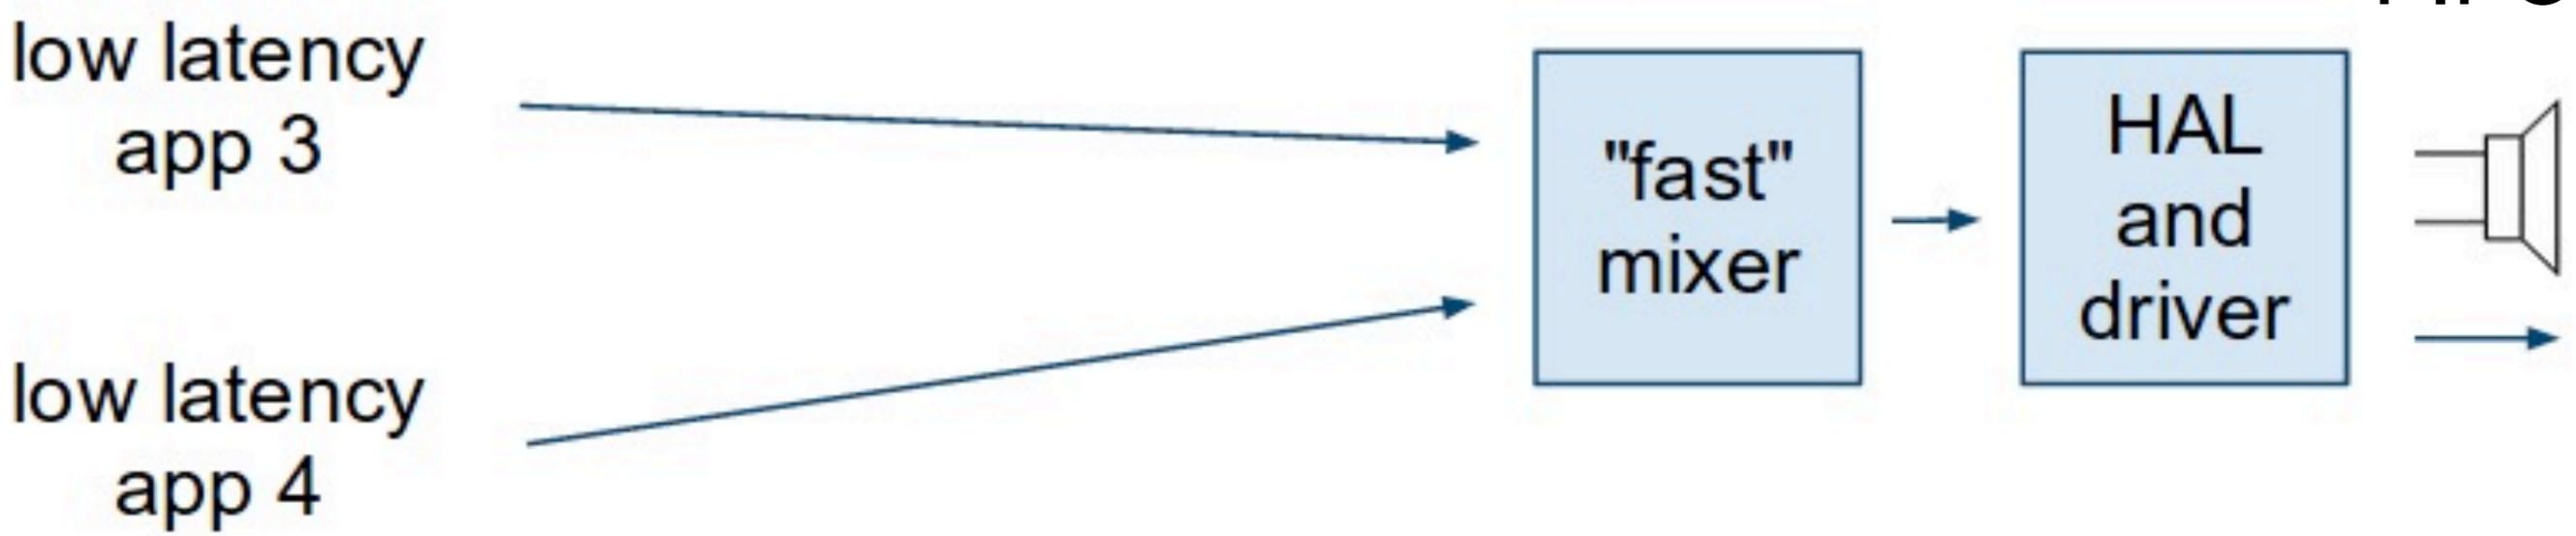
Control



CFS

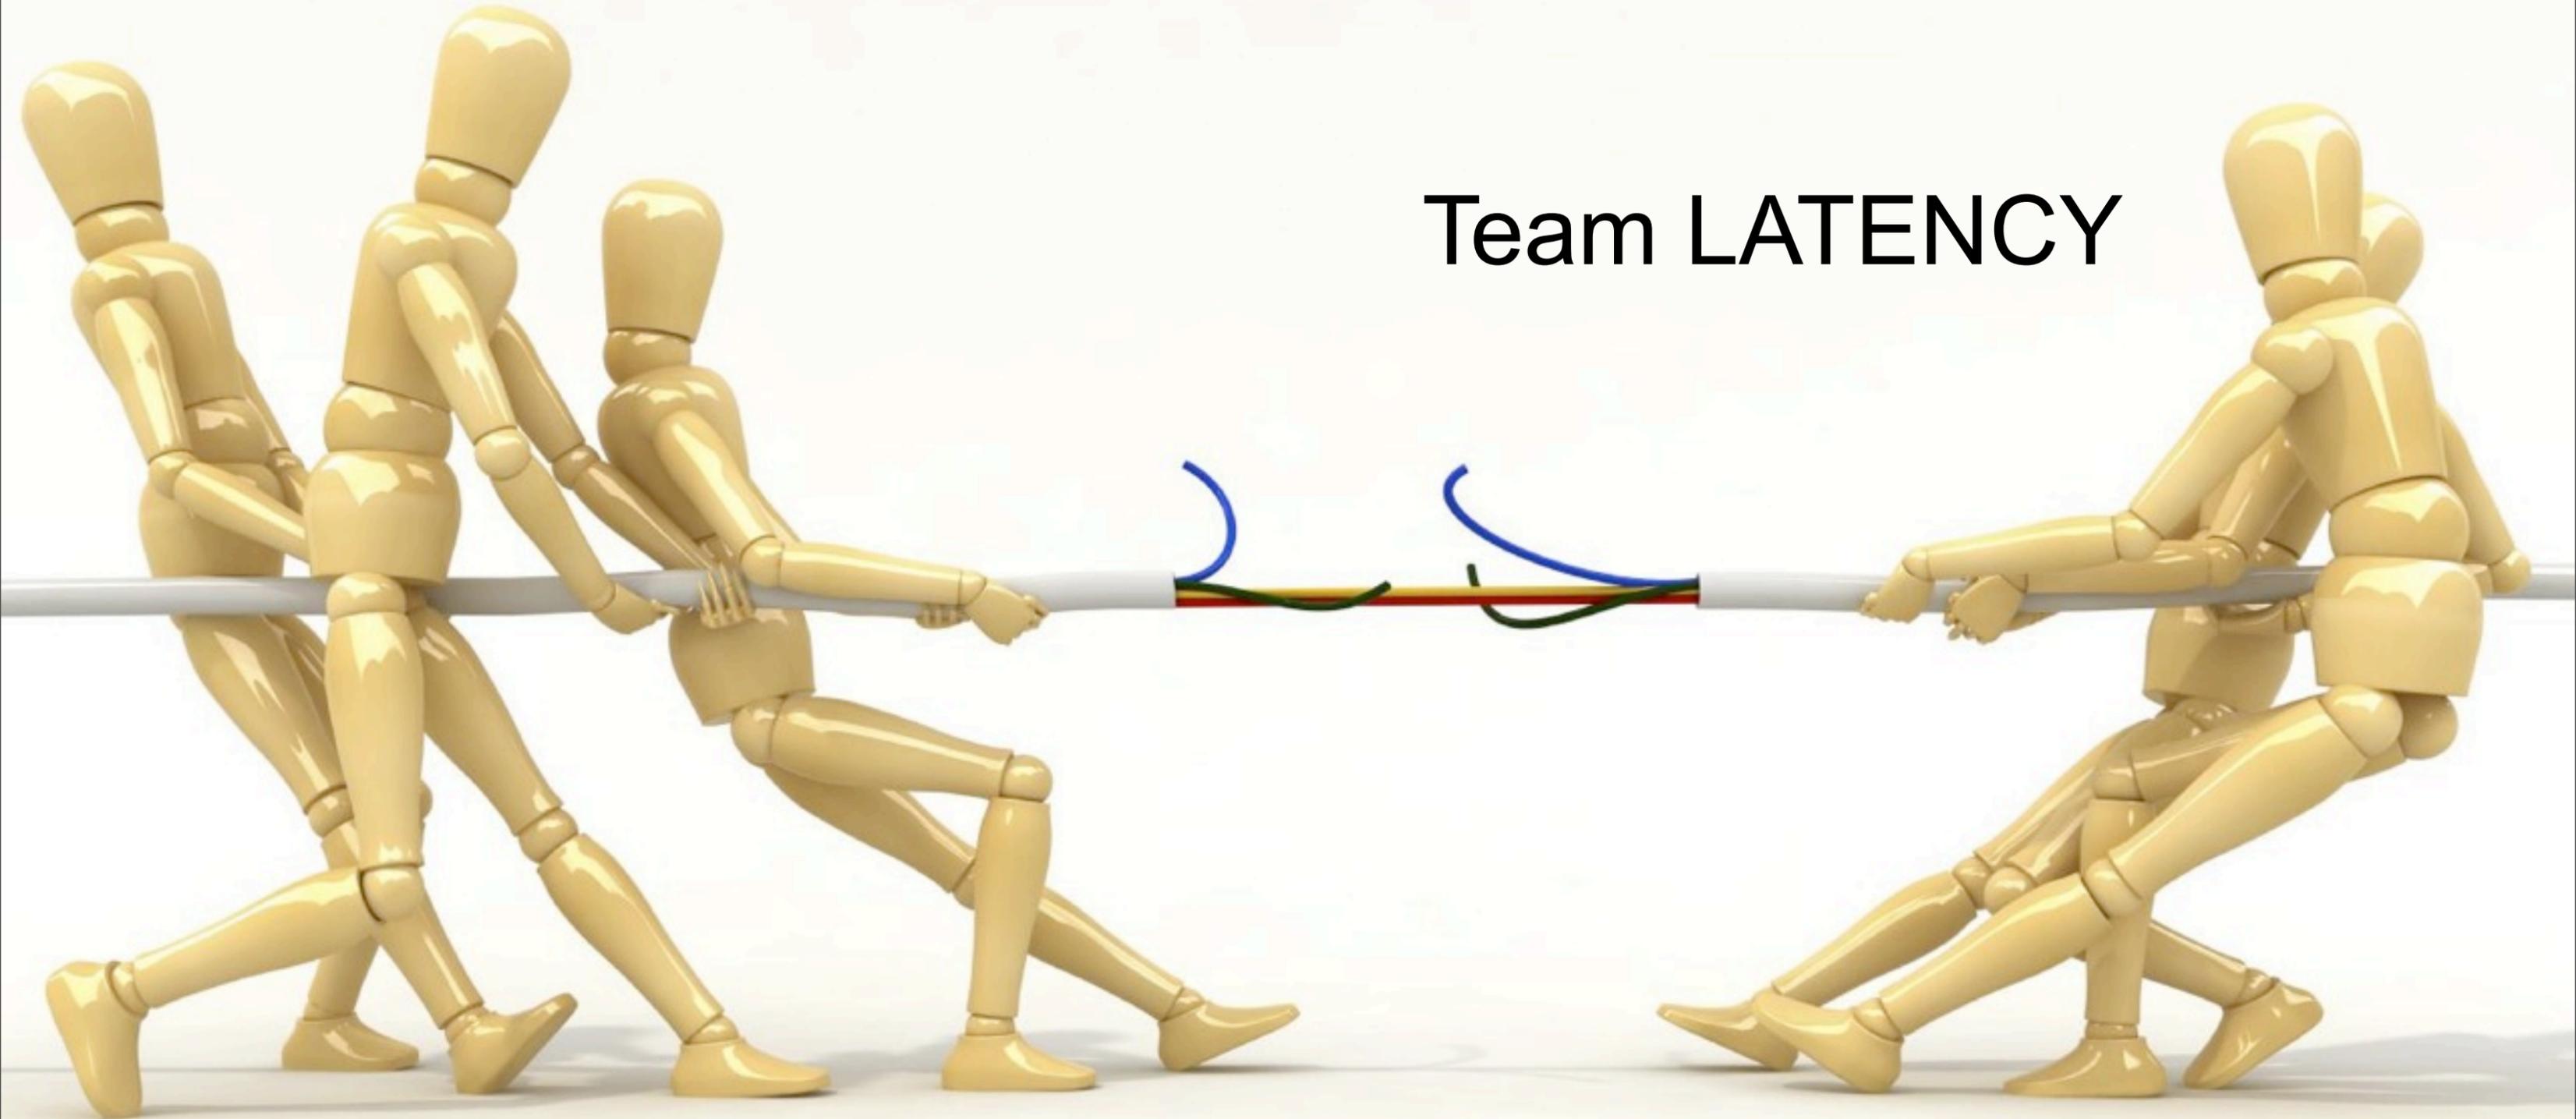


FIFO



# Team POWER

# Team LATENCY



	Sample Rate	Buffer Size (frames)	Buffer Size (ms)
Nexus S	44100	880	19.95
Galaxy Nexus	44100	144 (JB) 968 (ICS)	3.26 (JB) 21.95 (ICS)
Nexus 4	48000	240	5
Nexus 7	44100	512	11.6
Nexus 10	44100	256	5.8



