

Audio Toolbox™

User's Guide



MATLAB® & SIMULINK®

R2019b



How to Contact MathWorks



Latest news: www.mathworks.com
Sales and services: www.mathworks.com/sales_and_services
User community: www.mathworks.com/matlabcentral
Technical support: www.mathworks.com/support/contact_us



Phone: 508-647-7000



The MathWorks, Inc.
1 Apple Hill Drive
Natick, MA 01760-2098

Audio Toolbox™ User's Guide

© COPYRIGHT 2016 - 2019 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

Revision History

March 2016	Online only	New for Version 1.0 (Release 2016a)
September 2016	Online only	Revised for Version 1.1 (Release 2016b)
March 2017	Online only	Revised for Version 1.2 (Release 2017a)
September 2017	Online only	Revised for Version 1.3 (Release 2017b)
March 2018	Online only	Revised for Version 1.4 (Release 2018a)
September 2018	Online only	Revised for Version 1.5 (Release 2018b)
March 2019	Online only	Revised for Version 2.0 (Release 2019a)
September 2019	Online only	Revised for Version 2.1 (Release 2019b)

	Plugin GUI Design	
1	<hr/>	
	Design User Interface for Audio Plugin	1-2
	Use the Audio Labeler	
2	<hr/>	
	Label Audio Using Audio Labeler	2-2
	Load Unlabeled Data	2-2
	Define and Assign Labels	2-3
	Export Label Definitions	2-10
	Export Labeled Audio Data	2-11
	Prepare Audio Datastore for Deep Learning Workflow	2-13
	Speech2Text Walkthrough Chapter	
3	<hr/>	
	Speech-to-Text Transcription	3-2
	Measure Impulse Response of an Audio System	
4	<hr/>	
	Impulse Response Measurer Walkthrough	4-2
	Configure Audio I/O System	4-2
	Configure IR Acquisition Method	4-2

Acquire IR Measurements	4-4
Analyze and Manage IR Measurements	4-5
Export IR Measurements	4-9

Design and Play a MIDI Synthesizer

5

Design and Play a MIDI Synthesizer	5-2
Convert MIDI Note Messages to Sound Waves	5-2
Synthesize MIDI Messages	5-3
Synthesize Real-Time Note Messages from MIDI Device	5-4

MIDI Device Interface

6

MIDI Device Interface	6-2
MIDI	6-2
MIDI Devices	6-2
MIDI Messages	6-3

Dynamic Range Control

7

Dynamic Range Control	7-2
Linear to dB Conversion	7-3
Gain Computer	7-3
Gain Smoothing	7-5
Make-Up Gain	7-8
dB to Linear Conversion	7-9
Apply Calculated Gain	7-9
Example: Dynamic Range Limiter	7-9

MIDI Control for Audio Plugins

8

MIDI Control for Audio Plugins	8-2
MIDI and Plugins	8-2
Use MIDI with MATLAB Plugins	8-2

MIDI Control Surface Interface

9

MIDI Control Surface Interface	9-2
About MIDI	9-2
MIDI Control Surfaces	9-2
Use MIDI Control Surfaces with MATLAB and Simulink	9-3

Use the Audio Test Bench

10

Audio Test Bench Walkthrough	10-2
Choose Object Under Test	10-2
Run Audio Test Bench	10-3
Debug Source Code of Audio Plugin	10-4
Open Scopes	10-6
Configure Input to Audio Test Bench	10-6
Configure Output from Audio Test Bench	10-8
Call Custom Visualization of Audio Plugin	10-9
Synchronize Plugin Property with MIDI Control	10-10
Play the Audio and Save the Output File	10-11
Validate and Generate Audio Plugin	10-11
Generate MATLAB Script	10-12

Audio Plugin Example Gallery

11

Audio Plugin Example Gallery	11-2
Audio Effects	11-2
Filters	11-2
Gain Control	11-2
Spatial Audio	11-2
Communicate Between MATLAB and DAW	11-2
Music Information Retrieval	11-2
Speech Processing	11-2
Audio Plugin Examples	11-2

Equalization

12

Equalization	12-2
Equalization Design Using Audio Toolbox	12-2
EQ Filter Design	12-2
Lowpass and Highpass Filter Design	12-6
Shelving Filter Design	12-7

Deployment

13

Desktop Real-Time Audio Acceleration with MATLAB Coder	13-2
--	-------------

Audio I/O User Guide

14

Run Audio I/O Features Outside MATLAB and Simulink	14-2
---	-------------

Block Example Repository

15

Decrease Underrun	15-2
-------------------------	------

Block Example Repository

16

Extract Cepstral Coefficients	16-2
Tune Center Frequency Using Input Port	16-4
Gate Audio Signal Using VAD	16-6
Frequency-Domain Voice Activity Detection	16-8
Visualize Noise Power	16-9
Detect Presence of Speech	16-13
Perform Graphic Equalization	16-15
Split-Band De-Essing	16-17
Diminish Plosives from Speech	16-18
Suppress Loud Sounds	16-19
Attenuate Low-Level Noise	16-22
Suppress Volume of Loud Sounds	16-25
Gate Background Noise	16-28
Output Values from MIDI Control Surface	16-31
Apply Frequency Weighting	16-33
Compare Loudness Before and After Audio Processing	16-35

Two-Band Crossover Filtering for a Stereo Speaker System	16-37
Mimic Acoustic Environments	16-39
Perform Parametric Equalization	16-41
Perform Octave Filtering	16-43
Read from Microphone and Write to Speaker	16-45
Channel Mapping	16-48
Trigger Gain Control Based on Loudness Measurement ...	16-50

Real-Time Parameter Tuning

17

Real-Time Parameter Tuning	17-2
Programmatic Parameter Tuning	17-3

Sample Audio Files

18

Sample Audio Files	18-2
---------------------------------	-------------

Tips and Tricks for Plugin Authoring

19

Tips and Tricks for Plugin Authoring	19-2
Avoid Disrupting the Event Queue in MATLAB	19-2
Separate Code for Features Not Supported for Plugin Generation	19-6
Implement Reset Correctly	19-8

Implement Plugin Composition Correctly	19-8
Address "A set method for a non-Dependent property should not access another property" Warning in Plugin	19-11
Use System Object That Does Not Support Variable-Size Signals	19-13
Using Enumeration Parameter Mapping	19-16

Spectral Descriptors Chapter

20

Spectral Descriptors	20-2
-----------------------------------	-------------

Plugin GUI Design

Design User Interface for Audio Plugin

Audio plugins enable you to tune parameters of a processing algorithm while streaming audio in real time. To enhance usability, you can define a custom user interface (UI) that maps parameters to intuitively designed and positioned controls. You can use `audioPluginInterface`, `audioPluginParameter`, and `audioPluginGridLayout` to define the custom UI. You can interact with the custom UI in MATLAB® using `parameterTuner`, or deploy the plugin with a custom UI to a digital audio workstation (DAW). This tutorial walks through key design capabilities of audio plugins by sequentially enhancing a basic audio plugin UI.

To learn more about audio plugins in general, see “Audio Plugins in MATLAB”.

Default User Interface

The `equalizerV1` audio plugin enables you to tune the gains and center frequencies of a three-band equalizer, tune the overall volume, and toggle between enabled and disabled states.

```
classdef equalizerV1 < audioPlugin
    properties
        GainLow = 0
        FreqLow = sqrt(20*500)
        GainMid = 0
        FreqMid = sqrt(500*3e3)
        GainHigh = 0
        FreqHigh = sqrt(3e3*20e3)
        Volume = 1
        Enable = true
    end
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('GainLow', ...
                'Label','dB', ...
                'Mapping',{'lin',-20,20}), ...
            audioPluginParameter('FreqLow', ...
                'Label','Hz', ...
                'Mapping',{'log',20,500}), ...
            audioPluginParameter('GainMid', ...
                'Label','dB', ...
                'Mapping',{'lin',-20,20}), ...
            audioPluginParameter('FreqMid', ...
                'Label','Hz', ...
```

```

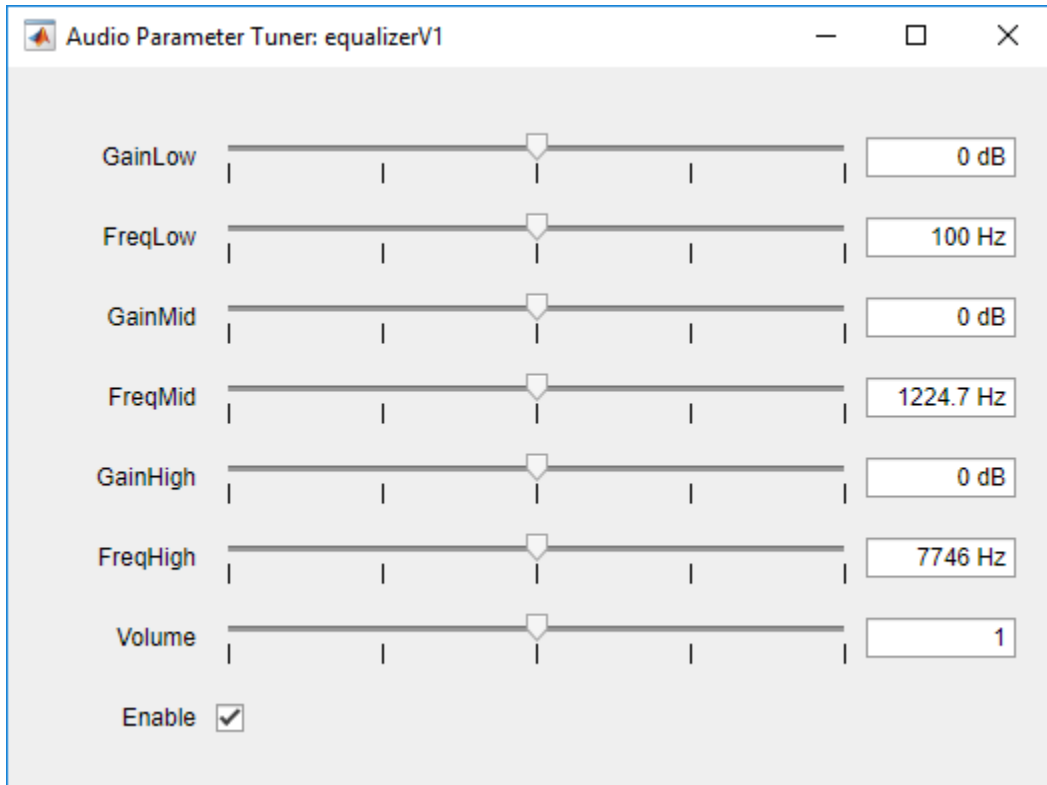
        'Mapping',{ 'log',500,3e3}), ...
    audioPluginParameter('GainHigh', ...
        'Label','dB', ...
        'Mapping',{ 'lin',-20,20}), ...
    audioPluginParameter('FreqHigh', ...
        'Label','Hz', ...
        'Mapping',{ 'log',3e3,20e3}), ...
    audioPluginParameter('Volume', ...
        'Mapping',{ 'lin',0,2}), ...
    audioPluginParameter('Enable'))
end
properties (Access = private)
    mPEQ
end
methods
    function obj = equalizerV1
        obj.mPEQ = multibandParametricEQ('HasHighpassFilter',false, ...
            'HasLowShelfFilter',false,'HasHighShelfFilter',false, ...
            'HasLowpassFilter',false,'Oversample',false,'NumEQBands',3, ...
            'EQOrder',2);
    end
    function y = process(obj, x)
        if obj.Enable
            y = step(obj.mPEQ,x);
            y = y*obj.Volume;
        else
            y = x;
        end
    end
    function reset(obj)
        obj.mPEQ.SampleRate = getSampleRate(obj);
        reset(obj.mPEQ);
    end
    function set.FreqLow(obj,val)
        obj.FreqLow = val;
        obj.mPEQ.Frequencies(1) = val; %#ok<*>MCSUP>
    end
    function set.GainLow(obj,val)
        obj.GainLow = val;
        obj.mPEQ.PeakGains(1) = val;
    end
    function set.FreqMid(obj,val)
        obj.FreqMid = val;
        obj.mPEQ.Frequencies(2) = val;
    end
end

```

```
end
function set.GainMid(obj,val)
    obj.GainMid = val;
    obj.mPEQ.PeakGains(2) = val;
end
function set.FreqHigh(obj,val)
    obj.FreqHigh = val;
    obj.mPEQ.Frequencies(3) = val;
end
function set.GainHigh(obj,val)
    obj.GainHigh = val;
    obj.mPEQ.PeakGains(3) = val;
end
end
end
```

Call `parameterTuner` to visualize the default UI of the audio plugin.

```
parameterTuner(equalizerV1)
```



Control Style and Layout

To define the UI grid, add `audioPluginGridLayout` to the `audioPluginInterface`. You can specify the number, size, spacing, and border of cells in the UI grid. In this example, specify "RowHeight" as `[20,20,160,20,100]` and "ColumnWidth" as `[100,100,100,50,150]`. This creates the following UI grid:



To define the UI control style, update the `audioPluginParameter` definition of each parameter to include the "Style" and "Layout" name-value pairs. `Style` defines the type of control (rotary knob, slider, or switch, for example). `Layout` defines which cells the controls occupy on the UI grid. You can specify `Layout` as the [row, column] of the grid to occupy, or as the [upper, left; lower, right] of the group of cells to occupy. By default, control display names are also displayed and occupy their own cells on the UI grid. The cells they occupy depend on the "DisplayNameLocation" name-value pair.

The commented arrows indicate the difference between `equalizerV1` and `equalzierV2`.

```
classdef equalizerV2 < audioPlugin
    ... % omitted for example purposes
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('GainLow', ...
```



```

        'Label', 'dB', ...
        'Mapping', {'lin', -20, 20}, ...
        'Style', 'vslider', ...
        'Layout', [2, 1; 4, 1], ...
        'DisplayName', 'Low', 'DisplayNameLocation', 'Above'), ...
audioPluginParameter('FreqLow', ...
    'Label', 'Hz', ...
    'Mapping', {'log', 20, 500}, ...
    'Style', 'rotaryknob', ...
    'Layout', [5, 1], ...
    'DisplayNameLocation', 'None'), ...
audioPluginParameter('GainMid', ...
    'Label', 'dB', ...
    'Mapping', {'lin', -20, 20}, ...
    'Style', 'vslider', ...
    'Layout', [2, 2; 4, 2], ...
    'DisplayNameLocation', 'None'), ...
audioPluginParameter('FreqMid', ...
    'Label', 'Hz', ...
    'Mapping', {'log', 500, 3e3}, ...
    'Style', 'rotaryknob', ...
    'Layout', [5, 2], ...
    'DisplayNameLocation', 'None'), ...
audioPluginParameter('GainHigh', ...
    'Label', 'dB', ...
    'Mapping', {'lin', -20, 20}, ...
    'Style', 'vslider', ...
    'Layout', [2, 3; 4, 3], ...
    'DisplayName', 'High', 'DisplayNameLocation', 'Above'), ...
audioPluginParameter('FreqHigh', ...
    'Label', 'Hz', ...
    'Mapping', {'log', 3e3, 20e3}, ...
    'Style', 'rotaryknob', ...
    'Layout', [5, 3], ...
    'DisplayNameLocation', 'None'), ...
audioPluginParameter('Volume', ...
    'Mapping', {'lin', 0, 2}, ...
    'Style', 'rotaryknob', ...
    'Layout', [3, 5], ...
    'DisplayNameLocation', 'Above'), ...
audioPluginParameter('Enable', ...
    'Style', 'vtoggle', ...
    'Layout', [5, 5], ...
    'DisplayNameLocation', 'None'), ...

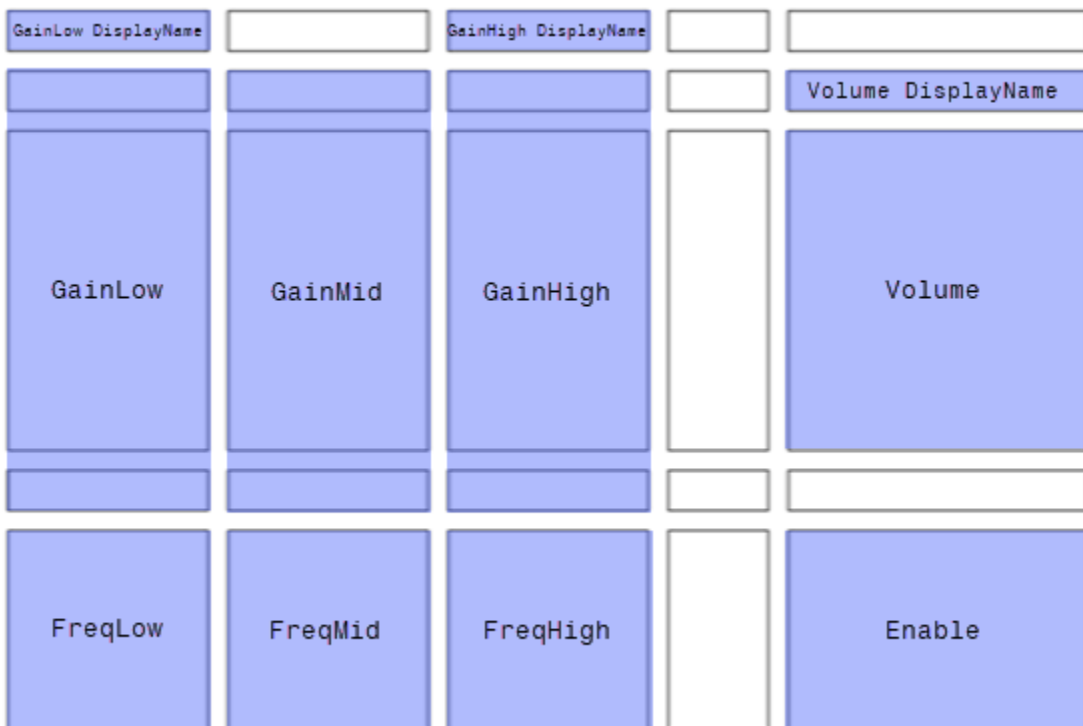
```

```

...
audioPluginGridLayout( ... %<--
    'RowHeight',[20,20,160,20,100], ... %<--
    'ColumnWidth',[100,100,100,50,150]) %<--
end
... % omitted for example purposes
end

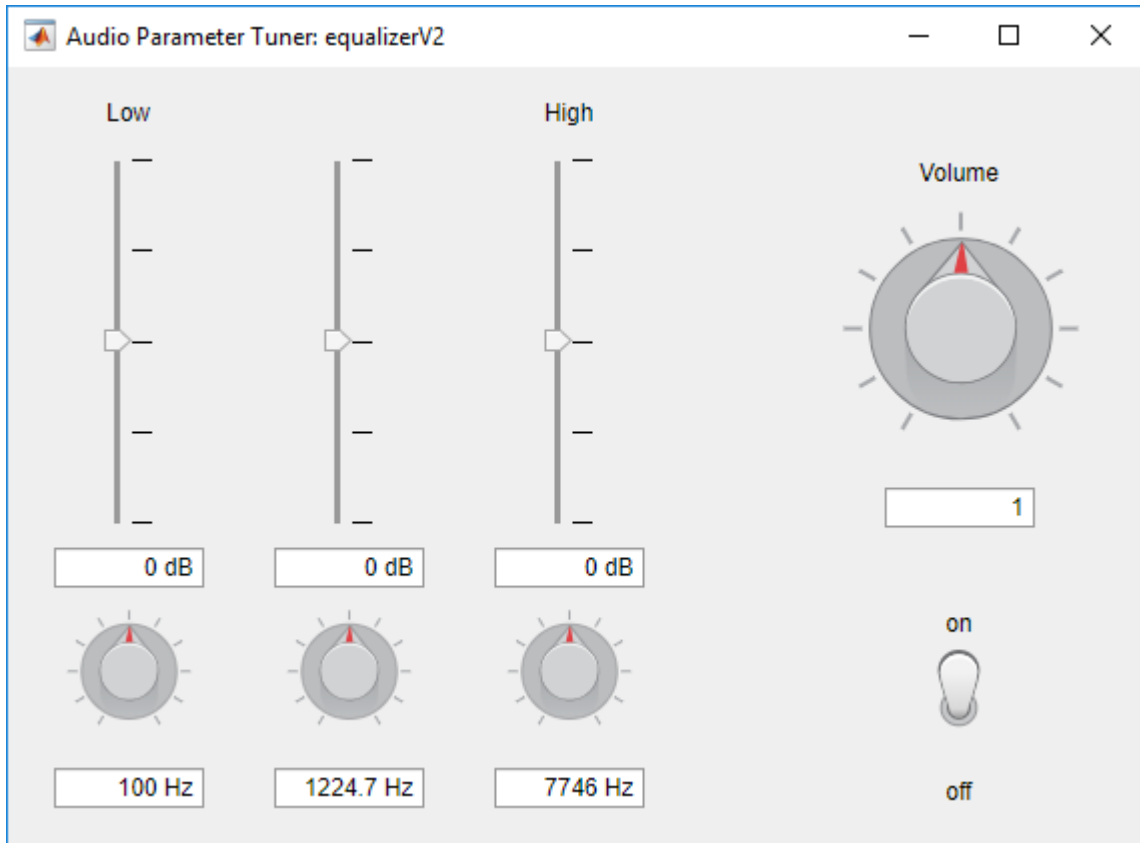
```

The `Layout` and `DisplayNameLocation` defined in the `audioPluginParameters` maps the respective parameters to the control grid as follows:



Call `parameterTuner` to visualize the UI of `equalizerV2`.

```
parameterTuner(equalizerV2)
```



Background Image and Color


To customize the background of your UI, specify “BackgroundImage” and “BackgroundColor” in `audioPluginInterface`.

The `BackgroundColor` can be specified as a short or long color name string or as an RGB triplet. See `ColorSpec` (`Color Specification`) for details. When you specify `BackgroundColor`, the color is applied to all space on the UI except space occupied by controls or a `BackgroundImage`. If the control or background image includes a transparency, then the background color shows through the transparency.

The `BackgroundImage` can be specified as a PNG, GIF, or JPG file. The image is applied to the UI grid by aligning the top left corners of the UI grid and image. If the image is

larger than the UI grid size defined in `audioPluginGridLayout`, then the image is clipped to the UI grid size. The background image is not resized. If the image is smaller than the UI grid, then unoccupied regions of the UI grid are treated as transparent.

In this example, you increase the padding around the perimeter of the grid to create space for the MathWorks® logo. You can calculate the total width of the UI grid as the sum of all column widths plus the left and right padding plus the column spacing (the default column spacing of 10 pixels is used in this example):
 $(100 + 100 + 100 + 50 + 150) + (20 + 20) + (4 \times 10) = 580$. The total height of the UI grid is the sum of all row heights plus the top and bottom padding plus the row spacing (the default row spacing of 10 pixels is used in this example):
 $(20 + 20 + 160 + 20 + 100) + (20 + 120) + (4 \times 10) = 500$. To locate the logo at the bottom of the UI grid, use a 580-by-500 image:



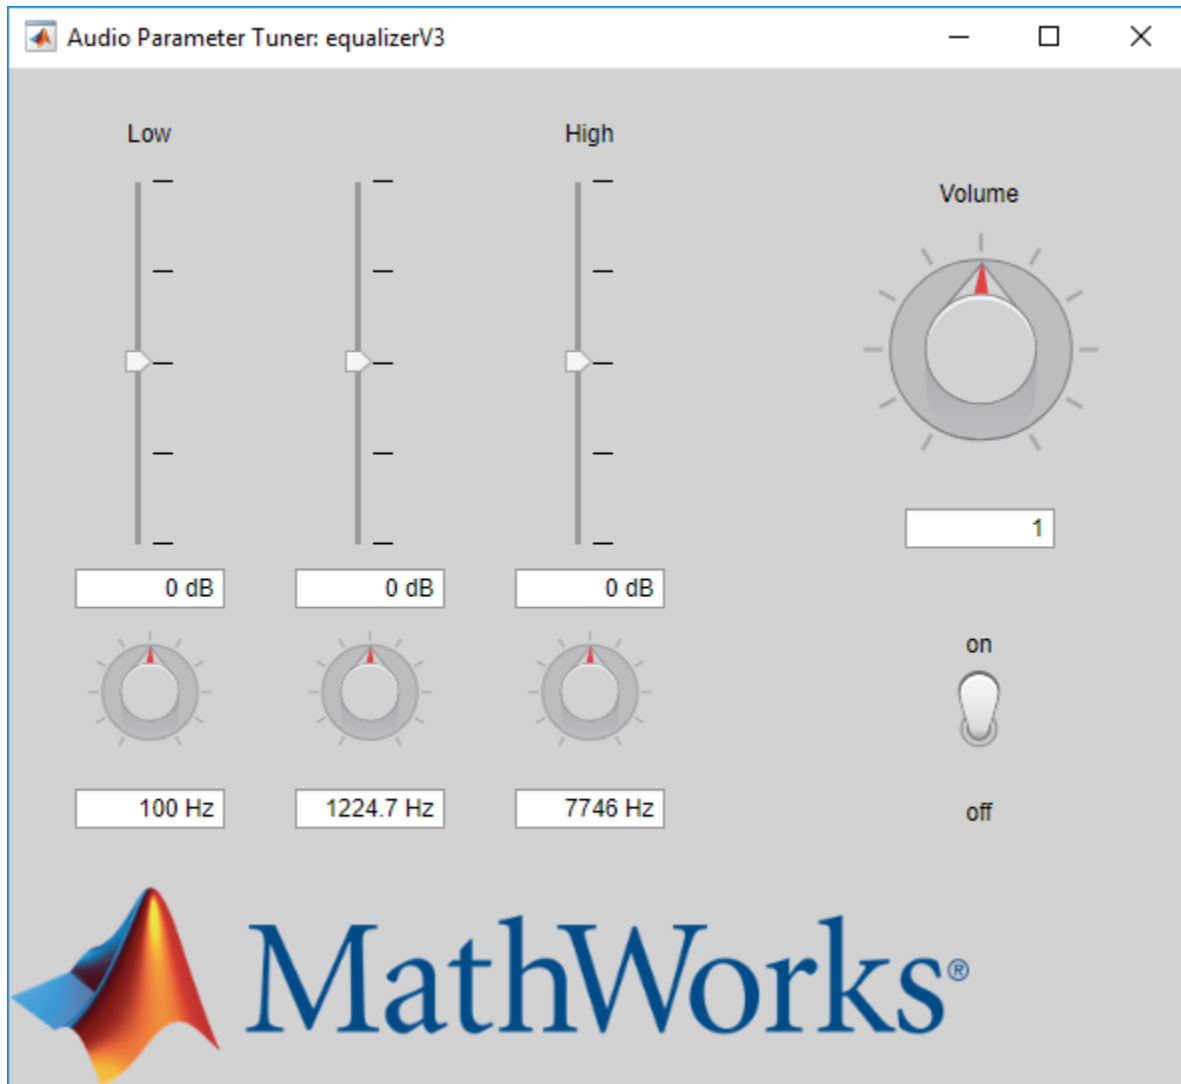
```
classdef equalizerV3 < audioPlugin
    ... % omitted for example purposes
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('GainLow', ...
                'Label','dB', ...
                'Mapping',{'lin',-20,20}, ...
                'Style','vslider', ...
```

```
        'Layout',[2,1;4,1], ...
        'DisplayName','Low','DisplayNameLocation','Above'), ...
audioPluginParameter('FreqLow', ...
    'Label','Hz', ...
    'Mapping',{'log',20,500}, ...
    'Style','rotaryknob', ...
    'Layout',[5,1], ...
    'DisplayNameLocation','None'), ...
audioPluginParameter('GainMid', ...
    'Label','dB', ...
    'Mapping',{'lin',-20,20}, ...
    'Style','vslider', ...
    'Layout',[2,2;4,2], ...
    'DisplayNameLocation','None'), ...
audioPluginParameter('FreqMid', ...
    'Label','Hz', ...
    'Mapping',{'log',500,3e3}, ...
    'Style','rotaryknob', ...
    'Layout',[5,2], ...
    'DisplayNameLocation','None'), ...
audioPluginParameter('GainHigh', ...
    'Label','dB', ...
    'Mapping',{'lin',-20,20}, ...
    'Style','vslider', ...
    'Layout',[2,3;4,3], ...
    'DisplayName','High','DisplayNameLocation','Above'), ...
audioPluginParameter('FreqHigh', ...
    'Label','Hz', ...
    'Mapping',{'log',3e3,20e3}, ...
    'Style','rotaryknob', ...
    'Layout',[5,3], ...
    'DisplayNameLocation','None'), ...
audioPluginParameter('Volume', ...
    'DisplayName','Volume', ...
    'Mapping',{'lin',0,2}, ...
    'Style','rotaryknob', ...
    'Layout',[3,5], ...
    'DisplayNameLocation','Above'), ...
audioPluginParameter('Enable', ...
    'Style','vtoggle', ...
    'Layout',[5,5], ...
    'DisplayNameLocation','None'), ...
...
audioPluginGridLayout( ...
```

```
        'RowHeight', [20,20,160,20,100], ...
        'ColumnWidth', [100,100,100,50,150], ...
        'Padding', [20,120,20,20]), ...
        ...
        'BackgroundImage', 'background.png', ...
        'BackgroundColor', [210/255,210/255,210/255])
    end
    ... % omitted for example purposes
end
```

Call `parameterTuner` to visualize the UI of `equalizerV3`.

```
parameterTuner(equalizerV3)
```



Custom Control Filmstrips

To use custom filmstrips, specify the "Filmstrip" and "FilmstripFrameSize" name-value pairs in `audioPluginParameter`. The filmstrip can be a PNG, GIF, or JPG file, and should consist of frames placed end-to-end either vertically or horizontally. The filmstrip is

mapped to the control's range so that the corresponding filmstrip frame is displayed on the plugin UI as you tune parameters. In this example, specify a two-frame filmstrip for the Enable parameter. As a best practice, the size of each frame of the film strip should equal the size of the region occupied by the parameter. The Enable parameter occupies one cell that is 150-by-100 pixels. To create a vertical filmstrip where each frame is 150-by-100, make the total filmstrip size 150-by-200 and set `FilmstripFrameSize` to `[150,100]`. The filmstrip used in this example contains the frame corresponding to the off position first, then the on position:



```
classdef equalizerV4 < audioPlugin
    ... % omitted for example purposes
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('GainLow', ...
                'Label', 'dB', ...
                'Mapping', {'lin', -20, 20}, ...
                'Style', 'vslider', ...
                'Layout', [2, 1; 4, 1], ...
                'DisplayName', 'Low', 'DisplayNameLocation', 'Above'), ...
            audioPluginParameter('FreqLow', ...
                'Label', 'Hz', ...
                'Mapping', {'log', 20, 500}, ...
                'Style', 'rotaryknob', ...
                'Layout', [5, 1], ...
                'DisplayNameLocation', 'None'), ...
            audioPluginParameter('GainMid', ...
                'Label', 'dB', ...
                'Mapping', {'lin', -20, 20}, ...
                'Style', 'vslider', ...
```

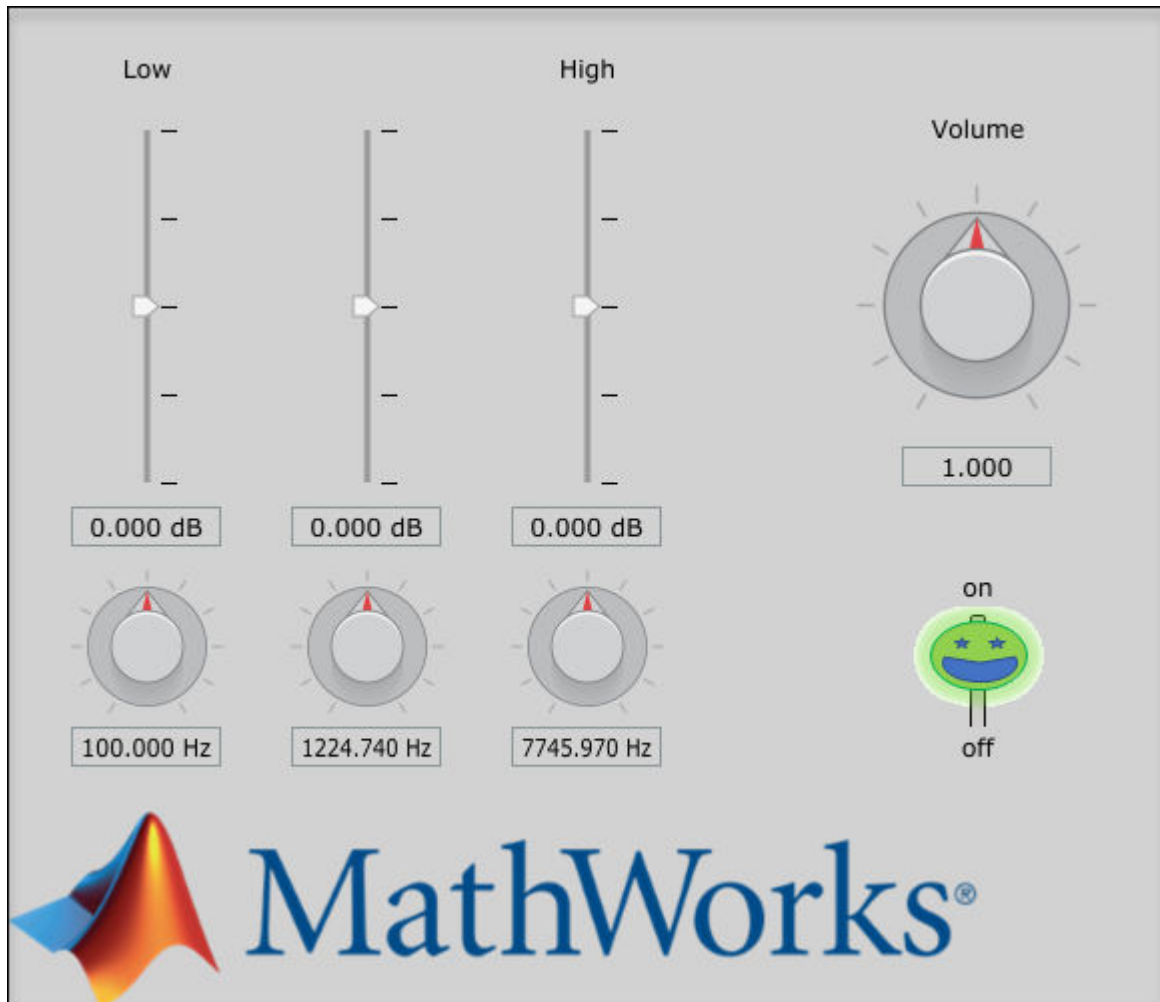
```
        'Layout',[2,2;4,2], ...
        'DisplayNameLocation','None'), ...
    audioPluginParameter('FreqMid', ...
        'Label','Hz', ...
        'Mapping',{'log',500,3e3}, ...
        'Style','rotaryknob', ...
        'Layout',[5,2], ...
        'DisplayNameLocation','None'), ...
    audioPluginParameter('GainHigh', ...
        'Label','dB', ...
        'Mapping',{'lin',-20,20}, ...
        'Style','vslider', ...
        'Layout',[2,3;4,3], ...
        'DisplayName','High','DisplayNameLocation','Above'), ...
    audioPluginParameter('FreqHigh', ...
        'Label','Hz', ...
        'Mapping',{'log',3e3,20e3}, ...
        'Style','rotaryknob', ...
        'Layout',[5,3], ...
        'DisplayNameLocation','None'), ...
    audioPluginParameter('Volume', ...
        'Mapping',{'lin',0,2}, ...
        'Style','rotaryknob', ...
        'Layout',[3,5], ...
        'DisplayNameLocation','Above'), ...
    audioPluginParameter('Enable', ...
        'Style','vtoggle', ...
        'Layout',[5,5], ...
        'DisplayNameLocation','None', ...
        'Filmstrip','vtoggle.png', ...
        'FilmstripFrameSize',[150,100]), ...
    ...
    audioPluginGridLayout( ...
        'RowHeight',[20,20,160,20,100], ...
        'ColumnWidth',[100,100,100,50,150], ...
        'Padding',[20,120,20,20]), ...
    ...
    'BackgroundImage','background.png', ...
    'BackgroundColor',[210/255,210/255,210/255])
end
... % omitted for example purposes
end
```

Filmstrips are not supported by `parameterTuner`. To see the custom plugin UI, you must deploy the plugin to a DAW. Use `generateAudioPlugin` to create a VST plugin.

```
generateAudioPlugin equalizerV4
```

```
.....
```

In this example, the plugin was opened in REAPER. A screenshot of the UI in REAPER is displayed below.



See Also

More About

- “Audio Plugins in MATLAB”

- [“Export a MATLAB Plugin to a DAW”](#)

See Also

[audioPlugin](#) | [audioPluginGridLayout](#) | [audioPluginInterface](#) | [audioPluginParameter](#) | [generateAudioPlugin](#) | [parameterTuner](#)

Use the Audio Labeler

Label Audio Using Audio Labeler

The **Audio Labeler** app enables you to interactively define and visualize ground-truth labels for audio data sets. This example shows how you can create label definitions and then interactively label a set of audio files. The example also shows how to export the labeled ground-truth data, which you can then use with `audioDatastore` to train a machine learning system.

Load Unlabeled Data

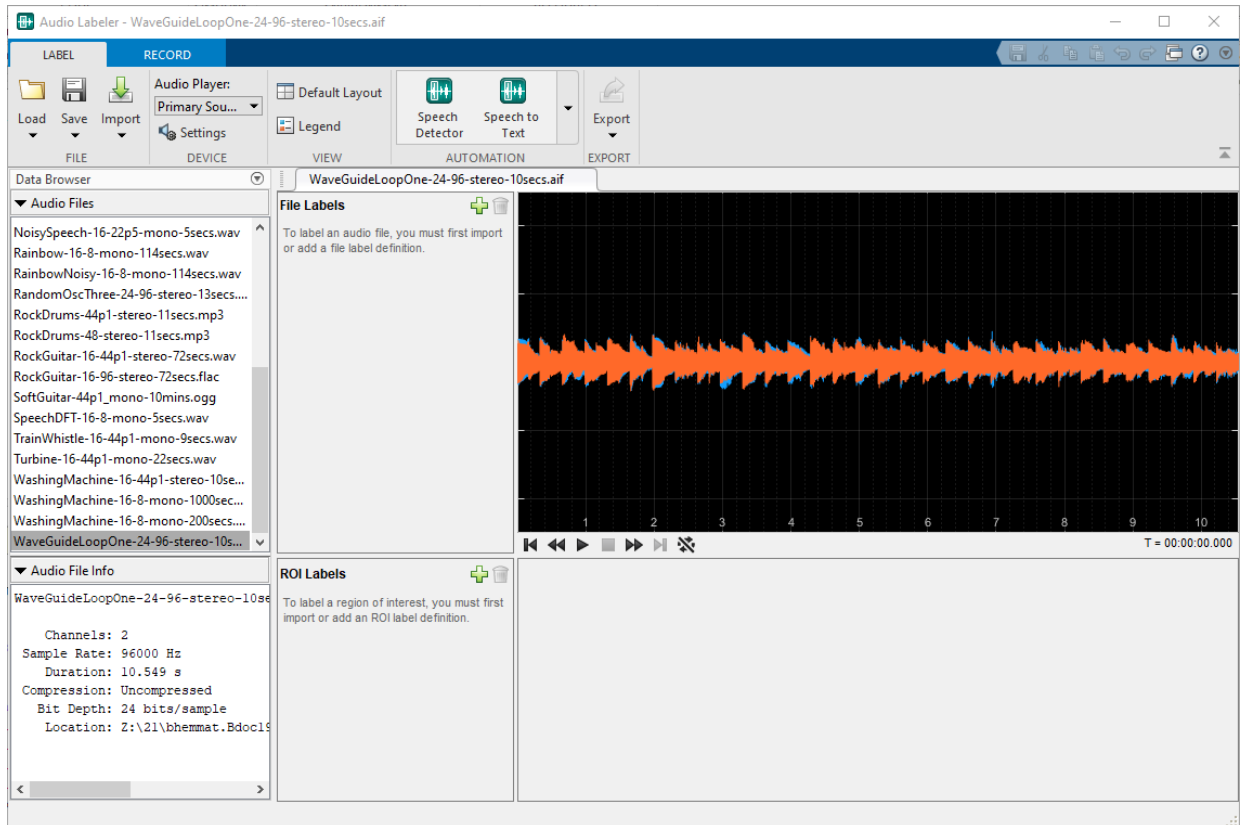
- 1 To open the **Audio Labeler**, at the MATLAB® command prompt, enter:

```
audioLabeler
```

- 2 This example uses the audio files included with Audio Toolbox. To locate the file path on your system, at the MATLAB command prompt, enter:


```
fullfile(matlabroot, 'toolbox', 'audio', 'samples')
```

To load audio from a file, click **Load > Audio Folders** and select the folder containing audio files you want to label.

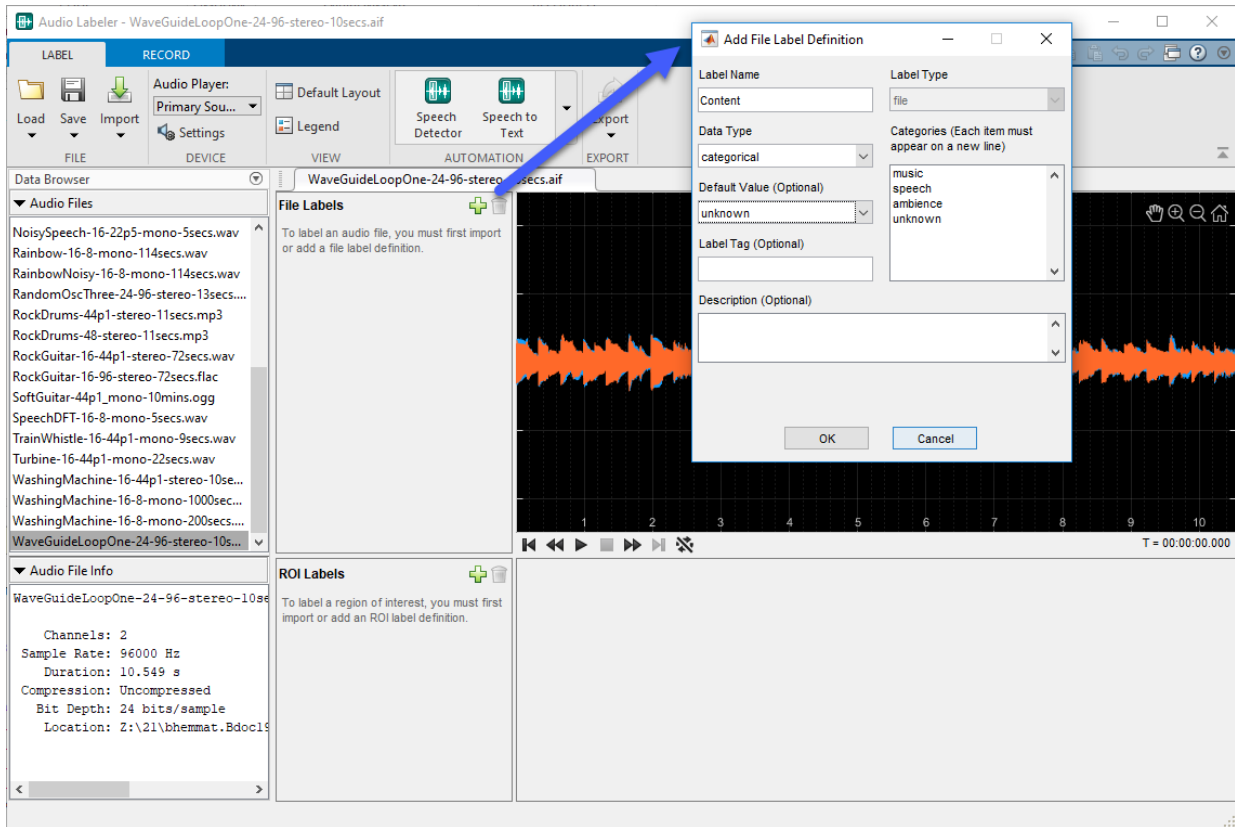



Define and Assign Labels

File-Level Labels

The audio samples include music, speech, and ambience. To create a file-level label that defines the contents of the audio file as music, speech, ambience, or unknown, click . Specify the **Label Name** as Content, the **Data Type** as categorical, and the **Categories** as music, speech, ambience, or unknown. Set the **Default Value** of the label definition to unknown.

2 Use the Audio Labeler



All audio files in the **Data Browser** are now associated with the **Content** label name. To listen to the audio file selected in the **Data Browser** and confirm that it is a music file, click . To set the value of the **Content** label, click **unknown** in the **File Labels** panel and select **music** from the drop-down menu.

The selected audio file now has the label name **Content** with value **music** assigned to it. You can continue setting the **Content** value for each file by selecting a file in the **Data Browser** and then selecting a value from the **File Labels** panel.

Region-Level Labels

You can define region-level labels manually or by using the provided automated algorithms. Audio Toolbox includes automatic labeling algorithms for speech detection and speech-to-text transcription.

Note To enable automatic speech-to-text transcription, you must download and set up the “Speech-to-Text Transcription” on page 3-2 functionality. Once you download and set up the speech-to-text transcription functionality, the **Speech to Text** automation algorithm appears as an option on the toolbar.

Select `Counting-16-44p1-mono-15secs.wav` from the **Data Browser**.

To create a region-level label that indicates if speech is detected, first select **Speech Detector** from the **AUTOMATION** section. You can control the speech detection algorithm using the **Window Length (s)** and **Merge Regions Within (s)** parameters. Use the default parameters for the speech detection algorithm. To create an ROI label and to label regions of the selected audio file, select **Run**.

The screenshot shows the Audio Labeler software interface. The main window is titled "Audio Labeler - Counting-16-44p1-mono-15secs.wav". The interface is divided into several sections:

- PARAMETERS:**
 - Window Length (s): 0.05
 - Merge Regions Within (s): 0.25
- SELECTION:**
 - Label: Create new label
 - Type ROI - Logical
 - Data: Selected audio files(s)
- AUTOMATE:**
 - Run
 - Undo Run
 - Close Speech Detector
 - CLOSE
- Data Browser:**
 - Counting-16-44p1-mono-15secs.wav (selected)
- File Labels:**

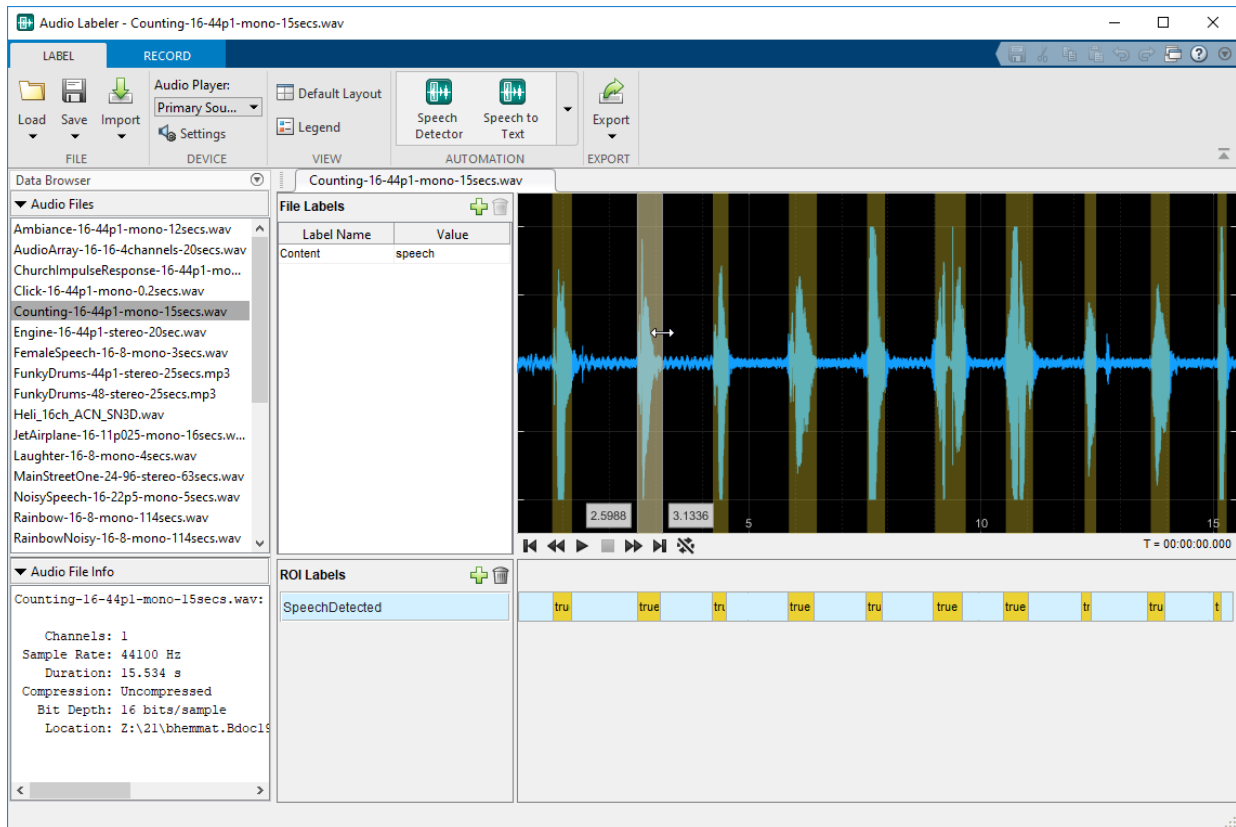
Label Name	Value
Content	speech
- ROI Labels:**

SpeechDetected
tru tru tru tru tru true true tru tru
- Audio File Info:**
 - Counting-16-44p1-mono-15secs.wav:
 - Channels: 1
 - Sample Rate: 44100 Hz
 - Duration: 15.534 s
 - Compression: Uncompressed
 - Bit Depth: 16 bits/sample
 - Location: Z:\21\bhemmat.Bdoc18

The main display area shows a waveform of the audio file with a blue line representing the signal. The x-axis is labeled with 5, 10, and 15 seconds. The y-axis represents amplitude. The waveform shows several distinct peaks, corresponding to the detected speech regions.

2 Use the Audio Labeler

Close the **Speech Detector** tab. You can correct or fine-tune the automatically generated **SpeechDetected** regions by selecting the ROI from the ROI bar, and then dragging the edges of the region. The ROI bar is directly to the right of the ROI label. When a region is selected, clicking ▶ plays only the selected region, enabling you to verify whether the selected region captures all relevant auditory information.



If you have set up a speech-to-text transcription service, select **Speech to Text** from the **Automation** section. You can control the speech-to-text transcription using name-value pair options specific to your selected service. This example uses the IBM® service and specifies no additional options.

The screenshot shows the Audio Labeler application window. The main interface is divided into several panels:

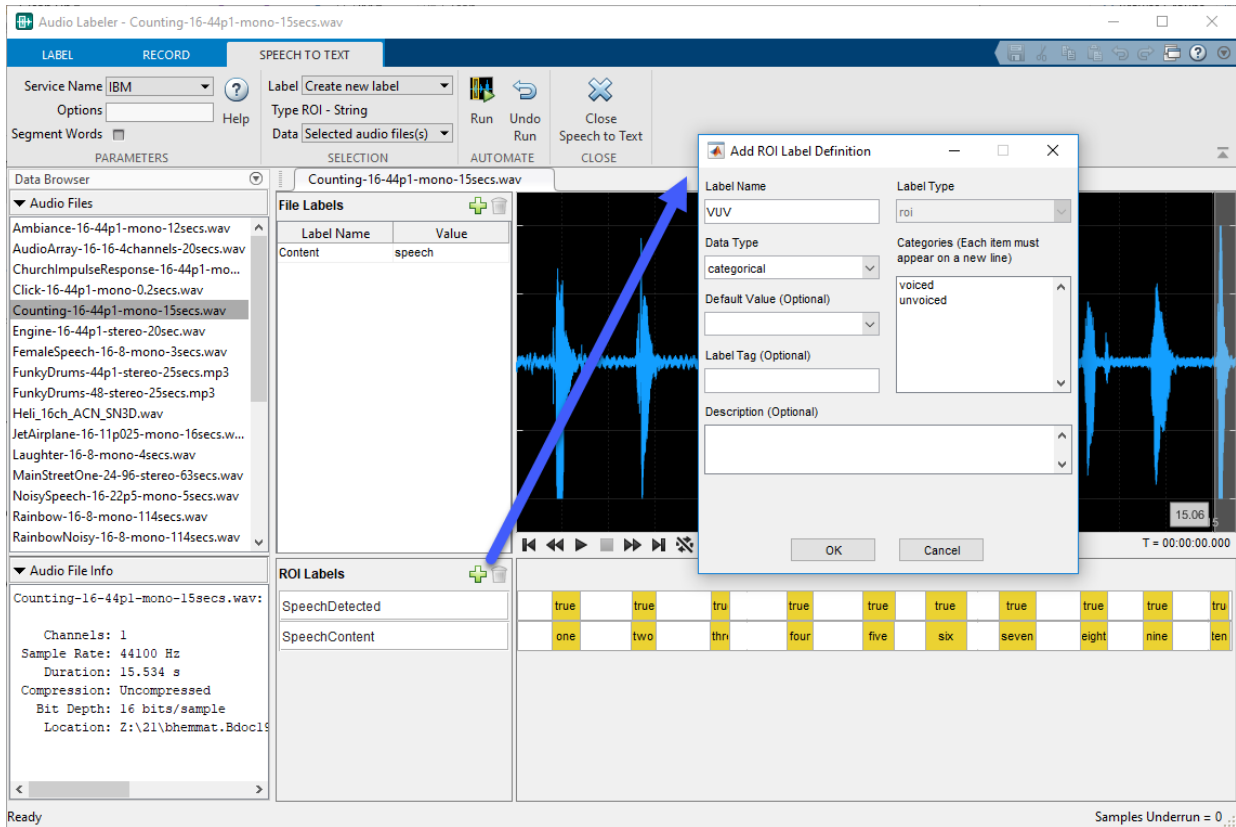
- Top Panel:** Includes tabs for LABEL, RECORD, and SPEECH TO TEXT. It features a 'Service Name' dropdown set to 'IBM', a 'Label' dropdown set to 'Create new label', and a 'Type ROI' dropdown set to 'String'. There are also buttons for 'Run', 'Undo Run', 'Close Speech to Text', and 'Close'.
- Data Browser:** A list of audio files, with 'Counting-16-44p1-mono-15secs.wav' selected.
- File Labels:** A table with columns 'Label Name' and 'Value'. It contains one entry: 'Content' with the value 'speech'.
- Audio File Info:** Displays metadata for the selected file: 'Counting-16-44p1-mono-15secs.wav', 'Channels: 1', 'Sample Rate: 44100 Hz', 'Duration: 15.534 s', 'Compression: Uncompressed', 'Bit Depth: 16 bits/sample', and 'Location: Z:\21\bhemmat.Bdoc15'.
- ROI Labels:** A table with columns 'Label Name' and 'Value'. It contains two entries: 'SpeechDetected' with values 'true' and 'SpeechContent' with values 'one', 'to', 'three', 'for', 'five', 'six', 'seven', 'eight', 'nine', 'an'.
- Waveform:** A blue waveform plot on a black background, showing the audio signal over time. The x-axis is labeled with '5', '10', and '15'.

The status bar at the bottom left shows 'Ready' and the bottom right shows 'Samples Underrun = 0'.

The ROI labels returned from the transcription service are strings with beginning and end points. The beginning and end points do not exactly correspond to the beginning and end points of the manually corrected speech detection regions. You can correct the endpoints of the **SpeechContent** ROI label by selecting the region and then dragging the edges of the region. The transcription service misclassified the words "two" as "to," "four" as "for," and "ten" as "then." You can correct the string by selecting the region and then entering a new string.

Create another region-level label by clicking  in the **ROI Labels** panel. Set **Label Name** to VUV, set **Data Type** to categorical, and **Categories** to voiced and unvoiced.

2 Use the Audio Labeler



By default, the waveform viewer shows the entire file. To display tools for zooming and panning, hover over the top right corner of the plot. Zoom in on the first five seconds of the audio file.

When you select a region in the plot and then hover over any of the two ROI bars, the shadow of the region appears. To assign the selected region to the category **voiced**, click **one** on the **SpeechContent** label bar. Hover over the **VUV** label bar and then click the shadow and choose **voiced**.

The screenshot shows the Audio Labeler software interface for the file "Counting-16-44p1-mono-15secs.wav". The interface is divided into several sections:

- Top Panel:** Includes tabs for "LABEL", "RECORD", and "SPEECH TO TEXT". It features a "Service Name" dropdown set to "IBM", a "Label" dropdown set to "Create new label", and a "Type ROI - String" dropdown set to "String". There are also buttons for "Run", "Undo Run", "Close Speech to Text", and "Close".
- Data Browser:** A list of audio files with "Counting-16-44p1-mono-15secs.wav" selected.
- File Labels:** A table with columns "Label Name" and "Value". It contains one entry: "Content" with the value "speech".
- Audio Waveform:** A blue waveform plot showing amplitude over time. The x-axis is labeled with time values: 0.76075, 1, 1.3992, 2, 2.5, 3, 3.5, 4, 4.5. The total time is shown as "T = 00:00:00.000".
- Audio File Info:** Metadata for the selected file:
 - Counting-16-44p1-mono-15secs.wav:
 - Channels: 1
 - Sample Rate: 44100 Hz
 - Duration: 15.534 s
 - Compression: Uncompressed
 - Bit Depth: 16 bits/sample
 - Location: Z:\21\bhemmat.Bdoc18
- ROI Labels:** A table for labeling regions of interest (ROIs) in the audio. It has columns for "SpeechDetected", "SpeechContent", and "VUV".

SpeechDetected	SpeechContent	VUV
true	one	voiced
true	two	voiced
true	three	unvoiced

At the bottom left, it says "Ready" and at the bottom right, "Samples Underrun = 0".

The next two words, "two" and "three," contain both voiced and unvoiced speech. Select each region of speech on the plot, hover over the VUV label bar, and select the correct category for that region.

2 Use the Audio Labeler

The screenshot displays the Audio Labeler software interface for the file "Counting-16-44p1-mono-15secs.wav". The interface is divided into several panels:

- Top Panel:** Includes tabs for "LABEL", "RECORD", and "SPEECH TO TEXT". It features a "Service Name" dropdown set to "IBM", "Options", "Segment Words" checkbox, and a "Help" button. The "SPEECH TO TEXT" section includes a "Label" dropdown set to "Create new label", "Type ROI - String", "Data" dropdown set to "Selected audio file(s)", and buttons for "Run", "Undo Run", "Close Speech to Text", and "Close".
- Data Browser:** A list of audio files with "Counting-16-44p1-mono-15secs.wav" selected.
- File Labels:** A table with columns "Label Name" and "Value". It contains one entry: "Content" with the value "speech".
- Audio Waveform:** A visual representation of the audio signal with a blue waveform and a yellow vertical line indicating the current time position at 4.6655 seconds.
- ROI Labels:** A table showing the results of the labeling process. It includes rows for "SpeechDetected", "SpeechContent", and "VUV".
- Audio File Info:** Technical details for the selected audio file, including Channels: 1, Sample Rate: 44100 Hz, Duration: 15.534 s, Compression: Uncompressed, Bit Depth: 16 bits/sample, and Location: Z:\21\bhemmat.Bdoc15.

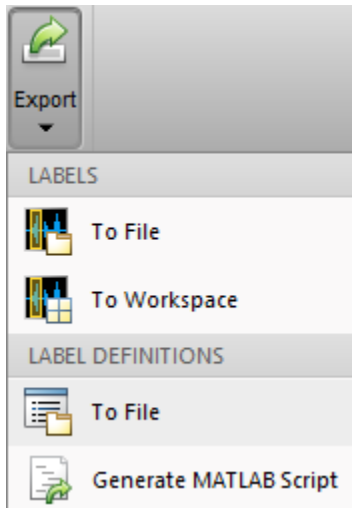
Label Name	Value
Content	speech

SpeechDetected	SpeechContent	VUV
true	one	voiced
true	two	unv, voiced
true	three	ur

Ready Samples Underrun = 0

Export Label Definitions

You can export label definitions as a MAT file or as a MATLAB script. Maintaining label definitions enables consistent labeling between users and sessions. Select **Export > Label Definitions > To File**.



The labels are saved as an array of `signalLabelDefinition` objects. In your next session, you can import the label definitions by selecting **Import > Label Definitions > From File**.

Export Labeled Audio Data

You can export the labeled signal set to a file or to your workspace. Select **Export > Labels > To Workspace**.

The **Audio Labeler** creates a `labeledSignalSet` object named `labeledSet_HHMMSS`, where *HHMMSS* is the time the object is created in hours, minutes, and seconds.

```
labeledSet_104620
```

```
labeledSet_104620 =
```

```
labeledSignalSet with properties:
```

```
    Source: {29x1 cell}
  NumMembers: 29
TimeInformation: "inherent"
    Labels: [29x4 table]
  Description: ""
```

Use `labelDefinitionsHierarchy` to see a list of labels and sublabels.
Use `setLabelValue` to add data to the set.

The labels you created are saved as a table to the `Labels` property.

`labeledSet_142356.Labels`

`ans =`

29×4 table

```
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Ambiance-16-44p1-mono-12secs.v
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\AudioArray-16-16-4channels-20
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\ChurchImpulseResponse-16-44p1
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Click-16-44p1-mono-0.2secs.wav
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Counting-16-44p1-mono-15secs.v
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Engine-16-44p1-stereo-20sec.wa
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FemaleSpeech-16-8-mono-3secs.v
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FunkyDrums-44p1-stereo-25secs
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FunkyDrums-48-stereo-25secs.mp
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Heli_16ch_ACN_SN3D.wav
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\JetAirplane-16-11p025-mono-16s
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Laughter-16-8-mono-4secs.wav
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples>MainStreetOne-24-96-stereo-63s
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\NoisySpeech-16-22p5-mono-5secs
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Rainbow-16-8-mono-114secs.wav
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RainbowNoisy-16-8-mono-114secs
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RandomOscThree-24-96-stereo-11
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockDrums-44p1-stereo-11secs.r
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockDrums-48-stereo-11secs.mp3
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockGuitar-16-44p1-stereo-72se
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockGuitar-16-96-stereo-72sec
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\SoftGuitar-44p1_mono-10mins.o
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\SpeechDFT-16-8-mono-5secs.wav
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\TrainWhistle-16-44p1-mono-9sec
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Turbine-16-44p1-mono-22secs.wa
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-44p1-stereo
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-8-mono-1000s
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-8-mono-200s
C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WaveGuideLoopOne-24-96-stereo
```

The file names associated with the labels are saved as a cell array to the Source property.

```
labeledSet_104620.Source
```

```
ans =
```

```
29x1 cell array
```

```
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Ambiance-16-44p1-mono-12secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\AudioArray-16-16-4channels-2secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\ChurchImpulseResponse-16-44p1-mono-12secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Click-16-44p1-mono-0.2secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Counting-16-44p1-mono-15secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Engine-16-44p1-stereo-20secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FemaleSpeech-16-8-mono-3secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FunkyDrums-44p1-stereo-25secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\FunkyDrums-48-stereo-25secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Heli_16ch_ACN_SN3D.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\JetAirplane-16-11p025-mono-12secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Laughter-16-8-mono-4secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples>MainStreetOne-24-96-stereo-4secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\NoisySpeech-16-22p5-mono-5secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Rainbow-16-8-mono-114secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RainbowNoisy-16-8-mono-114secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RandomOscThree-24-96-stereo-11secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockDrums-44p1-stereo-11secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockDrums-48-stereo-11secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockGuitar-16-44p1-stereo-72secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\RockGuitar-16-96-stereo-72secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\SoftGuitar-44p1_mono-10mins.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\SpeechDFT-16-8-mono-5secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\TrainWhistle-16-44p1-mono-9secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\Turbine-16-44p1-mono-22secs.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-44p1-stereo-10mins.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-8-mono-10mins.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WashingMachine-16-8-mono-20mins.wav'}
{'C:\Program Files\MATLAB\R2019b\toolbox\audio\samples\WaveGuideLoopOne-24-96-stereo-10mins.wav'}
```

Prepare Audio Datastore for Deep Learning Workflow

To continue on to a deep learning or machine learning workflow, use `audioDatastore`. Using an audio datastore enables you to apply capabilities that are common to machine

learning applications, such as `splitEachLabel`. `splitEachLabel` enables you split your data into train and test sets.

Create an audio datastore for your labeled signal set. Specify the location of the audio files as the first argument of `audioDatastore` and set the `Labels` property of `audioDatastore` to the `Labels` property of the labeled signal set.

```
ADS = audioDatastore(labeledSet_104620.Source, 'Labels', labeledSet_104620.Labels)
```

```
ADS =
```

```
audioDatastore with properties:
```

```
Files: {
    ' ...\toolbox\audio\samples\Ambiance-16-44p1-mono-12secs
    ' ...\toolbox\audio\samples\AudioArray-16-16-4channels-20
    ' ...\toolbox\audio\samples\ChurchImpulseResponse-16-44p
    ... and 26 more
}
Labels: 29-by-4 table
AlternateFileSystemRoots: {}
OutputDataType: 'double'
```

Call `countEachLabel` and specify the `Content` table variable to count the number of files that are labeled as `ambience`, `music`, `speech`, or `unknown`.

```
countEachLabel(ADS, 'TableVariable', 'Content')
```

```
ans =
```

```
4x2 table
```

Content	Count
ambience	13
music	9
speech	6
unknown	1

For examples of using labeled audio data in a machine learning or deep learning workflow, see:

- “Speech Command Recognition Using Deep Learning”
- “Speaker Identification Using Pitch and MFCC”
- “Denoise Speech Using Deep Learning Networks”
- “Classify Gender Using LSTM Networks”
- “Music Genre Classification Using Wavelet Time Scattering”

See Also

`audioDatastore` | `audioDeviceReader` | `audioDeviceWriter` | `labeledSignalSet`
| `signalLabelDefinition`

Speech2Text Walkthrough Chapter

Speech-to-Text Transcription

Audio Toolbox enables you to interface with third-party speech-to-text APIs from MATLAB.

```
ans =  
  
    "hello world"  
  
>>
```



To interface with third-party speech-to-text APIs, you must have the following:

- Audio Toolbox release R2017a or above
- Audio Toolbox extended functionality available from File Exchange
- One of the following APIs:
 - Google® Speech API
 - IBM Watson Speech API
 - Microsoft® Azure Speech API

The third-party APIs require you to generate keys for identification purposes. To begin, download the extended Audio Toolbox functionality from File Exchange. The File Exchange submission includes a tutorial to get you started. Once you have installed the speech-to-text functionality and set up your API keys, you can perform speech-to-text transcription programmatically or using the **Audio Labeler** app.

See Also

Audio Labeler | `voiceActivityDetector`

External Websites

- `speech2text` on File Exchange

Measure Impulse Response of an Audio System

Impulse Response Measurer Walkthrough

In this tutorial, explore key functionality of the **Impulse Response Measurer**. The **Impulse Response Measurer** app enables you to

- Configure your audio I/O system.
- Acquire impulse response (IR) measurements using either the Exponential Swept Sine or Maximum Length Sequences methods.
- View and manage captured IR data.
- Export the data to a file, workspace, or other app for further study.

To begin, open the **Impulse Response Measurer** app by selecting the  icon from the app gallery.

Configure Audio I/O System

The **Impulse Response Measurer** app enables you to specify an audio device, sample rate, player channel, and recorder channel. The audio device must be a real or virtual device enabled for simultaneous playback and recording (full-duplex mode) and must use a supported driver. Supported drivers are platform-specific:


- Windows® -- ASIO™
- Mac -- CoreAudio
- Linux® -- ALSA

Valid sample rates depend on your specified audio device.

You can use the level monitor to verify the configuration of your audio I/O system.

Configure IR Acquisition Method

To configure your IR acquisition method, use the **Method** and **Method Settings** sections of the toolstrip.



You can select the method to acquire IR measurements as either:

- Maximum Length Sequences (**MLS**)
- Exponential Swept Sine (**Exponential Swept Sine**)


Both methods for IR acquisition have the same basic settings, including:

- **Number of Runs** -- Number of times the excitation signal is sent within a single capture. Multiple runs are used to average individual impulse response captures to reduce measurement noise.
- **Duration per Run (s)** -- Total time of each run in seconds.
- **Excitation Level (dBFS)** -- The level of the excitation signal in dBFS.

Both methods for IR acquisition also have the same advanced run settings, including:

- **Wait before first run** -- Delay before starting first run. The delay allows time for any last-minute tasks, such as exiting a room before testing its acoustics.
- **Pause between runs** -- Duration of pause between runs. During a pause, the excitation signal is not sent, and audio is not recorded. When using the **Exponential Swept Sine** method, include a pause between runs to avoid buildup of reverberations. Pause between runs is always zero for the **MLS** method.
- **Number of warmup runs** -- Number of times to output the excitation signal before acquisition. The **MLS** method assumes the signal it acquires is a combination of the excitation signal and its impulse response. Use warmup runs to remove transients.

The total capture time is a sum of run durations, pauses, and the initial wait:




The **Exponential Swept Sine** method has additional **Advanced Settings** to control the excitation signal, including:

- **Sweep start frequency**
- **Sweep stop frequency**
- **Sweep duration**
- **End silence duration**

When using the **Exponential Swept Sine** method, the **Run Duration** is divided into **Sweep duration** and **End silence duration**. During the end silence, the app continues to record audio, enabling acquisition of the response over the entire range of the frequency sweep.

Acquire IR Measurements

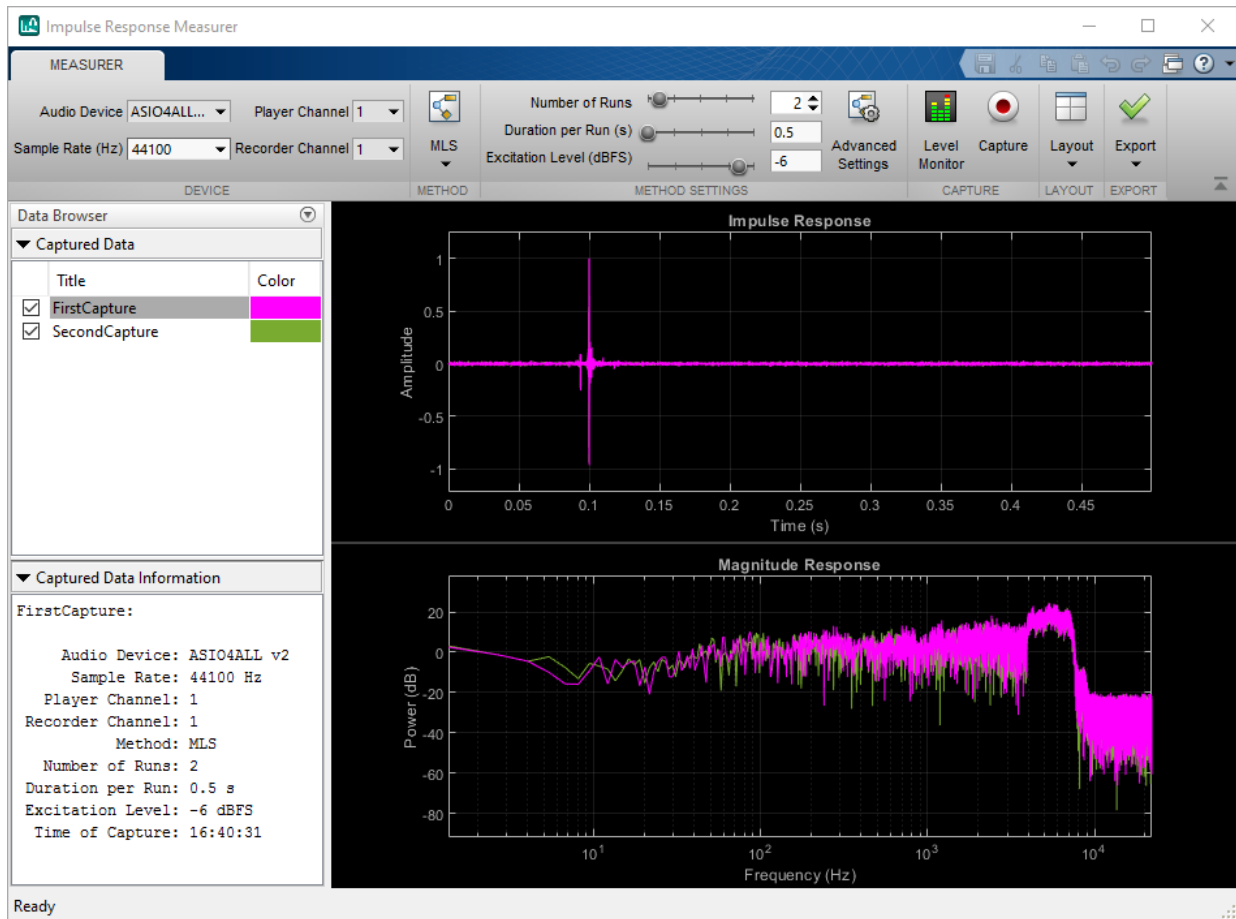
For this example, use the **MLS** method with default settings. Once you have your audio device set up, click **Capture**. A dialog box opens that displays the progress of your capture. IR Measurements are captured twice.



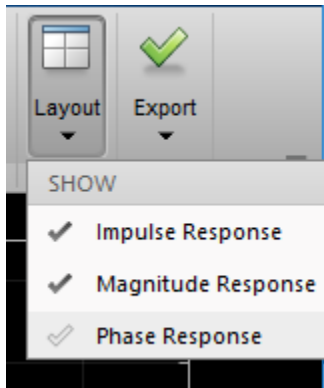
Analyze and Manage IR Measurements

After the capture, the **Impulse Response Measurer** app stores the captured data locally. The **Data Browser** displays the title of the captured data, the colors used for plotting, and information about the settings used to acquire the data. You can double-click a color in the **Data Browser** to choose which color you want associated with each impulse response. You can also double-click the title to rename your captured data. Rename your captures as *FirstCapture* and *SecondCapture*, and change the colors to pink and green. To make one impulse response plot appear on top of the other, select the title in the **Data Browser**. Select the capture you relabeled *FirstCapture*.

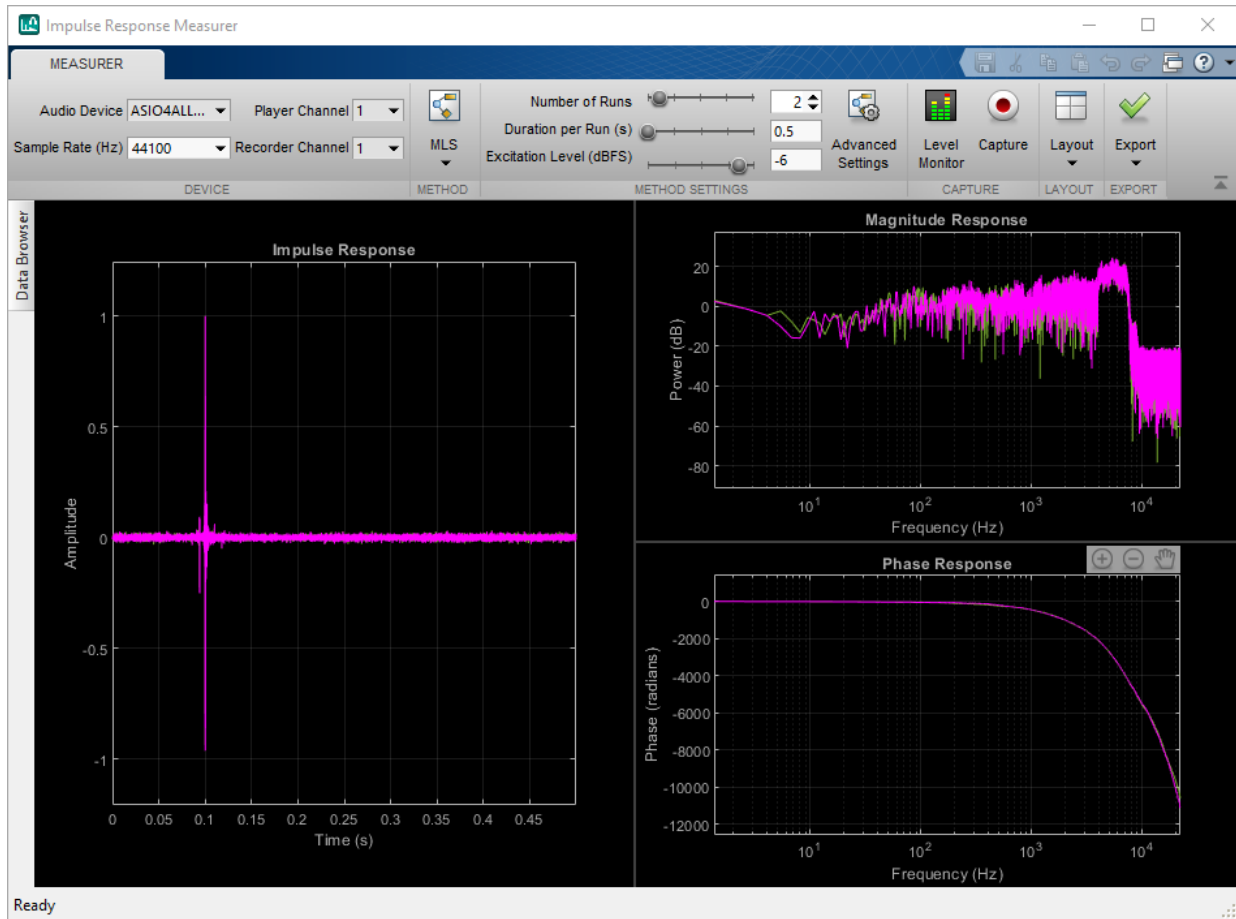
4 Measure Impulse Response of an Audio System




By default, the impulse response and magnitude response are plotted. You can view any combination of the impulse response, magnitude, and phase response using the **Layout** button. Minimize the **Data Browser**, then select the phase response plot for inclusion.



4 Measure Impulse Response of an Audio System

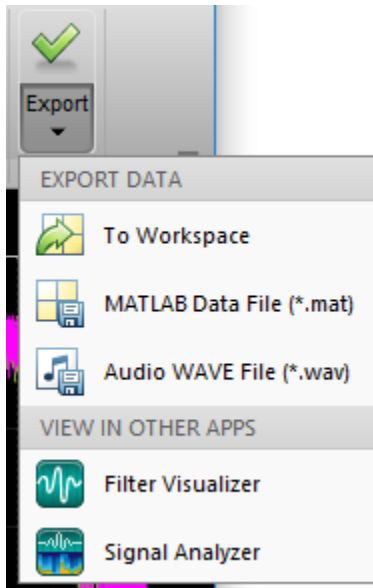


You can toggle the relative size of the plot by moving the dividers. You can zoom in and out by selecting the plus and minus icons on the UI. The icons appear when your pointer is over the plot. Zooming in and out of either the magnitude response or the phase response updates the other. Zoom in on the impulse response plot and in the range 100–1000 Hz of your frequency response plots.



Export IR Measurements

To view export options for further analysis or use, select the **Export** button.



Export the data to your workspace. The data is saved as a table. To inspect how the data is saved, display the table you exported.

```
irdata_172519
```

```
irdata_172519 =
```

```
2x14 table
```

	TimeOfCapture	ImpulseResponse	MagnitudeResponse	PhaseResponse
FirstCapture	21-Dec-2017 16:40:31 -0500	[1x1 struct]	[1x1 struct]	[1x1 struct]
SecondCapture	21-Dec-2017 16:42:18 -0500	[1x1 struct]	[1x1 struct]	[1x1 struct]

When you export the data as a MAT-file, the same table is saved as when you export to the workspace. When you select to export the data as a WAV file, each impulse response is saved as a separate WAV file. The title of the capture as the name of the WAV file. In this example, selecting to export data to audio WAV file places two WAV files in the specified folder, `FirstCapture.wav` and `SecondCapture.wav`.

To analyze your captured data further, view the data in the **Filter Visualization Tool** or **Signal Analyzer** app.

See Also

Impulse Response Measurer | `audioPlayerRecorder` | `reverberator` | `splMeter`

Related Examples

- “Measure Impulse Response of an Audio System”
- “Measure Frequency Response of an Audio Device”

Design and Play a MIDI Synthesizer

Design and Play a MIDI Synthesizer

The MIDI protocol enables you to send and receive information describing sound. A MIDI synthesizer is a device or software that synthesizes sound in response to incoming MIDI data. In its simplest form, a MIDI synthesizer converts MIDI note messages to an audio signal. More complicated synthesizers provide fine-tune control of the resulting sound, enabling you to mimic instruments. In this tutorial, you create a monophonic synthesizer that converts a stream of MIDI note messages to an audio signal in real time.

To learn about interfacing with MIDI devices in general, see “MIDI Device Interface” on page 6-2.

Convert MIDI Note Messages to Sound Waves

MIDI note information is packaged as a `NoteOn` or `NoteOff` `midimsg` object in Audio Toolbox. Both `NoteOn` and `NoteOff` `midimsg` objects have `Note` and `Velocity` properties:

- `Velocity` indicates how hard a note is played. By convention, Note On messages with velocity set to zero represent note off messages. Representing note off messages with note on messages is more efficient when using Running Status.
- `Note` indicates the frequency of the audio signal. The `Note` property takes a value between zero and 127, inclusive. The MIDI protocol specifies that 60 is Middle C, with all other notes relative to that note. Create a MIDI note on message that indicates to play Middle C.

```
channel = 1;  
note = 60;  
velocity = 64;  
msg = midimsg('NoteOn', channel, note, velocity)
```

```
msg =
```

```
MIDI message:  
NoteOn          Channel: 1 Note: 60 Velocity: 64 Timestamp: 0 [ 90 3C 40 ]
```

To interpret the note property as frequency, use the equal tempered scale and the A440 convention:

```
frequency = 440 * 2^((msg.Note-69)/12)
```



```
frequency =
    261.6256
```

Some MIDI synthesizers use an Attack Decay Sustain Release (ADSR) envelope to control the volume, or amplitude, of a note over time. For simplicity, use the note velocity to determine the amplitude. Conceptually, if a key is hit harder, the resulting sound is louder. The `Velocity` property takes a value between zero and 127, inclusive. Normalize the velocity and interpret as the note amplitude.

```
amplitude = msg(1).Velocity/127

amplitude =
    0.5039
```

To synthesize a sine wave, create an `audioOscillator System` object™. To play the sound to your computer's default audio output device, create an `audioDeviceWriter System` object. Step the objects for two seconds and listen to the note.

```
osc = audioOscillator('Frequency',frequency,'Amplitude',amplitude);
deviceWriter = audioDeviceWriter('SampleRate',osc.SampleRate);

tic
while toc < 2
    synthesizedAudio = osc();
    deviceWriter(synthesizedAudio);
end
```

Synthesize MIDI Messages

To play an array of `midimsg` objects with appropriate timing, create a loop.