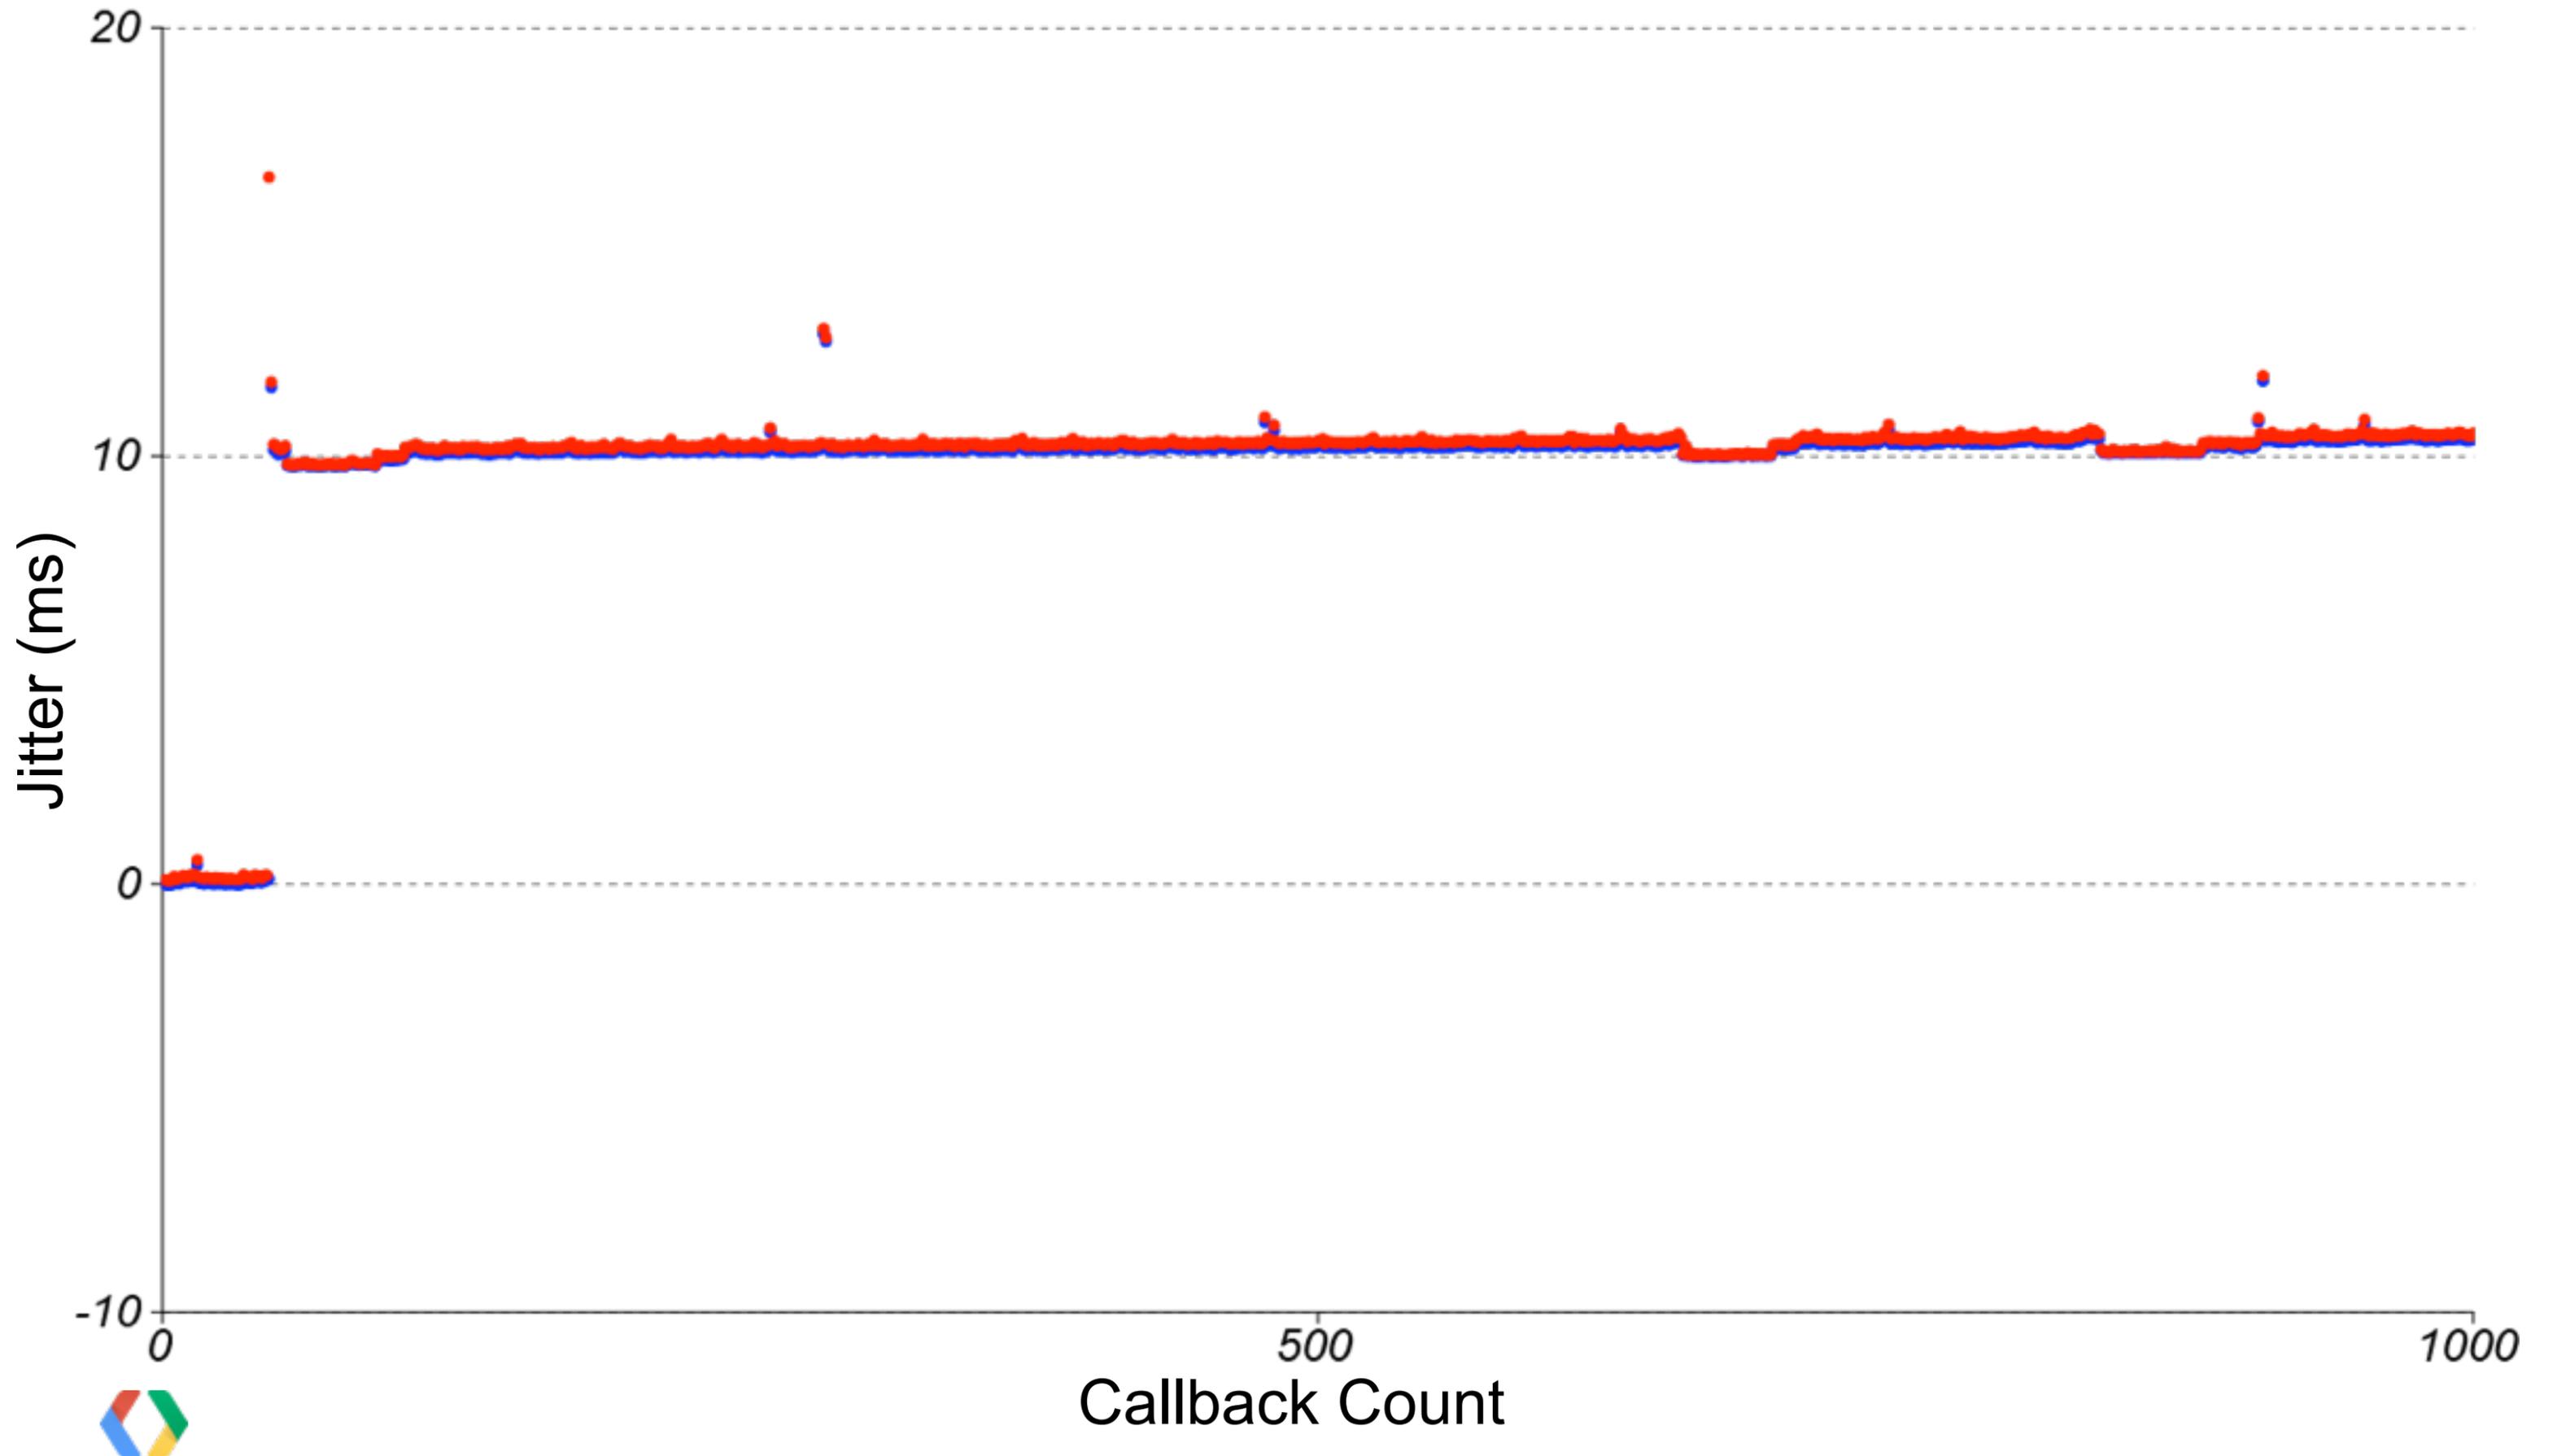
```
@TargetApi(Build.VERSION_CODES.JELLY_BEAN_MR1)
void getJbMr1Params(MyAudioParams params) {
    AudioManager audioManager =
        (AudioManager)this.getSystemService(Context.AUDIO_SERVICE);

    String rate = AudioManager.getProperty(
        AudioManager.PROPERTY_OUTPUT_SAMPLE_RATE);

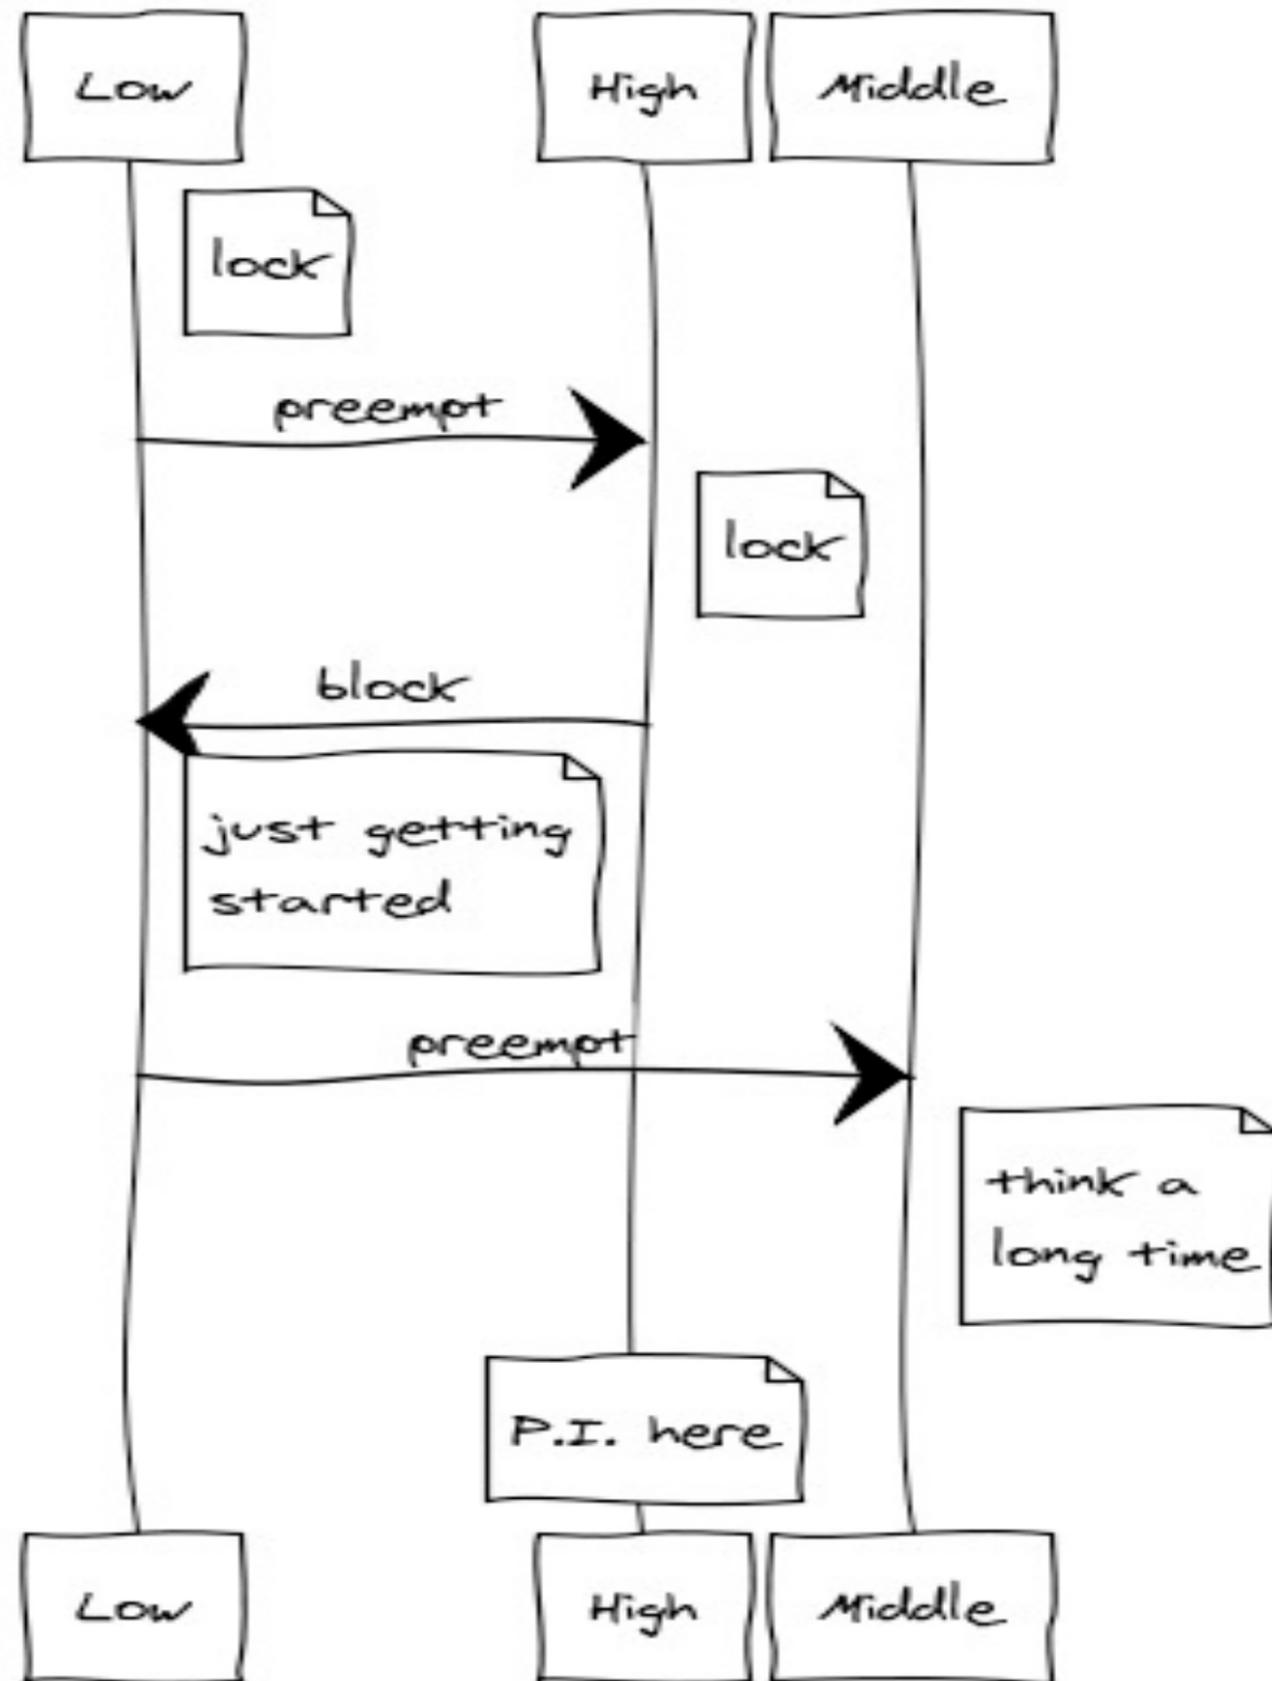
    String size = AudioManager.getProperty(
        AudioManager.PROPERTY_OUTPUT_FRAMES_PER_BUFFER);

    params.sampleRate = Integer.parseInt(rate);
    params.bufferSize = Integer.parseInt(size);
}
```