First, create an array of `midimsg` objects and cache the note on and note off times to the variable, `eventTimes`.

```
msgs = [midimsg('Note',channel,60,64,0.5,0), ...
        midimsg('Note',channel,62,64,0.5,.75), ...
        midimsg('Note',channel,57,40,0.5,1.5), ...
        midimsg('Note',channel,60,50,1,3)];
eventTimes = [msgs.Timestamp];
```

To mimic receiving notes in real time, create a for-loop that uses the `eventTimes` variable and `tic` and `toc` to play notes according to the MIDI message timestamps. Release your audio device after the loop is complete.

```
i = 1;
tic
while toc < max(eventTimes)
    if toc > eventTimes(i)
        msg = msgs(i);
        i = i+1;

        if msg.Velocity~= 0
            osc.Frequency = 440 * 2^((msg.Note-69)/12);
            osc.Amplitude = msg.Velocity/127;
        else
            osc.Amplitude = 0;
        end
    end
    deviceWriter(osc());
end
release(deviceWriter)
```

Synthesize Real-Time Note Messages from MIDI Device

To receive and synthesize note messages in real time, create an interface to a MIDI device. The `simplesynth` example function:

- receives MIDI note messages from a specified MIDI device
- synthesizes an audio signal
- plays them to your audio output device in real time

Save the `simplesynth` function to your current folder.

`simplesynth`

```
function simplesynth(midiDeviceName)

    midiInput = mididevice(midiDeviceName);
    osc = audioOscillator('square', 'Amplitude', 0);
    deviceWriter = audioDeviceWriter;
    deviceWriter.SupportVariableSizeInput = true;
    deviceWriter.BufferSize = 64; % small buffer keeps MIDI latency low

    while true
        msgs = midireceive(midiInput);
        for i = 1:numel(msgs)
            msg = msgs(i);
```

```

        if isNoteOn(msg)
            osc.Frequency = note2freq(msg.Note);
            osc.Amplitude = msg.Velocity/127;
        elseif isNoteOff(msg)
            if msg.Note == msg.Note
                osc.Amplitude = 0;
            end
        end
    end
    deviceWriter(osc());
end
end

function yes = isNoteOn(msg)
    yes = msg.Type == midimsgtype.NoteOn ...
        && msg.Velocity > 0;
end

function yes = isNoteOff(msg)
    yes = msg.Type == midimsgtype.NoteOff ...
        || (msg.Type == midimsgtype.NoteOn && msg.Velocity == 0);
end

function freq = note2freq(note)
    freqA = 440;
    noteA = 69;
    freq = freqA * 2.^((note-noteA)/12);
end

```

To query your system for your device name, use `mididevinfo`. To listen to your chosen device, call the `simplesynth` function with the device name. This example uses an M-Audio KeyRig 25 device, which registers with device name `USB 02` on the machine used in this example.

`mididevinfo`


```

MIDI devices available:
ID Direction Interface Name
0 output MMSystem 'Microsoft MIDI Mapper'
1 input MMSystem 'USB MIDI Interface '
2 input MMSystem 'USB 02'
3 output MMSystem 'Microsoft GS Wavetable Synth'
4 output MMSystem 'USB MIDI Interface '
5 output MMSystem 'USB 02'

```

Call the `simplesynth` function with your device name. The `simplesynth` function listens for note messages and plays them to your default audio output device. Play notes on your MIDI device and listen to the synthesized audio.

```
simplesynth('USB 02')
```



Use `Ctrl-C` to end the connection.

See Also

Classes

`mididevice` | `midimsg`

Functions

`mididevinfo` | `midireceive` | `midisend`

External Websites

- <https://www.midi.org>

MIDI Device Interface

MIDI Device Interface

MIDI

This tutorial introduces the Musical Instrument Digital Interface (MIDI) protocol and how you can use Audio Toolbox to interact with MIDI devices. The tools described here enable you to send and receive all MIDI messages as described by the MIDI protocol. If you are interested only in sending and receiving Control Change messages with a MIDI control surface, see “MIDI Control Surface Interface” on page 9-2. If you are interested in using MIDI to control your audio plugins, see “MIDI Control for Audio Plugins” on page 8-2. To learn more about MIDI in general, consult The MIDI Manufacturers Association.

MIDI is a technical standard for communication between electronic instruments, computers, and related devices. MIDI carries event messages specific to audio signals, such as pitch and velocity, as well as control signals for parameters and clock signals to synchronize tempo.

MIDI Devices

A MIDI device is any device capable of sending or receiving MIDI messages. MIDI devices have input ports, output ports, or both. The MIDI protocol defines messages as unidirectional. A MIDI device can be real-world or virtual.

Audio Toolbox enables you to create an interface to a MIDI device using `mididevice`. To create a MIDI interface to a specific device, use `mididevinfo` to query your system for available devices. Then create a `mididevice` object by specifying a MIDI device by name or ID.

```
mididevinfo
```

```
MIDI devices available:
```

ID	Direction	Interface	Name
0	output	MMSystem	'Microsoft MIDI Mapper'
1	input	MMSystem	'USB MIDI Interface '
2	output	MMSystem	'Microsoft GS Wavetable Synth'
3	output	MMSystem	'USB MIDI Interface '

```
device = mididevice('USB MIDI Interface ')
```

```
device =
```

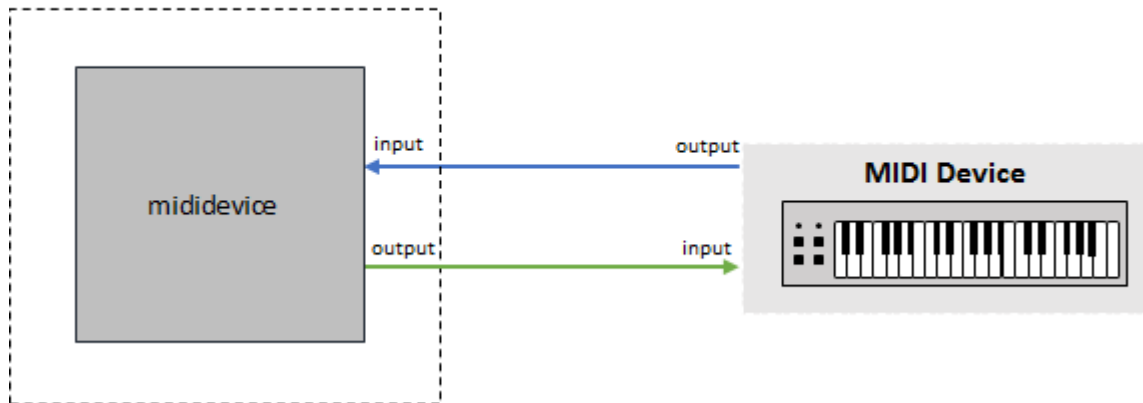
```

mididevice connected to
  Input: 'USB MIDI Interface ' (1)
  Output: 'USB MIDI Interface ' (3)

```

You can specify a `mididevice` object to listen for input messages, send output messages, or both. In this example, the `mididevice` object receives MIDI messages at the input port named 'USB MIDI Interface ', and sends MIDI messages from the output port named 'USB MIDI Interface '.

MATLAB



MIDI Messages

A MIDI message contains information that describes an audio-related action. For example, when you press a key on a keyboard, the corresponding MIDI message contains 3 bytes:

- 1 The first byte describes the kind of action and the channel. The first byte is referred to as the Status Byte.
- 2 The second byte describes which key is pressed. The second byte is referred to as a Data Byte.
- 3 The third byte describes how hard the key is played. The third byte is also a Data Byte.


This message is a Note On message. Note On is referred to as the message name, command, or type.

In MATLAB, a MIDI message is packaged as a `midimsg` object and can be manipulated as scalars or arrays. To create a MIDI message, call `midimsg` with a message type and then specify the required parameters for the specific message type. For example, to create a note on message, specify the `midimsg` Type as 'NoteOn' and then specify the required inputs: channel, note, and velocity.

```
channel = 1;
note = 60;
velocity = 64;
msg = midimsg('NoteOn',channel,note,velocity)

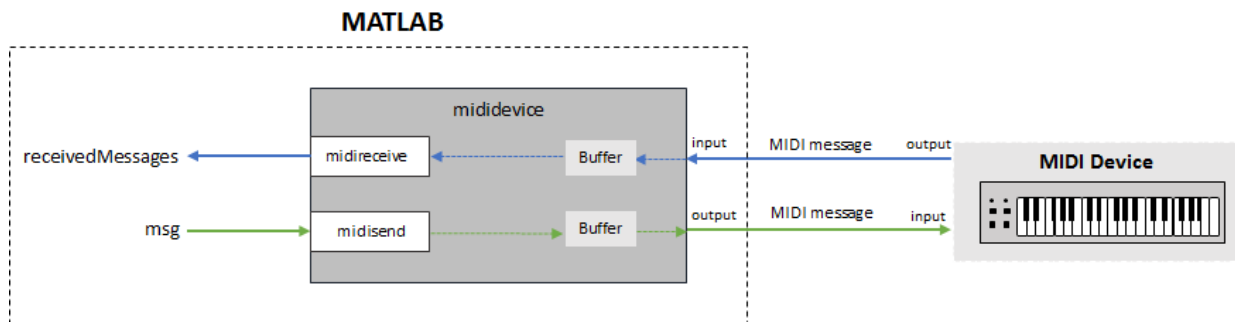
msg =
    MIDI message:
    NoteOn          Channel: 1 Note: 60 Velocity: 64 Timestamp: 0 [ 90 3C 40 ]
```

For convenience, `midimsg` displays the message type, channel, additional parameters, timestamp, and the constructed message in hexadecimal form. Hexadecimal is the preferred form because it has a straightforward interpretation:



Sending and Receiving MIDI Messages

To send and receive MIDI messages, use the `mididevice` object functions `midisend` and `midireceive`. When you create a `mididevice` object, it begins receiving data at its input and placing it in a buffer.



To retrieve MIDI messages from the buffer, call `midireceive`.

```
receivedMessages = midireceive(device)
```

```
receivedMessages =
```

```
  MIDI message:
    NoteOn           Channel: 1 Note: 36 Velocity: 64 Timestamp: 15861.9 [ 90 24 40 ]
    NoteOn           Channel: 1 Note: 36 Velocity: 0   Timestamp: 15862.1 [ 90 24 00 ]
```

The MIDI messages are returned as an array of `midimsg` objects. In this example, a MIDI keyboard key is pressed.

To send MIDI messages to a MIDI device, call `midisend`.

```
midisend(device,msg)
```

MIDI Message Types

The type of MIDI message you create is defined as a character vector or string. To create a MIDI message, specify it by its type and the required property values. For example, create a Channel Pressure MIDI message by entering the following at the command prompt:

```
channelPressureMessage = midimsg('ChannelPressure',1,20)
```

```
channelPressureMessage =
```

```
  MIDI message:
    ChannelPressure Channel: 1 ChannelPressure: 20 Timestamp: 0 [ D0 14 ]
```

After you create a MIDI message, you can modify the properties, but you cannot modify the type.

```
channelPressureMessage.ChannelPressure = 37
```

```
channelPressureMessage =
```

```
  MIDI message:
    ChannelPressure Channel: 1 ChannelPressure: 37 Timestamp: 0 [ D0 25 ]
```

The table summarizes valid MIDI message types.

MIDI Spec Description	midimsg(Type, PropertyValue1,...,PropertyValueN)			
	Type	properties		
Channel Voice Messages				
Note Off event	NoteOff	Channel	Note	Velocity
Note On event	NoteOn	Channel	Note	Velocity
Polyphonic Key Pressure (Aftertouch)	PolyKeyPressure	Channel	Note	KeyPressure
Control Change	ControlChange	Channel	CCNumber	CCValue
Program Change	ProgramChange	Channel	Program	
Channel Pressure (Aftertouch)	ChannelPressure	Channel	ChannelPressure	
Pitch Bend Change	PitchBend	Channel	PitchChange	
Channel Mode Messages				
All Sound Off	AllSoundOff	Channel		
Reset All Controllers	ResetAllControllers	Channel		
Local Control	LocalControl	Channel	LocalControl	
All Notes Off	AllNotesOff	Channel		
Omni Mode Off	OmniOff	Channel		
Omni Mode On	OmniOn	Channel		
Mono Mode On (Poly Off)	MonoOn	Channel	MonoChannels	
Poly Mode On (Mono Off)	PolyOn	Channel		
System Common Messages				
System Exclusive	SystemExclusive			
MIDI Time Code Quarter Frame	MIDITimeCodeQuarterFrame	TimeCodeSequence	TimeCodeValue	
Song Position Pointer	Song	SongPosition		
Song Select	SongSelect	Song		
Tune Request	TuneRequest			
End of Exclusive	EOX			
System Real-Time Messages				
Timing Clock	TimingClock			
Start	Start			
Continue	Continue			
Stop	Stop			
Active Sensing	ActiveSensing			
6-6 Reset	SystemReset			

The Audio Toolbox provides convenience syntaxes to create multiple MIDI messages used in sequence and to create arrays of MIDI messages. See `midimsg` for a complete list of syntaxes.

MIDI Message Timing


The MIDI protocol does not define message timing and assumes that messages are acted on immediately. Many applications require timing information for queuing and batch processing. For convenience, the Audio Toolbox packages timing information with MIDI messages into a single `midimsg` object. All `midimsg` objects have a `Timestamp` property, which is set during creation as an optional last argument or after creation. The default `Timestamp` is zero.

The interpretation of the `Timestamp` property depends on how a MIDI message is created and used:

- When receiving MIDI messages using `midireceive`, the underlying infrastructure assigns a timestamp when receiving MIDI messages. Conceptually, the timing clock starts when a `mididevice` object is created and attached as a listener to a given MIDI input port. If another `mididevice` is attached to the same input port, it receives timestamps from the same timing clock as the first object.
- When sending MIDI messages using `midisend`, timestamps are interpreted as when to send the message.

If there have been no recent calls to `midisend`, then `midisend` interprets timestamps as relative to the current real-world time. A message with a timestamp of zero is sent immediately. If there has been a recent call to `midisend`, then `midisend` interprets timestamps as relative to the largest timestamp of the last call to `midisend`. The timestamp clock for `midisend` is specific to the MIDI output port that `mididevice` is connected to.

Consider a pair of MIDI messages that turn a note on and off. The messages specify that the note starts after one second and is sustained for one second.

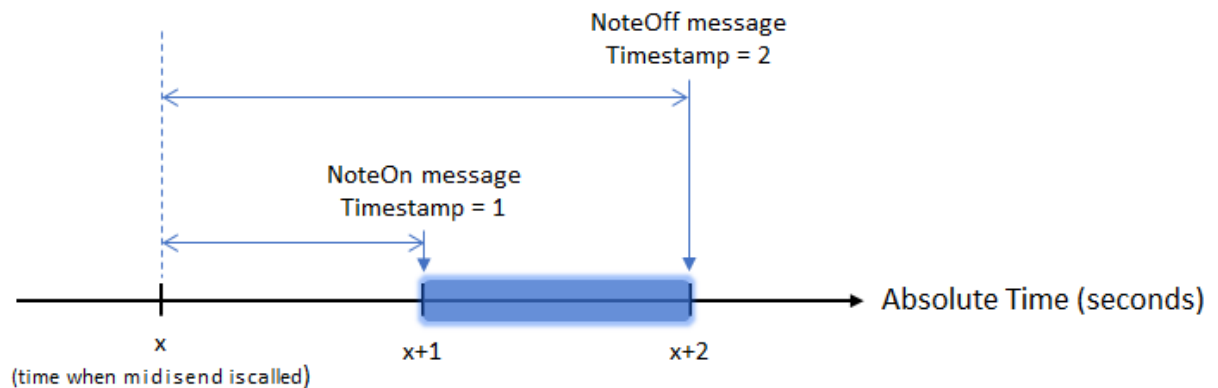


Create Note On and Note Off messages.

```
OnMsg = midimsg('NoteOn',1,59,64);
OffMsg = midimsg('NoteOn',1,59,0);
```

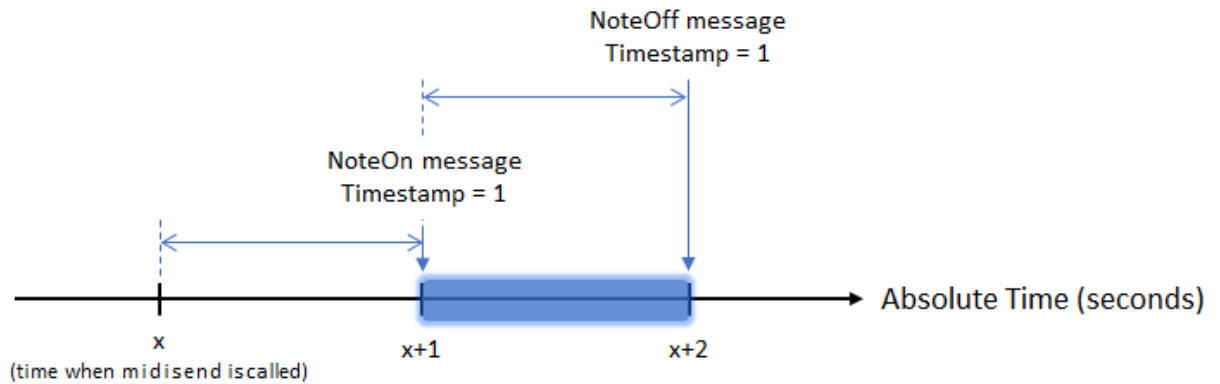
To send on and off messages using a single call to `midisend`, specify the timestamps of the messages relative to the same start time.

```
OnMsg.Timestamp = 1;
OffMsg.Timestamp = 2;
midisend(device,[OnMsg;OffMsg])
```




To send the Note Off message separately, specify the timestamp of the Note Off message relative to the largest timestamp of the previous call to `midisend`.

```
OnMsg.Timestamp = 1;
OffMsg.Timestamp = 1;
midisend(device,OnMsg)
midisend(device,OffMsg)
```



The "start" time, or reference time, for `midisend` is the max between the absolute time and the largest timestamp in the last call to `midisend`. For example, consider that x , the arbitrary start time, is equal to the current absolute time. If there is a 1.5-second pause between sending the note on and note off messages, the resulting note duration is 1.5 seconds.

```
OnMsg.Timestamp = 1;
OffMsg.Timestamp = 1;
midisend(device, OnMsg)
pause(1.5)
midisend(device, OffMsg)
```



Usually, MIDI messages are sent faster than or at real-time speeds so there is no need to track the absolute time.

For live performances or to enable interrupts in a MIDI stream, you can set timestamps to zero and then call `midisend` at appropriate real-world time intervals. Depending on your use case, you can divide your MIDI stream into small repeatable time chunks.

See Also

Classes

`mididevice` | `midimsg`

Functions

`mididevinfo` | `midireceive` | `midisend`

Related Examples

- “Design and Play a MIDI Synthesizer” on page 5-2

External Websites

- [MIDI Manufacturers Association](#)
- [Summary of MIDI Messages](#)

Dynamic Range Control

Dynamic Range Control

Dynamic range control is the adaptive adjustment of the dynamic range of a signal. The dynamic range of a signal is the logarithmic ratio of maximum to minimum signal amplitude specified in dB.

You can use dynamic range control to:


- Match an audio signal level to its environment
- Protect AD converters from overload
- Optimize information
- Suppress low-level noise

Types of dynamic range control include:

- Dynamic range compressor -- Attenuates the volume of loud sounds that cross a given threshold. They are often used in recording systems to protect hardware and to increase overall loudness.
- Dynamic range limiter -- A type of compressor that brickwalls sound above a given threshold.
- Dynamic range expander -- Attenuates the volume of quiet sounds below a given threshold. They are often used to make quiet sounds even quieter.
- Noise gate -- A type of expander that brickwalls sound below a given threshold.

This tutorial shows how to implement dynamic range control systems using the compressor, expander, limiter, and noiseGate System objects from Audio Toolbox. The tutorial also provides an illustrated example of dynamic range limiting at various stages of a dynamic range limiting system.

The diagram depicts a general dynamic range control system.



In a dynamic range control system, a gain signal is calculated in a sidechain and then applied to the input audio signal. The sidechain consists of:

- Linear to dB conversion: $x \rightarrow x_{dB}$
- Gain computation, by passing the dB signal through a static characteristic equation, and then taking the difference: $g_c = x_{sc} - x_{dB}$
- Gain smoothing over time: $g_c \rightarrow g_s$
- Addition of make-up gain (for compressors and limiters only): $g_s \rightarrow g_m$
- dB to linear conversion: $g_m \rightarrow g_{lin}$
- Application of the calculated gain signal to the original audio signal: $y = g_{lin} \times x$


Linear to dB Conversion

The gain signal used in dynamic range control is processed on a dB scale for all dynamic range controllers. There is no reference for the dB output; it is a straight conversion: $x_{dB} = 20\log_{10}(x)$. You might need to adjust the output of a dynamic range control system to the range of your system.


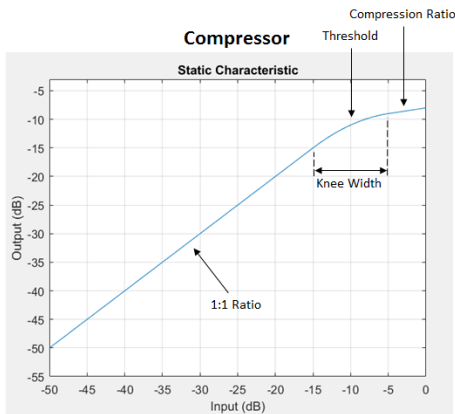

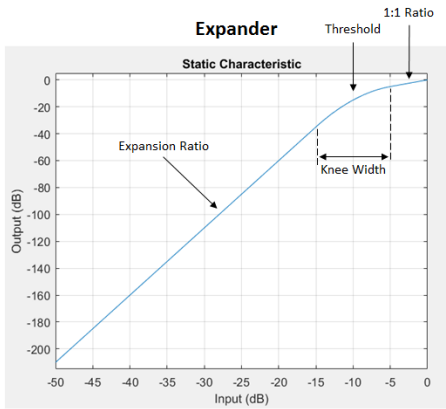
Gain Computer

The gain computer provides the first rough estimate of a gain signal for dynamic range control. The principal component of the gain computer is the static characteristic. Each type of dynamic range control has a different static characteristic with different tunable properties:

- **Threshold** -- All static characteristics have a threshold. On one side of the threshold, the input is given to the output with no modification. On the other side of the threshold, compression, expansion, brickwall limiting, or brickwall gating is applied.
- **Ratio** -- Expanders and compressors enable you to adjust the input-to-output ratio of the static characteristic above or below a given threshold.
- **KneeWidth** -- Expanders, compressors, and limiters enable you to adjust the knee width of the static characteristic. The knee of a static characteristic is centered at the threshold. An increase in knee width creates a smoother transition around the threshold. A knee width of zero provides no smoothing and is known as a hard knee. A knee width greater than zero is known as a soft knee.



In these static characteristic plots, the expander, limiter, and compressor each have a 10 dB knee width.



Gain Smoothing

All dynamic range controllers provide gain smoothing over time. Gain smoothing diminishes sharp jumps in the applied gain, which can result in artifacts and an unnatural sound. You can conceptualize gain smoothing as the addition of impedance to your gain signal.

The `expander` and `noiseGate` objects have the same smoothing equation, because a noise gate is a type of expander. The `limiter` and `compressor` objects have the same smoothing equation, because a limiter is a type of compressor.

The type of gain smoothing is specified by a combination of attack time, release time, and hold time coefficients. Attack time and release time correspond to the time it takes the gain signal to go from 10% to 90% of its final value. Hold time is a delay period before gain is applied. See the algorithms of individual dynamic range controller pages for more detailed explanations.

Smoothing Equations

expander and noiseGate

$$g_s[n] = \begin{cases} \alpha_A g_s[n-1] + (1 - \alpha_A) g_c[n] & \text{if } (C_A > k) \ \& \ (g_c[n] \leq g_s[n-1]) \\ g_s[n-1] & \text{if } C_A \leq k \\ \alpha_R g_s[n-1] + (1 - \alpha_R) g_c[n] & \text{if } (C_R > k) \ \& \ (g_c[n] > g_s[n-1]) \\ g_s[n-1] & \text{if } C_R \leq k \end{cases}$$

- α_A and α_R are determined by the sample rate and specified attack and release time:

$$\alpha_A = \exp\left(\frac{-\log(9)}{FS \times T_A}\right), \quad \alpha_R = \exp\left(\frac{-\log(9)}{FS \times T_R}\right)$$

- k is the specified hold time in samples.
- C_A and C_R are hold counters for attack and release, respectively.

compressor and limiter


$$g_s[n] = \begin{cases} \alpha_A g_s[n-1] + (1 - \alpha_A) g_c[n] & \text{if } g_c[n] \leq g_s[n-1] \\ \alpha_R g_s[n-1] + (1 - \alpha_R) g_c[n] & \text{if } g_c[n] > g_s[n-1] \end{cases}$$

- α_A and α_R are determined by the sample rate and specified attack and release time:

$$\alpha_A = \exp\left(\frac{-\log(9)}{FS \times T_A}\right), \quad \alpha_R = \exp\left(\frac{-\log(9)}{FS \times T_R}\right)$$



Gain Smoothing Example

Examine a trivial case of dynamic range compression for a two-step input signal. In this example, the compressor has a threshold of -10 dB, a compression ratio of 5, and a hard knee.





Several variations of gain smoothing are shown. On the top, a smoothed gain curve is shown for different attack time values, with release time set to zero seconds. In the middle, release time is varied and attack time is held constant at zero seconds. On the bottom, both attack and release time are specified by nonzero values.



Vary Attack Time
(Release Time = 0 s)



Vary Release Time
(Attack Time = 0 s)



Vary Release and Attack Time



Make-Up Gain

Make-up gain applies for compressors and limiters, where higher dB portions of a signal are attenuated or brickwalled. The dB reduction can significantly reduce total signal power. In these cases, make-up gain is applied after gain smoothing to increase the signal power. In Audio Toolbox, you can specify a set amount of make-up gain or specify the make-up gain mode as 'auto'.

The 'auto' make-up gain ensures that a 0 dB input results in a 0 dB output. For example, assume a static characteristic of a compressor with a soft knee:

$$x_{sc}(x_{dB}) = \begin{cases} x_{dB} & x_{dB} < \left(T - \frac{W}{2}\right) \\ x_{dB} + \frac{\left(\frac{1}{R} - 1\right)\left(x_{dB} - T + \frac{W}{2}\right)^2}{2W} & \left(T - \frac{W}{2}\right) \leq x_{dB} \leq \left(T + \frac{W}{2}\right) \\ T + \frac{(x_{dB} - T)}{R} & x_{dB} > \left(T + \frac{W}{2}\right) \end{cases}$$

T is the threshold, W is the knee width, and R is the compression ratio. The calculated auto make-up gain is the negative of the static characteristic equation evaluated at 0 dB:

$$\text{MAKE-UP GAIN} = -x_{sc}(0) = \begin{cases} 0 & \frac{W}{2} < T \\ -\frac{\left(\frac{1}{R} - 1\right)\left(T - \frac{W}{2}\right)^2}{2W} & -\frac{W}{2} \leq T \leq \frac{W}{2} \\ -T + \frac{T}{R} & -\frac{W}{2} > T \end{cases}$$

dB to Linear Conversion

Once the gain signal is determined in dB, it is transferred to the linear domain:

$$g_{lin} = 10^{g_m/20}.$$

Apply Calculated Gain

The final step in a dynamic control system is to apply the calculated gain by multiplication in the linear domain.

Example: Dynamic Range Limiter

The audio signal described in this example is a 0.5 second interval of a drum track. The limiter properties are:

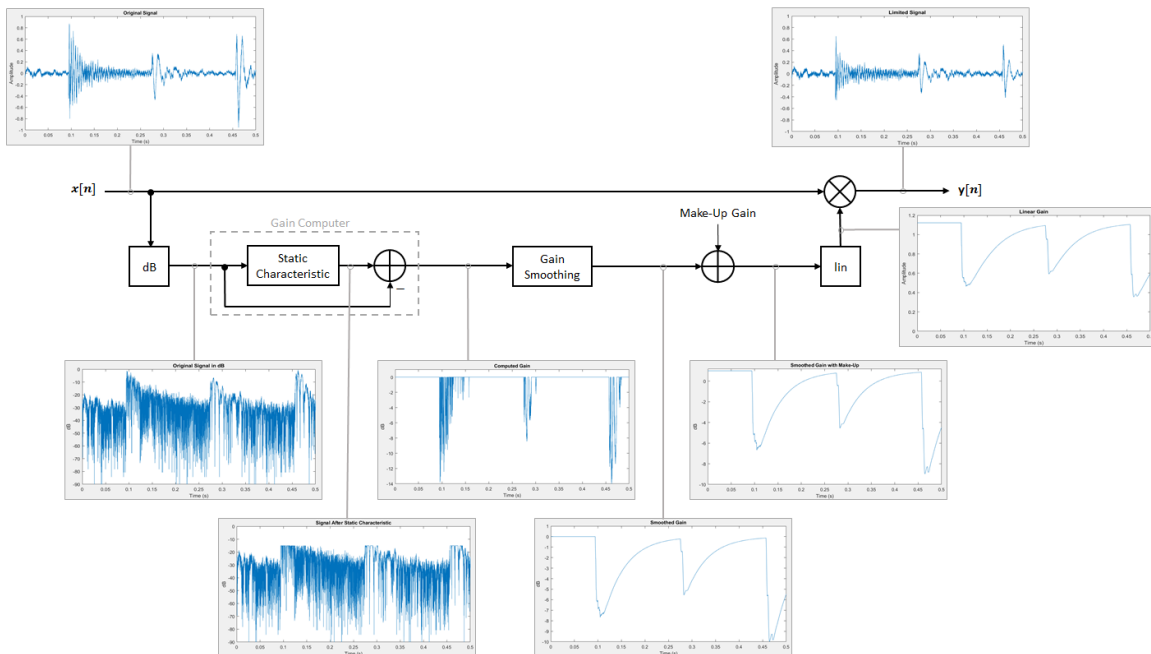
- Threshold = -15 dB
- Knee width = 0 (hard knee)
- Attack time = 0.004 seconds

- Release time = 0.1 seconds
- Make-up gain = 1 dB

To create a limiter System object with these properties, at the MATLAB command prompt, enter:


```
dRL = limiter('Threshold',-15,...
             'KneeWidth',0,...
             'AttackTime',0.004,...
             'ReleaseTime',0.1,...
             'MakeUpGainMode','property',...
             'MakeUpGain',1);
```

This example provides a visual walkthrough of the various stages of the dynamic range limiter system.




Linear to dB Conversion

The input signal is converted to a dB scale element by element.





Gain Computer

The static characteristic brickwall limits the dB signal at -15 dB. To determine the dB gain that results in this limiting, the gain computer subtracts the original dB signal from the dB signal processed by the static characteristic.




Gain Smoothing

The relatively short attack time specification results in a steep curve when the applied gain is suddenly increased. The relatively long release time results in a gradual diminishing of the applied gain.




Make-Up Gain

Assume a limiter with a 1 dB make-up gain value. The make-up gain is added to the smoothed gain signal.




$$+ (1 \text{ dB}) =$$




dB to Linear Conversion

The gain in dB is converted to a linear scale element by element.



Apply Calculated Gain

The original signal is multiplied by the linear gain.



References

- [1] Zolzer, Udo. "Dynamic Range Control." *Digital Audio Signal Processing*. 2nd ed. Chichester, UK: Wiley, 2008.
- [2] Giannoulis, Dimitrios, Michael Massberg, and Joshua D. Reiss. "Digital Dynamic Range Compressor Design -- A Tutorial And Analysis." *Journal of Audio Engineering Society*. Vol. 60, Issue 6, 2012, pp. 399-408.

See Also

Compressor | Expander | Limiter | Noise Gate | compressor | expander | limiter | noiseGate

More About

- “Dynamic Range Compression Using Overlap-Add Reconstruction”

MIDI Control for Audio Plugins

MIDI Control for Audio Plugins

MIDI and Plugins

MIDI control surfaces are commonly used in conjunction with audio plugins in digital audio workstation (DAW) environments. Synchronizing MIDI controls with plugin parameters provides a tangible interface for audio processing and is an efficient approach to parameter tuning.

In the MATLAB environment, audio plugins are defined as any valid class that derives from the `audioPlugin` base class or the `audioPluginSource` base class. For more information about how audio plugins are defined in the MATLAB environment, see “Audio Plugins in MATLAB”.


Use MIDI with MATLAB Plugins

The Audio Toolbox product provides three functions for enabling the interface between MIDI control surfaces and audio plugins:

- `configureMIDI` -- Configure MIDI connections between audio plugin and MIDI controller.
- `getMIDIConnections` -- Get MIDI connections of audio plugin.
- `disconnectMIDI` -- Disconnect MIDI controls from audio plugin.

These functions combine the abilities of general MIDI functions into a streamlined and user-friendly interface suited to audio plugins in MATLAB. For a tutorial on the general functions and the MIDI protocol, see “MIDI Control Surface Interface” on page 9-2.

This tutorial walks you through the MIDI functions for audio plugins in MATLAB.



1. Connect MIDI Device and Then Start MATLAB

Before starting MATLAB, connect your MIDI control surface to your computer and turn it on. For connection instructions, see the instructions for your MIDI device. If you start

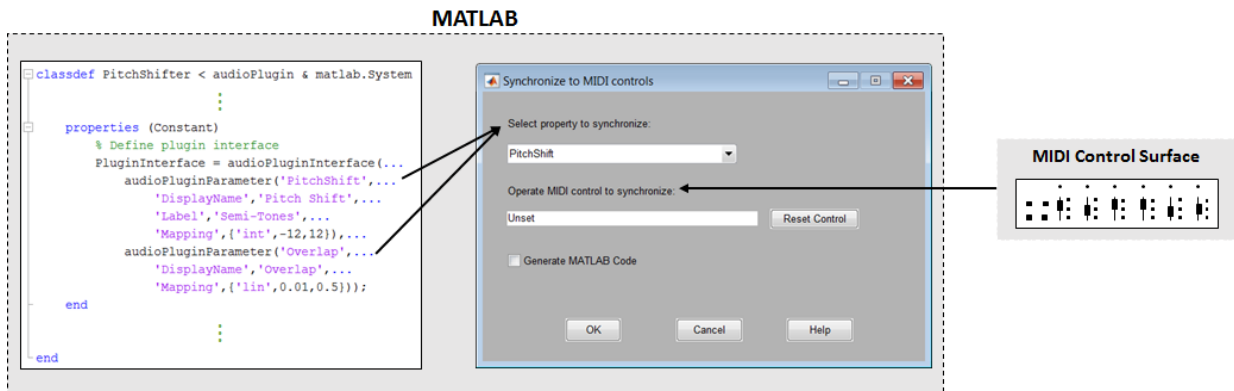
MATLAB before connecting your device, MATLAB might not recognize your device when you connect it. To correct the problem, restart MATLAB with the device already connected.

2. Establish MIDI Connections

Use `configureMIDI` to establish MIDI connections between your default MIDI device and an audio plugin. You can use `configureMIDI` programmatically, or you can open a user interface (UI) to guide you through the process. The `configureMIDI` UI reads from your audio plugin and populates a drop-down list of tunable plugin properties. You are then prompted to move individual controls on your MIDI control surface to associate the position of each control with the normalized value of each property you select. For example, create an object of `audiopluginexample.PitchShifter` and then call `configureMIDI` with the object as the argument:

```
ctrlPitch = audiopluginexample.PitchShifter;
configureMIDI(ctrlPitch)
```

The Synchronize to MIDI controls dialog box opens with the tunable properties of your plugin automatically populated. When you operate a MIDI control, its identification is entered into the **Operate MIDI control to synchronize** box. After you synchronize tunable properties with MIDI controls, click **OK** to complete the configuration. If your MIDI control surface is bidirectional, it automatically shifts the position of the synchronized controls to the initial property values specified by your plugin.



To open a MATLAB function with the programmatic equivalent of your actions in the UI, select the **Generate MATLAB Code** check box. Saving this function enables you to reuse your settings and quickly establish the configuration in future sessions.

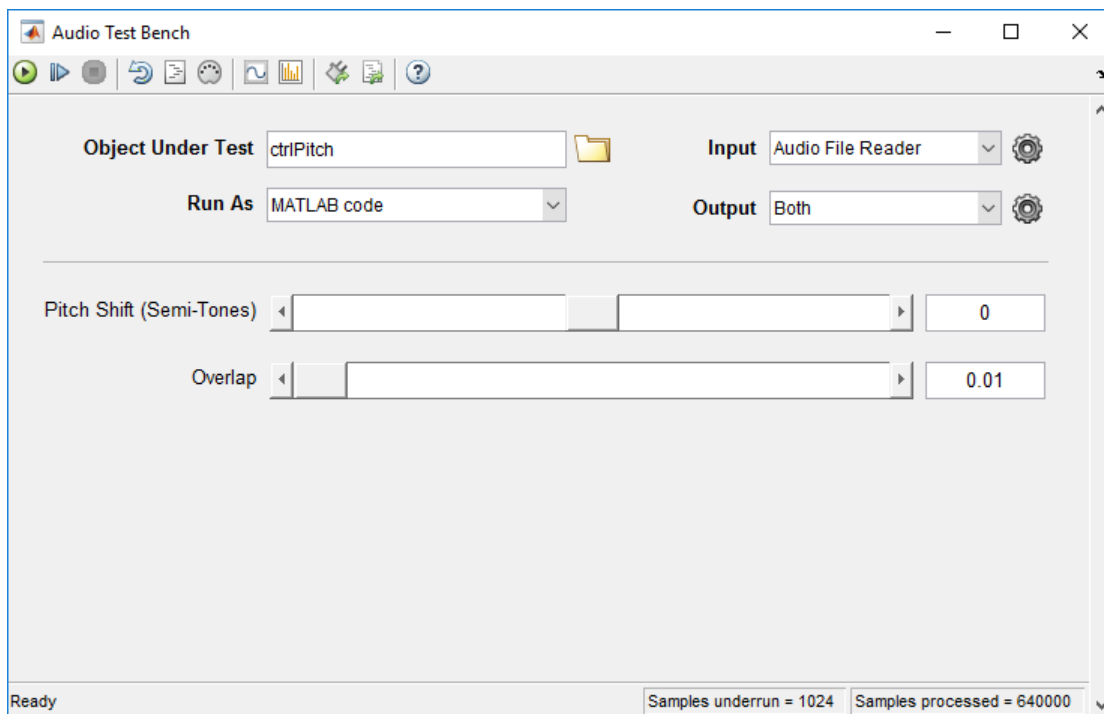
3. Tune Plugin Parameters Using MIDI


After you establish connections between plugin properties and MIDI controls, you can tune the properties in real time using your MIDI control surface.


Audio Toolbox provides an all-in-one app for running and testing your audio plugin. The test bench mimics how a DAW interacts with plugins.

Open the **Audio Test Bench** for your `ctrlPitch` object.

```
audioTestBench(ctrlPitch)
```



When you adjust the controls on your MIDI surface, the corresponding plugin parameter sliders move. Click  to run the plugin. Move the controls on your MIDI surface to hear the effect of tuning the plugin parameters.

To establish MIDI connections and modify existing ones, click the Synchronize to MIDI Controls  button to open a configureMIDI UI.

Alternatively, you can use the MIDI connections you established in a script or function. For example, run the following code and move your synchronized MIDI controls to hear the pitch-shifting effect:

```
fileReader = dsp.AudioFileReader(...
    'Filename', 'Counting-16-44p1-mono-15secs.wav');
deviceWriter = audioDeviceWriter;

% Audio stream loop
while ~isDone(fileReader)
    input = fileReader();
    output = ctrlPitch(input);
    deviceWriter(output);
    drawnow limitrate; % Process callback immediately
end

release(fileReader);
release(deviceWriter);
```

4. Get Current MIDI Connections

To query the MIDI connections established with your audio plugin, use the `getMIDIConnections` function. `getMIDIConnections` returns a structure with fields corresponding to the tunable properties of your plugin. The corresponding values are nested structures containing information about the mapping between your plugin property and the specified MIDI control.

```
connectionInfo = getMIDIConnections(ctrlPitch)

connectionInfo =

    struct with fields:

        PitchShift: [1x1 struct]
        Overlap: [1x1 struct]

connectionInfo.PitchShift

ans =

    struct with fields:
```

```
        Law: 'int'  
        Min: -12  
        Max: 12  
MIDIControl: 'control 1081 on 'BCF2000''
```

5. Disconnect MIDI Surface

As a best practice, release external devices such as MIDI control surfaces when you are done.

```
disconnectMIDI(ctrlPitch)
```

See Also

Apps

Audio Test Bench

Classes

audioPlugin | audioPluginSource

Functions

configureMIDI | disconnectMIDI | getMIDIConnections

More About

- “What Are DAWs, Audio Plugins, and MIDI Controllers?”
- “MIDI Control Surface Interface” on page 9-2
- “Audio Plugins in MATLAB”
- “Host External Audio Plugins”

External Websites

- <https://www.midi.org>

MIDI Control Surface Interface

MIDI Control Surface Interface

In this section...

“About MIDI” on page 9-2

“MIDI Control Surfaces” on page 9-2

“Use MIDI Control Surfaces with MATLAB and Simulink” on page 9-3


About MIDI

Musical Instrument Digital Interface (MIDI) was originally developed to interconnect electronic musical instruments. This interface is flexible and has uses in applications far beyond musical instruments. Its simple unidirectional messaging protocol supports many different kinds of messaging. One kind of MIDI message is the Control Change message, which is used to communicate changes in controls, such as knobs, sliders, and buttons.

MIDI Control Surfaces


A MIDI control surface is a device with controls that sends MIDI Control Change messages when you turn a knob, move a slider, or push a button on its surface. Each Control Change message indicates which control changed and what its new position is.