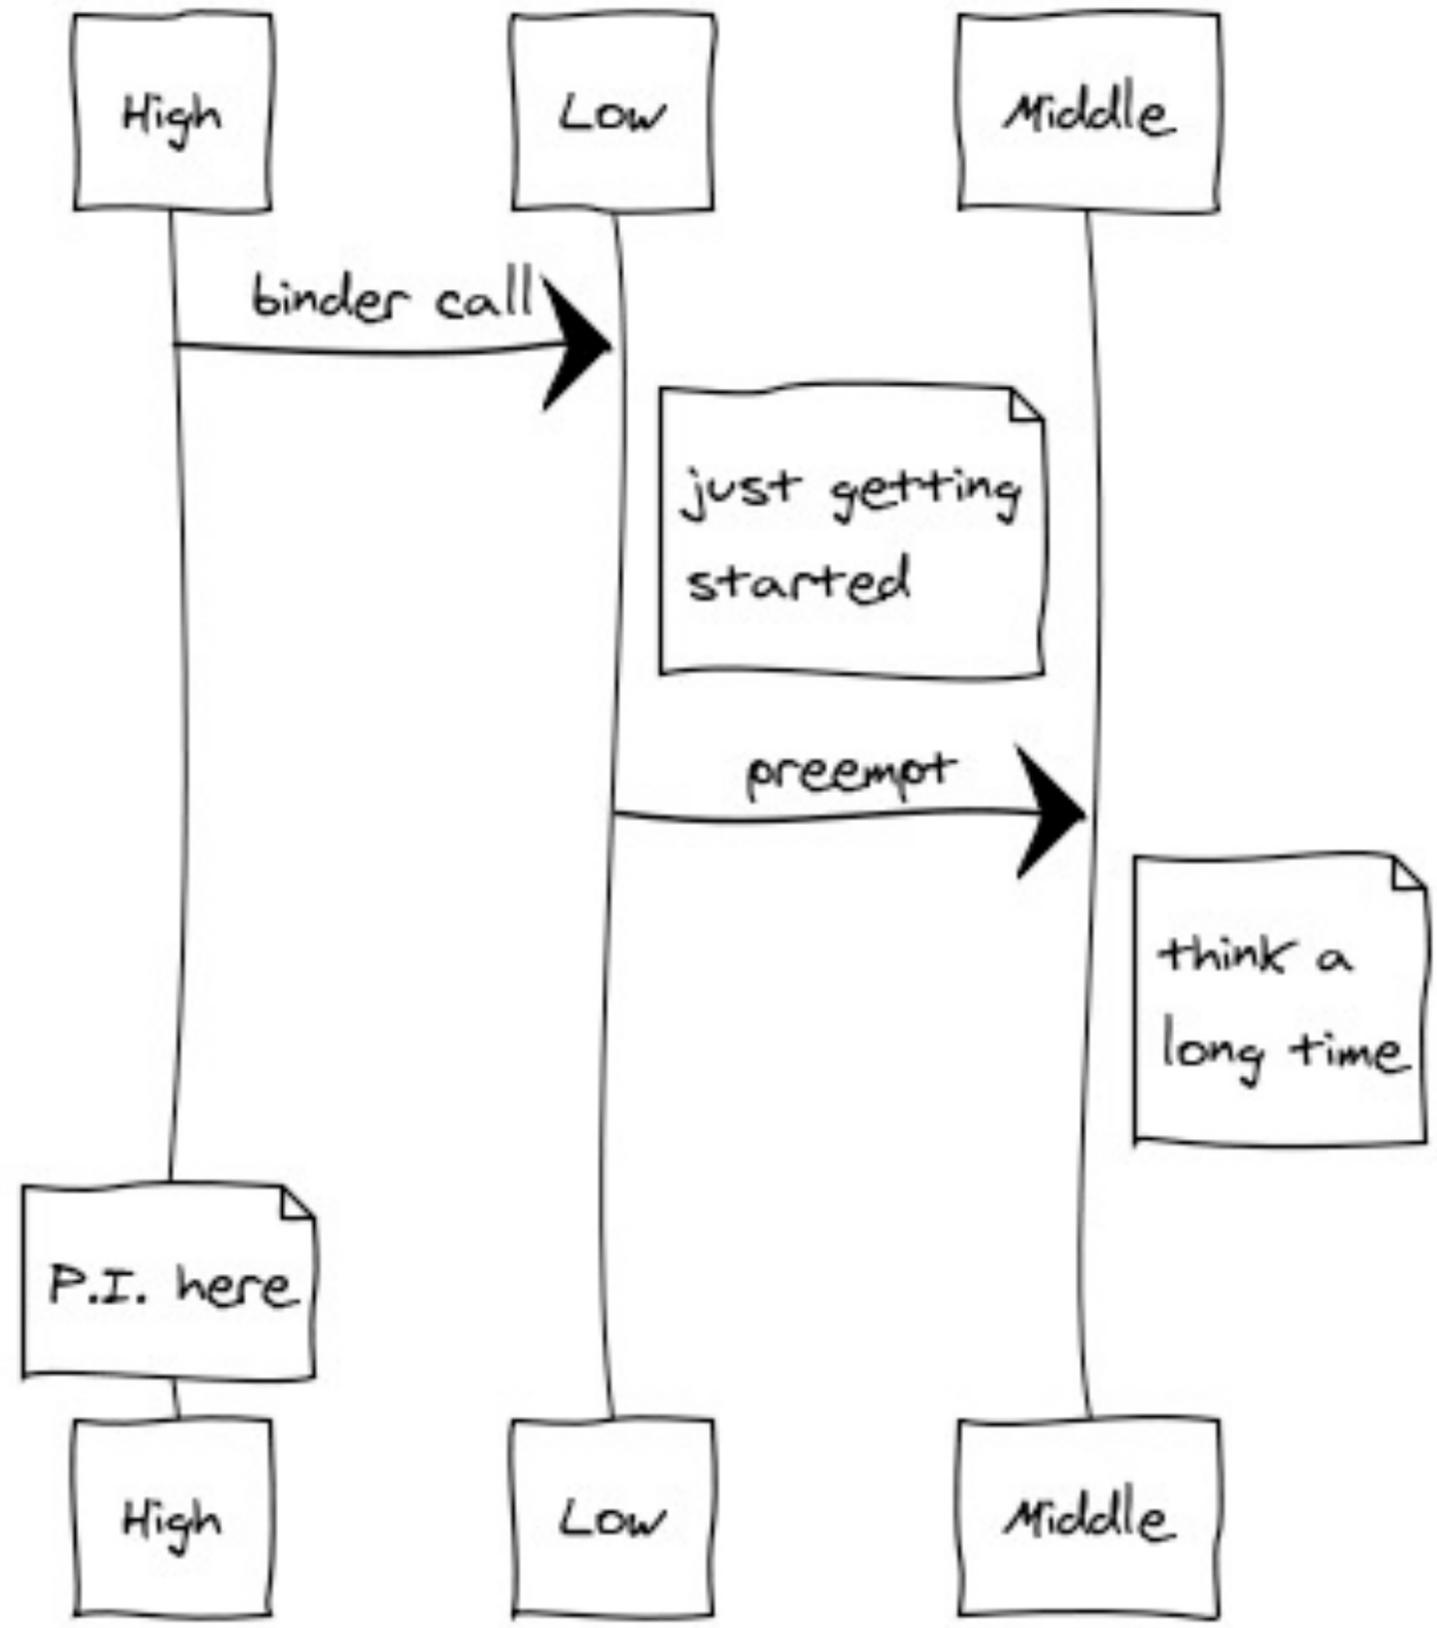


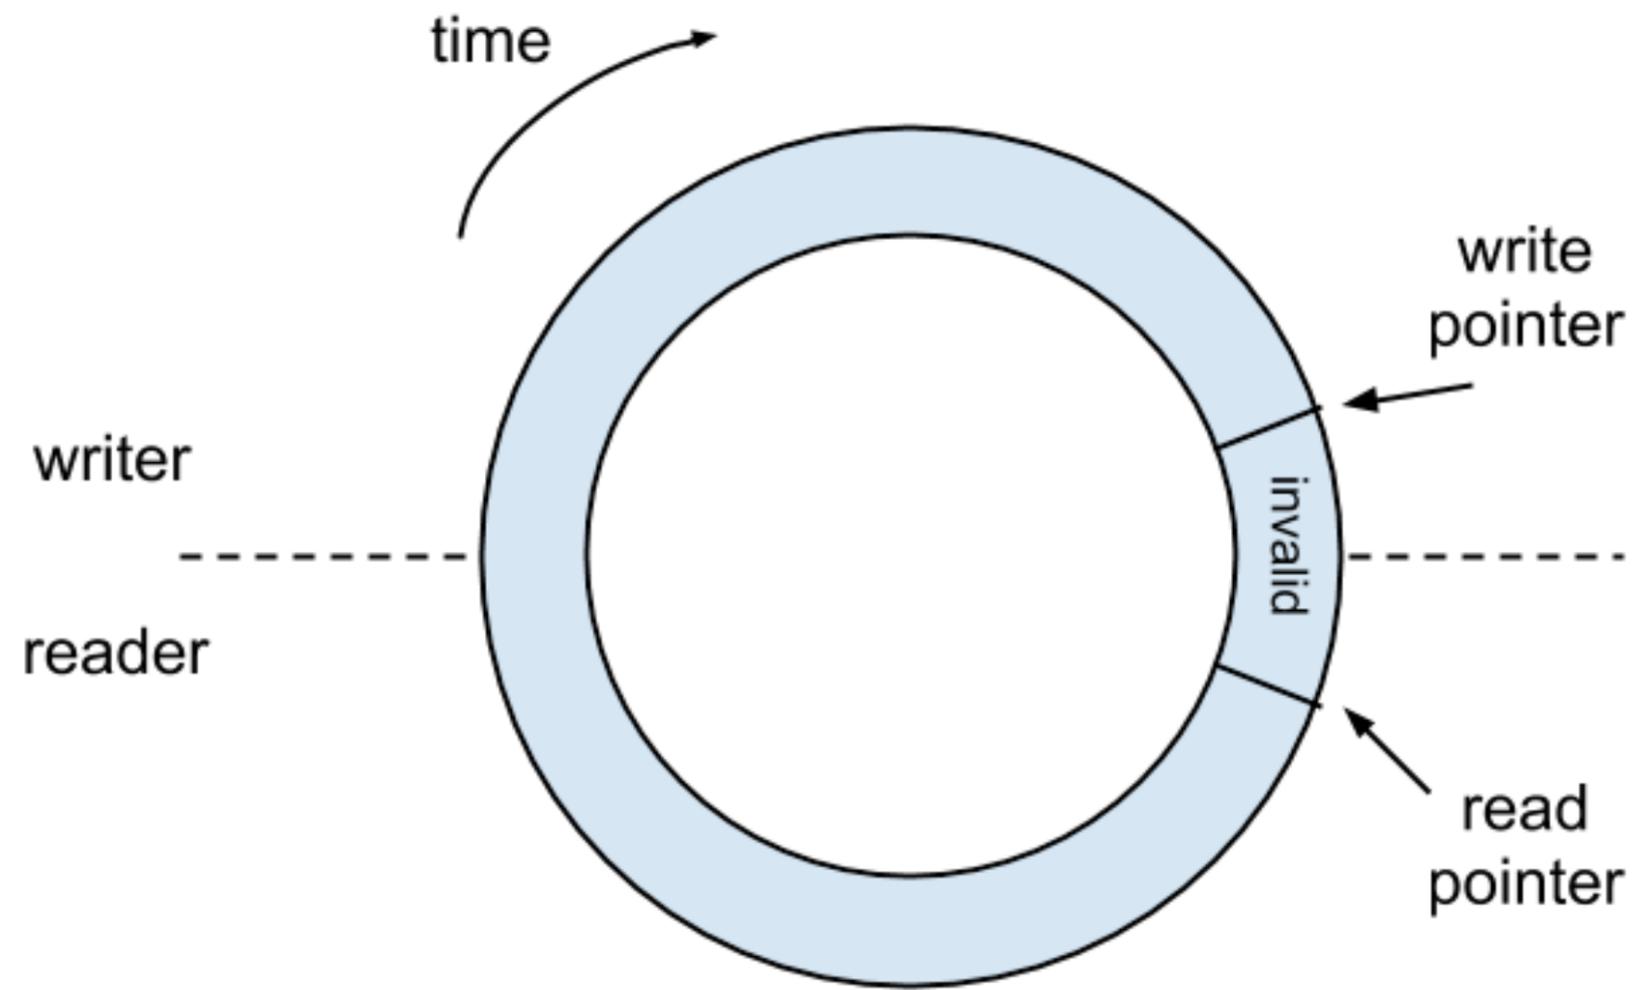


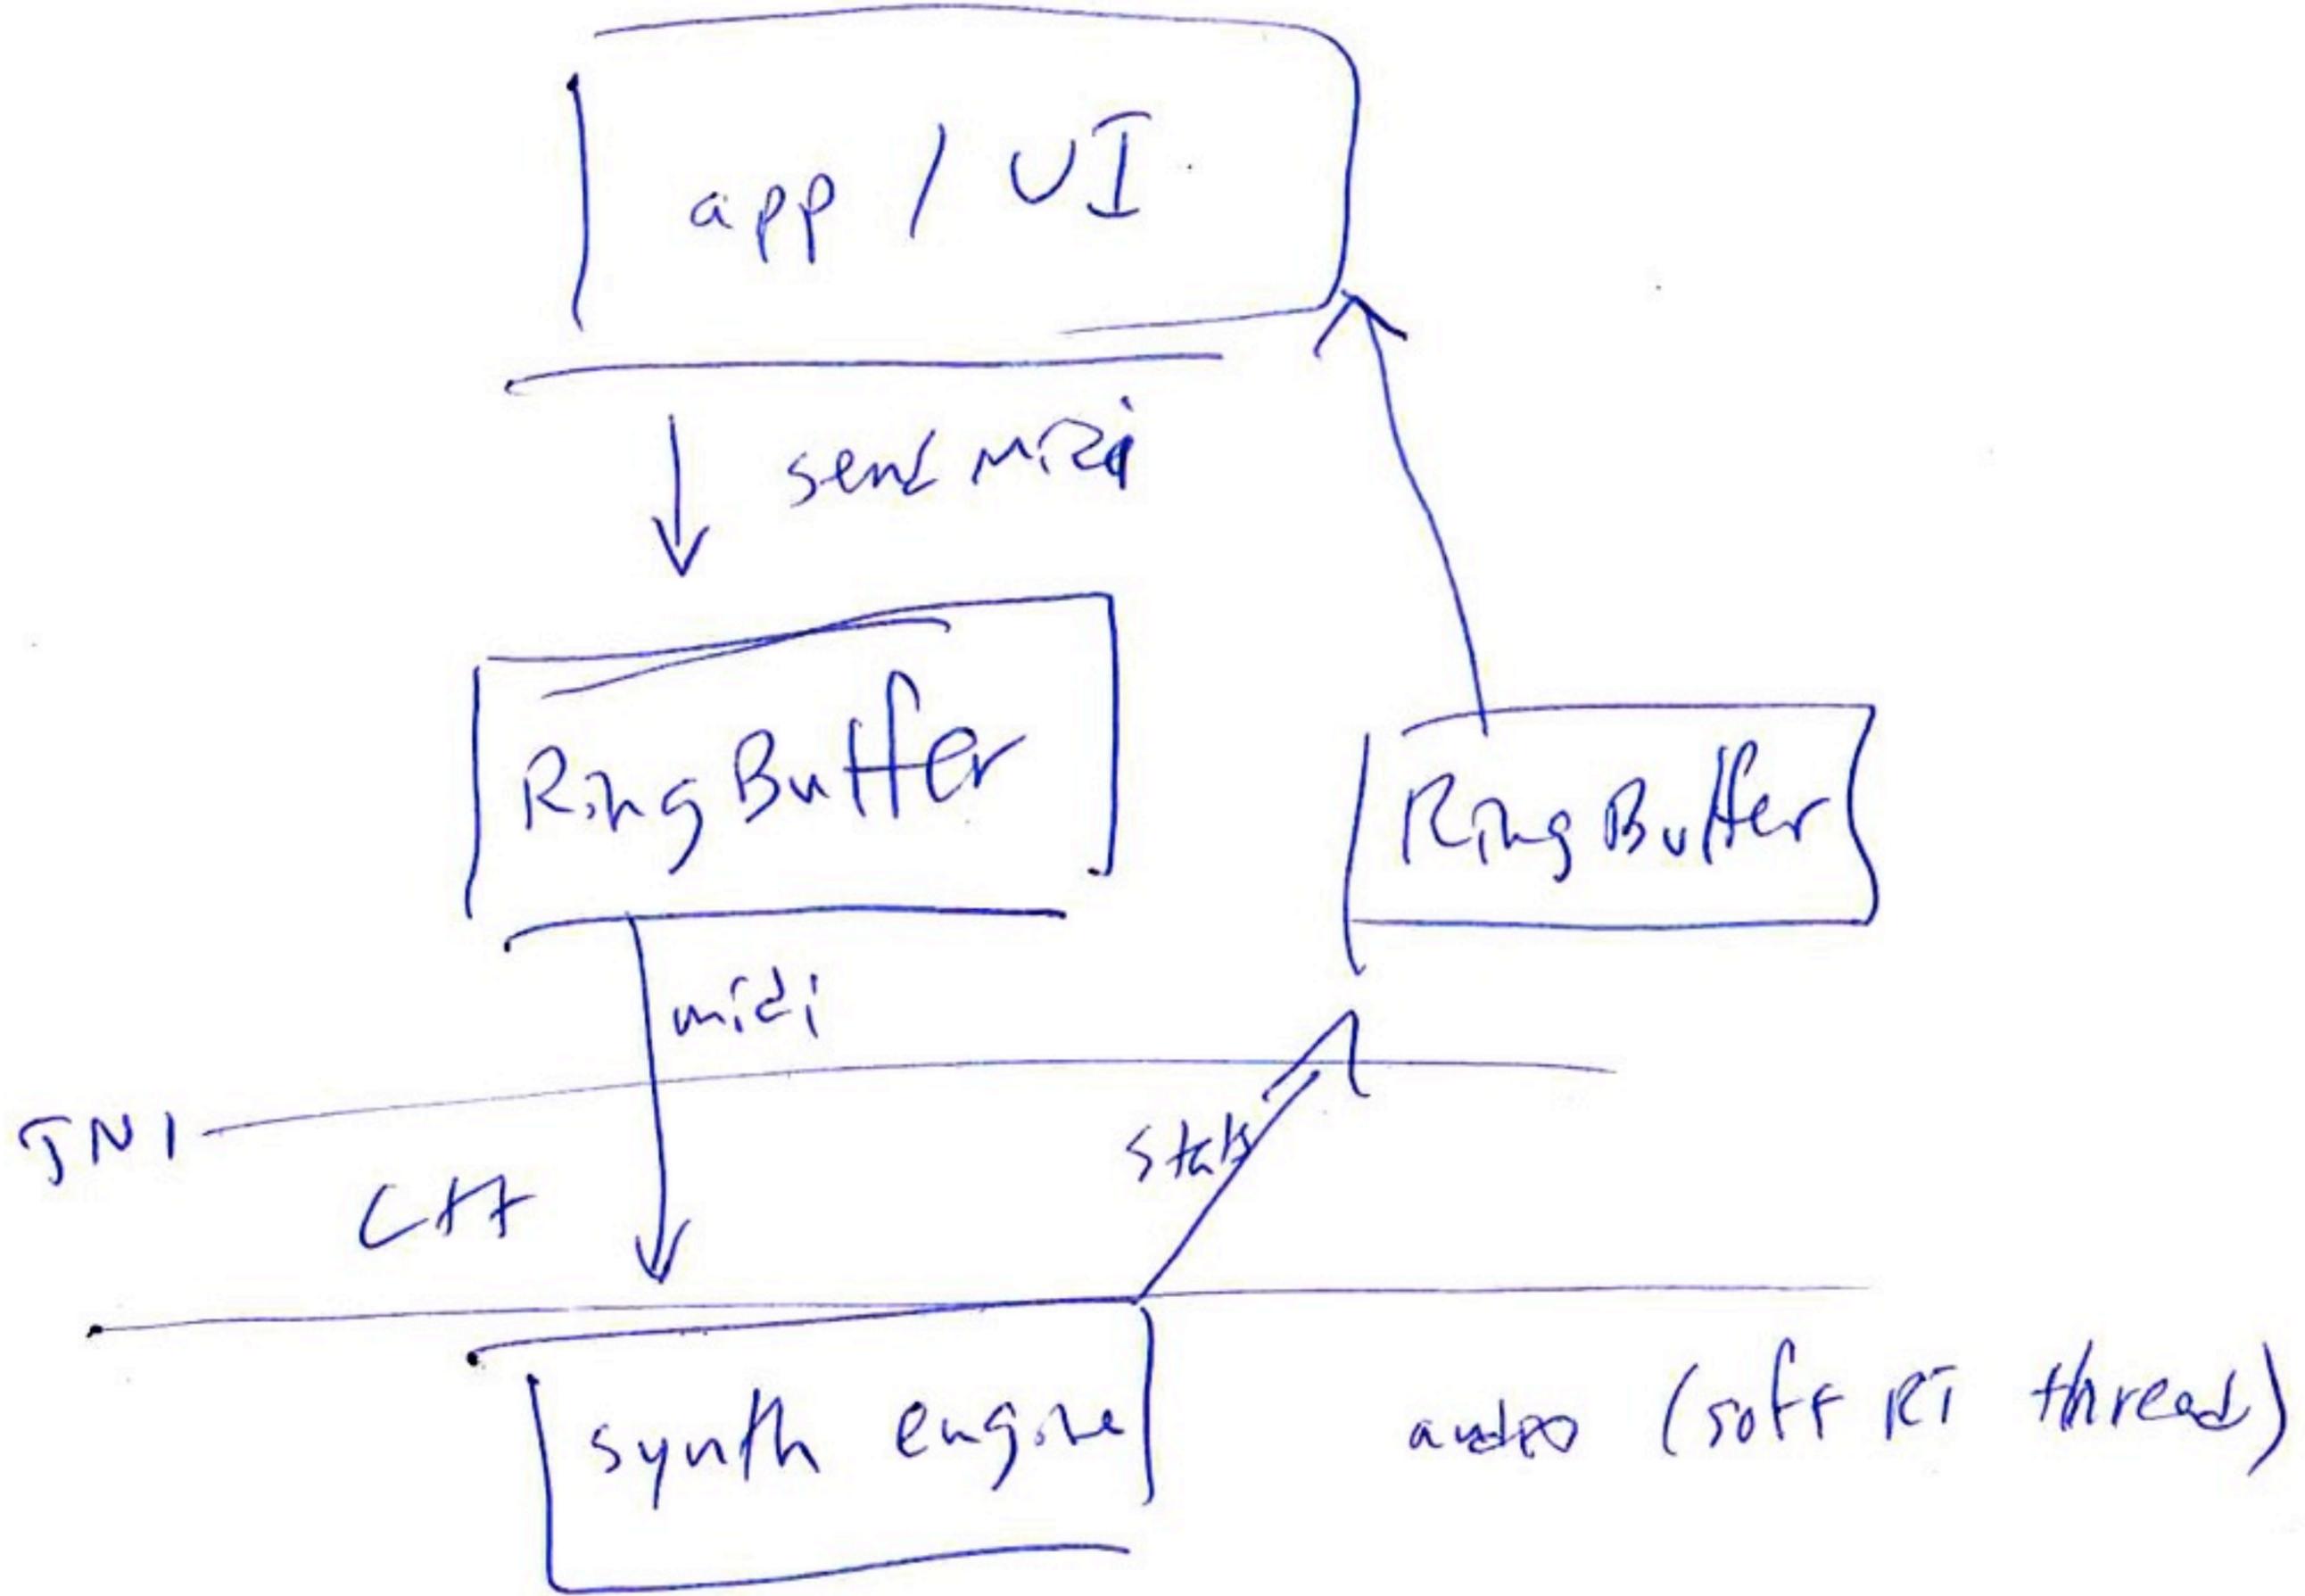
# Classic Priority Inversion