Because the MIDI messaging protocol is unidirectional, determining a particular controller position requires that the receiver listen for Control Change messages that the controller sends. The protocol does not support querying the MIDI controller for its position.



The simplest MIDI control surfaces are unidirectional: They send MIDI Control Change messages but do not receive them. More sophisticated control surfaces are bidirectional: They can both send and receive Control Change messages. These control surfaces have

knobs or sliders that can operate automatically. For example, a control surface can have motorized sliders or knobs. When it receives a Control Change message, the appropriate control moves to the position in the message.




Use MIDI Control Surfaces with MATLAB and Simulink

Audio Toolbox enables you to use MIDI control surfaces to control MATLAB programs and Simulink® models by providing the capability to listen to Control Change messages. The toolbox also provides a limited capability to send Control Change messages to support synchronizing MIDI controls. This tutorial covers general MIDI functions. For functions specific to audio plugins in MATLAB, see “MIDI Control for Audio Plugins” on page 8-2. The Audio Toolbox general interface to MIDI control surfaces includes five functions and one block:

- `midid` -- Interactively identify MIDI control.
- `midicontrols` -- Open group of MIDI controls for reading.
- `midiread` -- Return most recent value of MIDI controls.
- `midisync` -- Send values to MIDI controls for synchronization.
- `midicallback` -- Call function handle when MIDI controls change value.
- MIDI Controls (block) -- Output values from controls on MIDI control surface. The MIDI Controls block combines functionality of the general MIDI functions into one block for the Simulink environment.

This diagram shows a typical workflow involving general MIDI functions in MATLAB. For the Simulink environment, follow steps 1 and 2, and then use the MIDI Controls block for a user-interface guided workflow.



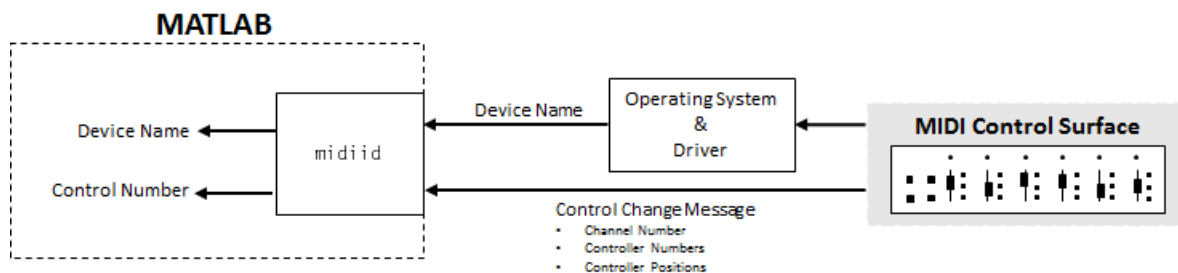
1. Connect MIDI Device and Then Start MATLAB

Before starting MATLAB, connect your MIDI control surface to your computer and turn it on. For connection instructions, see the instructions for your MIDI device. If you start MATLAB before connecting your device, MATLAB might not recognize your device when you connect it. To correct the problem, restart MATLAB with the device already connected.

2. Determine Device Name and Control Numbers

Use the `midid` function to determine the device name and control numbers of your MIDI control surface. After you call `midid`, it continues to listen until it receives a Control Change message. When it receives a Control Change message, it returns the control number associated with the MIDI controller number that you manipulated, and optionally returns the device name of your MIDI control surface. The manufacturer and host operating system determine the device name. See “Control Numbers” on page 9-10 for an explanation of how MATLAB calculates the control number.

To set a default device name, see “Set Default MIDI Device” on page 9-9.




View Example

Call `midid` with two outputs and then move a controller on your MIDI device. `midid` returns the control number specific to the controller you moved and the device name of the MIDI control surface.

```
[controlNumber,deviceName] = midid;
```

3. Create Listener for Control Change Messages

Use the `midicontrols` function to create an object that listens for Control Change messages and caches the most recent values corresponding to specified controllers. When you create a `midicontrols` object, you specify a MIDI control surface by its device name and specific controllers on the surface by their associated control numbers. Because the `midicontrols` object cannot query the MIDI control surface for initial values, consider setting initial values when creating the object.



View Example

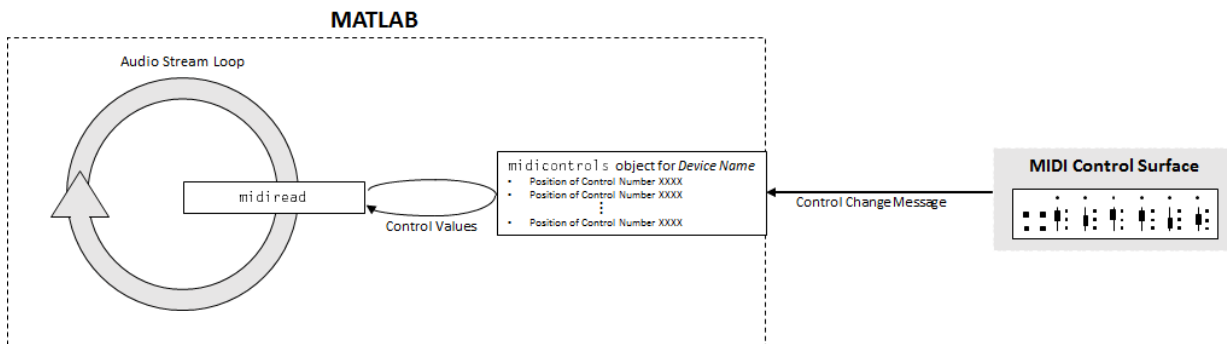
Identify two control numbers on your MIDI control surface. Choose initial control values for the controls you identified. Create a `midicontrols` object that listens to Control Change messages that arrive from the controllers you identified on the device you identified. When you create your `midicontrols` object, also specify initial control values. These initial control values work as default values until a Control Change message is received.

```
controlNum1 = midid;
[controlNum2,deviceName] = midid;
initialControlValues = [0.1,0.9];

midicontrolsObject = midicontrols([controlNum1,controlNum2], ...
    initialControlValues, ...
    'MIDIDevice',deviceName);
```

4. Get Current Control Values

Use the `midiread` function to query your `midicontrols` object for current control values. `midiread` returns a matrix with values corresponding to all controllers the `midicontrols` object is listening to. Generally, you want to place `midiread` in an audio stream loop for continuous updating.



View Example

Place `midiread` in an audio stream loop to return the current control value of a specified controller. Use the control value to apply gain to an audio signal.

```
[controlNumber, deviceName] = midiid;
initialControlValue = 1;
midicontrolsObject = midicontrols(controlNumber, initialControlValue, 'MIDIDevice', deviceName);

% Create a dsp.AudioFileReader System object™ with default settings. Create
% an audioDeviceWriter System object and specify the sample rate.
fileReader = dsp.AudioFileReader('RockDrums-44p1-stereo-11secs.mp3');
deviceWriter = audioDeviceWriter(...
    'SampleRate', fileReader.SampleRate);

% In an audio stream loop, read an audio signal frame from the file, apply
% gain specified by the control on your MIDI device, and then write the
% frame to your audio output device. By default, the control value returned
% by midiread is normalized.
while ~isDone(fileReader)
    audioData = step(fileReader);

    controlValue = midiread(midicontrolsObject);

    gain = controlValue*2;
    audioDataWithGain = audioData*gain;

    play(deviceWriter, audioDataWithGain);
end

% Close the input file and release your output device.
```



```
release(fileReader);
release(deviceWriter);
```

5. Synchronize Bidirectional MIDI Control Surfaces

You can use `midisync` to send Control Change messages to your MIDI control surface. If the MIDI control surface is bidirectional, it adjusts the specified controllers. One important use of `midisync` is to set the controller positions on your MIDI control surface to initial values.

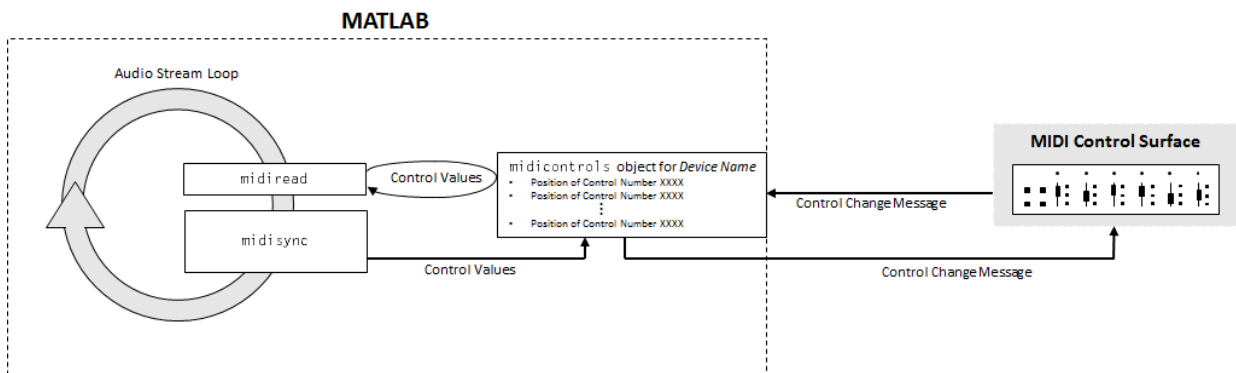
View Example

In this example, you initialize a controller on your MIDI control surface to a specified position. Calling `midisync(midicontrolsObject)` sends a Control Change message to your MIDI control surface, using the initial control values specified when you created the `midicontrols` object.

```
[controlNumber,deviceName] = midiid;
initialControlValue = 0.5;
midicontrolsObject = midicontrols(controlNumber,initialControlValue,'MIDIDevice',deviceName);

midisync(midicontrolsObject);
```

Another important use of `midisync` is to update your MIDI control surface if control values are adjusted in an audio stream loop. In this case, you call `midisync` with both your `midicontrols` object and the updated control values.



View Example

In this example, you check the normalized output volume in an audio stream loop. If the volume is above a given threshold, `midisync` is called and the MIDI controller that controls the applied gain is reduced.

```
[controlNumber, deviceName] = midiid;
initialControlValue = 0.5;
midicontrolsObject = midicontrols(controlNumber,initialControlValue);
fileReader = dsp.AudioFileReader('Ambiance-16-44p1-mono-12secs.wav');
deviceWriter = audioDeviceWriter(...
    'SampleRate',fileReader.SampleRate);

% Synchronize specified initial value with the MIDI control surface.
midisync(midicontrolsObject);

while ~isDone(fileReader)
    audioData = step(fileReader);
    controlValue = midiread(midicontrolsObject);
    gain = controlValue*2;
    audioDataWithGain = audioData*gain;

    % Check if max output is above a given threshold.
    if max(audioDataWithGain) > 0.7

        % Force new control value to be nonnegative.
        newControlValue = max(0,controlValue-0.5);

        % Send a Control Change message to the MIDI control surface.
        midisync(midicontrolsObject,newControlValue)
    end

    play(deviceWriter,audioDataWithGain);
end


release(fileReader);
release(deviceWriter);
```

`midisync` is also a powerful tool in systems that also involve user interfaces (UIs), so that when one control is changed, the other control tracks it. Typically, you implement such tracking by setting callback functions on both the `midicontrols` object (using `midicallback`) and the UI control. The callback for the `midicontrols` object sends new values to the UI control. The UI uses `midisync` to send new values to the `midicontrols` object and MIDI control surface. See `midisync` for examples.

Alternative to Stream Processing

You can use `midicallback` as an alternative to placing `midiread` in an audio stream loop. If a `midicontrols` object receives a Control Change message, `midicallback` automatically calls a specified function handle. The callback function typically calls

`midiread` to determine the new value of the MIDI controls. You can use this callback when you want a MIDI controller to trigger an action, such as updating a UI. Using this approach prevents having a MATLAB program continuously running in the command window.



Set Default MIDI Device

You can set the default MIDI device in the MATLAB environment by using the `setpref` function. Use `midiid` to determine the name of the device, and then use `setpref` to set the preference.

```
[~,deviceName] = midiid
```

```
Move the control you wish to identify; type ^C to abort.
Waiting for control message... done
```

```
deviceName =
```

```
BCF2000
```

```
setpref('midi', 'DefaultDevice', deviceName)
```

This preference persists across MATLAB sessions, so you only have to set it once, unless you want to change devices.

If you do not set this preference, MATLAB and the host operating system choose a device for you. However, such autoselection can cause unpredictable results because many computers have "virtual" (software) MIDI devices installed that you might not be aware of. For predictable behavior, set the preference.

Control Numbers

MATLAB defines control numbers as $(MIDI\ channel\ number) \times 1000 + (MIDI\ controller\ number)$.

- MIDI channel number is the transmission channel that your device uses to send messages. This value is in the range 1-16.
- MIDI controller number is a number assigned to an individual control on your MIDI device. This value is in the range 1-127.

Your MIDI device determines the values of *MIDI channel number* and *MIDI controller number*.

See Also

Blocks

MIDI Controls

Functions

`midicallback` | `midicontrols` | `midiid` | `midiread` | `midisync`

More About

- “What Are DAWs, Audio Plugins, and MIDI Controllers?”
- “Real-Time Audio in MATLAB”
- “MIDI Device Interface” on page 6-2
- “MIDI Control for Audio Plugins” on page 8-2

External Websites

- <https://www.midi.org>

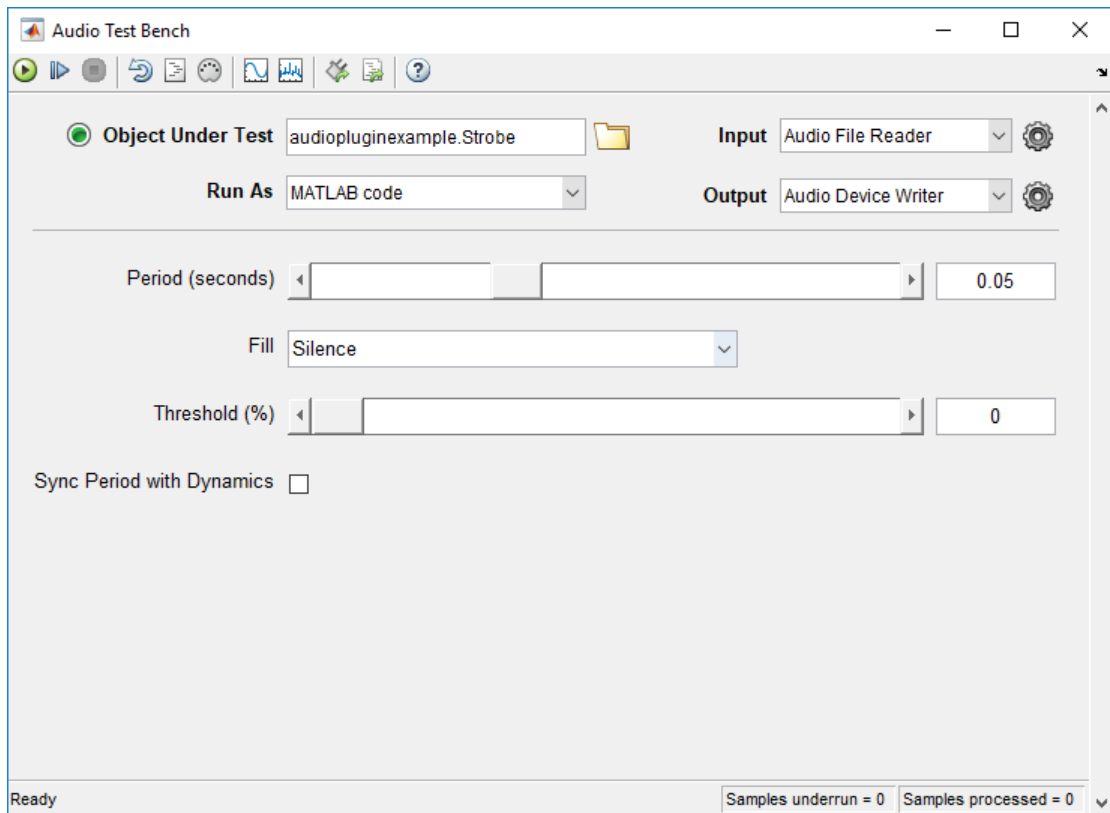
Use the Audio Test Bench

Audio Test Bench Walkthrough

In this tutorial, explore key functionality of the **Audio Test Bench**. The **Audio Test Bench** app enables you to debug, visualize, and configure audio plugins.

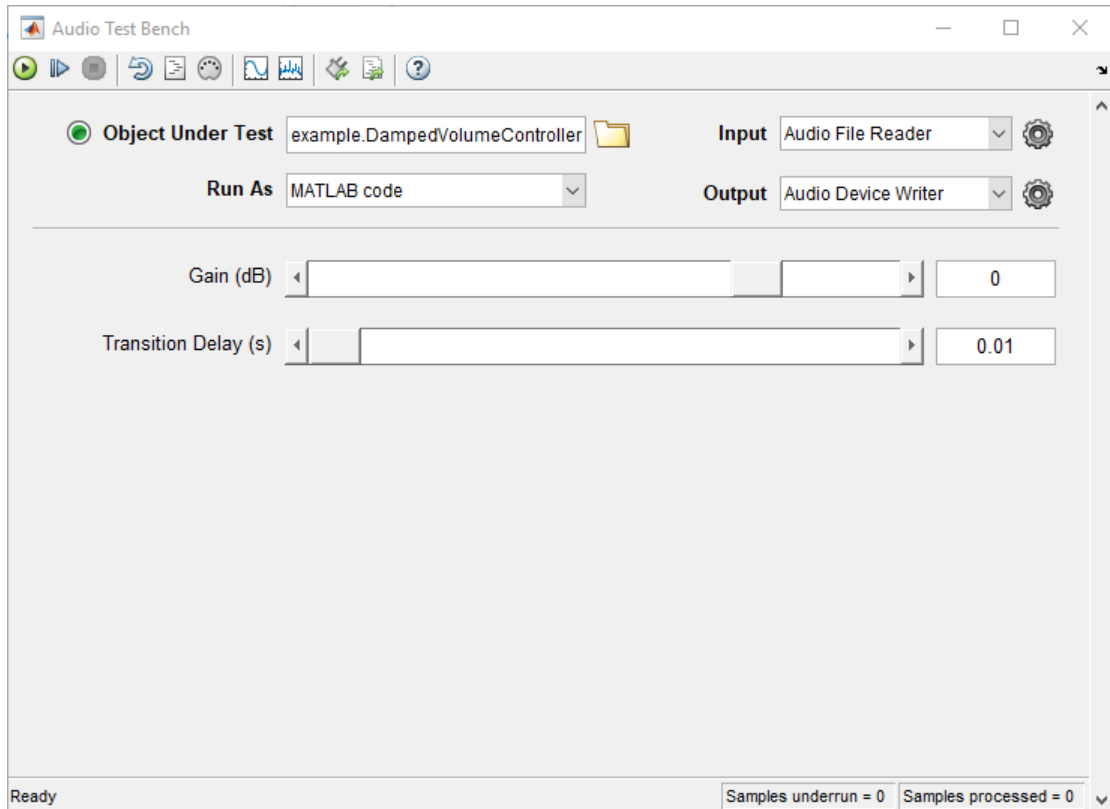
Choose Object Under Test

- 1 To open the **Audio Test Bench**, at the MATLAB command prompt, enter:
`audioTestBench`
- 2 In the **Object Under Test** box, enter `audiopluginexample.Strobe` and press **Enter**. The **Audio Test Bench** automatically displays the tunable parameters of the `audiopluginexample.Strobe` audio plugin.





The mapping between the tunable parameters of your object and the UI widgets on the **Audio Test Bench** is determined by `audioPluginInterface` and `audioPluginParameter` in the class definition of your object.


- 3 In the **Object Under Test** box, enter `audiopluginexample.DampedVolumeController` and press **Enter**. The **Audio Test Bench** automatically displays the tunable parameters of the `audiopluginexample.DampedVolumeController` audio plugin.



Run Audio Test Bench

To run the **Audio Test Bench** for your plugin with default settings, click . Move the sliders to modify the **Gain (dB)** and **Transition Delay (s)** parameters while streaming.


To stop the audio stream loop, click . The MATLAB command line and objects used by the test bench are now released.

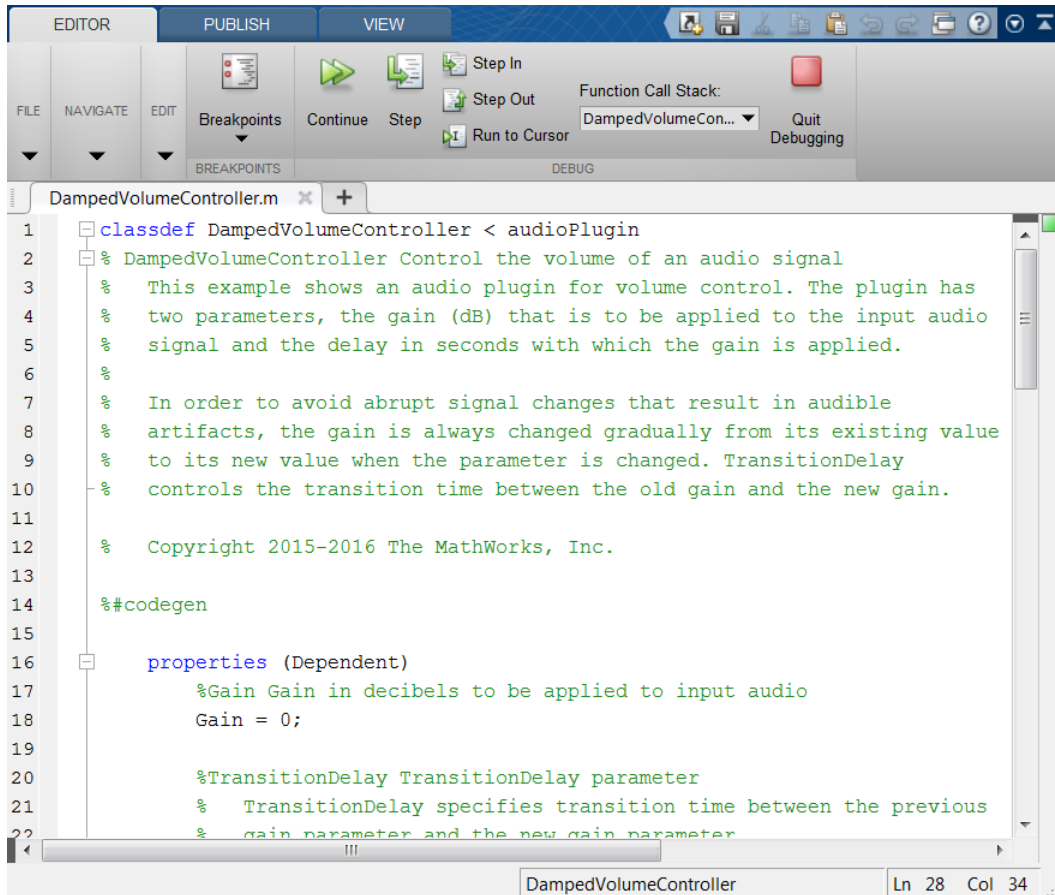
To reset internal states of your audio plugin and return the sliders to their initial positions, click .

Click  to run the **Audio Test Bench** again.

Debug Source Code of Audio Plugin

To pause the **Audio Test Bench**, click .


To open the source file of your audio plugin, click .

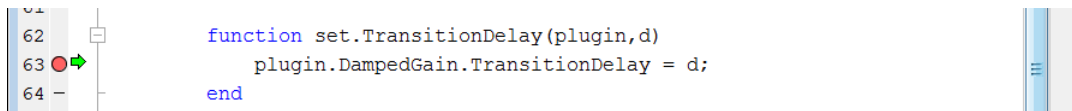


```

1  classdef DampedVolumeController < audioPlugin
2  % DampedVolumeController Control the volume of an audio signal
3  % This example shows an audio plugin for volume control. The plugin has
4  % two parameters, the gain (dB) that is to be applied to the input audio
5  % signal and the delay in seconds with which the gain is applied.
6  %
7  % In order to avoid abrupt signal changes that result in audible
8  % artifacts, the gain is always changed gradually from its existing value
9  % to its new value when the parameter is changed. TransitionDelay
10 % controls the transition time between the old gain and the new gain.
11
12 % Copyright 2015-2016 The MathWorks, Inc.
13
14 %#codegen
15
16 properties (Dependent)
17     %Gain Gain in decibels to be applied to input audio
18     Gain = 0;
19
20     %TransitionDelay TransitionDelay parameter
21     % TransitionDelay specifies transition time between the previous
22     % gain parameter and the new gain parameter

```

You can inspect the source code of your audio plugin, set breakpoints on it, and modify the code. Set a breakpoint at line 63 and then click  on the **Audio Test Bench**.





```

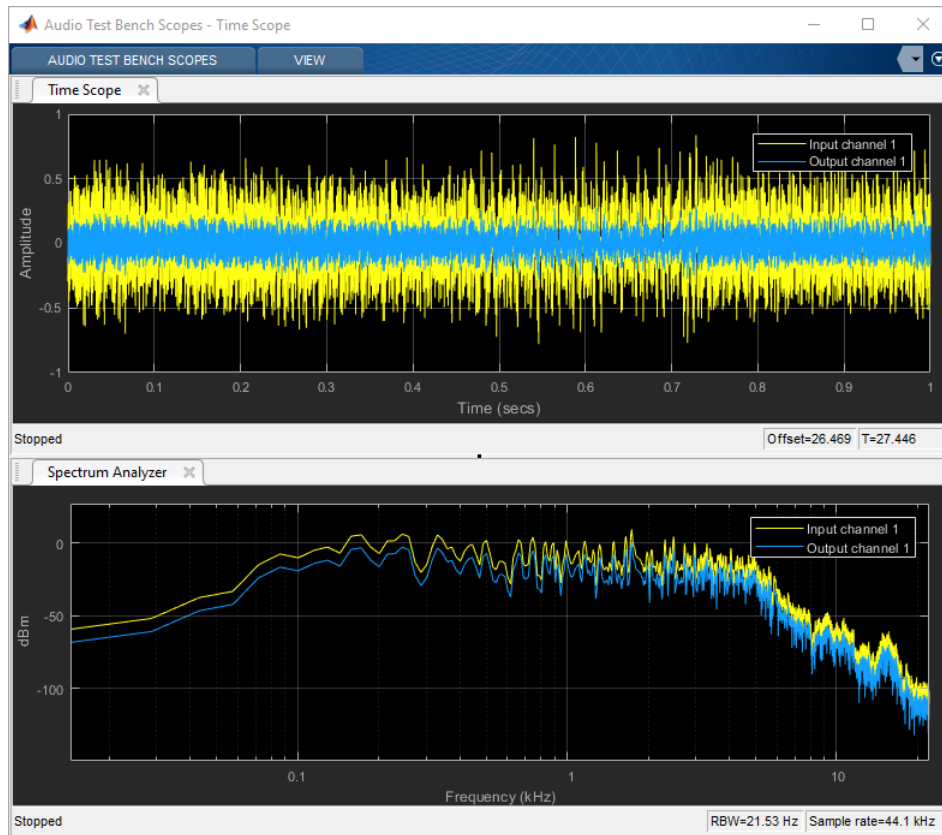
62     function set.TransitionDelay(plugin,d)
63         plugin.DampedGain.TransitionDelay = d;
64     end


```

The **Audio Test Bench** runs your plugin until it reaches the breakpoint. To reach the breakpoint, move the **Transition Delay (s)** slider. To quit debugging, remove the breakpoint. In the MATLAB editor, click **Quit Debugging**.

Open Scopes

To open a time scope to visualize the time-domain input and output for your audio plugin, click . To open a spectrum analyzer to visualize the frequency-domain input and output, click .




To release objects and stop the audio stream loop, click .

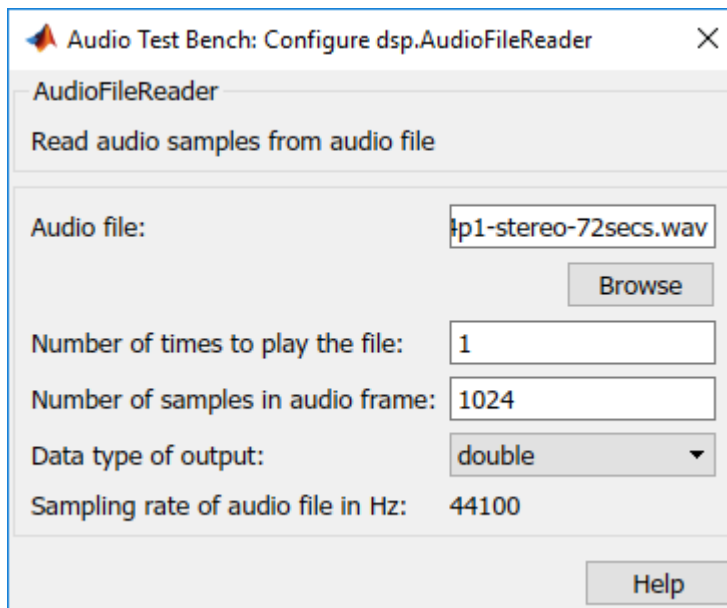
Configure Input to Audio Test Bench

The **Input** list contains these options:

- Audio File Reader -- `dsp.AudioFileReader`
- Audio Device Reader -- `audioDeviceReader`
- Audio Oscillator -- `audioOscillator`
- Wavetable Synthesizer -- `wavetableSynthesizer`
- Chirp Signal -- `dsp.Chirp`
- Colored Noise -- `dsp.ColoredNoise`

1 Select Audio File Reader.


2 Click  to open a dialog for Audio File Reader configuration.



You can enter any file name included on the MATLAB path. To specify a file that is not on the MATLAB path, specify the full file path.

3 In the **Audio file** box, enter: `RockDrums-44p1-stereo-11secs.mp3`


Press **Enter**, and then exit the Audio File Reader configuration dialog. To run the audio test bench with your new input, click .

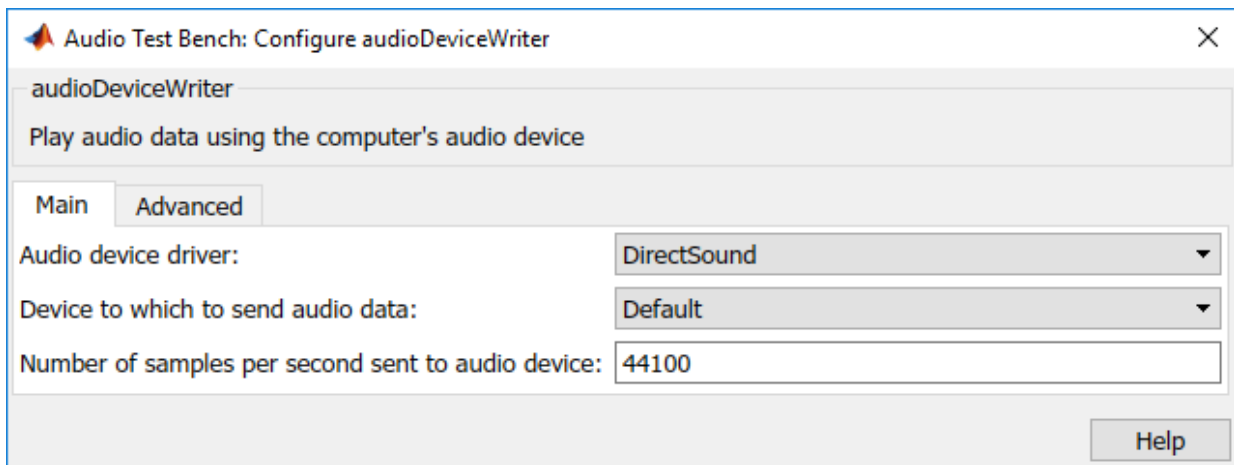
To release your output object and stop the audio stream loop, click .

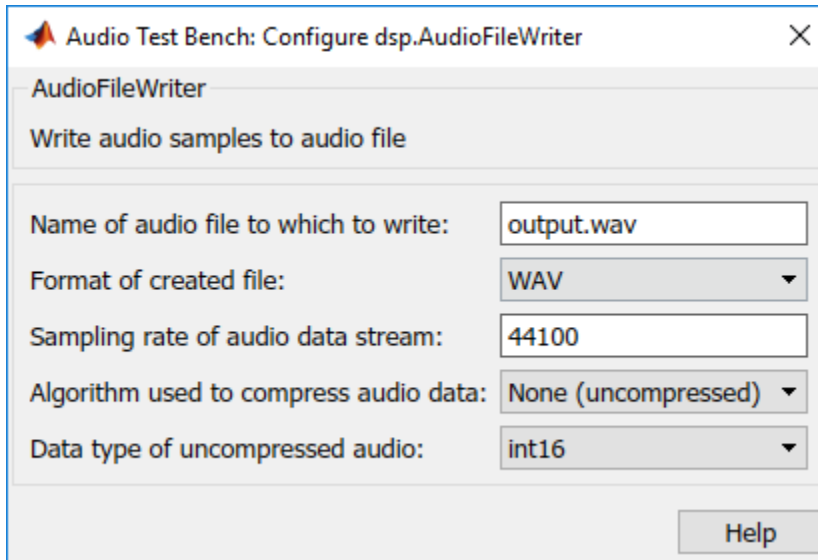
Configure Output from Audio Test Bench

The **Output** list contains these options:



- Audio Device Writer -- `audioDeviceWriter`
- Audio File Writer -- `dsp.AudioFileWriter`
- Both -- `audioDeviceWriter` and `dsp.AudioFileWriter`
- None -- The audio signal is not routed to a file or device. Use this option if you are only interested in using the visualization and tuning capabilities of the test bench.


- 1 Choose to output to device and file by selecting **Both** from the **Output** menu.
- 2 To open a dialog for Audio Device Writer and Audio File Writer configuration, click .






Call Custom Visualization of Audio Plugin

If your audio plugin has a custom visualization method, the  button appears on the **Audio Test Bench**. In the **Object Under Test** box, enter `audiopluginexample.VarSlopeBandpassFilter` and press **Enter**. To open the custom visualization of `audiopluginexample.VarSlopeBandpassFilter`, click . The custom visualization plots the frequency response of the filter. Tune the plugin parameters and observe the plot update in real time.




Custom visualizations are MATLAB-only features. Custom visualizations are not available for generated plugins.

Synchronize Plugin Property with MIDI Control


If you have a MIDI device connected to your computer, you can synchronize plugin properties with MIDI controls. To open a MIDI configuration UI, click . Synchronize the `LowCutoff` and `HighCutoff` properties with MIDI controls you choose. Click **OK**.

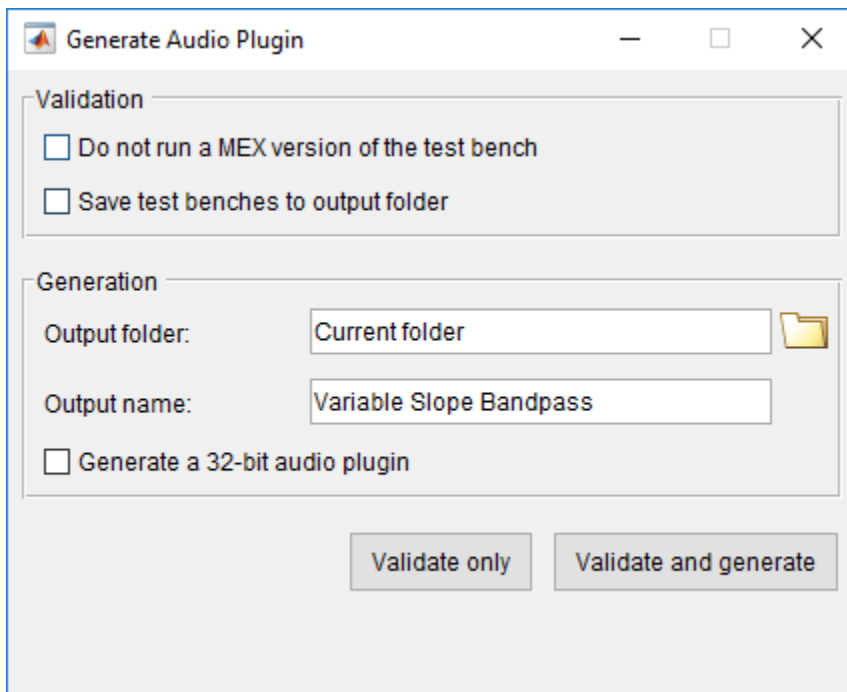
See `configureMIDI` for more information.

Play the Audio and Save the Output File

To run your audio plugin, click . Adjust your plugin properties in real time using your synchronized MIDI controls and sliders. Your processed audio file is saved to the current folder.

Validate and Generate Audio Plugin

To open the validation and generation dialog box, click .



You can validate only, or validate and generate your MATLAB audio plugin code in VST 2 plugin format. The **Generate a 32-bit audio plugin** check box is available only on win64 machines. See `validateAudioPlugin` and `generateAudioPlugin` for more information.

Generate MATLAB Script

To generate a MATLAB script that implements a test bench for your audio plugin, click




```

Untitled* X +
1 % Test bench script for 'audiopluginexample.VarSlopeBandpassFilter'.
2 % Generated by Audio Test Bench on 10-Jul-2018 11:48:59 -0400.
3
4 % Create test bench input and output(s)
5 fileReader = dsp.AudioFileReader('Filename','C:\Program Files\MATLAB\R2018a
6 deviceWriter = audioDeviceWriter('SampleRate',fileReader.SampleRate);
7 fileWriter = dsp.AudioFileWriter;
8
9 % Set up the system under test
10 sut = audiopluginexample.VarSlopeBandpassFilter;
11 setSampleRate(sut,fileReader.SampleRate);
12 tuneParameters(sut);
13
14 % Stream processing loop
15 nUnderruns = 0;
16 while ~isDone(fileReader)
17     % Read from input, process, and write to outputs
18     in = fileReader();
19     out = sut(in);
20     nUnderruns = nUnderruns + deviceWriter(out);
21     fileWriter(out);
22
23     % Visualize input and output data in scopes
24     timeScope([in(:,1),out(:,1)]);
25     specScope([in(:,1),out(:,1)]);
26
27     % To tune parameters during streaming, uncomment the code below and
28     % customize tuneParameters function
29     % tuneParameters(sut);
30 end
31
32 % Clean up
33 release(sut);
34 release(fileReader);
35 release(deviceWriter);
36 release(fileWriter);
37
38
39 function tuneParameters(sut)
40 %tuneParameters Tune the parameter values on the system under test
41 sut.LowCutoff = 72.5972;
42 sut.HighCutoff = 730.639;
43 sut.LowSlope = '42';

```

You can modify the code for complete control over the test bench environment, including the ability to create processing chains by placing plugins in cascade.

See Also

Audio Test Bench | `audioPlugin` | `generateAudioPlugin` | `validateAudioPlugin`

More About

- “Audio Plugins in MATLAB”
- “Audio Plugin Example Gallery” on page 11-2
- “Export a MATLAB Plugin to a DAW”

Audio Plugin Example Gallery

Audio Plugin Example Gallery

Use these Audio Toolbox plugin examples as building blocks in larger systems, as models for design patterns, or as benchmarks for comparison. Search the plugin descriptions to find an example that meets your needs.

Audio Effects

Filters

Gain Control

Spatial Audio

Communicate Between MATLAB and DAW

Music Information Retrieval

Speech Processing

Audio Plugin Examples

For a list of available audio plugins, see the online documentation.

See Also

Audio Test Bench | `audioPlugin` | `audioPluginInterface` | `audioPluginParameter` | `audioPluginSource`

More About

- “Audio Test Bench Walkthrough” on page 10-2
- “Audio Plugins in MATLAB”

Equalization

Equalization

Equalization (EQ) is the process of weighting the frequency spectrum of an audio signal.

You can use equalization to:

- Enhance audio recordings
- Analyze spectral content

Types of equalization include:

- Lowpass and highpass filters -- Attenuate high frequency and low frequency content, respectively.
- Low-shelf and high-shelf equalizers -- Boost or cut frequencies equally above or below a desired cutoff point.
- Parametric equalizers -- Selectively boost or cut frequency bands. Also known as peaking filters.
- Graphic equalizers -- Selectively boost or cut octave or fractional octave frequency bands. The bands have standards-based center frequencies. Graphic equalizers are a special case of parametric equalizers.

This tutorial describes how Audio Toolbox implements the design functions: `designParamEQ`, `designShelvingEQ`, and `designVarSlopeFilter`. The `multibandParametricEQ` System object combines the filter design functions into a multiband parametric equalizer. The `graphicEQ` System object combines the filter design functions and the `octaveFilter` System object for standards-based graphic equalization. For a tutorial focused on using the design functions in MATLAB, see “Parametric Equalizer Design”.

Equalization Design Using Audio Toolbox


EQ Filter Design

Audio Toolbox design functions use the bilinear transform method of digital filter design to determine your equalizer coefficients. In the bilinear transform method, you:

- 1 Choose an analog prototype.
- 2 Specify filter design parameters.

3 Perform the bilinear transformation.

Analog Low-Shelf Prototype



Audio Toolbox uses the high-order parametric equalizer design presented in [1]. In this design method, the analog prototype is taken to be a low-shelf Butterworth filter:

$$H_a(s) = \left[\frac{g\beta + s}{\beta + s} \right]^r \prod_{i=1}^L \left[\frac{g^2\beta^2 + 2gs_i\beta s + s^2}{\beta^2 + 2s_i\beta s + s^2} \right]$$

- L = Number of analog SOS sections

- N = Analog filter order

- $r = \begin{cases} 0 & N \text{ even} \\ 1 & N \text{ odd} \end{cases}$

- $g = G^{1/N}$

- $\beta = \Omega_B \times \left(\sqrt{\frac{G^2 - G_B^2}{G_B^2 - 1}} \right)^{-1/N} = \tan\left(\pi \frac{\Delta\omega}{2}\right) \times \left(\sqrt{\frac{G^2 - G_B^2}{G_B^2 - 1}} \right)^{-1/N}$, where $\Delta\omega$ is the

desired digital bandwidth

$$\bullet \quad s_i = \sin\left(\frac{(2i-1)\pi}{2N}\right), \quad i = 1, 2, \dots, L$$

For parametric equalizers, the analog prototype is reduced by setting the bandwidth gain to the square root of the peak gain ($G_B = \text{sqrt}(G)$).

After the design parameters are specified, the analog prototype is transformed directly to the desired digital equalizer by a bandpass bilinear transformation:

$$s = \frac{1 - 2\cos(\omega_0)z^{-1} + z^{-2}}{1 - z^{-2}}$$

ω_0 is the desired digital center frequency.

This transformation doubles the filter order. Every first-order analog section becomes a second-order digital section. Every second-order analog section becomes a fourth-order digital section. Audio Toolbox always calculates fourth-order digital sections, which means that returning second-order sections requires the computation of roots, and is less efficient.

Digital Coefficients

The digital transfer function is implemented as a cascade of second-order and fourth-order sections.

$$H(z) = \left[\frac{b_{00} + b_{01}z^{-1} + b_{02}z^{-2}}{1 + a_{01}z^{-1} + a_{02}z^{-2}} \right]^r \prod_{i=1}^L \left[\frac{b_{i0} + b_{i1}z^{-1} + b_{i2}z^{-2} + b_{i3}z^{-3} + b_{i4}z^{-4}}{1 + a_{i1}z^{-1} + a_{i2}z^{-2} + a_{i3}z^{-3} + a_{i4}z^{-4}} \right]$$

The coefficients are given by performing the bandpass bilinear transformation on the analog prototype design.

Second-Order Section Coefficients	Fourth-Order Section Coefficients
$D_0 = \beta + 1$	$D_i = \beta^2 + 2s_i\beta + 1$
$b_{00} = (1 + g\beta)/D_0$	$b_{i0} = (g^2\beta^2 + 2gs_i\beta + 1)/D_i$
$b_{01} = -2\cos(\omega_0)/D_0$	$b_{i1} = -4c_0(1 + gs_i\beta)/D_i$
$b_{02} = (1 - g\beta)/D_0$	$b_{i2} = 2(1 + 2\cos^2(\omega_0) - g^2\beta^2)/D_i$
$a_{01} = -2\cos(\omega_0)/D_0$	$b_{i3} = -4c_0(1 - gs_i\beta)/D_i$
$a_{02} = (1 - \beta)/D_0$	$b_{i4} = (g^2\beta^2 - 2gs_i\beta + 1)/D_i$
	$a_{i1} = -4c_0(1 + s_i\beta)/D_i$
	$a_{i2} = 2(1 + 2\cos^2(\omega_0) - \beta^2)/D_i$
	$a_{i3} = -4\cos(\omega_0)(1 - s_i\beta)/D_i$
	$a_{i4} = (\beta^2 - 2s_i\beta + 1)/D_i$

Biquadratic Case

In the biquadratic case, when $N = 1$, the coefficients reduce to:

$$D_0 = \frac{\Omega_B}{\sqrt{G}} + 1$$

$$b_{00} = (1 + \Omega_B\sqrt{G})/D_0, \quad b_{01} = -2\cos(\omega_0)/D_0, \quad b_{02} = (1 - \Omega_B\sqrt{G})/D_0$$

$$a_{01} = -2\cos(\omega_0)/D_0, \quad a_{02} = \left(1 - \frac{\Omega_B}{\sqrt{G}}\right)/D_0$$

Denormalizing the a_{00} coefficient, and making substitutions of $A = \text{sqrt}(G)$, $\Omega_B \cong \alpha$ yields the familiar peaking EQ coefficients described in [2].

Orfanidis notes the approximate equivalence of Ω_B and α in [1].

By using trigonometric identities,

$$\Omega_B = \tan\left(\frac{\Delta\omega}{2}\right) = \sin(\omega_0)\sinh\left(\frac{\ln 2}{2}B\right),$$

where B plays the role of an equivalent octave bandwidth.

Bristow-Johnson obtained an approximate solution for B in [4]:

$$B = \frac{\omega_0}{\sin\omega_0} \times BW$$

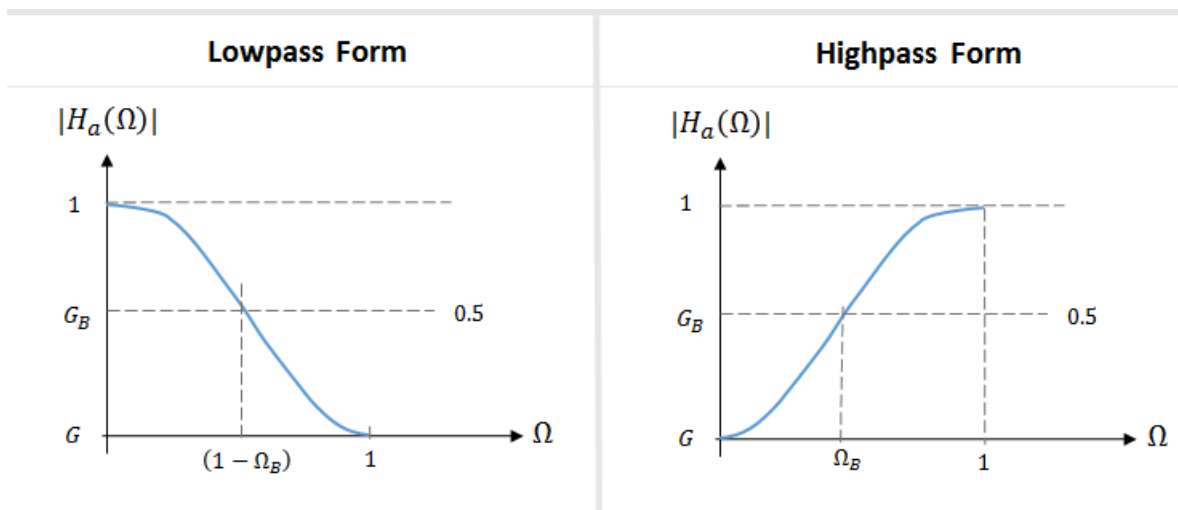
Substituting the approximation for B into the Ω_B equation yields the definition of α in [2]:

$$\alpha = \sin(\omega_0)\sinh\left(\frac{\ln 2}{2} \times \frac{\omega_0}{\sin\omega_0} \times BW\right)$$

Lowpass and Highpass Filter Design

Analog Low-Shelf Prototype

To design lowpass and highpass filters, Audio Toolbox uses a special case of the filter design for parametric equalizers. In this design, the peak gain, G , is set to 0, and G_B^2 is set to 0.5 (-3 dB cutoff). The cutoff frequency of the lowpass filter corresponds to $1 - \Omega_B$. The cutoff frequency of the highpass filter corresponds to Ω_B .



Digital Coefficients

The table summarizes the results of the bandpass bilinear transformation. The digital center frequency, ω_0 , is set to π for lowpass filters and 0 for highpass filters.

Second Order Section Coefficients	Fourth Order Section Coefficients
$D_0 = 1 + \tan\left(\pi\frac{\Delta\omega}{2}\right)$	$D_i = \tan^2\left(\pi\frac{\Delta\omega}{2}\right) + 2s_i\tan\left(\pi\frac{\Delta\omega}{2}\right) + 1$
$b_{00} = 1/D_0$	$b_{i0} = 1/D_i$
$b_{01} = -2\cos(\omega_0)/D_0$	$b_{i1} = -4\cos(\omega_0)/D_i$
$b_{02} = 1/D_0$	$b_{i2} = 2(1 + 2\cos^2(\omega_0))/D_i$
$a_{01} = -2\cos(\omega_0)/D_0$	$b_{i3} = -4\cos(\omega_0)/D_i$
$a_{02} = \left(1 - \tan\left(\pi\frac{\Delta\omega}{2}\right)\right)/D_0$	$b_{i4} = 1/D_i$
	$a_{i1} = -4\cos(\omega_0)\left(1 + s_i\tan\left(\pi\frac{\Delta\omega}{2}\right)\right)/D_i$
	$a_{i2} = 2\left(1 + 2\cos^2(\omega_0) - \tan^2\left(\pi\frac{\Delta\omega}{2}\right)\right)/D_i$
	$a_{i3} = -4\cos(\omega_0)\left(1 - s_i\tan\left(\pi\frac{\Delta\omega}{2}\right)\right)/D_i$
	$a_{i4} = \left(\tan^2\left(\pi\frac{\Delta\omega}{2}\right) - 2s_i\tan\left(\pi\frac{\Delta\omega}{2}\right) + 1\right)/D_i$

Shelving Filter Design

Analog Prototype

Audio Toolbox implements the shelving filter design presented in [2]. In this design, the high-shelf and low-shelf analog prototypes are presented separately:

$$H_L(s) = A \left(\frac{As^2 + (\sqrt{A}/Q)s + 1}{s^2 + (\sqrt{A}/Q)s + A} \right) \quad H_H(s) = A \left(\frac{s^2 + (\sqrt{A}/Q)s + A}{As^2 + (\sqrt{A}/Q)s + 1} \right)$$

For compactness, the analog filters are presented with variables A and Q . You can convert A and Q to available Audio Toolbox design parameters:

$$A = 10^{G/40}$$

$$\frac{1}{Q} = \sqrt{(A + 1/A)(1/slope - 1) + 2}$$

After you specify the design parameters, the analog prototype is transformed to the desired digital shelving filter by a bilinear transformation with prewarping:

$$s = \left(\frac{z-1}{z+1} \right) \times \left(\frac{1}{\tan\left(\frac{\omega_0}{2}\right)} \right)$$

Digital Coefficients

The table summarizes the results of the bilinear transformation with prewarping.

Low-Shelf	$b_0 = A((A + 1) - (A - 1)\cos(\omega_0) + 2\alpha\sqrt{A})$ $b_1 = 2A((A - 1) - (A + 1)\cos(\omega_0))$ $b_2 = A((A + 1) - (A - 1)\cos(\omega_0) - 2\alpha\sqrt{A})$ $a_0 = (A + 1) + (A - 1)\cos(\omega_0) + 2\alpha\sqrt{A}$ $a_1 = -2((A - 1) + (A + 1)\cos(\omega_0))$ $a_2 = (A + 1) + (A - 1)\cos(\omega_0) - 2\alpha\sqrt{A}$
High-Shelf	$b_0 = A((A + 1) + (A - 1)\cos(\omega_0) + 2\alpha\sqrt{A})$ $b_1 = -2A((A - 1) + (A + 1)\cos(\omega_0))$ $b_2 = A((A + 1) + (A - 1)\cos(\omega_0) - 2\alpha\sqrt{A})$ $a_0 = (A + 1) - (A - 1)\cos(\omega_0) + 2\alpha\sqrt{A}$ $a_1 = 2((A - 1) + (A + 1)\cos(\omega_0))$ $a_2 = (A + 1) - (A - 1)\cos(\omega_0) - 2\alpha\sqrt{A}$
Intermediate Variables	$\alpha = \frac{\sin(\omega_0)}{2} \sqrt{\left(A + \frac{1}{A}\right) \left(\frac{1}{\text{slope}} - 1\right) + 2A}$ $\omega_0 = 2\pi \frac{\text{Cutoff Frequency}}{F_s}$

References

- [1] Orfanidis, Sophocles J. "High-Order Digital Parametric Equalizer Design." *Journal of the Audio Engineering Society*. Vol. 53, November 2005, pp. 1026-1046.
- [2] Bristow-Johnson, Robert. "Cookbook Formulae for Audio EQ Biquad Filter Coefficients." Accessed March 02, 2016. <http://www.musicdsp.org/files/Audio-EQ-Cookbook.txt>.

[3] Orfanidis, Sophocles J. *Introduction to Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 2010.

[4] Bristow-Johnson, Robert. "The Equivalence of Various Methods of Computing Biquad Coefficients for Audio Parametric Equalizers." Presented at the 97th Convention of the AES, San Francisco, November 1994, AES Preprint 3906.

See Also

[designParamEQ](#) | [designShelvingEQ](#) | [designVarSlopeFilter](#) | [graphicEQ](#) | [multibandParametricEQ](#)

More About

- "Parametric Equalizer Design"
- "Graphic Equalization"
- "Octave-Band and Fractional Octave-Band Filters"
- "Audio Weighting Filters"

Deployment

Desktop Real-Time Audio Acceleration with MATLAB Coder

This example shows how to accelerate a real-time audio application using C code generation with MATLAB® Coder™. You must have the MATLAB Coder™ software installed to run this example.

Introduction

Replacing parts of your MATLAB code with an automatically generated MATLAB executable (MEX-function) can speed up simulation. Using MATLAB Coder, you can generate readable and portable C code and compile it into a MEX-function that replaces the equivalent section of your MATLAB algorithm.

This example showcases code generation using an audio notch filtering application.

Notch Filtering

A notch filter is used to eliminate a specific frequency from a signal. Typical filter design parameters for notch filters are the notch center frequency and the 3 dB bandwidth. The center frequency is the frequency at which the filter has a linear gain of zero. The 3 dB bandwidth measures the frequency width of the notch of the filter computed at the half-power or 3 dB attenuation point.

The helper function used in this example is `helperAudioToneRemoval`. The function reads an audio signal corrupted by a 250 Hz sinusoidal tone from a file. `helperAudioToneRemoval` uses a notch filter to remove the interfering tone and writes the filtered signal to a file.

You can visualize the corrupted audio signal using a spectrum analyzer.

```
reader = dsp.AudioFileReader('guitar_plus_tone.ogg');


scope = dsp.SpectrumAnalyzer('SampleRate',reader.SampleRate, ...
    'RBWSource','Property','RBW',5, ...
    'PlotAsTwoSidedSpectrum',false, ...
    'SpectralAverages',10, ...
    'FrequencySpan','Start and stop frequencies', ...
    'StartFrequency',20, ...
    'StopFrequency',1000, ...
    'Title','Audio signal corrupted by 250 Hz tone');
```



```

while ~isDone(reader)
    audio = reader();
    scope(audio(:,1));
end

```



C Code Generation Speedup

Measure the time it takes to read the audio file, filter out the interfering tone, and write the filtered output using MATLAB code.

```

tic
helperAudioToneRemoval
t1 = toc;

fprintf('MATLAB Simulation Time: %d\n',t1)

```

```
MATLAB Simulation Time: 5.336342e+00
```

Next, generate a MEX-function from `helperAudioToneRemoval` using the MATLAB Coder function, `codegen`.

```
codegen helperAudioToneRemoval
```

Measure the time it takes to execute the MEX-function and calculate the speedup gain with a compiled function.

```
tic
helperAudioToneRemoval_mex
t2 = toc;

fprintf('Code Generation Simulation Time: %d\n',t2)
Code Generation Simulation Time: 4.645470e+00
fprintf('Speedup factor: %6.2f\n',t1/t2)
Speedup factor: 1.15
```

See Also

Related Examples

- “Generate Standalone Executable for Parametric Audio Equalizer”
- “Deploy Audio Applications with MATLAB Compiler”

Audio I/O User Guide

Run Audio I/O Features Outside MATLAB and Simulink

You can deploy these audio input and output features outside the MATLAB and Simulink environments:

System Objects

- `audioPlayerRecorder`
- `audioDeviceReader`
- `audioDeviceWriter`
- `dsp.AudioFileReader`
- `dsp.AudioFileWriter`

Blocks

- Audio Device Reader
- Audio Device Writer
- From Multimedia File
- To Multimedia File

The generated code for the audio I/O features relies on prebuilt dynamic library files included with MATLAB. You must account for these extra files when you run audio I/O features outside the MATLAB and Simulink environments. To run a standalone executable generated from a model or code containing the audio I/O features, set your system environment using commands specific to your platform.

Platform	Command
Mac	<pre>setenv DYLD_LIBRARY_PATH "\$ {DYLD_LIBRARY_PATH}:\$MATLABROOT/b in/maci64" (csh/tcsh) export DYLD_LIBRARY_PATH= \$LD_LIBRARY_PATH:\$MATLABROOT/bin/ maci64 (Bash)</pre>

Platform	Command
Linux	<pre>setenv LD_LIBRARY_PATH \$ {LD_LIBRARY_PATH}:\$MATLABROOT/bin /glnxa64 (csh/tcsh) export LD_LIBRARY_PATH= \$LD_LIBRARY_PATH:\$MATLABROOT/bin/ glnxa64 (Bash)</pre>
Windows	<pre>set PATH=%PATH%;%MATLABROOT%\bin \win64</pre>

The path in these commands is valid only on systems that have MATLAB installed. If you run the standalone app on a machine with only MCR, and no MATLAB installed, replace `$MATLABROOT/bin/...` with the path to the MCR.

To run the code generated from the above System objects and blocks on a machine does not have MCR or MATLAB installed, use the `packNGo` function. The `packNGo` function packages all relevant files in a compressed zip file so that you can relocate, unpack, and rebuild your project in another development environment with no MATLAB installed.

You can use the `packNGo` function at the command line or the **Package** option in the MATLAB Coder app. The files are packaged in a compressed file that you can relocate and unpack using a standard zip utility. For more details on how to pack the code generated from MATLAB code, see “Package Code for Other Development Environments” (MATLAB Coder). For more details on how to pack the code generated from Simulink blocks, see the `packNGo` function.

See Also


More About

- “MATLAB Programming for Code Generation” (MATLAB Coder)

Block Example Repository

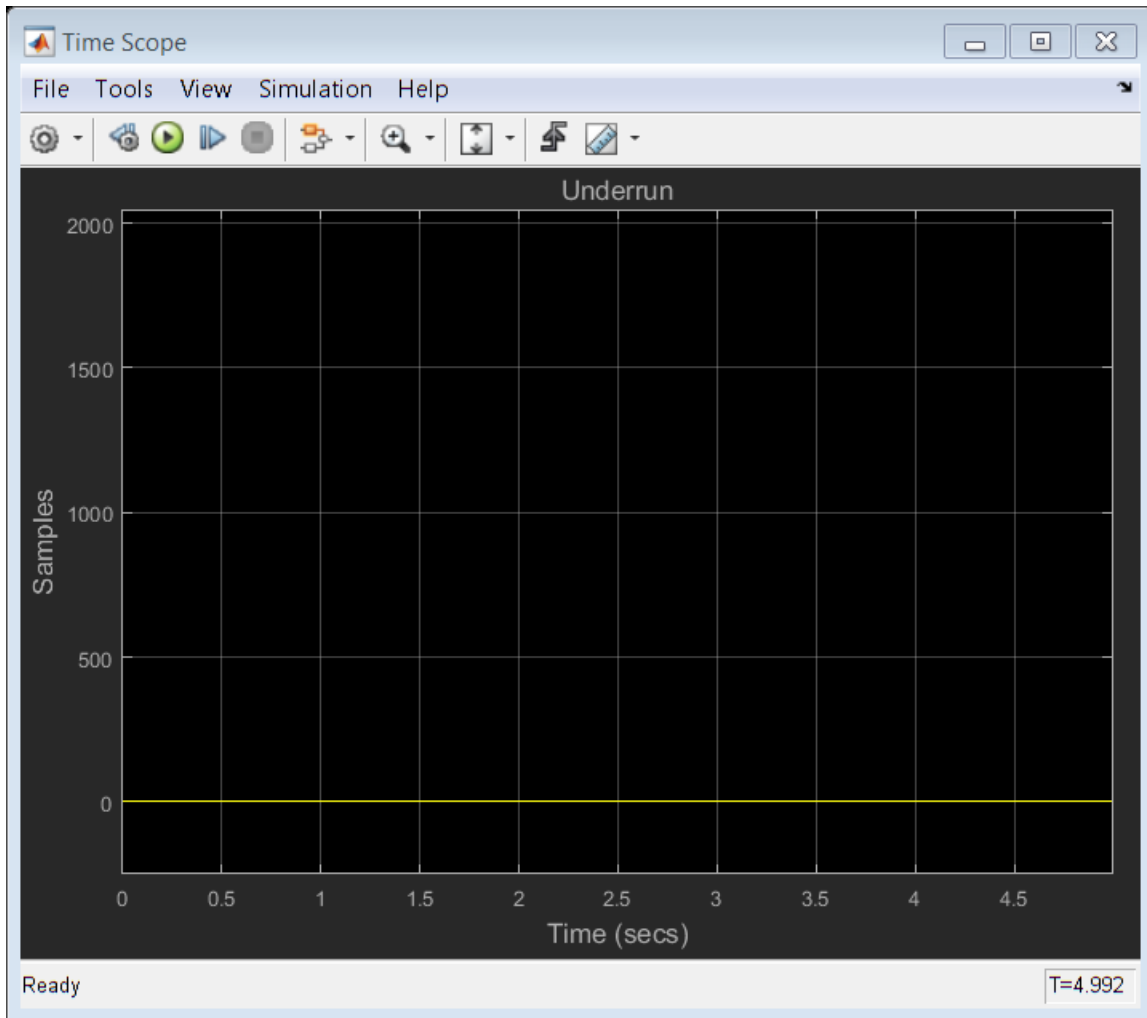
Decrease Underrun

Examine the Audio Device Writer block in a Simulink® model, determine underrun, and decrease underrun.



Copyright 2016 The MathWorks, Inc.

1. Run the model. The Audio Device Writer sends an audio stream to your computer's default audio output device. The Audio Device Writer block sends the number of samples underrun to your Time Scope.



2. Uncomment the Artificial Load block. This block performs computations that slow the simulation.
3. Run the model. If your device writer is dropping samples:
 - a. Stop the simulation.
 - b. Open the From Multimedia File block.

- c. Set the **Samples per frame** parameter to 1024.
- d. Close the block and run the simulation.

If your model continues to drop samples, increase the frame size again. The increased frame size increases the buffer size used by the sound card. A larger buffer size decreases the possibility of underruns at the cost of higher audio latency.

See Also

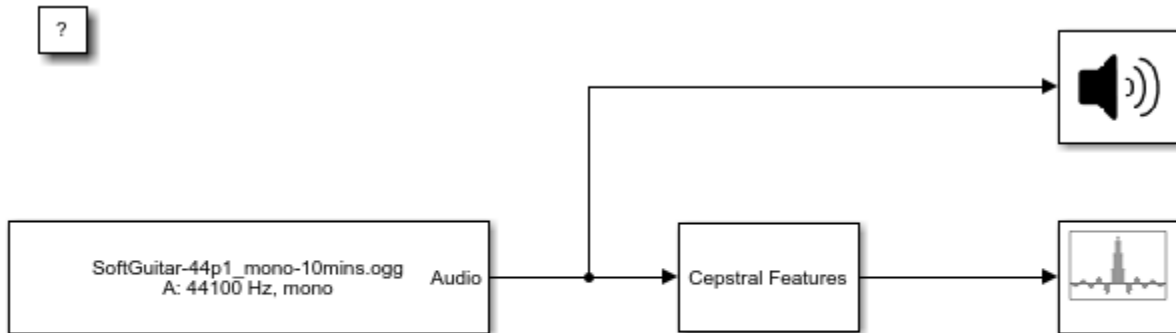
From Multimedia File | Time Scope

Block Example Repository

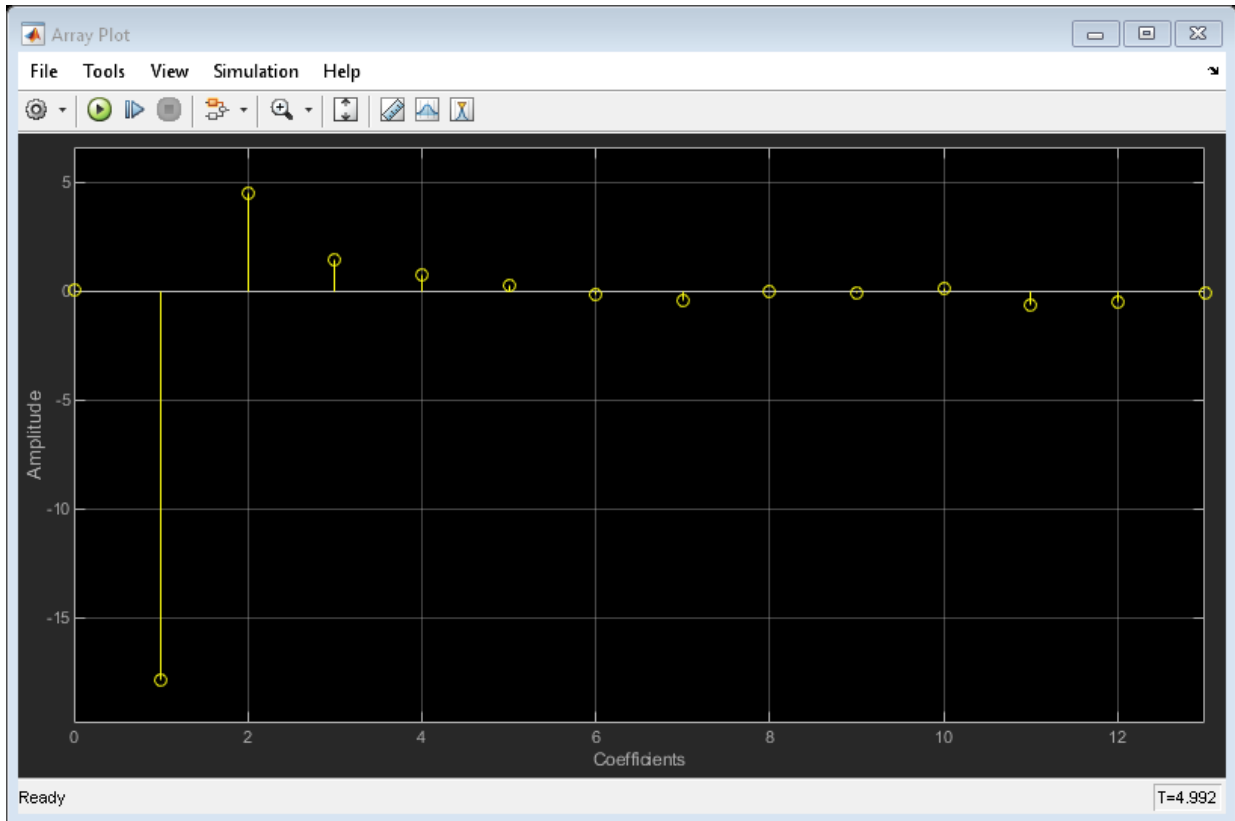
- “Extract Cepstral Coefficients” on page 16-2
- “Tune Center Frequency Using Input Port” on page 16-4
- “Gate Audio Signal Using VAD” on page 16-6
- “Frequency-Domain Voice Activity Detection” on page 16-8
- “Visualize Noise Power” on page 16-9
- “Detect Presence of Speech” on page 16-13
- “Perform Graphic Equalization” on page 16-15
- “Split-Band De-Essing” on page 16-17
- “Diminish Plosives from Speech” on page 16-18
- “Suppress Loud Sounds” on page 16-19
- “Attenuate Low-Level Noise” on page 16-22
- “Suppress Volume of Loud Sounds” on page 16-25
- “Gate Background Noise” on page 16-28
- “Output Values from MIDI Control Surface” on page 16-31
- “Apply Frequency Weighting” on page 16-33
- “Compare Loudness Before and After Audio Processing” on page 16-35
- “Two-Band Crossover Filtering for a Stereo Speaker System” on page 16-37
- “Mimic Acoustic Environments” on page 16-39
- “Perform Parametric Equalization” on page 16-41
- “Perform Octave Filtering” on page 16-43
- “Read from Microphone and Write to Speaker” on page 16-45
- “Channel Mapping” on page 16-48
- “Trigger Gain Control Based on Loudness Measurement” on page 16-50

Extract Cepstral Coefficients

Use the Cepstral Feature Extractor block to extract and visualize cepstral coefficients from an audio file.



Copyright 2018 The MathWorks, Inc.




See Also

Array Plot | Audio Device Writer | Cepstral Feature Extractor | From Multimedia File | cepstralFeatureExtractor | gtcc | mfcc

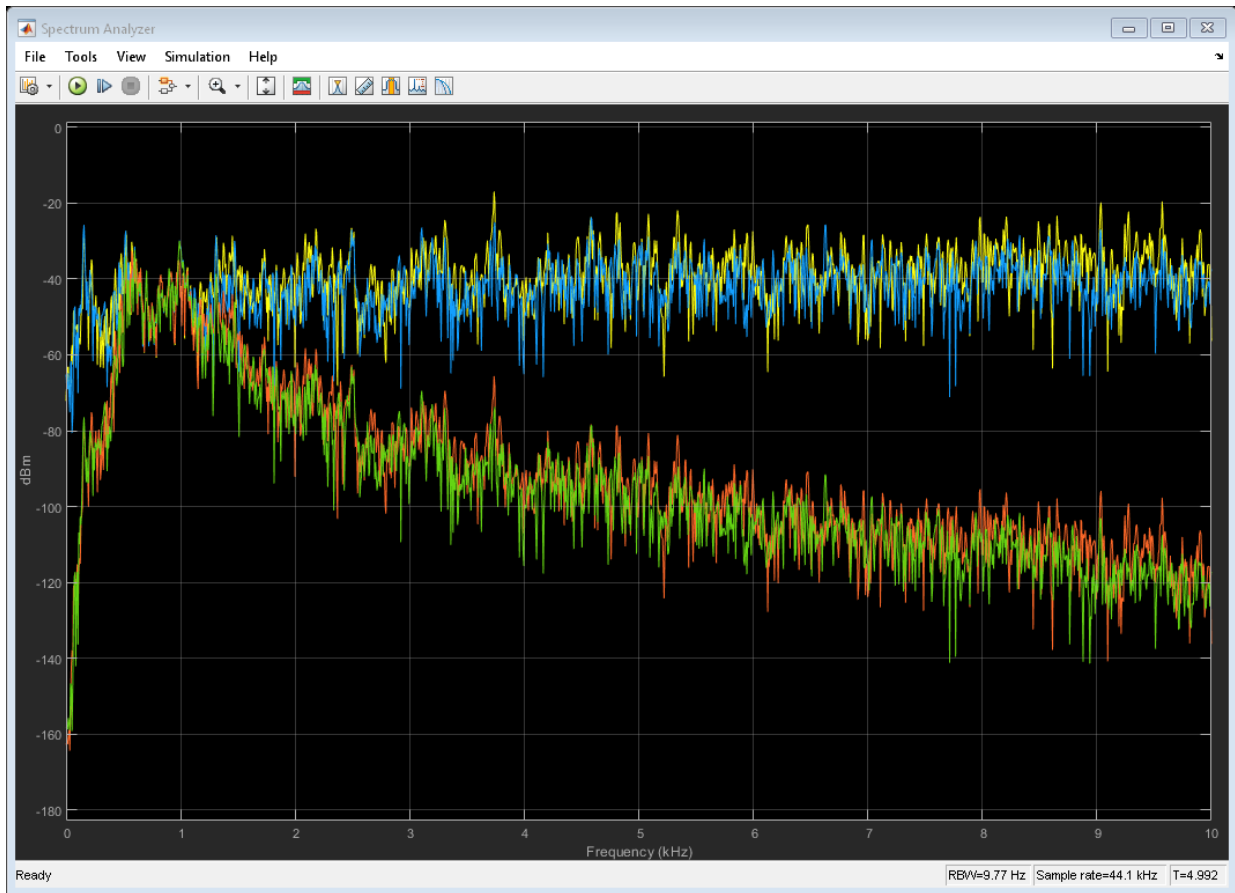
Tune Center Frequency Using Input Port

Tune the center frequency of an Octave Filter block in Simulink® using the optional input port.



Copyright 2018 The MathWorks Inc.

1. Run the simulation. The From Multimedia File block sends a stereo audio stream to the Octave Filter block. The center frequency of the Octave Filter block can be tuned using the manual switches routed into the optional input port. The filtered audio is sent to your computer's default audio device. The filtered audio and unfiltered audio are sent to a Spectrum Analyzer block for visualization.
2. Tune the center frequency by toggling manual switches routing constant values. The constant value routed from the left is multiplied with the constant value routed from the right. The center frequency of the Octave Filter block can be set at 400, 800, 4000, and 8000 Hz.
3. Observe the Spectrum Analyzer as you tune the center frequency. Note how the center frequency changes as you toggle the manual switches.



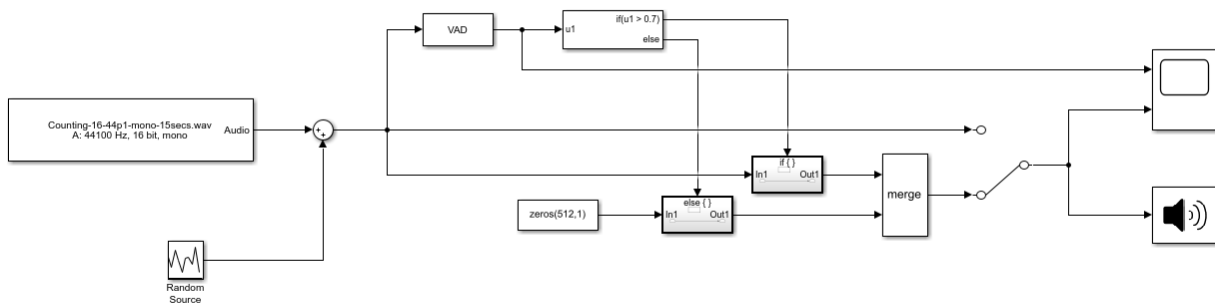
See Also

[Audio Device Writer](#) | [From Multimedia File](#) | [Manual Switch](#) | [Octave Filter](#) | [Time Scope](#)

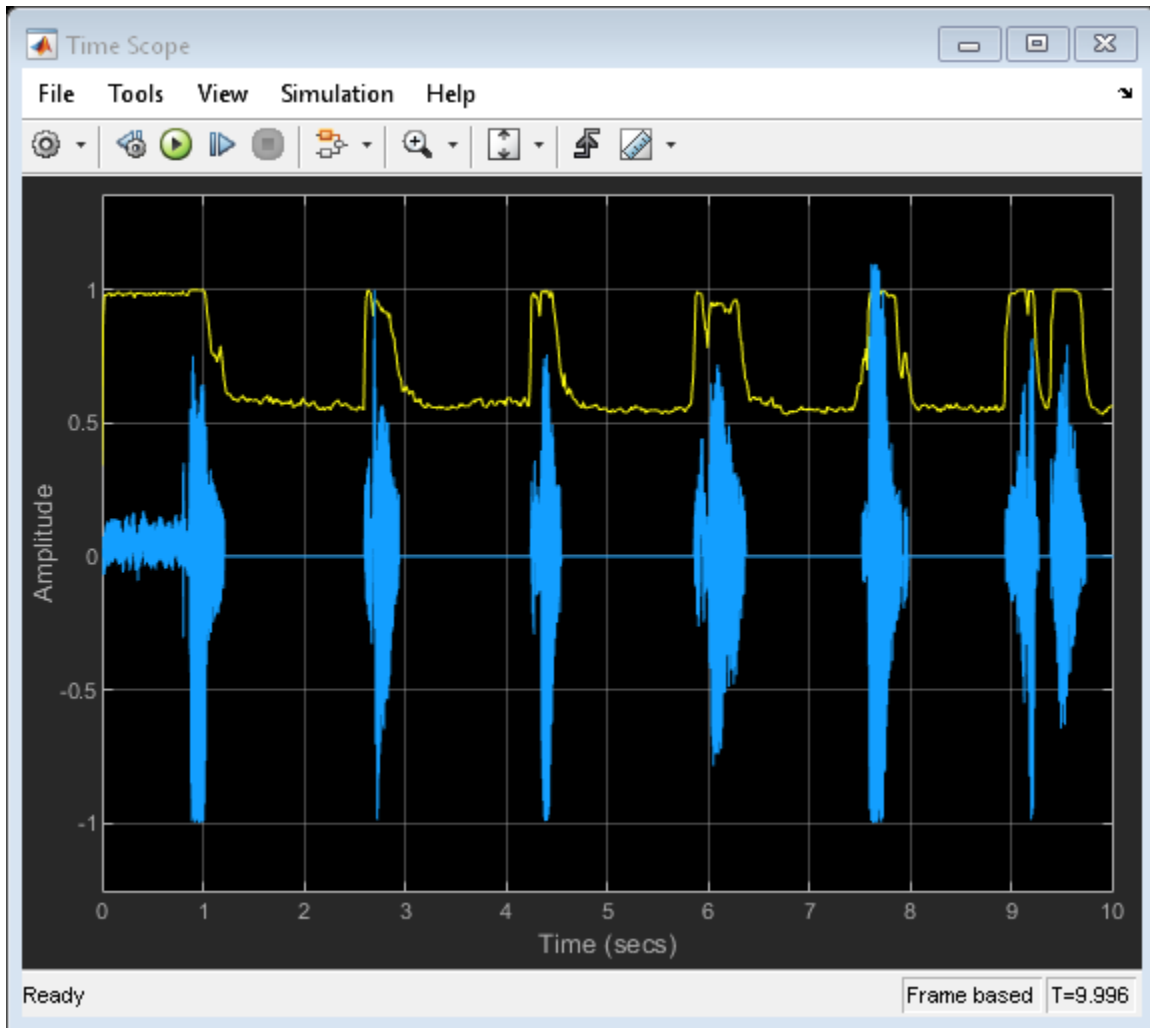
Gate Audio Signal Using VAD

This model uses if-else block signal routing to replace regions of no speech with zeros.

To explore this model, tune the **Probability of transition from a silence frame to a speech frame** and **Probability of transition from a speech frame to a silence frame** parameters of the Voice Activity Detector (VAD) and observe the effect on the speech presence probability. Toggle between the gated and original audio signal to assess the quality of your system.



Copyright 2017 The MathWorks, Inc.



See Also

Audio Device Writer | From Multimedia File | If | If Action Subsystem | Manual Switch | Random Source | Time Scope | Voice Activity Detector

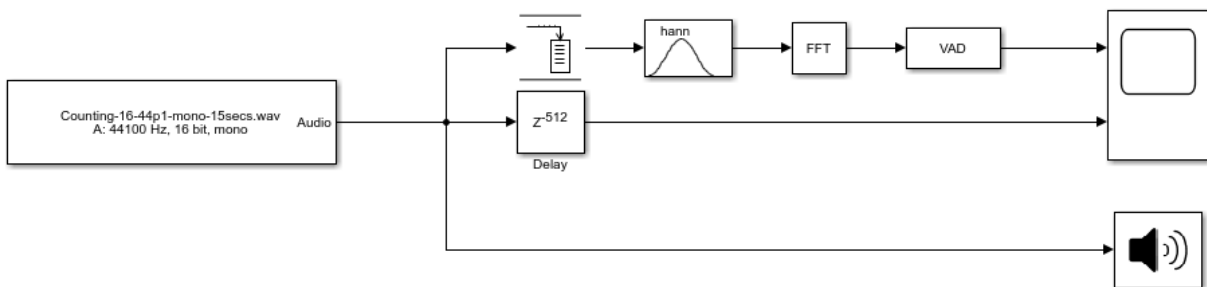
Frequency-Domain Voice Activity Detection

This model detects voice activity using a frequency-domain audio signal.

Voice Activity Detection is often used as an indication whether further processing or analysis of a signal is required. Many processing and analysis techniques require a frequency-domain representation of the signal. For example, the voice activity detection algorithm operates in the frequency domain. To save computation, you can convert the audio signal to the frequency domain once, and then feed the frequency-domain signal to downstream analysis and processing.

This model additionally buffers the signal so that the VAD operates on half-overlapped frames. Overlapping the input frames to the VAD increases the accuracy and resolution in time of the probability of speech.

?




Copyright 2017 The MathWorks, Inc.