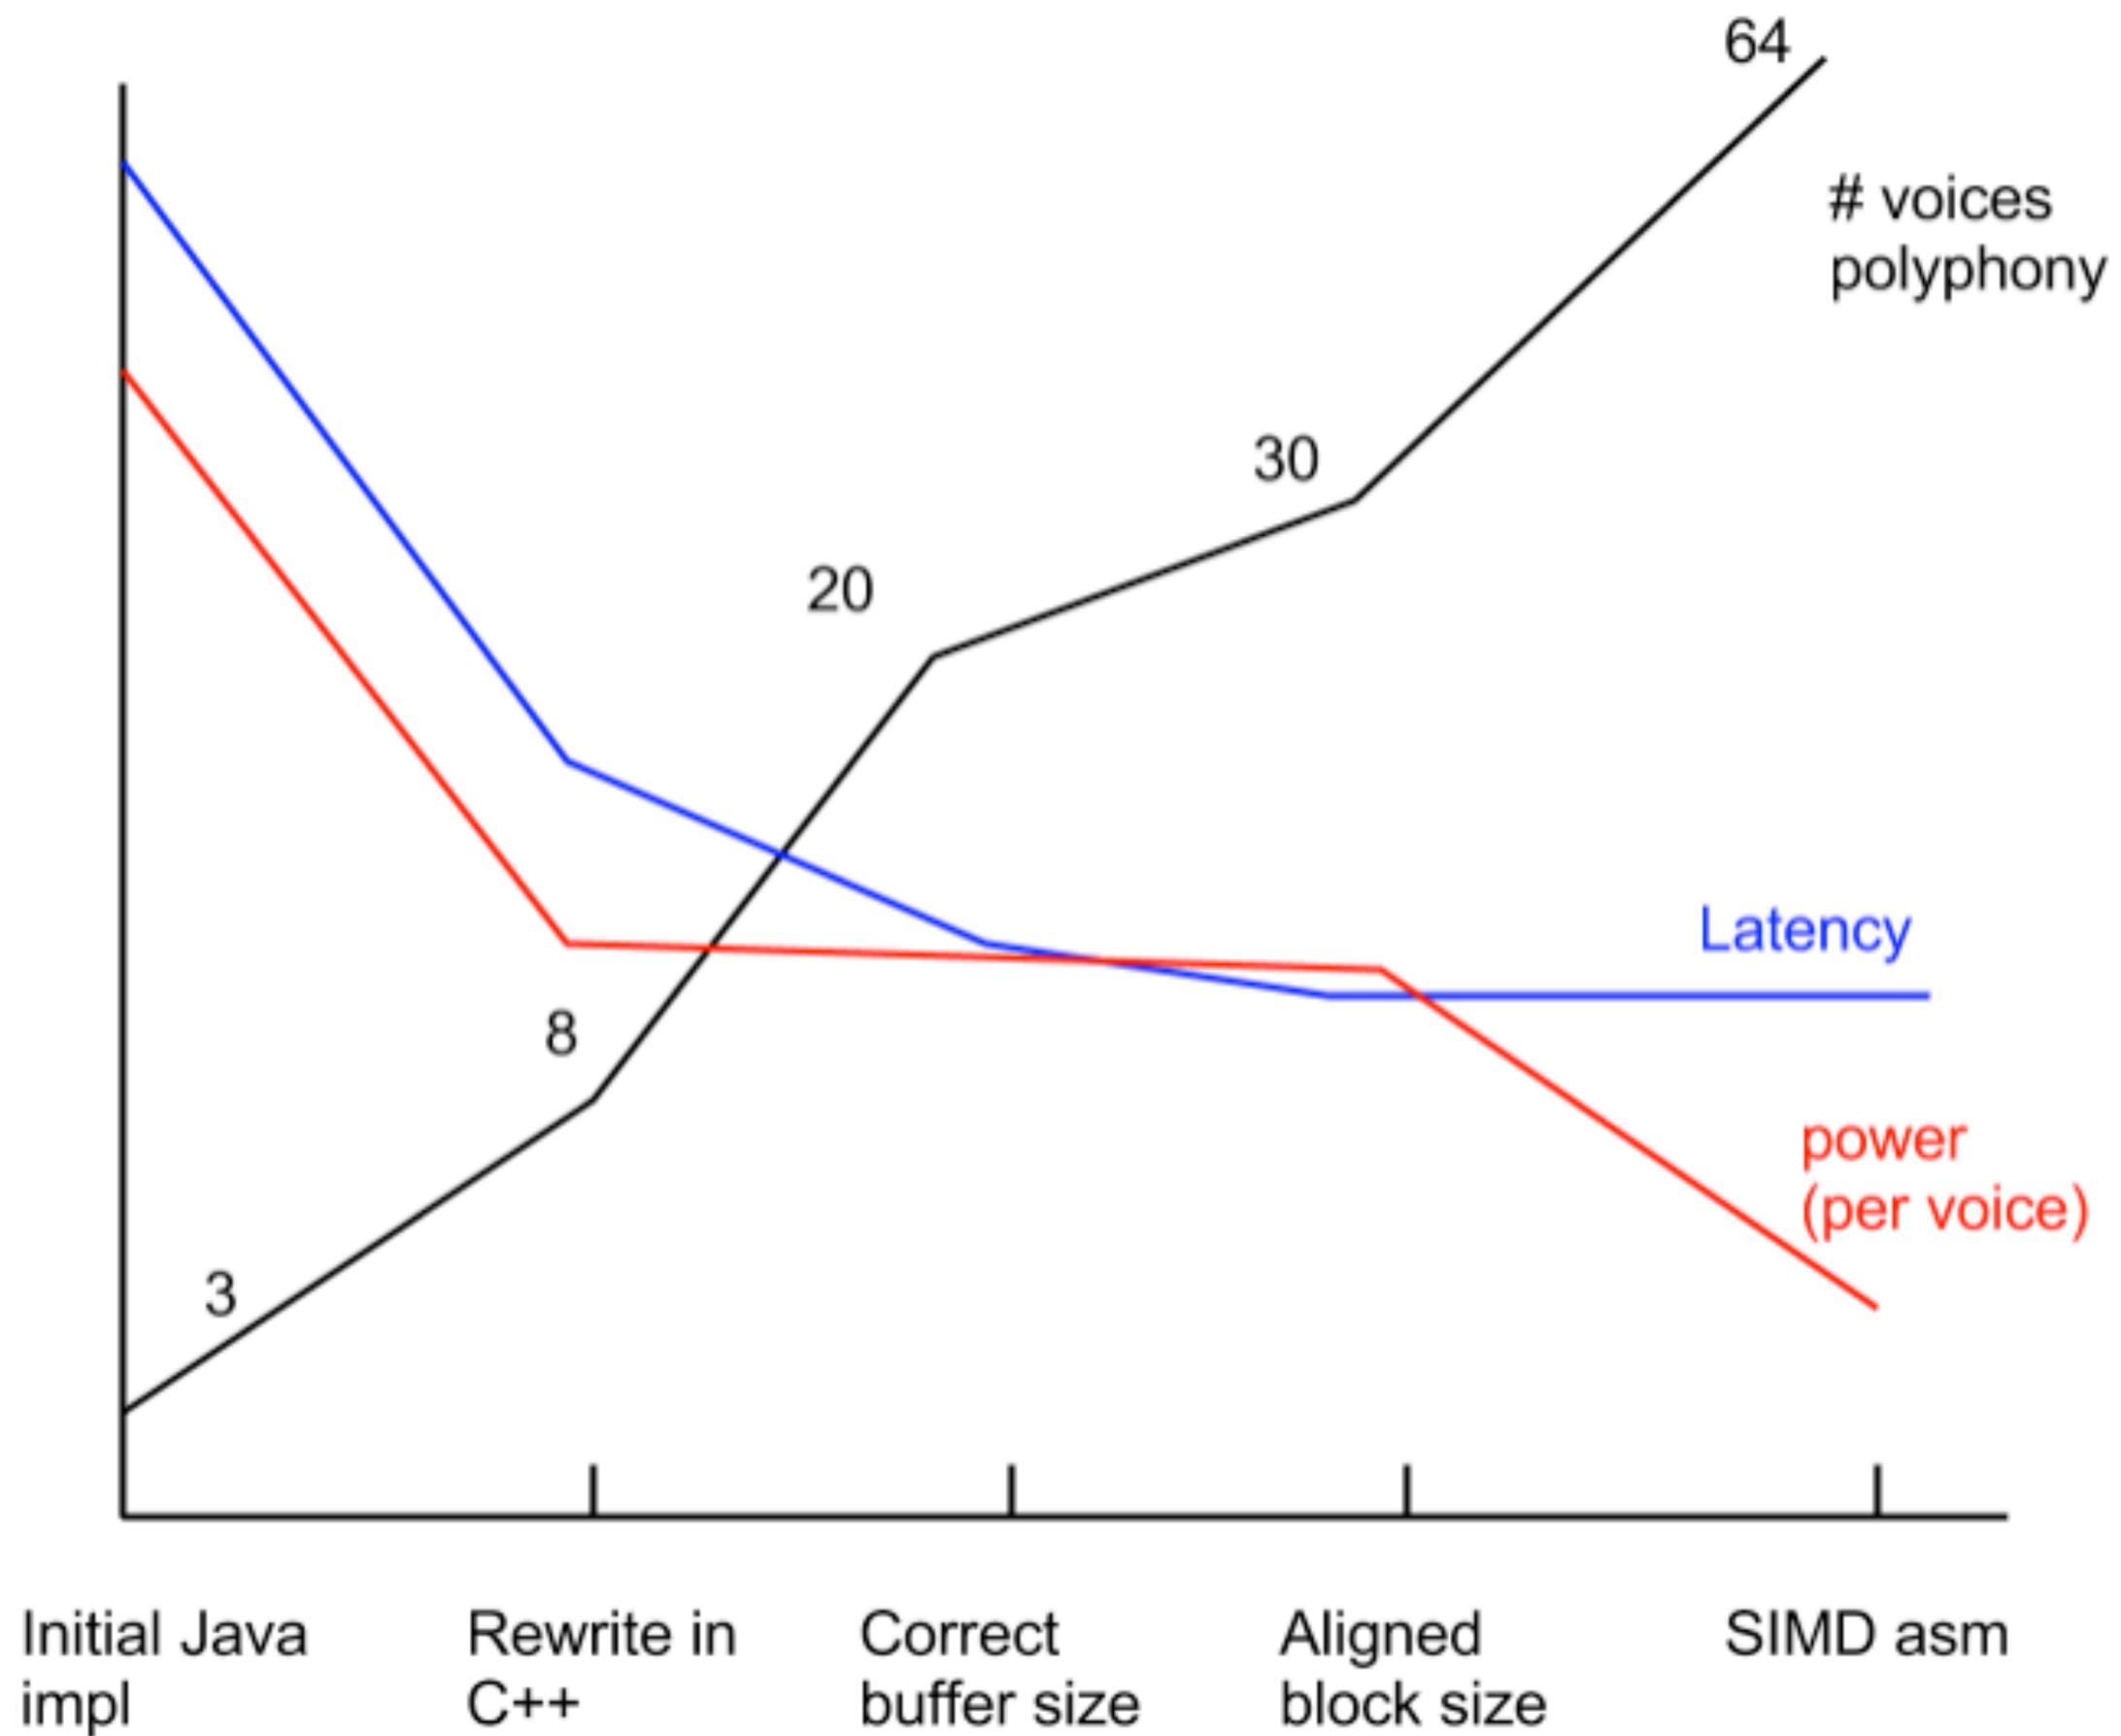


# Priority inversion during IPC









# To Do: Android

1. Measure everything!
2. Use SCHED\_FIFO
3. Implement fast path to system mixer
4. Add API to discover sample rate & buffer size
5. Keep improving!



# To Do: You

1. Measure everything!
2. Process audio in OpenSL callback
3. Match device sample rate & buffer size
4. Eliminate blocking



Google  
Developers

[code.google.com/p/high-performance-audio/](https://code.google.com/p/high-performance-audio/)