See Also

[Buffer](#) | [Audio Device Writer](#) | [Delay](#) | [FFT](#) | [From Multimedia File](#) | [Time Scope](#) | [Voice Activity Detector](#) | [Window Function](#)


Visualize Noise Power

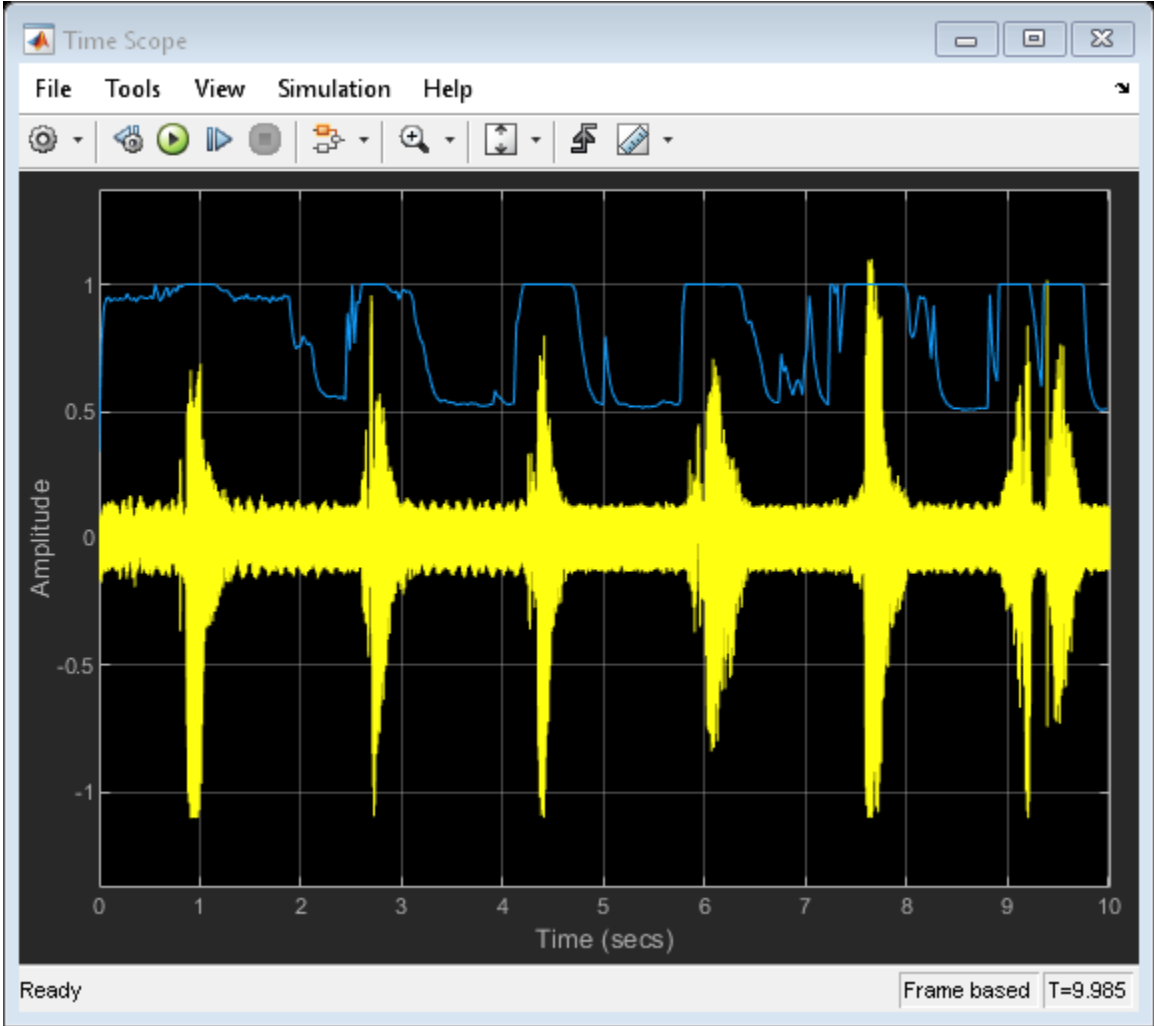
This model plots the noise power estimated by the Voice Activity Detector.



Copyright 2017 The MathWorks, Inc.

To explore this model, tune the **Frequency (Hz)** parameter of the Sine Wave block and observe the noise power estimate updated on the Array Plot block.





Zoom in on the Array Plot to verify that the Voice Activity Detector outputs a good estimate of the noise tone.

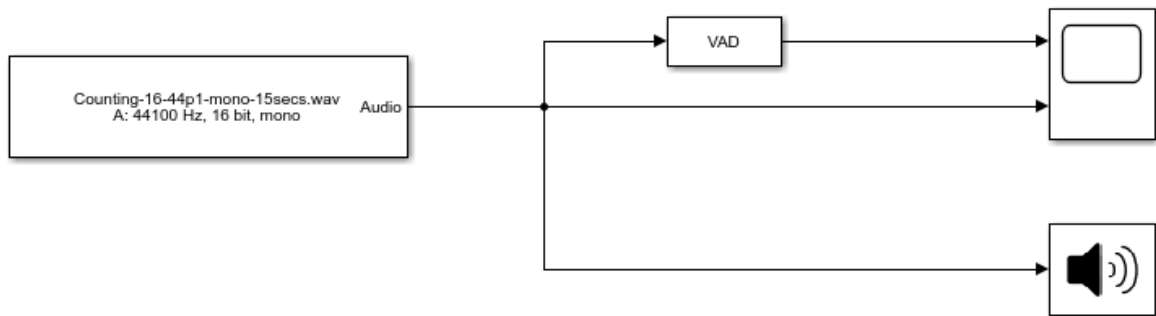
See Also

[Array Plot](#) | [Audio Device Writer](#) | [From Multimedia File](#) | [Sine Wave](#) | [Time Scope](#) | [Voice Activity Detector](#)

Detect Presence of Speech

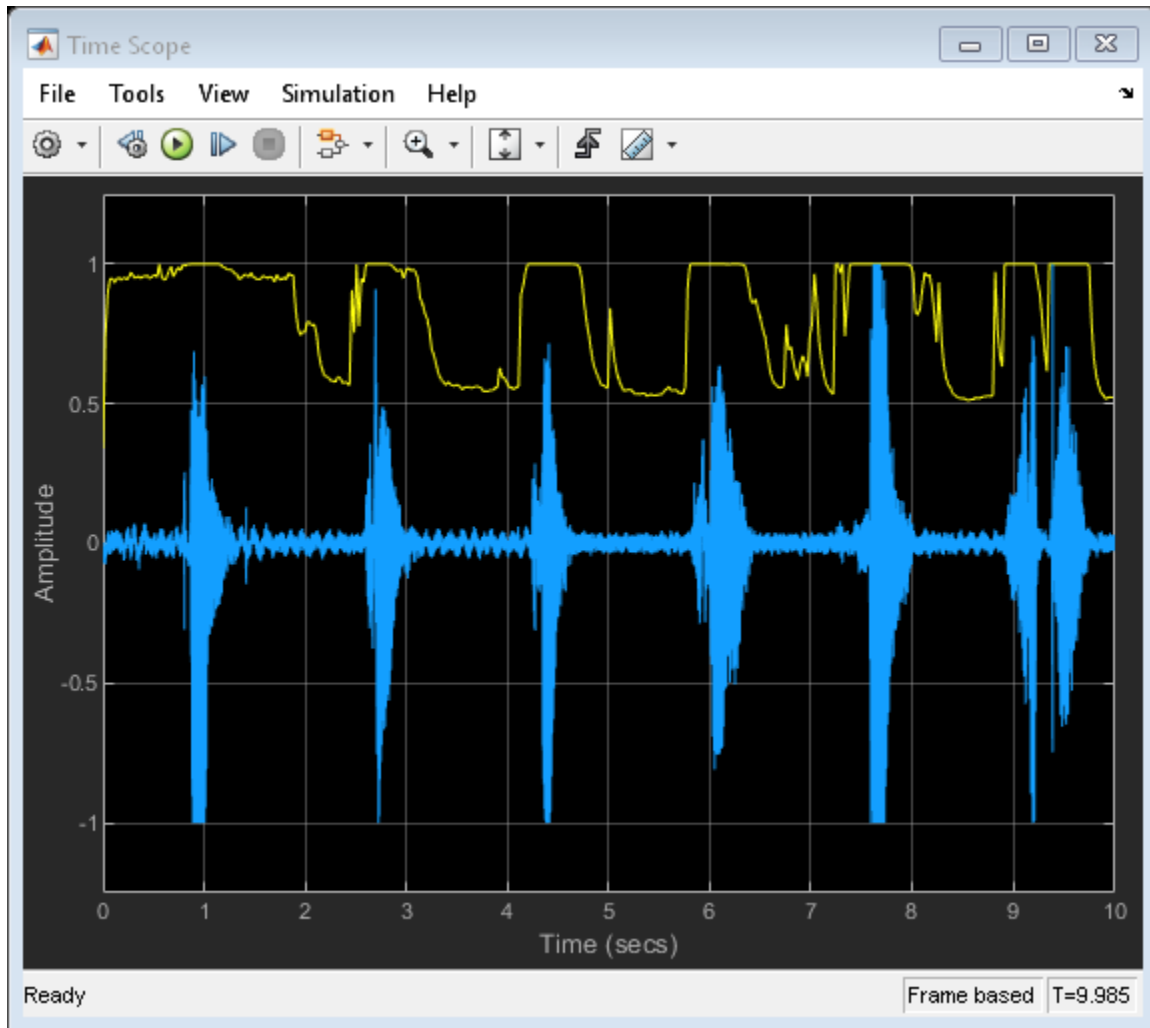
This model uses the Voice Activity Detector block to visualize the probability of speech presence in an audio signal.

To explore this model, tune the **Probability of transition from a silence frame to a speech frame** and **Probability of transition from a speech frame to a silence frame** parameters of the Voice Activity Detector (VAD) and observe the effect on the speech presence probability.



Copyright 2017 The MathWorks, Inc.

The Time Scope blocks plots the audio signal and associated voice activity probability.




See Also

[Audio Device Writer](#) | [From Multimedia File](#) | [Time Scope](#) | [Voice Activity Detector](#)

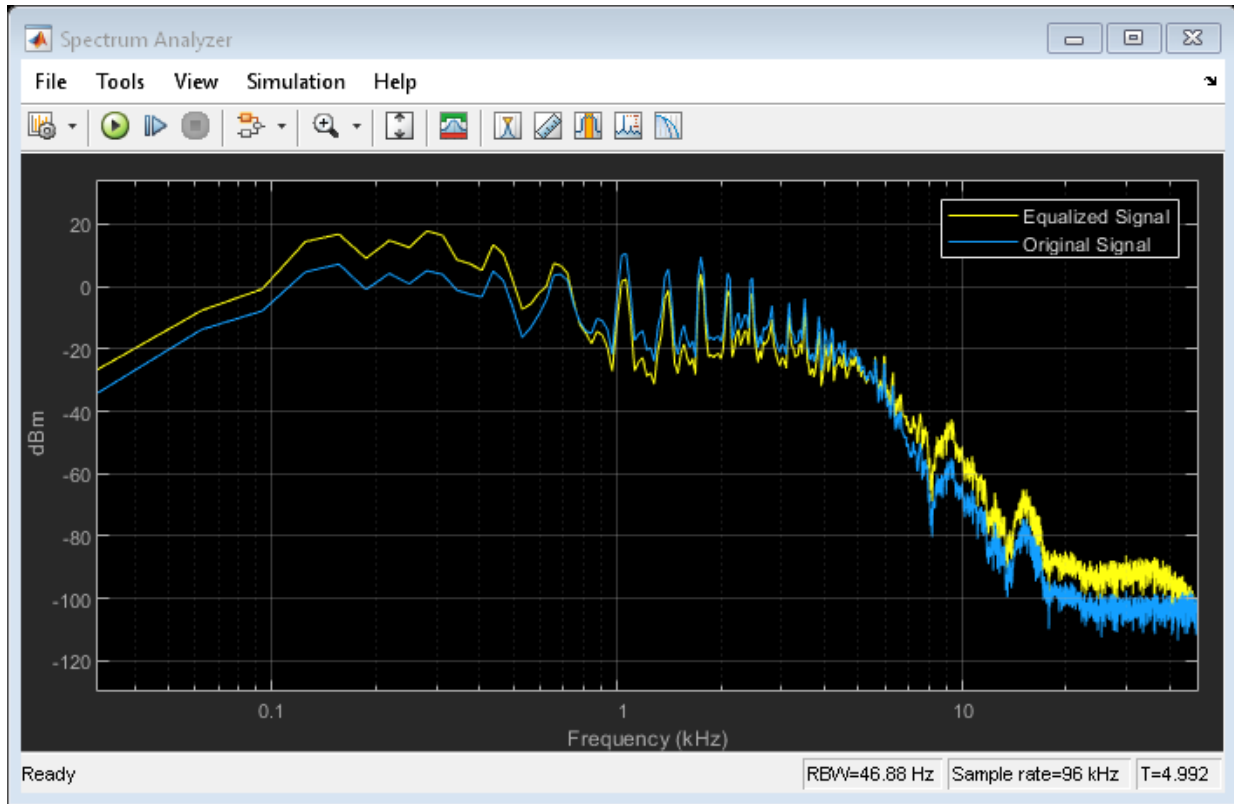
Perform Graphic Equalization

Examine the Graphic EQ block in a Simulink® model and tune parameters.



Copyright 2017 The MathWorks, Inc.

1. Open the Spectrum Analyzer and Graphic EQ blocks.
2. In the Graphic EQ block, click **Visualize equalizer response**. Modify gains of the graphic equalizer and see the magnitude response plot update automatically.
3. Run the model. Tune gains on the Graphic EQ to listen to the effect on your audio device and see the effect on the Spectrum Analyzer display. Double-click the Manual Switch block to toggle between the original and equalized signal as output.



See Also


[Audio Device Writer](#) | [From Multimedia File](#) | [Graphic EQ](#) | [Spectrum Analyzer](#)

Split-Band De-Essing

This model implements split-band de-essing by separating a speech signal into high and low frequencies, applying dynamic range expansion to diminish the sibilant frequencies, and then remixing the channels.

De-essing is the process of diminishing sibilant sounds in an audio signal. Sibilance refers to the *s*, *z*, and *sh* sounds in speech, which can be disproportionately emphasized during recording. *es* sounds fall under the category of unvoiced speech with all consonants, and have a higher frequency than voiced speech.

To explore the model, tune the parameters of the Expander and Crossover Filter blocks. To switch between listening to the processed and unprocessed speech signal, double-click the Manual Switch block. To view the effect of the processing, double-click the Time Scope block.



Copyright 2017 The MathWorks, Inc.

See Also

Audio Device Writer | Crossover Filter | Expander | From Multimedia File | Time Scope


Diminish Plosives from Speech

This model minimizes the plosives of a speech signal by applying highpass filtering and low-band compression.

Plosives are consonant sounds resulting from a sudden release of airflow. They are most pronounced in *p*, *d*, and *g* words. Plosives can be emphasized by the recording process and are often displeasurable to hear.

To explore this model, tune the highpass filter cutoff and the parameters on the Compressor and Crossover Filter blocks. Switch between listening to the original and processed signals by double-clicking the Manual Switch block.

?



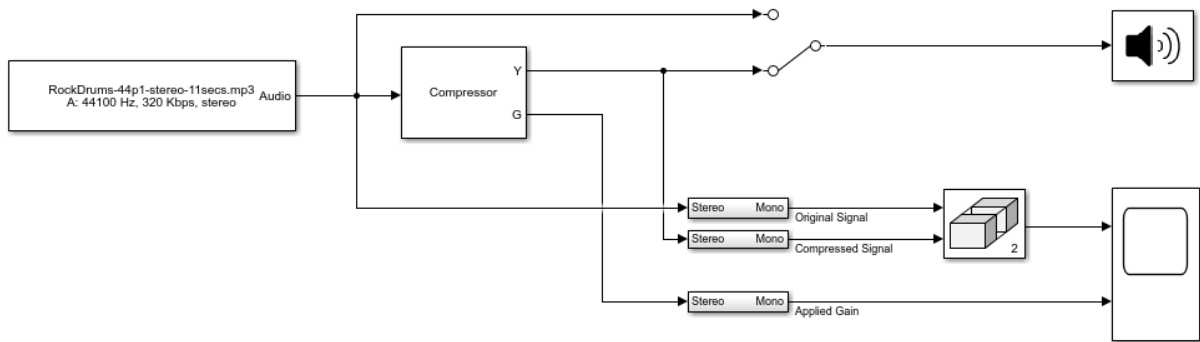
Copyright 2017 The MathWorks, Inc.

See Also

[Audio Device Writer](#) | [Compressor](#) | [Crossover Filter](#) | [From Multimedia File](#)

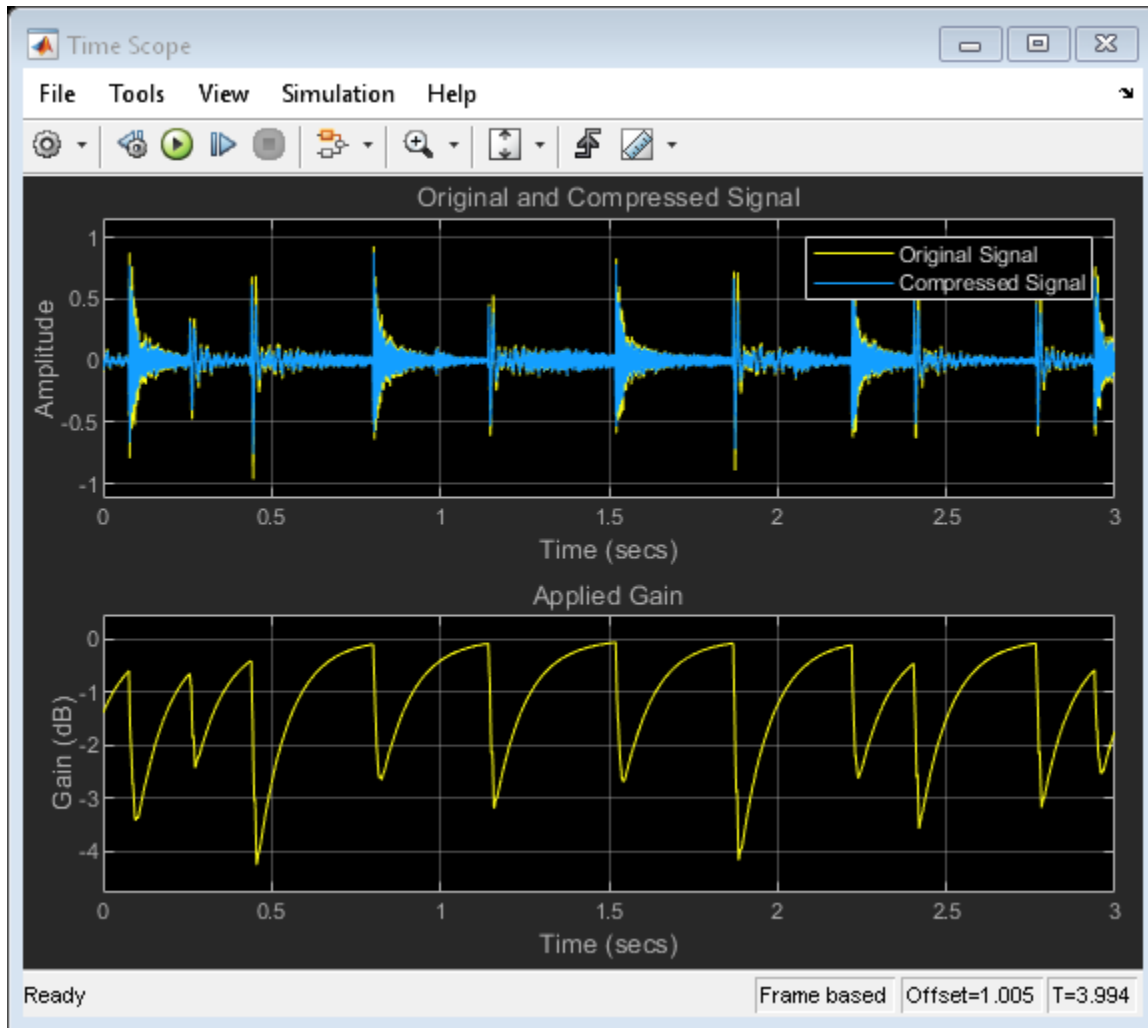
Suppress Loud Sounds

Use the Compressor block to suppress loud sounds and visualize the applied compression gain.



Copyright 2016-2017 The MathWorks, Inc.

1. Open the Time Scope and Compressor blocks.
2. Run the model. To switch between listening to the compressed signal and the original signal, double-click the Manual Switch block.
3. Observe how the applied gain depends on compression parameters and input signal dynamics by tuning the Compressor block parameters and viewing the results on the Time Scope.



See Also

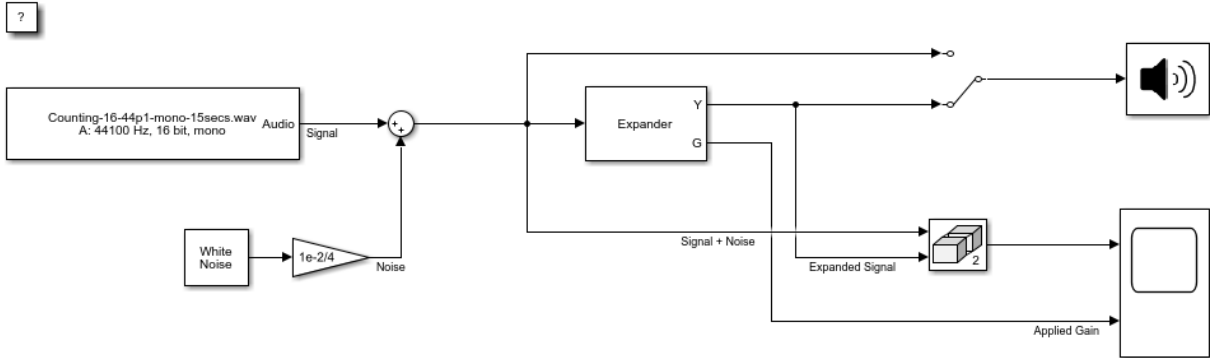
Audio Device Writer | Compressor | From Multimedia File | Time Scope | Vector Concatenate, Matrix Concatenate

More About

- “Dynamic Range Control” on page 7-2

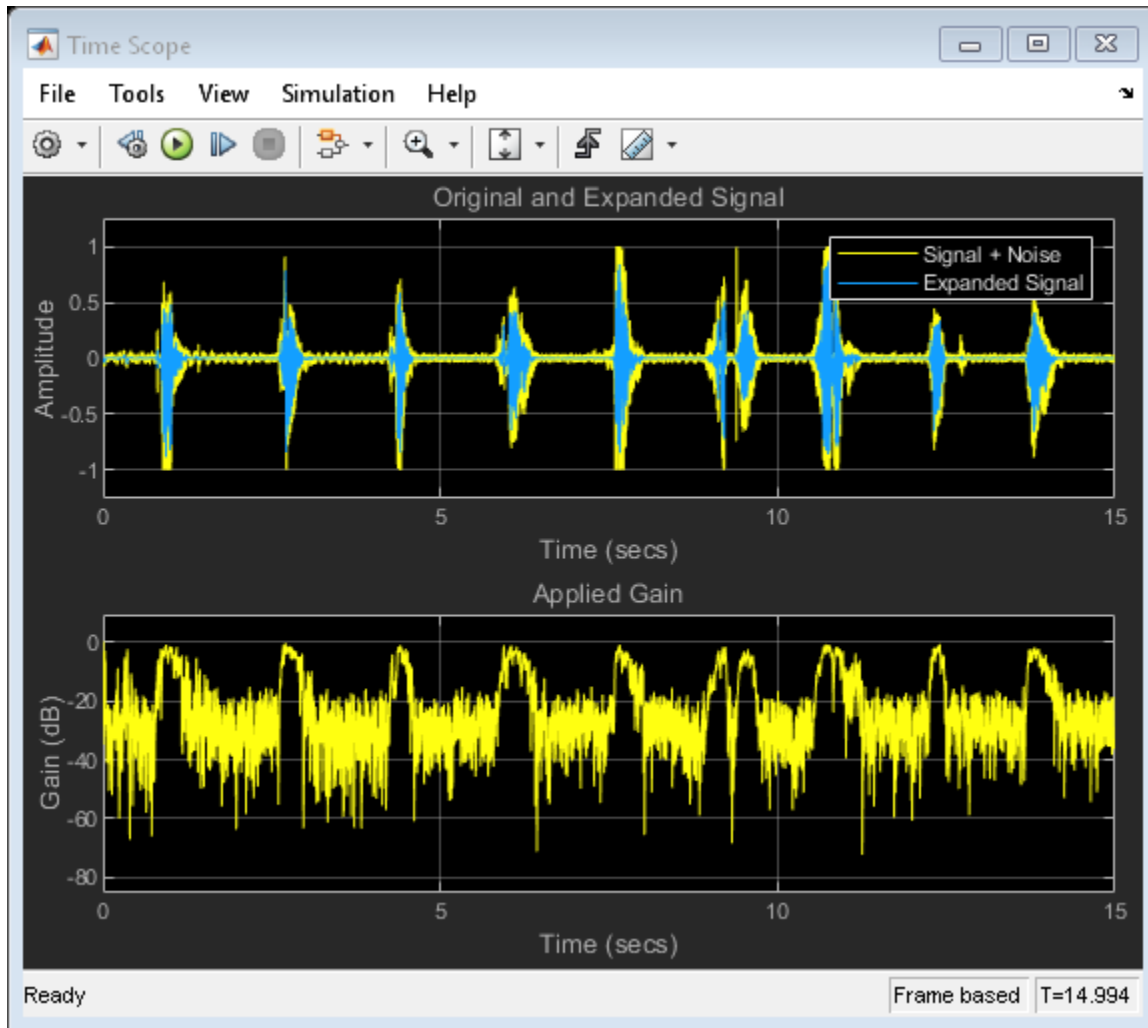
Attenuate Low-Level Noise

Use the Expander block to attenuate low-level noise and visualize the applied dynamic range control gain.



Copyright 2016 The MathWorks, Inc.

1. Open the Time Scope and Expander blocks.
2. Run the model. To switch between listening to the expanded signal and the original signal, double-click the Manual Switch block.
3. Observe how the applied gain depends on expansion parameters and input signal dynamics by tuning the Expander block parameters and viewing the results on the Time Scope.



See Also

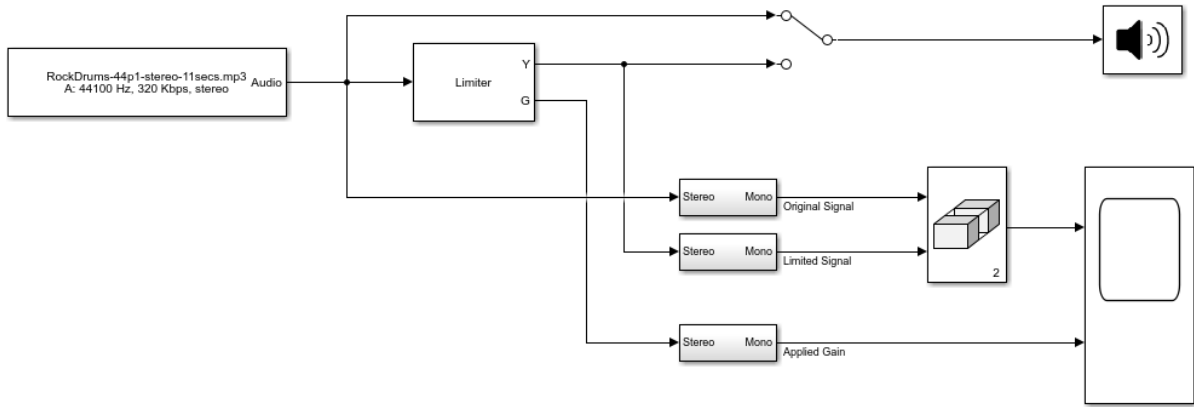
Audio Device Writer | Colored Noise | Expander | From Multimedia File | Time Scope |
Vector Concatenate, Matrix Concatenate

More About

- “Dynamic Range Control” on page 7-2

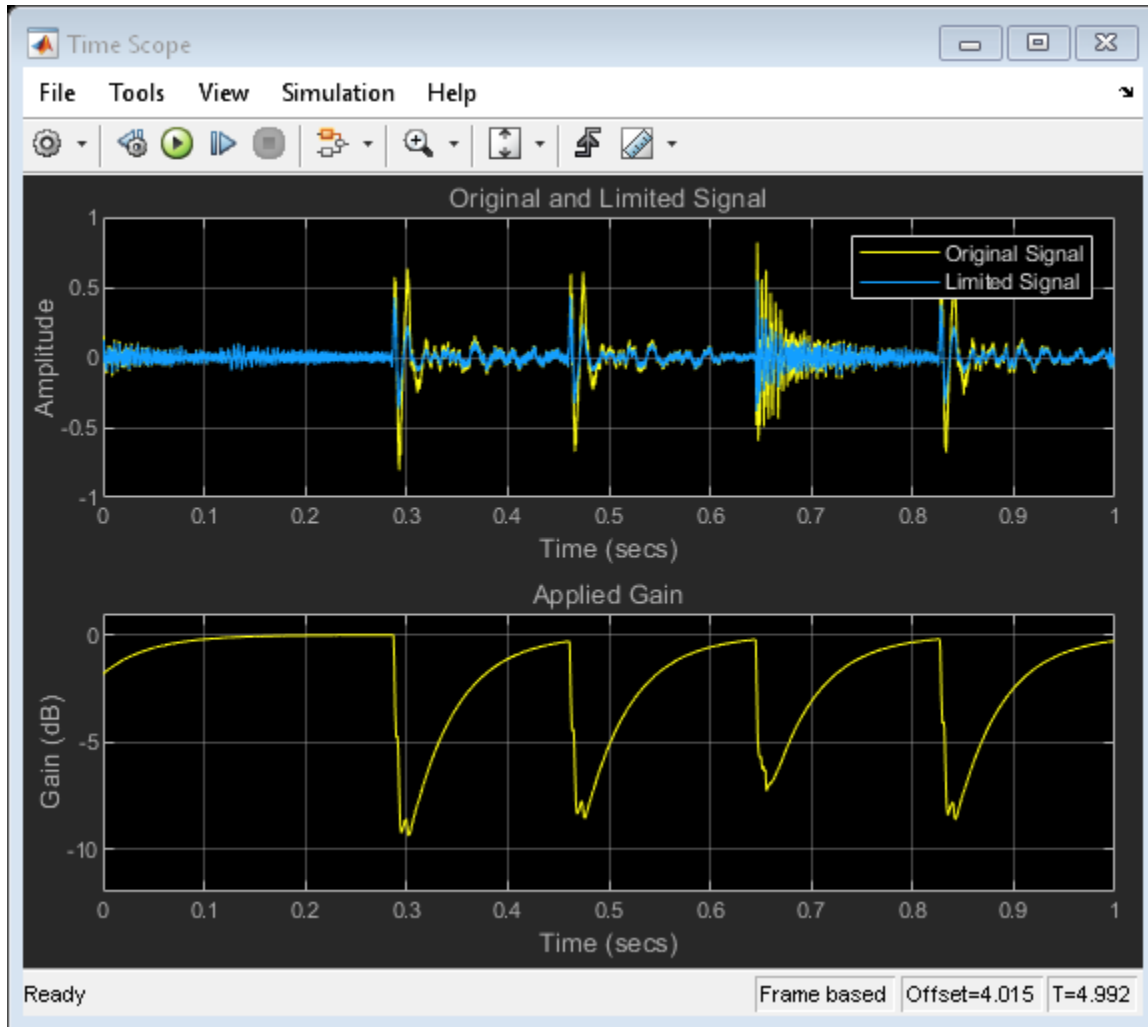
Suppress Volume of Loud Sounds

Suppress the volume of loud sounds and visualize the applied dynamic range control gain.



Copyright 2016 The Mathworks, Inc.

1. Open the Time Scope and Limiter blocks.
2. Run the model. To switch between listening to the gated signal and the original signal, double-click the Manual Switch block.
3. Observe how the applied gain depends on dynamic range limiting parameters and input signal dynamics by tuning Limiter block parameters and viewing the results on the Time Scope.



See Also

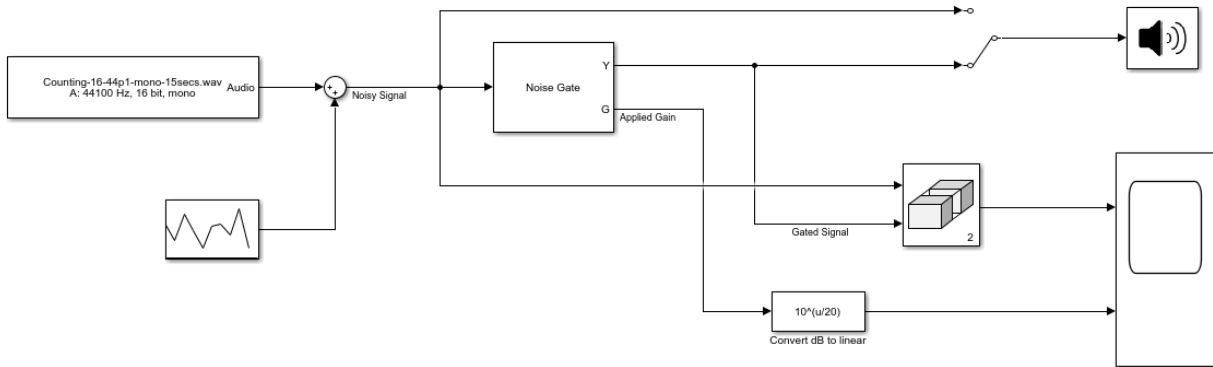
Audio Device Writer | From Multimedia File | Limiter | Time Scope | Vector Concatenate, Matrix Concatenate

More About

- “Dynamic Range Control” on page 7-2

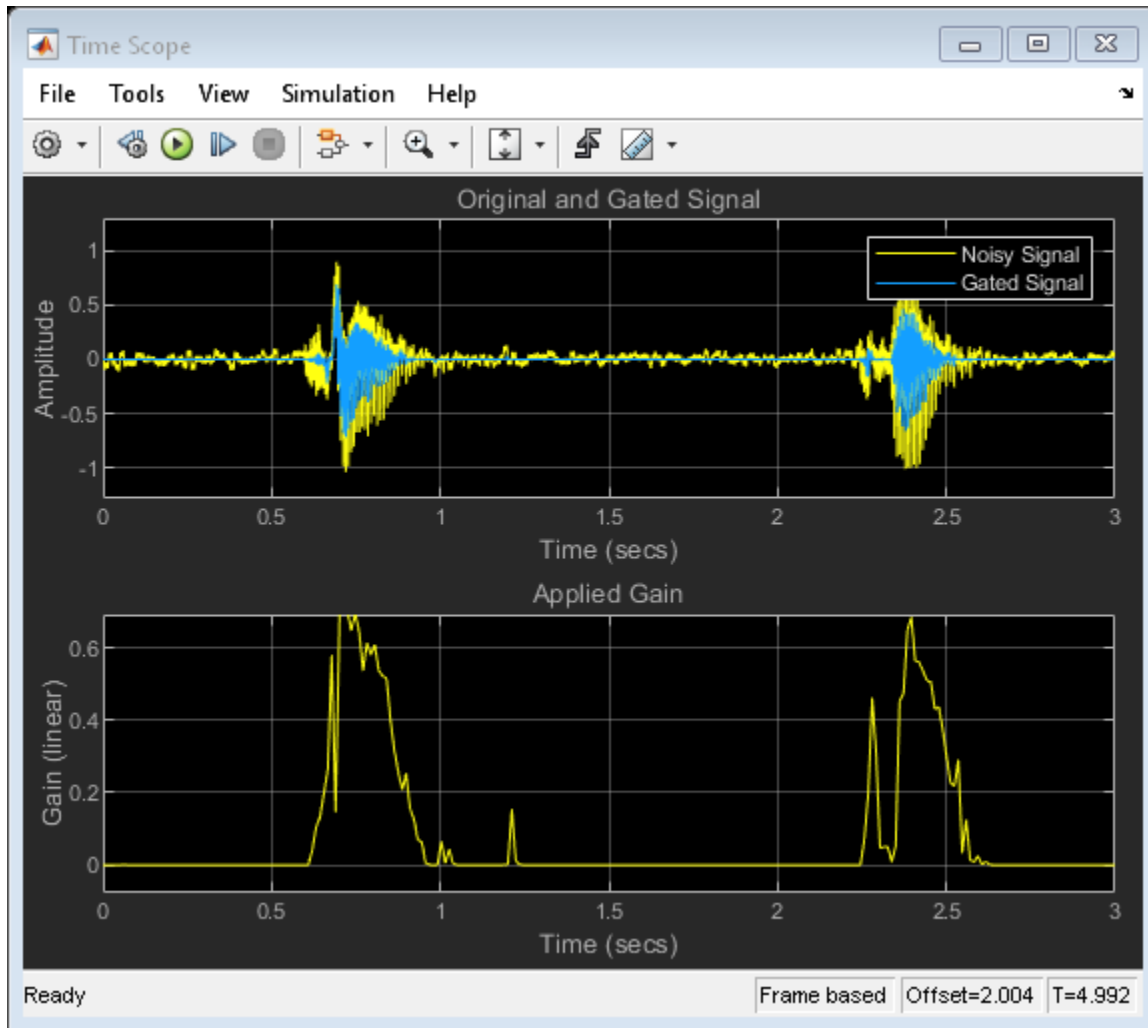
Gate Background Noise

Apply dynamic range gating to remove low-level noise from an audio file.



Copyright 2016 The MathWorks, Inc.

1. Open the Time Scope and Noise Gate blocks.
2. Run the model. To switch between listening to the gated signal and the original signal, double-click the Manual Switch block.
3. Observe how the applied gain depends on noise gate parameters and input signal dynamics by tuning Noise Gate block parameters and viewing the results on the Time Scope.



See Also

Audio Device Writer | From Multimedia File | Noise Gate | Random Source | Time Scope | Vector Concatenate, Matrix Concatenate

More About

- “Dynamic Range Control” on page 7-2

Output Values from MIDI Control Surface

The example shows how to set the MIDI Controls block parameters to output control values from your MIDI device.


1. Connect a MIDI device to your computer and then open the model.



Copyright 2016 The MathWorks, Inc.

2. Run the model with default settings. Move any controller on your default MIDI device to update the Display block.
3. Stop the simulation.
4. At the MATLAB™ command line, use `midid` to determine the name of your MIDI device and two control numbers associated with your device.
5. In the MIDI Control block dialog box, set **MIDI device** to Specify other and enter the name of your MIDI device.
6. Set **MIDI controls** to Respond to specified controls and enter the control numbers determined using `midid`.
7. Specify initial values as a vector the same size as **MIDI control numbers**. The initial values you specify are quantized according to the MIDI protocol and your particular MIDI surface.

The dialog box shows sample values for a 'BCF2000' MIDI device with control numbers 1081 and 1083.



8. Click **OK**, and then run the model. Verify that the Display block shows initial values and updates when you move the specified controls.

See Also


Audio Device Writer | From Multimedia File | MIDI Controls | Time Scope | Vector Concatenate, Matrix Concatenate

More About

- “MIDI Control Surface Interface” on page 9-2

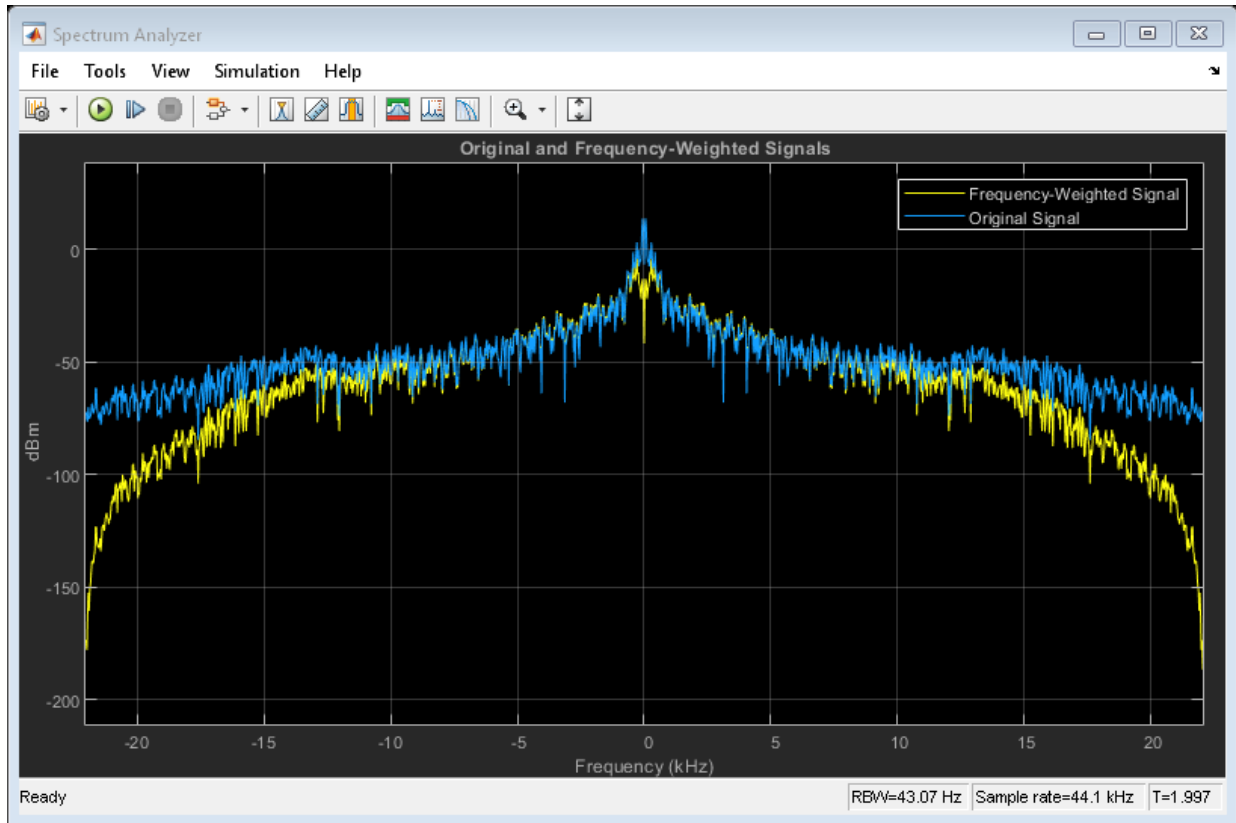
Apply Frequency Weighting

Examine the Weighting Filter block in a Simulink® model and tune parameters.



Copyright 2016 The MathWorks, Inc.

1. Open the Spectrum Analyzer block.
2. Run the model. Switch between listening to the frequency-weighted signal and the original signal by double-clicking the Manual Switch block.
3. Stop the model. Open the Weighting Filter block and choose a different weighting method. Observe the difference in simulation.

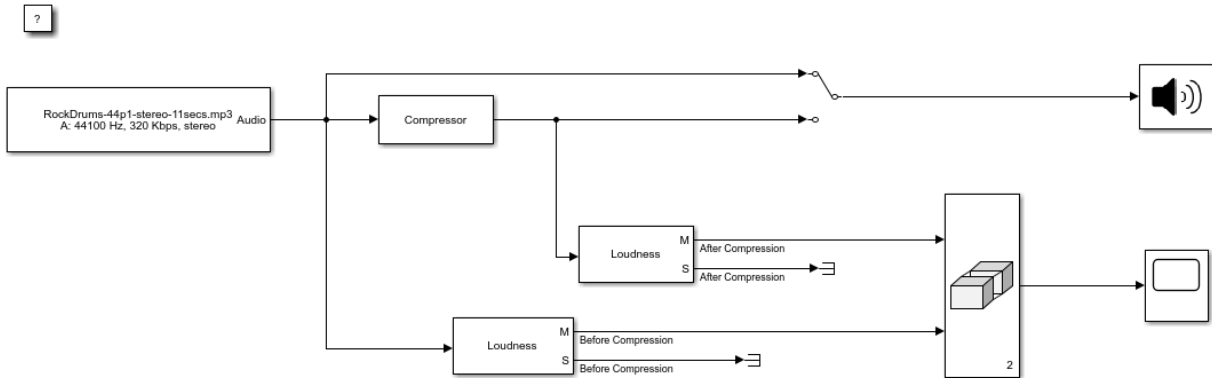


See Also

[Audio Device Writer](#) | [From Multimedia File](#) | [Spectrum Analyzer](#) | [Weighting Filter](#)

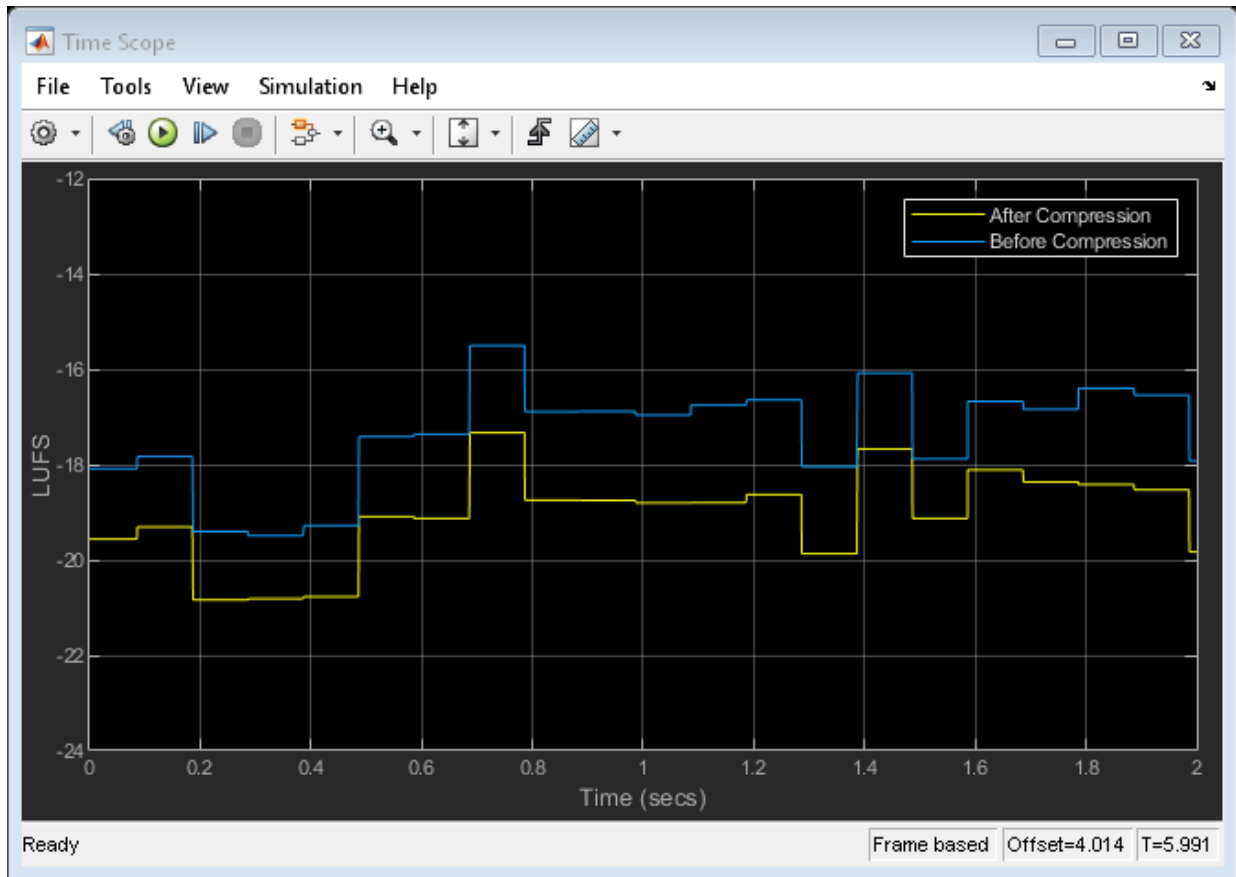
Compare Loudness Before and After Audio Processing

Measure momentary and short-term loudness before and after compression of a streaming audio signal in Simulink®.



Copyright 2016 The MathWorks, Inc.

1. Open the Time Scope and Compressor blocks.
2. Run the model. To switch between listening to the compressed signal and the original signal, double-click the switch.
3. Observe the effect of compression on loudness by tuning the Compressor block parameters and viewing the momentary loudness on the Time Scope block.




4. Stop the model. For both Loudness blocks, replace momentary loudness with short-term loudness as input to the Matrix Concatenate block. Run the model again and observe the effect of compression on short-term loudness.

See Also

Audio Device Writer | Compressor | From Multimedia File | Loudness Meter | Time Scope | Vector Concatenate, Matrix Concatenate

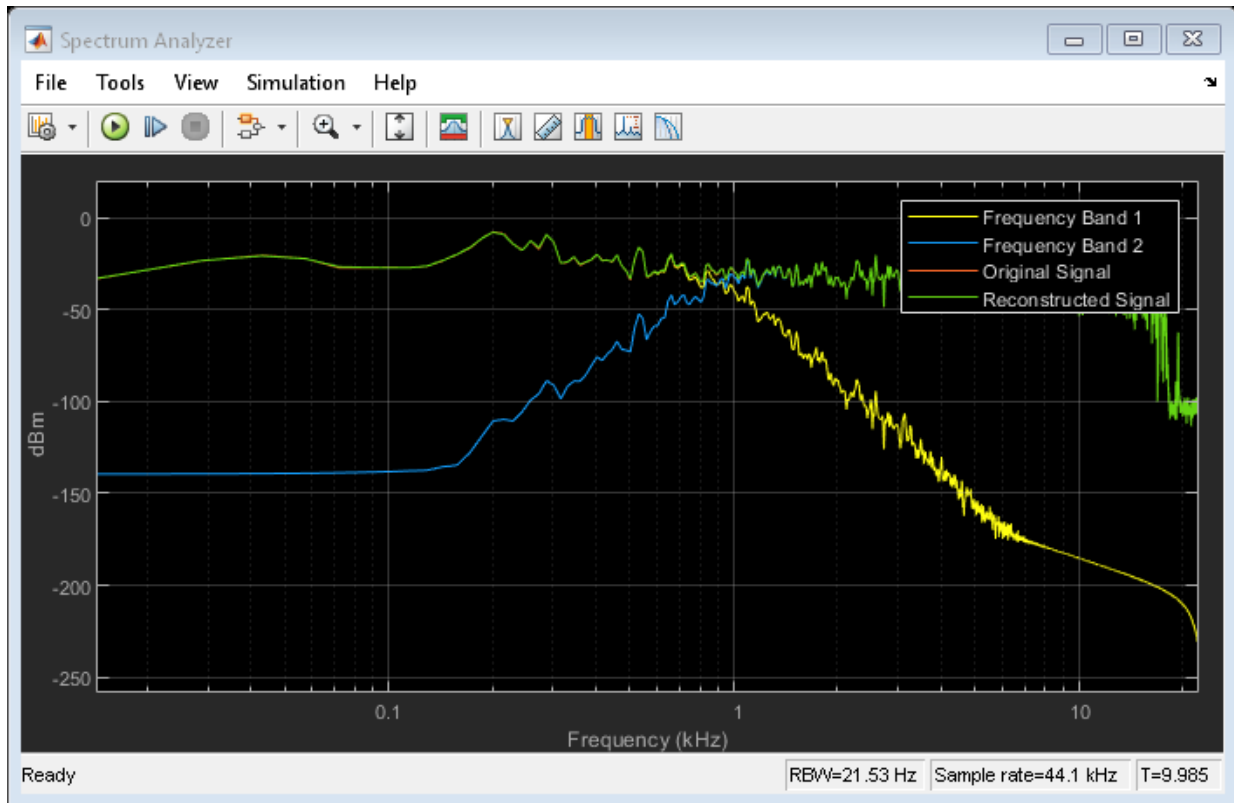
Two-Band Crossover Filtering for a Stereo Speaker System

Divide a mono signal into a stereo signal with distinct frequency bands. To hear the full effect of this simulation, use a stereo speaker system, such as headphones.



Copyright 2016 The MathWorks, Inc.

1. Open the Spectrum Analyzer and Crossover Filter blocks.
2. Run the model. To switch between listening to the filtered and original signal, double-click the Manual Switch block.
3. Tune the crossover frequency on the Crossover Filter block to listen to the effect on your speakers and view the effect on the Spectrum Analyzer block.

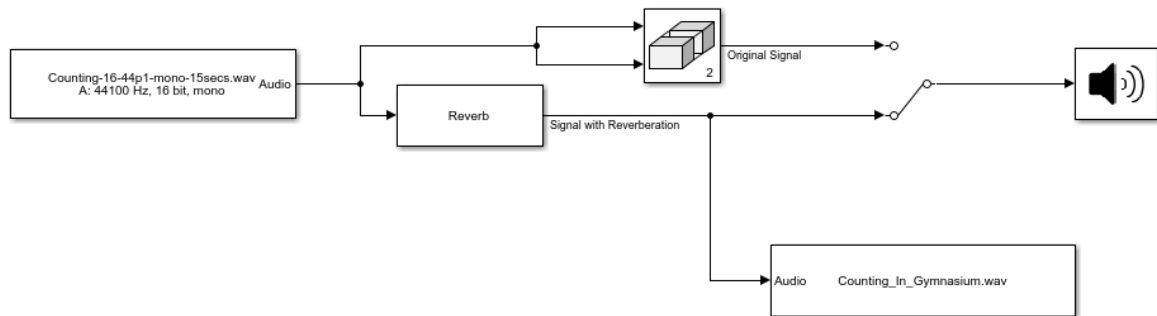


See Also

[Audio Device Writer](#) | [Crossover Filter](#) | [From Multimedia File](#) | [Spectrum Analyzer](#) | [Vector Concatenate](#), [Matrix Concatenate](#)

Mimic Acoustic Environments

Examine the Reverberator block in a Simulink® model and tune parameters. The reverberation parameters in this model mimic a large room with hard walls, such as a gymnasium.



Copyright 2016 The Mathworks, Inc.

1. Run the simulation. Listen to the audio signal with and without reverberation by double-clicking the Manual Switch block.
2. Stop the simulation.
3. Disconnect the From Multimedia File block so that you can run the model without recording.
4. Open the Reverberator block.
5. Run the simulation and tune the parameters of the Reverberator block.
6. After you are satisfied with the reverberation environment, stop the simulation.

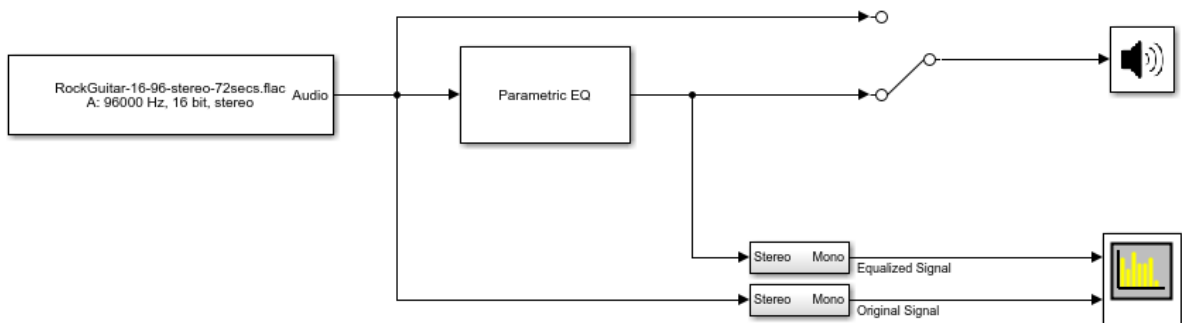
7. Reconnect the To Multimedia File block. Rename the output file with a description to match your reverberation environment, and rerun the model.

See Also

Audio Device Writer | From Multimedia File | Reverberator | To Multimedia File | Vector Concatenate, Matrix Concatenate

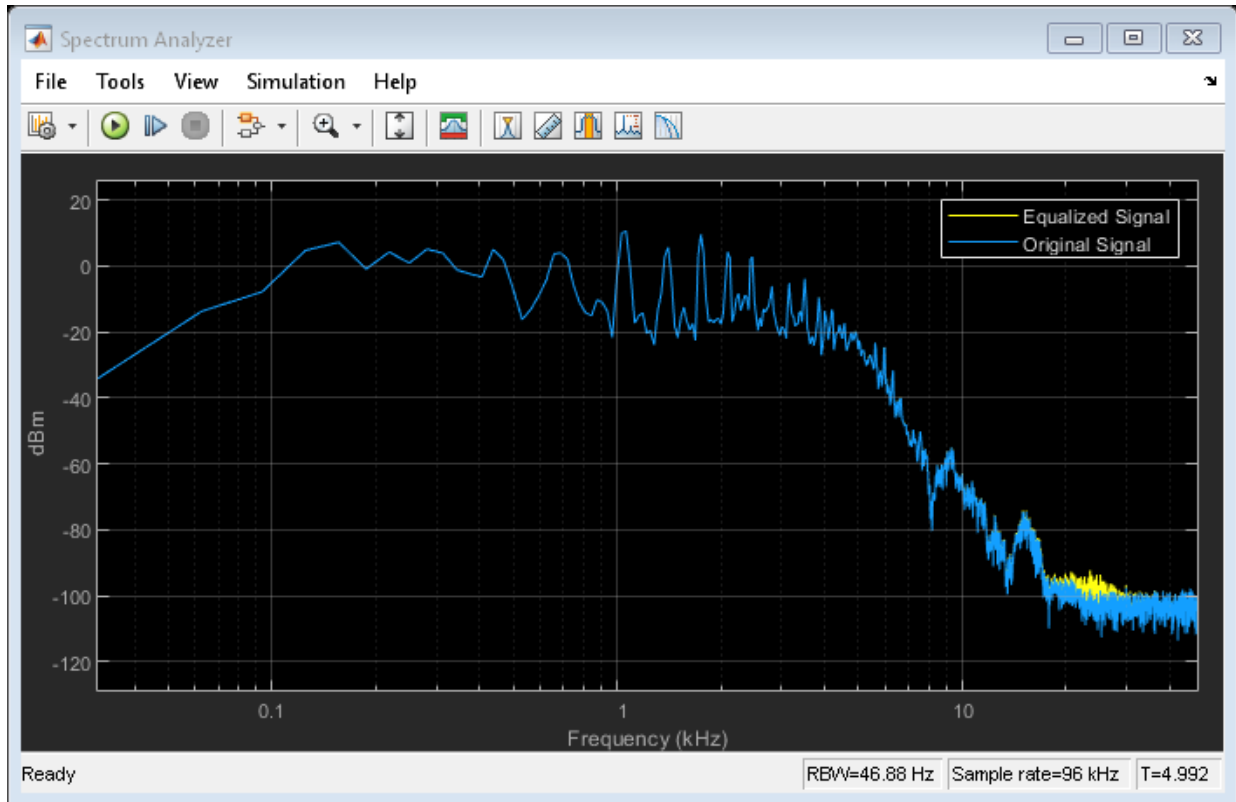
Perform Parametric Equalization

Examine the Parametric EQ block in a Simulink® model and tune parameters.



Copyright 2016-2019 The MathWorks, Inc.

1. Open the Spectrum Analyzer and Parametric EQ blocks.
2. In the Parametric EQ block, click **View Filter Response**. Modify parameters of the parametric equalizer and see the magnitude response plot update automatically.
3. Run the model. Tune parameters on the Parametric EQ to listen to the effect on your audio device and see the effect on the Spectrum Analyzer display. Double-click the Manual Switch block to toggle between the original and equalized signal as output.




See Also

[Audio Device Writer](#) | [From Multimedia File](#) | [Parametric EQ](#) | [Spectrum Analyzer](#) | [Vector Concatenate](#), [Matrix Concatenate](#)

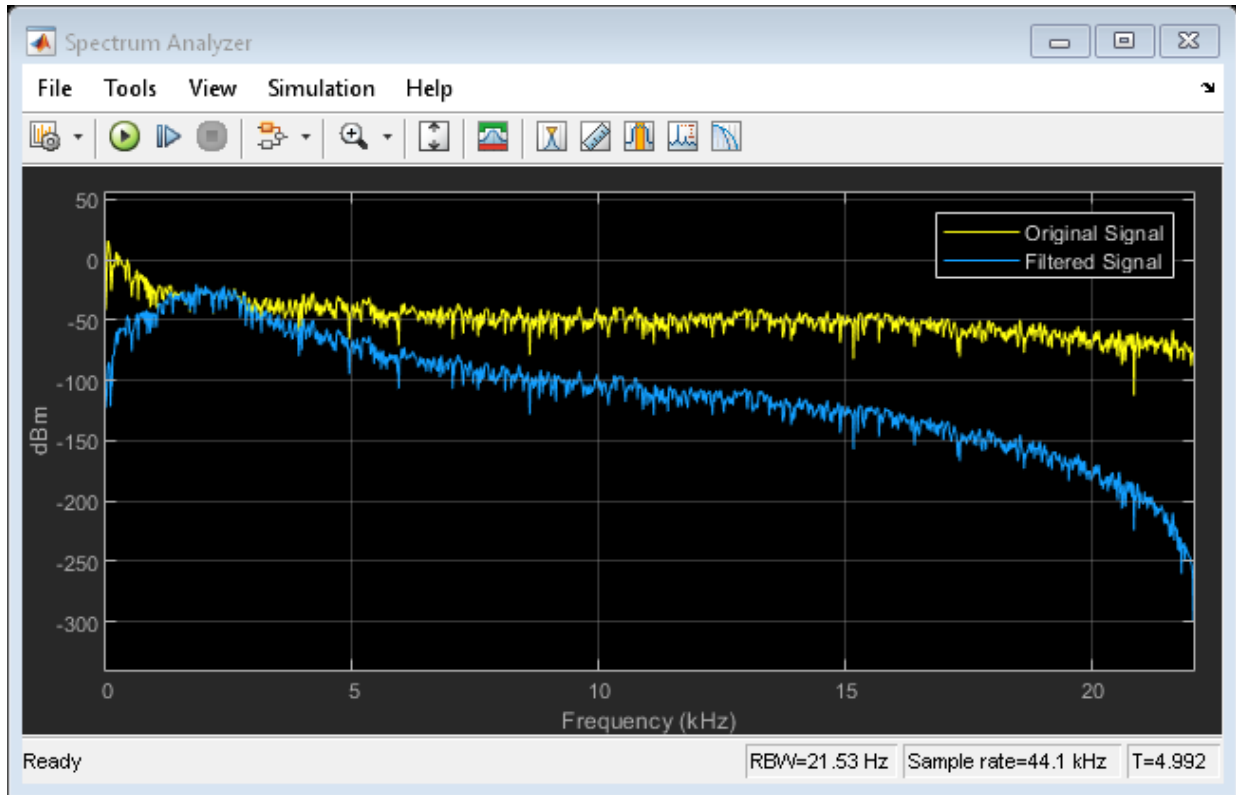
Perform Octave Filtering

Examine the Octave Filter block in a Simulink® model and tune parameters.



Copyright 2016 The MathWorks, Inc.

1. Open the Octave Filter block and click **Visualize filter response**. Tune parameters on the Octave Filter dialog. The filter response visualization updates automatically. If you break compliance with the ANSI S1.11-2004 standard, the filter mask is drawn in red.
2. Run the model. Open the Spectrum Analyzer block. Tune parameters on the Octave Filter block to listen to the effect on your audio device and see the effect on the Spectrum Analyzer display. Switch between listening to the filtered and unfiltered audio by double-clicking the Manual Switch block.




See Also

[Audio Device Writer](#) | [From Multimedia File](#) | [Octave Filter](#) | [Spectrum Analyzer](#)

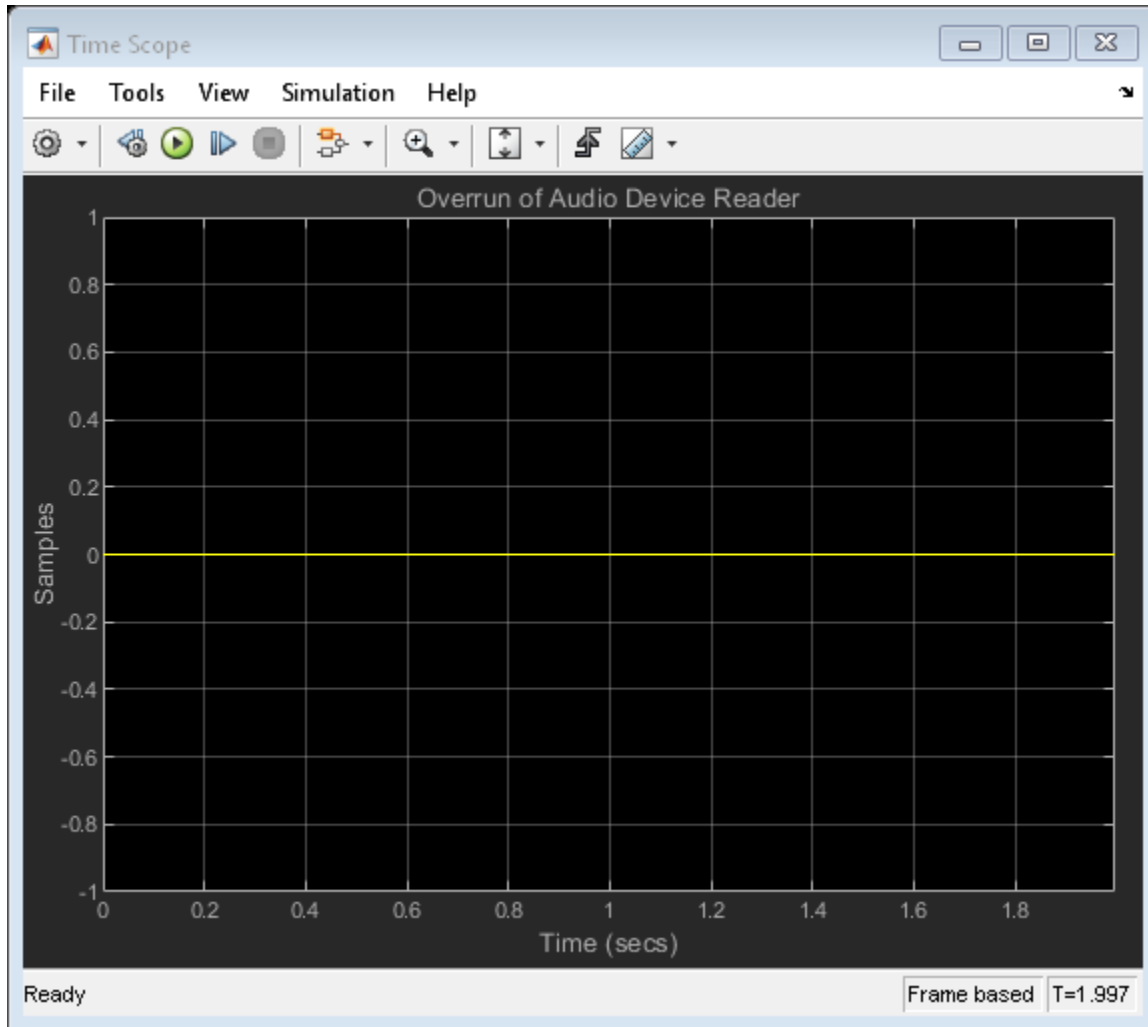
Read from Microphone and Write to Speaker

Examine the Audio Device Reader block in a Simulink® model, modify parameters, and explore overrun.



Copyright 2016 The MathWorks, Inc.

1. Run the model. The Audio Device Reader records an audio stream from your computer's default audio input device. The Reverberator block processes your input audio. The Audio Device Writer block sends the processed audio to your default audio output device.



2. Stop the model. Open the Audio Device Reader block and lower the **Samples per frame** parameter. Open the Time Scope block to view overrun.

3. Run the model again. Lowering the **Samples per frame** decreases the buffer size of your Audio Device Reader block. A smaller buffer size decreases audio latency while increasing the likelihood of overruns.

See Also


[Audio Device Reader](#) | [Audio Device Writer](#) | [Reverberator](#) | [Time Scope](#)

More About

- [“Audio I/O: Buffering, Latency, and Throughput”](#)

Channel Mapping

Examine the Audio Device Writer block in a Simulink® model and specify a nondefault channel mapping.



Copyright 2016 The MathWorks, Inc.

1. Run the simulation. The Audio Device Writer sends a stereo audio stream to your computer's default audio output device. If you are using a stereo audio output device, such as headphones, you can hear a tone from one speaker and noise from the other speaker.
2. Specify a nondefault channel mapping:
 - a. Stop the simulation.
 - b. Open the Audio Device Writer block to modify parameters.
 - c. On the **Advanced** tab, clear the **Use default channel mapping** parameter.

d. Specify the **Device output channels** in reverse order: [2, 1]. If you are using a stereo output device, such as headphones, you hear that the noise and tone have switched speakers.

See Also

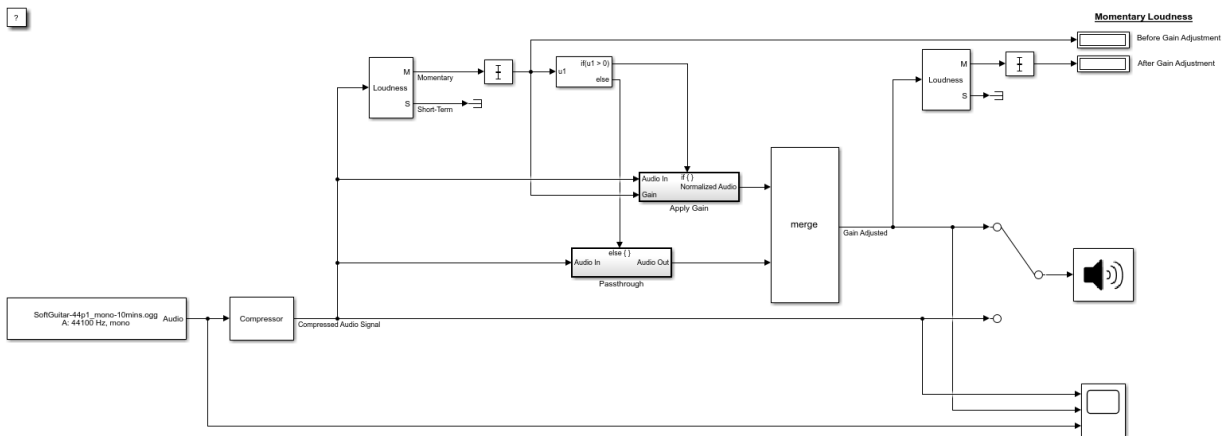
Audio Device Writer | Random Source | Sine Wave | Vector Concatenate, Matrix Concatenate

More About

- “Audio I/O: Buffering, Latency, and Throughput”

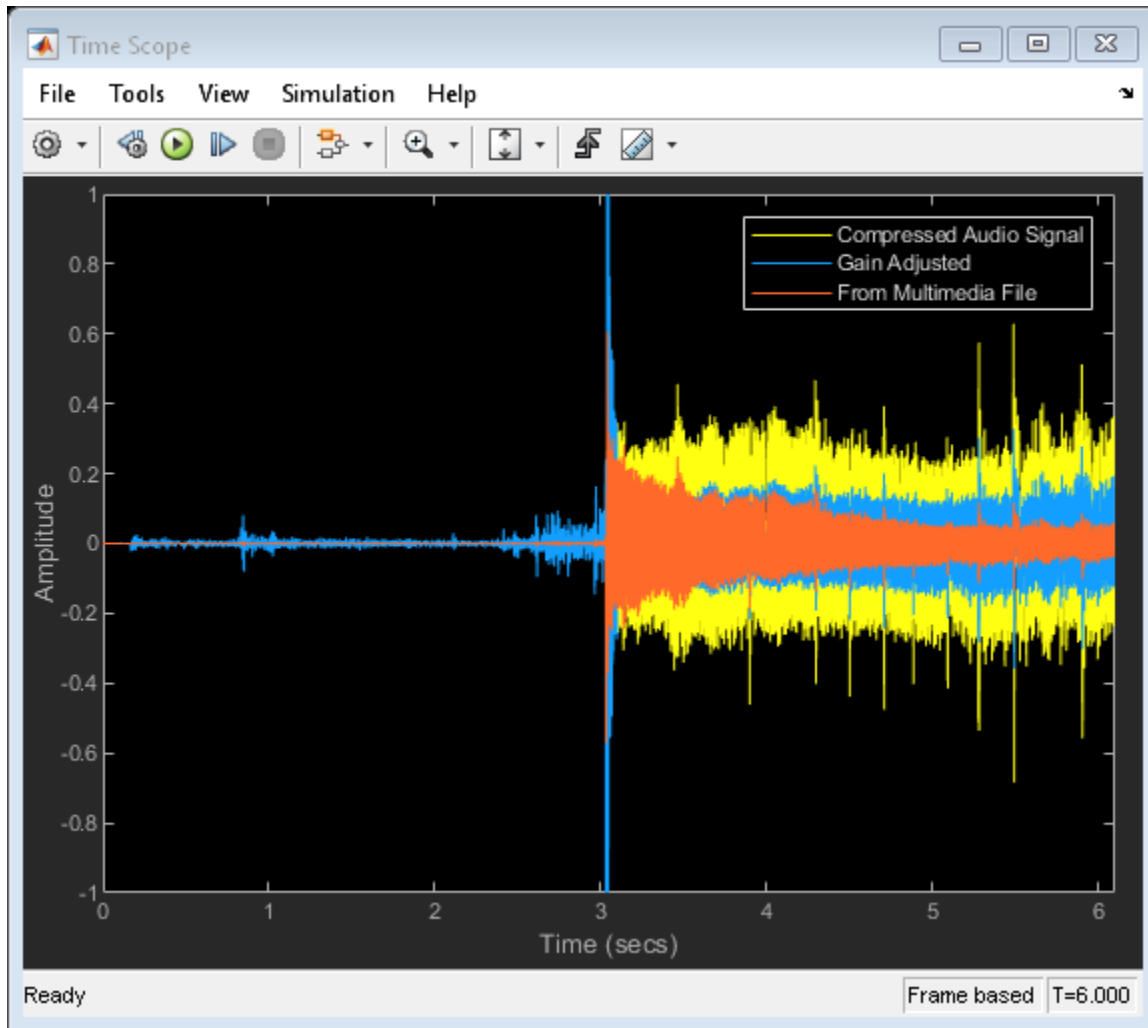
Trigger Gain Control Based on Loudness Measurement

This model enables you to apply dynamic range compression to an audio signal while staying inside a preset loudness range. In this model, a Compressor block increases the loudness and decreases the dynamic range of an audio signal. A Loudness Meter block calculates the momentary loudness of the compressed audio signal. If momentary loudness crosses a -23 LUFS threshold, an enabled subsystem applies gain to lower the corresponding level of the audio signal.



Copyright 2016 The MathWorks, Inc.

1. Open the Time Scope and Compressor blocks.
2. Run the model. To switch between listening to the compressed signal with and without gain adjustment, double-click the switch.
3. To observe the effect of compression on loudness, tune the Compressor block parameters and view the compressed audio signal on the Time Scope block.



See Also

Blocks

Audio Device Writer | Compressor | From Multimedia File | Loudness Meter | Time Scope

Objects

LoudnessMeter

Functions

integratedLoudness

More About

- “Loudness Normalization in Accordance with EBU R 128 Standard”

Real-Time Parameter Tuning

Real-Time Parameter Tuning

Parameter tuning is the ability to modify parameters of your audio system in real time while streaming an audio signal. In algorithm development, tunable parameters enable you to quickly prototype and test various parameter configurations. In deployed applications, tunable parameters enable users to fine-tune general algorithms for specific purposes, and to react to changing dynamics.

Audio Toolbox is optimized for parameter tuning in a real-time audio stream. The System objects, blocks, and audio plugins provide various tunable parameters, including sample rate and frame size, making them robust tools when used in an audio stream loop.

To optimize your use of Audio Toolbox, package your audio processing algorithm as an audio plugin. Packaging your audio algorithm as an audio plugin enables you to graphically tune your algorithm using `parameterTuner` or **Audio Test Bench**:

- **Audio Test Bench** -- Creates a user interface (UI) for tunable parameters, enables you to specify input and output from your audio stream loop, and provides access to analysis tools such as the time scope and spectrum analyzer. Packaging your code as an audio plugin also enables you to quickly synchronize your parameters with MIDI controls.
- `parameterTuner` -- Creates a UI for tunable parameters that can be used from any MATLAB programmatic environment. You can customize your parameter controls to render as knobs, sliders, rocker switches, toggle switches, check boxes, or drop-downs. You can also define a custom background color, background image, or both. You can then place your audio plugin in an audio processing loop in a programmatic environment such as a script, and then tune parameters while the loop executes.

For more information, see “Audio Plugins in MATLAB”.

Other methods to create UIs in MATLAB include:

- App Designer -- Development environment for a large set of interactive controls with support for 2-D plots. See “Create and Run a Simple App Using App Designer” (MATLAB) for more information.
- GUIDE -- Drag-and-drop environment for laying out user interfaces with support for any type of plot. See “Create a Simple App Using GUIDE” (MATLAB) for more information.
- Programmatic workflow -- Use MATLAB functions to define your app element-by-element. This tutorial uses a programmatic approach.

See “Ways to Build Apps” (MATLAB) for a more detailed list of the costs and benefits of the different approaches to parameter tuning.

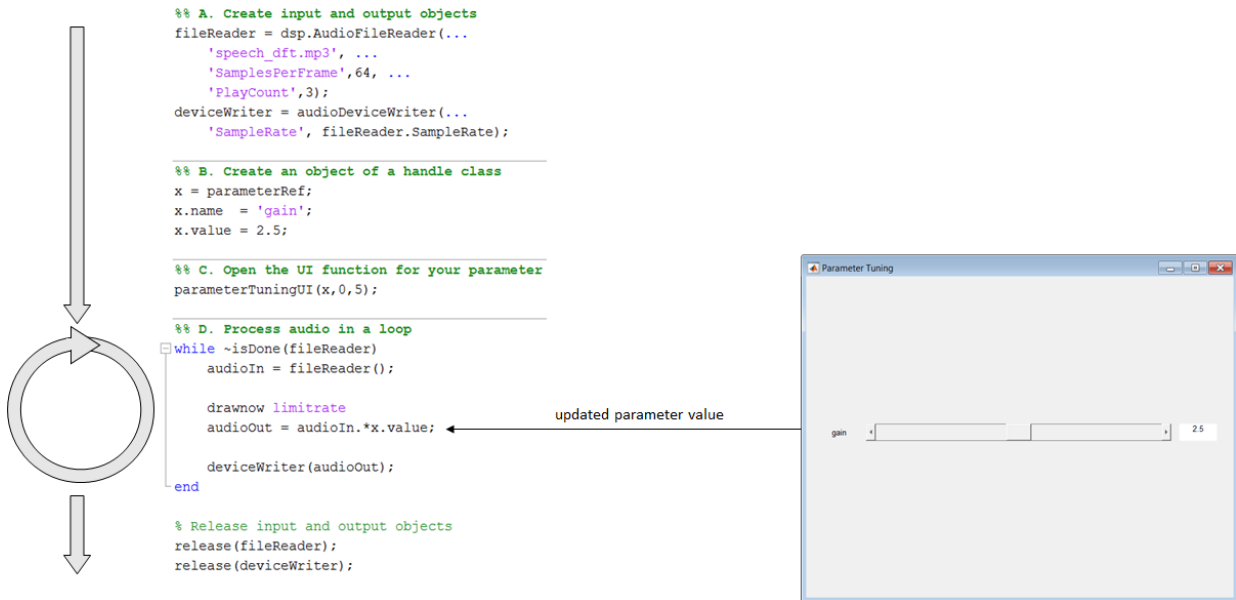
Programmatic Parameter Tuning

If you can not package your algorithm as an audio plugin, you can create a tuning UI using basic MATLAB techniques.

This tutorial contains three files:

- 1 parameterRef -- Class definition that contains tunable parameters
- 2 parameterTuningUI -- Function that creates a UI for parameter tuning
- 3 AudioProcessingScript -- Script for audio processing

Inspect the diagram for an overview of how real-time parameter tuning is implemented. To implement real-time parameter tuning, walk through the example for explanations and step-by-step instructions.



1. Create Class with Tunable Parameters

To tune a parameter in an audio stream loop using a UI, you need to associate the parameter with the position of a UI widget. To associate a parameter with a UI widget, make the parameter an object of a handle class. Objects of handle classes are passed by reference, meaning that you can modify the value of the object in one place and use the updated value in another. For example, you can modify the value of the object using a slider on a figure and use the updated value in an audio processing loop.

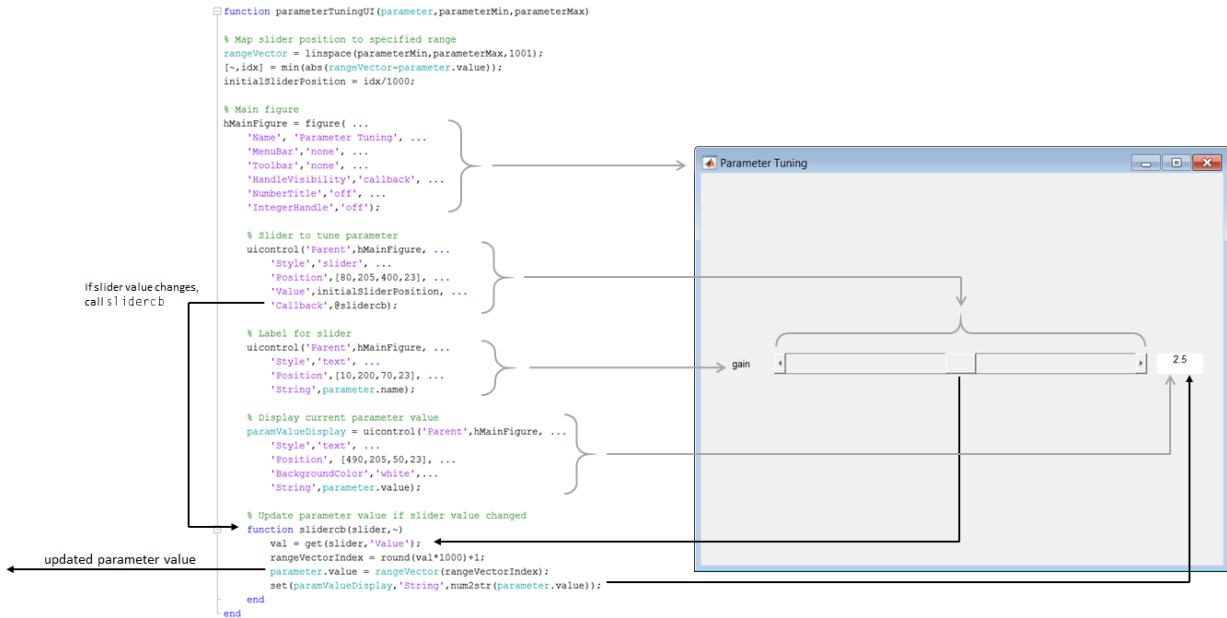
Save the `parameterRef` class definition file to your current folder.

```
classdef parameterRef < handle
    properties
        name
        value
    end
end
```

Objects of the `parameterRef` class have a `name` and `value`. The `name` is for display purposes on the UI. You use the `value` for tuning.

2. Create Function to Generate a UI

The `parameterTuningUI` function accepts your parameter, specified as an object handle, and the desired range. The function creates a figure with a slider associated with your parameter. The nested function, `slidercb`, is called whenever the slider position changes. The slider callback function maps the position of the slider to the parameter range, updates the value of the parameter, and updates the text on the UI. You can easily modify this function to tune multiple parameters in the same UI.



Save parameterTuningUI to Current Folder

Open parameterTuningUI and save it to your current folder.

```
function parameterTuningUI(parameter,parameterMin,parameterMax)
```

```
% Map slider position to specified range
rangeVector = linspace(parameterMin,parameterMax,1001);
[~,idx] = min(abs(rangeVector-parameter.value));
initialSliderPosition = idx/1000;
```

```
% Main figure
hMainFigure = figure( ...
    'Name', 'Parameter Tuning', ...
    'MenuBar', 'none', ...
    'ToolBar', 'none', ...
    'HandleVisibility', 'callback', ...
    'NumberTitle', 'off', ...
    'IntegerHandle', 'off');
```

```
% Slider to tune parameter
uicontrol('Parent',hMainFigure, ...
```

```
        'Style','slider', ...
        'Position',[80,205,400,23], ...
        'Value',initialSliderPosition, ...
        'Callback',@slidercb);

% Label for slider
uicontrol('Parent',hMainFigure, ...
        'Style','text', ...
        'Position',[10,200,70,23], ...
        'String',parameter.name);

% Display current parameter value
paramValueDisplay = uicontrol('Parent',hMainFigure, ...
        'Style','text', ...
        'Position', [490,205,50,23], ...
        'BackgroundColor','white', ...
        'String',parameter.value);

% Update parameter value if slider value changed
function slidercb(slider,~)
    val = get(slider,'Value');
    rangeVectorIndex = round(val*1000)+1;
    parameter.value = rangeVector(rangeVectorIndex);
    set(paramValueDisplay,'String',num2str(parameter.value));
end
end
```

3. Create Script for Audio Processing

The audio processing script:

- A** Creates input and output objects for an audio stream loop.
- B** Creates an object of the handle class, `parameterRef`, that stores your parameter name and value.
- C** Calls the tuning UI function, `parameterTuningUI`, with your parameter and the parameter range.
- D** Processes the audio in a loop. You can tune your parameter, `x`, in the audio stream loop.

Run AudioProcessingScript

Open `AudioProcessingScript`, save it to your current folder, and then run the file.

```
%% A. Create input and output objects
fileReader = dsp.AudioFileReader(...
    'speech_dft.mp3', ...
    'SamplesPerFrame',64, ...
    'PlayCount',3);
deviceWriter = audioDeviceWriter(...
    'SampleRate', fileReader.SampleRate);

%% B. Create an object of a handle class
x = parameterRef;
x.name = 'gain';
x.value = 2.5;

%% C. Open the UI function for your parameter
parameterTuningUI(x,0,5);

%% D. Process audio in a loop
while ~isDone(fileReader)
    audioIn = fileReader();

    drawnow limitrate
    audioOut = audioIn.*x.value;

    deviceWriter(audioOut);
end

% Release input and output objects
release(fileReader)
release(deviceWriter)
```

While the script runs, move the position of the slider to update your parameter value and hear the result.

See Also

Audio Test Bench | parameterTuner

More About

- “Real-Time Audio in MATLAB”
- “Audio Plugins in MATLAB”
- “Audio Test Bench Walkthrough” on page 10-2

- “Create and Run a Simple App Using App Designer” (MATLAB)
- “Create a Simple App Using GUIDE” (MATLAB)
- “Ways to Build Apps” (MATLAB)

Sample Audio Files

Sample Audio Files

Use these audio files as input to your audio system.

See Also

From Multimedia File | `audioread` | `dsp.AudioFileReader`

More About

- “Audio I/O: Buffering, Latency, and Throughput”

Tips and Tricks for Plugin Authoring

Tips and Tricks for Plugin Authoring

To author your algorithm as an audio plugin, you must conform to the audio plugin API. When authoring audio plugins in the MATLAB environment, keep these common pitfalls and best practices in mind.


To learn more about audio plugins in general, see “Audio Plugins in MATLAB”.

Avoid Disrupting the Event Queue in MATLAB

When the **Audio Test Bench** runs an audio plugin, it sequentially:

- 1 Calls the reset method
- 2 Sets tunable properties associated with parameters
- 3 Calls the process method

While running, the **Audio Test Bench** calls in a loop the process method and then the set methods for tuned properties. The plugin API does not specify the order that the tuned properties are set.



It is possible to disrupt the normal methods timing by interrupting the event queue. Common ways to accidentally interrupt the event queue include using a `plot` or `drawnow` function.

Note `plot` and `drawnow` are only available in the MATLAB environment. `plot` and `drawnow` cannot be included in generated plugins. See “Separate Code for Features Not Supported for Plugin Generation” on page 19-6 for more information.

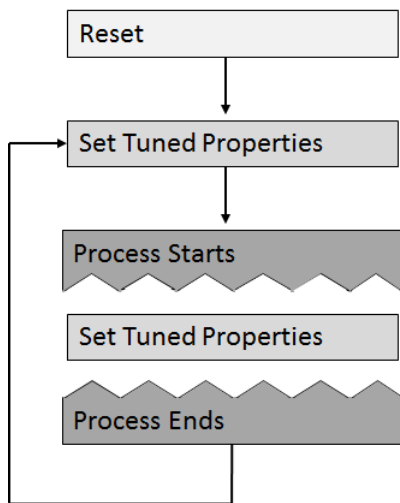
In the following code snippet, the gain applied to the left and right channels is not the same if the associated `Gain` parameter is tuned during the call to `process`:

```
...  
L = plugin.Gain*in(:,1);  
drawnow  
R = plugin.Gain*in(:,2);  
out = [L,R];  
...
```

See Full Code

```
classdef badPlugin < audioPlugin  
    properties  
        Gain = 0.5;  
    end  
    properties (Constant)  
        PluginInterface = audioPluginInterface(audioPluginParameter('Gain'));  
    end  
    methods  
        function out = process(plugin,in)  
  
            L = plugin.Gain*in(:,1);  
  
            drawnow  
  
            R = plugin.Gain*in(:,2);  
  
            out = [L,R];  
        end  
        function set.Gain(plugin,val)  
            plugin.Gain = val;  
        end  
    end  
end
```

The author interrupts the event queue in the code snippet, causing the `set` methods of properties associated with parameters to be called while the `process` method is in the middle of execution.



Depending on your processing algorithm, interrupting the event queue can lead to inconsistent and buggy behavior. Also, the `set` method might not be explicit, which can make the issue difficult to track down. Possible fixes for the problem of event queue disruption include saving properties to local variables, and moving the queue disruption to the beginning or end of the process method.

Save Properties to Local Variables

You can save tunable property values to local variables at the start of your processing. This technique ensures that the values used during the process method are not updated within a single call to process. Because accessing the value of a local variable is cheaper than accessing the value of a property, saving properties to local variables that are accessed multiple times is a best practice.

```

...
gain = plugin.Gain;
L = gain*in(:,1);
drawnow
R = gain*in(:,2);
out = [L,R];
...

```

See Full Code

```

classdef goodPlugin < audioPlugin
    properties

```

```

    Gain = 0.5;
end
properties (Constant)
    PluginInterface = audioPluginInterface(audioPluginParameter('Gain'));
end
methods
    function out = process(plugin,in)
        gain = plugin.Gain;

        L = gain*in(:,1);

        drawnow

        R = gain*in(:,2);

        out = [L,R];
    end
    function set.Gain(plugin,val)
        plugin.Gain = val;
    end
end
end
end

```

Move Queue Disruption to Bottom or Top of Process Method

You can move the disruption to the event queue to the bottom or top of the process method. This technique ensures that property values are not updated in the middle of the call.

```

...
L = plugin.Gain*in(:,1);
R = plugin.Gain*in(:,2);
out = [L,R];
drawnow
...

```

See Full Code

```

classdef goodPlugin < audioPlugin
    properties
        Gain = 0.5;
    end
    properties (Constant)
        PluginInterface = audioPluginInterface(audioPluginParameter('Gain'));
    end
    methods
        function out = process(plugin,in)

            L = plugin.Gain*in(:,1);

            R = plugin.Gain*in(:,2);

            out = [L,R];

            drawnow
        end
        function set.Gain(plugin,val)

```

```
        plugin.Gain = val;
    end
end
end
```

Separate Code for Features Not Supported for Plugin Generation

The MATLAB environment offers functionality not supported for plugin generation. You can mark code to ignore during plugin generation by placing it inside a conditional statement by using `coder.target`.

```
...
    if coder.target('MATLAB')
        ...
    end
...

```

If you generate the plugin using `generateAudioPlugin`, code inside the statement `if coder.target('MATLAB')` is ignored.

For example, `dsp.TimeScope` is not enabled for code generation. If you run the following plugin in MATLAB, you can use the `visualize` function to open a time scope that plots the input and output power per frame.

See Full Example Code

```
classdef pluginWithMATLABOnlyFeatures < audioPlugin
    properties
        Threshold = -10;
    end
    properties (Access = private)
        aCompressor
        aScope
        SamplesPerFrame = 1;
    end
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('Threshold','Mapping',{'lin',-60,20}));
    end
    methods
        function plugin = pluginWithMATLABOnlyFeatures
            plugin.aCompressor = compressor;
            setup(plugin.aCompressor,[0,0])
        end
        function out = process(plugin,in)
            out = plugin.aCompressor(in);

            % The contents of this if-statement are ignored during plugin
            % generation.
            if coder.target('MATLAB')
                if ~isempty(plugin.aScope) && isvalid(plugin.aScope)
```

```

        numSamples = size(in,1);

        % The time scope object is not enabled for
        % variable-size signals. Call release if the samples
        % per frame is changed.
        % Because this code is intended for use in MATLAB only,
        % it is okay to call release on the time scope object.
        % Do not call release on a System object in generated
        % code.
        if plugin.SamplesPerFrame(1) ~= numSamples
            release(plugin.aScope)
            plugin.SamplesPerFrame = numSamples;
        end

        power = 20*log10(mean(var(in)))*ones(numSamples,1);
        adjustedPower = 20*log10(mean(var(out)))*ones(numSamples,1);
        plugin.aScope([power,adjustedPower]);
    end
end
end
function reset(plugin)
    fs = getSampleRate(plugin);
    plugin.aCompressor.SampleRate = fs;
    reset(plugin.aCompressor)

    % The contents of this if-statement are ignored during plugin
    % generation.
    if coder.target('MATLAB')
        if ~isempty(plugin.aScope)
            % Because this code is intended for use in MATLAB only,
            % it is okay to call release on the time scope object.
            % Do not call release on a System object in generated
            % code.
            release(plugin.aScope)
            plugin.aScope.SampleRate = fs;
            plugin.aScope.BufferLength = 2*fs;
        end
    end
end
function visualize(plugin)
    % Visualization function. This function is public in the MATLAB
    % environment. Because the plugin does not call this function
    % directly, the function is not part of the code generated by
    % generateAudioPlugin.

    % Create a time scope object for visualization in the MATLAB
    % environment.
    plugin.aScope = dsp.TimeScope( ...
        'SampleRate',getSampleRate(plugin), ...
        'TimeSpan',1, ...
        'YLimits',[-40,0], ...
        'BufferLength',2*getSampleRate(plugin), ...
        'TimeSpanOverrunAction','Scroll', ...
        'YLabel','Power (dB)');
    show(plugin.aScope)
end
function set.Threshold(plugin,val)
    plugin.Threshold = val;
    plugin.aCompressor.Threshold = val;
end
end
end
end

```

Implement Reset Correctly

A common error in audio plugin authoring is misusing the reset method. Valid uses of the reset method include:

- Clearing state
- Passing down calls to reset to component objects
- Updating properties which depend on sample rate

Invalid use of the reset method includes setting the value of any properties associated with parameters. Do not use your reset method to set properties associated with parameters to their initial conditions. Directly setting a property associated with a parameter causes the property to be out of sync with the parameter. For example, the following plugin is an example of incorrect use of the reset method.

```
classdef badReset < audioPlugin
    properties
        Gain = 1;
    end
    properties (Constant)
        PluginInterface = audioPluginInterface(audioPluginParameter('Gain'));
    end
    methods
        function out = process(plugin,in)
            out = in*plugin.Gain;
        end
        function reset(plugin) % <-- Incorrect use of reset method.
            plugin.Gain = 1; % <-- Never set values of a property that is
                            % associated with a plugin parameter.
        end
    end
end
```

Implement Plugin Composition Correctly

If your plugin is composed of other plugins, then you must pass down the sample rate and calls to reset to the component plugins. Call `setSampleRate` in the reset method to pass down the sample rate to the component plugins. To tune parameters of the component plugins, create an audio plugin interface in the composite plugin for tunable parameters of the component plugins. Then pass down the values in the `set` methods for the associated properties. The following is an example of plugin composition that was constructed using best practices.

Plugin Composition Using Basic Plugins

```
classdef compositePlugin < audioPlugin
    properties
        PhaserQ = 1.6;
        EchoGain = 0.5;
    end
end
```



```

end
properties (Access = private)
    aEcho
    aPhaser
end
properties (Constant)
    PluginInterface = audioPluginInterface( ...
        audioPluginParameter('PhaserQ', ...
            'DisplayName','Phaser Q', ...
            'Mapping',{'lin',0.5, 25}), ...
        audioPluginParameter('EchoGain', ...
            'DisplayName','Gain'));
end
methods
function plugin = compositePlugin
    % Construct your component plugins in the composite plugin's
    % constructor.
    plugin.aPhaser = audiopluginexample.Phaser;
    plugin.aEcho   = audiopluginexample.Echo;
end
function out = process(plugin,in)
    % Call the process method of your component plugins inside the
    % call to the process method of your composite plugin.
    x = process(plugin.aPhaser,in);
    y = process(plugin.aEcho,x);
    out = y;
end
function reset(plugin)
    % Use the setSampleRate method to set the sample rate of
    % component plugins and pass the call to reset down.
    fs = getSampleRate(plugin);

    setSampleRate(plugin.aPhaser, fs)
    setSampleRate(plugin.aEcho, fs)

    reset(plugin.aPhaser)
    reset(plugin.aEcho);
end
% Use the set method of your properties to pass down property
% values to your component plugins.
function set.PhaserQ(plugin,val)
    plugin.PhaserQ = val;
    plugin.aPhaser.QualityFactor = val;
end
function set.EchoGain(plugin,val)
    plugin.EchoGain = val;
    plugin.aEcho.Gain = val;
end
end
end
end

```

Plugin composition using System objects has these key differences from plugin composition using basic plugins.

- Immediately call `setup` on your component System object after it is constructed. Construction and setup of the component object occurs inside the constructor of the composite plugin.

- If your component System object requires sample rate information, then it has a sample rate property. Set the sample rate property in the reset method.

Plugin Composition Using System Objects

```

classdef compositePluginWithSystemObjects < audioPlugin
    properties
        CrossoverFrequency = 100;
        CompressorThreshold = -40;
    end
    properties (Access = private)
        aCrossoverFilter
        aCompressor
    end
    properties (Constant)
        PluginInterface = audioPluginInterface( ...
            audioPluginParameter('CrossoverFrequency', ...
                'DisplayName','Crossover Frequency', ...
                'Mapping',{'lin',50, 200}), ...
            audioPluginParameter('CompressorThreshold', ...
                'DisplayName','Compressor Threshold', ...
                'Mapping',{'lin',-100,0}));
    end
    methods
        function plugin = compositePluginWithSystemObjects
            % Construct your component System objects within the composite
            % plugin's constructor. Call setup immediately after
            % construction.
            %
            % The audio plugin API requires plugins to declare the number
            % of input and output channels in the plugin interface. This
            % plugin uses the default 2-in 2-out configuration. Call setup
            % with a sample input that has the same number of channels as
            % defined in the plugin interface.
            %
            sampleInput = zeros(1,2);

            plugin.aCrossoverFilter = crossoverFilter;
            setup(plugin.aCrossoverFilter,sampleInput)

            plugin.aCompressor = compressor;
            setup(plugin.aCompressor,sampleInput)
        end
        function out = process(plugin,in)
            % Call your component System objects inside the call to
            % process of your composite plugin.
            [band1,band2] = plugin.aCrossoverFilter(in);
            band1Compressed = plugin.aCompressor(band1);
            out = band1Compressed + band2;
        end
        function reset(plugin)
            % Set the sample rate properties of your component System
            % objects.
            fs = getSampleRate(plugin);

            plugin.aCrossoverFilter.SampleRate = fs;
            plugin.aCompressor.SampleRate = fs;

            reset(plugin.aCrossoverFilter)
            reset(plugin.aCompressor);
        end
    end
end

```

```

% Use the set method of your properties to pass down property
% values to your component System objects.
function set.CrossoverFrequency(plugin,val)
    plugin.CrossoverFrequency = val;
    plugin.aCrossoverFilter.CrossoverFrequencies = val;
end
function set.CompressorThreshold(plugin,val)
    plugin.CompressorThreshold = val;
    plugin.aCompressor.Threshold = val;
end
end
end

```

Address "A set method for a non-Dependent property should not access another property" Warning in Plugin

It is recommended that you suppress the warning when authoring audio plugins.

The following code snippet follows the plugin authoring best practice for processing changes in parameter property Cutoff.

```

...
properties (Constant)
    PluginInterface = audioPluginInterface( ...
        audioPluginParameter('Cutoff', ...
            'Label','Hz',...
            'Mapping',{'log',20,2000}));
end
methods
    function y = process(plugin,x)
        [y,plugin.State] = filter(plugin.B,plugin.A,x,plugin.State);
    end

    function set.Cutoff(plugin,val)
        plugin.Cutoff = val;
        [plugin.B,plugin.A] = highpassCoeffs(plugin,val,getSampleRate(plugin)); % <<<< warning occurs here
    end
end
...

```

See Full Code Example

```

classdef highpassFilter < audioPlugin
    %-----
    % Public Properties - End user interacts with these
    %-----
    properties
        Cutoff = 20;
    end

    %-----
    % Private Properties - Used for internal storage
    %-----
    properties (Access = private)
        State = zeros(2);
        B = zeros(1,3);
        A = zeros(1,3);
    end
end

```

```

end

%-----
% Constant Properties - Used to define plugin interface
%-----
properties (Constant)
    PluginInterface = audioPluginInterface( ...
        audioPluginParameter( 'Cutoff', ...
            'Label','Hz', ...
            'Mapping',{'log',20,2000});
end

methods

%-----
% Main processing function
%-----
function y = process(plugin,x)
    [y,plugin.State] = filter(plugin.B,plugin.A,x,plugin.State);
end

%-----
% Set Method
%-----
function set.Cutoff(plugin,val)
    plugin.Cutoff = val;
    [plugin.B,plugin.A] = highpassCoeffs(plugin,val,getSampleRate(plugin)); % <<<< warning occurs here
end

%-----
% Reset Method
%-----
function reset(plugin)
    plugin.State = zeros(2);
    [plugin.B,plugin.A] = highpassCoeffs(plugin,plugin.Cutoff,getSampleRate(plugin));
end
end
methods (Access = private)
%-----
% Calculate Filter Coefficients
%-----
function [B,A] = highpassCoeffs(~,fc,fs)
    w0 = 2*pi*fc/fs;
    alpha = sin(w0)/sqrt(2);
    cosw0 = cos(w0);
    norm = 1/(1+alpha);
    B = (1 + cosw0)*norm * [.5 -1 .5];
    A = [1 -2*cosw0*norm (1 - alpha)*norm];
end
end
end
end

```

The `highpassCoeffs` function might be expensive, and should be called only when necessary. You do not want to call `highpassCoeffs` in the `process` method, which runs in the real-time audio processing loop. The logical place to call `highpassCoeffs` is in `set.Cutoff`. However, `mlint` shows a warning for this practice. The warning is intended to help you avoid initialization order issues when saving and loading classes. See “Avoid Property Initialization Order Dependency” (MATLAB) for more details. The solution recommended by the warning is to create a dependent property with a `get` method and

compute the value there. However, following the recommendation complicates the design and pushes the computation back into the real-time processing method, which you are trying to avoid.

You might also incur the warning when correctly implementing plugin composition. For an example of a correct implementation of composition, see “Implement Plugin Composition Correctly” on page 19-8.

Use System Object That Does Not Support Variable-Size Signals

The audio plugin API requires audio plugins to support variable-size inputs and outputs. For a partial list of System objects that support variable-size signals, see “Variable-Size Signal Support DSP System Objects” (DSP System Toolbox). You might encounter issues if you attempt to use objects that do not support variable-size signals in your plugin.

For example, `dsp.AnalyticSignal` does not support variable-size signals. The `BrokenAnalyticSignalTransformer` plugin uses a `dsp.AnalyticSignal` object incorrectly and fails the `validateAudioPlugin` test bench:

```
validateAudioPlugin BrokenAnalyticSignalTransformer
```

```
Checking plug-in class 'BrokenAnalyticSignalTransformer'... passed.
Generating testbench file 'testbench_BrokenAnalyticSignalTransformer.m'... done.
Running testbench...
Error using dsp.AnalyticSignal/parenReference
Changing the size on input 1 is not allowed without first calling the release() method.

Error in BrokenAnalyticSignalTransformer/process (line 13)
    analyticSignal = plugin.Transformer(in);

Error in testbench_BrokenAnalyticSignalTransformer (line 61)
    o1 = process(plugin, in(:,1));

Error in validateAudioPlugin
```

See BrokenAnalyticSignalTransformer Code

```
classdef BrokenAnalyticSignalTransformer < audioPlugin
    properties (Access = private)
        Transformer
    end
    properties (Constant)
        PluginInterface = audioPluginInterface('InputChannels',1,'OutputChannels',2);
    end
    methods
        function plugin = BrokenAnalyticSignalTransformer
            plugin.Transformer = dsp.AnalyticSignal;
        end
    end
end
```

```


function out = process(plugin,in)
    analyticSignal = plugin.Transformer(in);
    realPart = real(analyticSignal);
    imaginaryPart = imag(analyticSignal);
    out = [realPart,imaginaryPart];
end
end
end

```

If you want to use the functionality of a System object that does not support variable-size signals, you can buffer the input and output of the System object, or always call the object with one sample.

Always Call the Object with One Sample

You can create a loop around your call to an object. The loop iterates for the number of samples in your variable frame size. The call to the object inside the loop is always a single sample.



See Full Code Example

```

classdef ExpensiveAnalyticSignalTransformer < audioPlugin
    properties (Access = private)
        Transformer
    end
    properties (Constant)
        PluginInterface = audioPluginInterface('InputChannels',1,'OutputChannels',2);
    end
    methods
        function plugin = ExpensiveAnalyticSignalTransformer
            plugin.Transformer = dsp.AnalyticSignal;
        end
        function out = process(plugin,in)
            analyticSignal = complex(zeros(size(in,1),1),0);
            for i = 1:size(in,1)
                analyticSignal(i,:) = plugin.Transformer(in(i,1));
            end
        end
    end
end

```

```

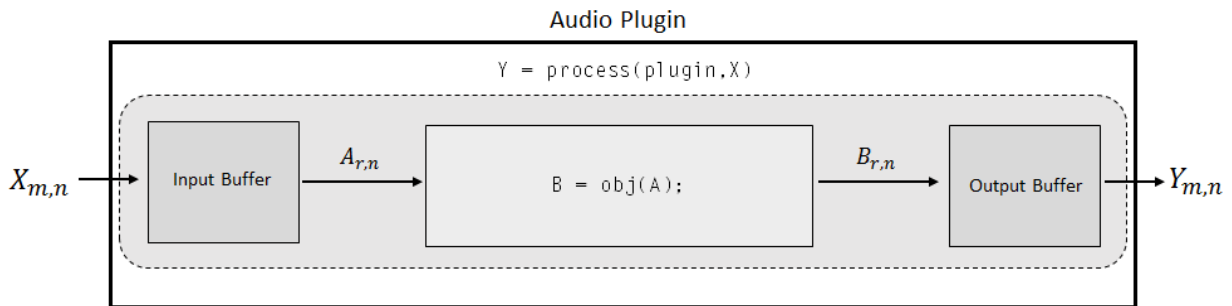
        end
        out = [real(analyticSignal), imag(analyticSignal)];
    end
end
end

```

Note Depending on your implementation and the particular object, calling an object sample by sample in a loop might result in significant computational cost.

Buffer Input and Output of Object

You can buffer the input to your object to a consistent frame size, and then buffer the output of your object back to the original frame size. The `dsp.AsyncBuffer` System object is well-suited for this task.



See Full Code Example

```

classdef DelayedAnalyticSignalTransformer < audioPlugin
    properties (Access = private)
        Transformer
        InputBuffer
        OutputBuffer
    end
    properties (Constant)
        PluginInterface = audioPluginInterface('InputChannels',1,'OutputChannels',2);
        MinSampleDelay = 256;
    end
    methods
        function plugin = DelayedAnalyticSignalTransformer
            plugin.Transformer = dsp.AnalyticSignal;
            setup(plugin.Transformer, ones(plugin.MinSampleDelay,1));

            plugin.InputBuffer = dsp.AsyncBuffer;
            setup(plugin.InputBuffer, 1);

            plugin.OutputBuffer = dsp.AsyncBuffer;
            setup(plugin.OutputBuffer, [1,1]);
        end
    end
end

```

```
function out = process(plugin,in)
    write(plugin.InputBuffer,in);

    while plugin.InputBuffer.NumUnreadSamples >= plugin.MinSampleDelay
        x = read(plugin.InputBuffer,plugin.MinSampleDelay);
        analyticSignal = plugin.Transformer(x(1:plugin.MinSampleDelay,:));
        write(plugin.OutputBuffer,[real(analyticSignal),imag(analyticSignal)]);
    end

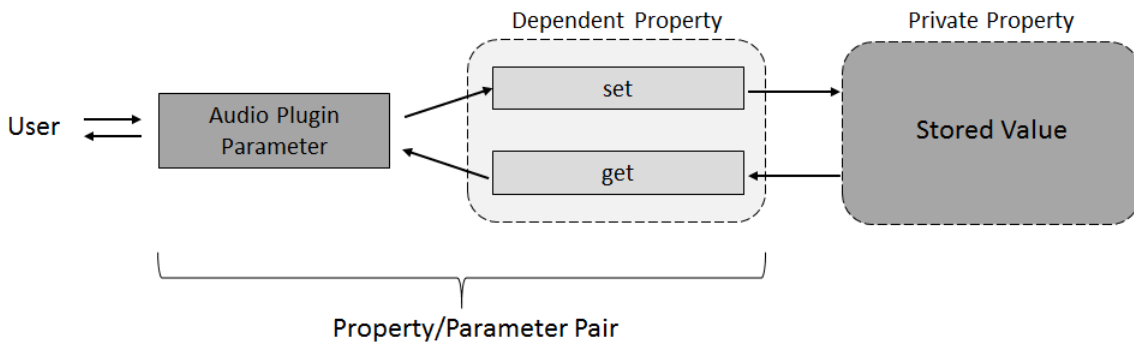
    if plugin.OutputBuffer.NumUnreadSamples >= size(in,1)
        out = read(plugin.OutputBuffer,size(in,1));
    else
        out = zeros(size(in,1),2);
    end
end
function reset(plugin)
    reset(plugin.Transformer)
    reset(plugin.InputBuffer)
    reset(plugin.OutputBuffer)
end
end
end
```

Note Use of the asynchronous buffering object forces a minimum latency of your specified frame size.

Using Enumeration Parameter Mapping

It is often useful to associate a property with a set of strings or character vectors. However, restrictions on plugin generation require cached values, such as property values, to have a static size. To work around this issue, you can use a separate enumeration class that maps the strings to the enumerations, as described in the `audioPluginParameter` documentation.

Alternatively, if you want to avoid writing an enumeration class and keep all your code in one file, you can use a dependent property to map your parameter names to a set of values. In this scenario, you map your enumeration value to a value that you can cache.



See Full Code Example

```

classdef pluginWithEnumMapping < audioPlugin
    properties (Dependent)
        Mode = '+6 dB';
    end
    properties (Access = private)
        pMode = 1; % '+6 dB'
    end
    properties (Constant)
        PluginInterface = audioPluginInterface(...
            audioPluginParameter('Mode',...
                'Mapping',{ 'enum', '+6 dB', '-6 dB', 'silence', 'white noise'}));
    end
    methods
        function out = process(plugin,in)
            switch (plugin.pMode)
                case 1
                    out = in * 2;
                case 2
                    out = in / 2;
                case 3
                    out = zeros(size(in));
                otherwise % case 4
                    out = rand(size(in)) - 0.5;
            end
        end
        function set.Mode(plugin,val)
            validatestring(val,{' +6 dB', '-6 dB', 'silence', 'white noise'},'set.Mode','Mode');
            switch val
                case '+6 dB'
                    plugin.pMode = 1;
                case '-6 dB'
                    plugin.pMode = 2;
                case 'silence'
                    plugin.pMode = 3;
                otherwise % 'white noise'
                    plugin.pMode = 4;
            end
        end
        function out = get.Mode(plugin)
            switch plugin.pMode
                case 1
    
```

```
        out = '+6 dB';  
    case 2  
        out = '-6 dB';  
    case 3  
        out = 'silence';  
    otherwise % case 4  
        out = 'white noise';  
    end  
end  
end  
end
```

See Also

More About

- “Audio Plugins in MATLAB”
- “Audio Plugin Example Gallery” on page 11-2
- “Export a MATLAB Plugin to a DAW”

Spectral Descriptors Chapter

Spectral Descriptors

Audio Toolbox™ provides a suite of functions that describe the shape, sometimes referred to as *timbre*, of audio. This example defines the equations used to determine the spectral features, cites common uses of each feature, and provides examples so that you can gain intuition about what the spectral descriptors are describing.

Spectral descriptors are widely used in machine and deep learning applications, and perceptual analysis. Spectral descriptors have been applied to a range of applications, including:

- Speaker identification and recognition [21 on page 20-0]
- Acoustic scene recognition [11 on page 20-0] [17 on page 20-0]
- Instrument recognition [22 on page 20-0]
- Music genre classification [16 on page 20-0] [18 on page 20-0]
- Mood recognition [19 on page 20-0] [20 on page 20-0]
- Voice activity detection [5 on page 20-0] [7 on page 20-0] [8 on page 20-0] [10 on page 20-0] [12 on page 20-0] [13 on page 20-0]

Spectral Centroid

The spectral centroid (`spectralCentroid`) is the frequency-weighted sum normalized by the unweighted sum [1 on page 20-0]:

$$\mu_1 = \frac{\sum_{k=b_1}^{b_2} f_k s_k}{\sum_{k=b_1}^{b_2} s_k}$$

where

- f_k is the frequency in Hz corresponding to bin k .
- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral centroid.

The spectral centroid represents the "center of gravity" of the spectrum. It is used as an indication of *brightness* [2 on page 20-0] and is commonly used in music analysis and


genre classification. For example, observe the jumps in the centroid corresponding to high hat hits in the audio file.

```
[audio,fs] = audioread('FunkyDrums-44p1-stereo-25secs.mp3');  
audio = sum(audio,2)/2;
```

```
centroid = spectralCentroid(audio,fs);
```

```
subplot(2,1,1)  
t = linspace(0,size(audio,1)/fs,size(audio,1));  
plot(t,audio)  
ylabel('Amplitude')
```

```
subplot(2,1,2)  
t = linspace(0,size(audio,1)/fs,size(audio,1));  
plot(t,centroid)  
xlabel('Time (s)')  
ylabel('Centroid (Hz)')
```




The spectral centroid is also commonly used to classify speech as voiced or unvoiced [3 on page 20-0]. For example, the centroid jumps in regions of unvoiced speech.

```
[audio,fs] = audioread('Counting-16-44p1-mono-15secs.wav');
centroid = spectralCentroid(audio,fs);

subplot(2,1,1)
t = linspace(0,size(audio,1)/fs,size(audio,1));
plot(t,audio)
ylabel('Amplitude')

subplot(2,1,2)
t = linspace(0,size(audio,1)/fs,size(centroid,1));
```

```
plot(t,centroid)
xlabel('Time (s)')
ylabel('Centroid (Hz)')
```



Spectral Spread

Spectral spread (`spectralSpread`) is the standard deviation around the spectral centroid [1 on page 20-0]:

$$\mu_2 = \sqrt{\frac{\sum_{k=b_1}^{b_2} (f_k - \mu_1)^2 s_k}{\sum_{k=b_1}^{b_2} s_k}}$$

where


- f_k is the frequency in Hz corresponding to bin k .
- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral spread.
- μ_1 is the spectral centroid.

The spectral spread represents the "instantaneous bandwidth" of the spectrum. It is used as an indication of the dominance of a tone. For example, the spread increases as the tones diverge and decreases as the tones converge.

```
fs = 16e3;
tone = audioOscillator('SampleRate',fs,'NumTones',2,'SamplesPerFrame',512,'Frequency',
duration = 5;
numLoops = floor(duration*fs/tone.SamplesPerFrame);
signal = [];
for i = 1:numLoops
    signal = [signal;tone()];
    if i<numLoops/2
        tone.Frequency = tone.Frequency + [0,50];
    else
        tone.Frequency = tone.Frequency - [0,50];
    end
end
spread = spectralSpread(signal,fs);

subplot(2,1,1)
spectrogram(signal,round(fs*0.05),round(fs*0.04),2048,fs,'yaxis')

subplot(2,1,2)
t = linspace(0,size(signal,1)/fs,size(spread,1));
plot(t,spread)
xlabel('Time (s)')
ylabel('Spread')
```

Spectral Skewness

Spectral skewness (`spectralSkewness`) is computed from the third order moment [1 on page 20-0]:

$$\mu_3 = \frac{\sum_{k=b_1}^{b_2} (f_k - \mu_1)^3 s_k}{(\mu_2)^3 \sum_{k=b_1}^{b_2} s_k}$$

where

- f_k is the frequency in Hz corresponding to bin k .

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral skewness.
- μ_1 is the spectral centroid.
- μ_2 is the spectral spread.

The spectral skewness measures symmetry around the centroid. In phonetics, spectral skewness is often referred to as *spectral tilt* and is used with other spectral moments to distinguish the place of articulation [4 on page 20-0]. For harmonic signals, it indicates the relative strength of higher and lower harmonics. For example, in the four-tone signal, there is a positive skew when the lower tone is dominant and a negative skew when the upper tone is dominant.

```

fs = 16e3;
duration = 99;
tone = audioOscillator('SampleRate',fs,'NumTones',4,'SamplesPerFrame',fs,'Frequency',[
    1000,2000,3000,4000]);


signal = [];
for i = 1:duration
    signal = [signal;tone()];
    tone.Amplitude = tone.Amplitude + [0.01,0,0,-0.01];
end

skewness = spectralSkewness(signal,fs);
t = linspace(0,size(signal,1)/fs,size(skewness,1))/60;

subplot(2,1,1)
spectrogram(signal,round(fs*0.05),round(fs*0.04),round(fs*0.05),fs,'yaxis','power')
view([-58 33])

subplot(2,1,2)
plot(t,skewness)
xlabel('Time (minutes)')
ylabel('Skewness')

```



Spectral Kurtosis

Spectral kurtosis (`spectralKurtosis`) is computed from the fourth order moment [1 on page 20-0]:

$$\mu_4 = \frac{\sum_{k=b_1}^{b_2} (f_k - \mu_1)^4 s_k}{(\mu_2)^4 \sum_{k=b_1}^{b_2} s_k}$$

where

- f_k is the frequency in Hz corresponding to bin k .

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral kurtosis.
- μ_1 is the spectral centroid.
- μ_2 is the spectral spread.

The spectral kurtosis measures the flatness, or non-Gaussianity, of the spectrum around its centroid. Conversely, it is used to indicate the peakiness of a spectrum. For example, as the white noise is increased on the speech signal, the kurtosis decreases, indicating a less peaky spectrum.

```
[audioIn,fs] = audioread('Counting-16-44p1-mono-15secs.wav');

noiseGenerator = dsp.ColoredNoise('Color','white','SamplesPerFrame',size(audioIn,1));


noise = noiseGenerator();
noise = noise/max(abs(noise));
ramp = linspace(0,.25,numel(noise))';
noise = noise.*ramp;

audioIn = audioIn + noise;

kurtosis = spectralKurtosis(audioIn,fs);

t = linspace(0,size(audioIn,1)/fs,size(audioIn,1));
subplot(2,1,1)
plot(t,audioIn)
ylabel('Amplitude')

t = linspace(0,size(audioIn,1)/fs,size(kurtosis,1));
subplot(2,1,2)
plot(t,kurtosis)
xlabel('Time (s)')
ylabel('Kurtosis')
```



Spectral Entropy

Spectral entropy (`spectralEntropy`) measures the peakiness of the spectrum [6 on page 20-0]:

$$\text{entropy} = \frac{-\sum_{k=b_1}^{b_2} s_k \log(s_k)}{\log(b_2 - b_1)}$$


where

- f_k is the frequency in Hz corresponding to bin k .

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral entropy.

Spectral entropy has been used successfully in voiced/unvoiced decisions for automatic speech recognition [6 on page 20-0]. Because entropy is a measure of disorder, regions of voiced speech have lower entropy compared to regions of unvoiced speech.

```
[audioIn,fs] = audioread('Counting-16-44p1-mono-15secs.wav');  
  
entropy = spectralEntropy(audioIn,fs);  
  
t = linspace(0,size(audioIn,1)/fs,size(audioIn,1));  
subplot(2,1,1)  
plot(t,audioIn)  
ylabel('Amplitude')  
  
t = linspace(0,size(audioIn,1)/fs,size(entropy,1));  
subplot(2,1,2)  
plot(t,entropy)  
xlabel('Time (s)')  
ylabel('Entropy')
```



Spectral entropy has also been used to discriminate between speech and music [7 on page 20-0] [8 on page 20-0]. For example, compare histograms of entropy for speech, music, and background audio files.

```
fs = 8000;
[speech,speechFs] = audioread('Rainbow-16-8-mono-114secs.wav');
speech = resample(speech,fs,speechFs);
speech = speech./max(speech);

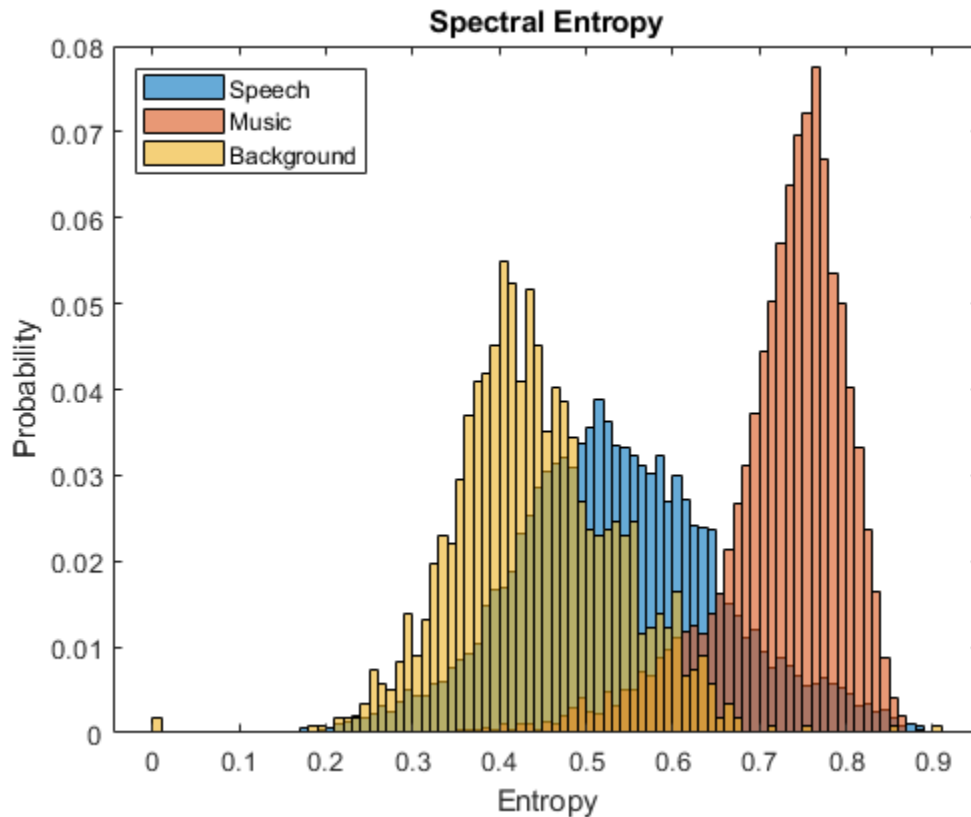
[music,musicFs] = audioread('RockGuitar-16-96-stereo-72secs.flac');
music = sum(music,2)/2;
music = resample(music,fs,musicFs);
music = music./max(music);
```

```
[background,backgroundFs] = audioread('Ambiance-16-44p1-mono-12secs.wav');
background = resample(background,fs,backgroundFs);
background = background./max(background);

speechEntropy      = spectralEntropy(speech,fs);
musicEntropy       = spectralEntropy(music,fs);
backgroundEntropy = spectralEntropy(background,fs);

figure
h1 = histogram(speechEntropy);
hold on
h2 = histogram(musicEntropy);
h3 = histogram(backgroundEntropy);

h1.Normalization = 'probability';
h2.Normalization = 'probability';
h3.Normalization = 'probability';
h1.BinWidth = 0.01;
h2.BinWidth = 0.01;
h3.BinWidth = 0.01;
title('Spectral Entropy')
legend('Speech','Music','Background','Location','northwest')
xlabel('Entropy')
ylabel('Probability')
hold off
```

Spectral Flatness

Spectral flatness (`spectralFlatness`) measures the ratio of the geometric mean of the spectrum to the arithmetic mean of the spectrum [9 on page 20-0]:

$$\text{flatness} = \frac{\left(\prod_{k=b_1}^{b_2} s_k \right)^{\frac{1}{b_2 - b_1}}}{\frac{1}{b_2 - b_1} \sum_{k=b_1}^{b_2} s_k}$$

where

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral flatness.

Spectral flatness is an indication of the peakiness of the spectrum. A higher spectral flatness indicates noise, while a lower spectral flatness indicates tonality.

```
[audio,fs] = audioread('WaveGuideLoopOne-24-96-stereo-10secs.aif');
audio = sum(audio,2)/2;


noise = (2*rand(numel(audio),1)-1).*linspace(0,0.05,numel(audio))';

audio = audio + noise;

flatness = spectralFlatness(audio,fs);

subplot(2,1,1)
t = linspace(0,size(audio,1)/fs,size(audio,1));
plot(t,audio)
ylabel('Amplitude')

subplot(2,1,2)
t = linspace(0,size(audio,1)/fs,size(flatness,1));
plot(t,flatness)
ylabel('Flatness')
xlabel('Time (s)')
```



Spectral flatness has also been applied successfully to singing voice detection [10 on page 20-0] and to audio scene recognition [11 on page 20-0].

Spectral Crest

Spectral crest (`spectralCrest`) measures the ratio of the maximum of the spectrum to the arithmetic mean of the spectrum [1 on page 20-0]:

$$\text{crest} = \frac{\max(S_{k \in [b_1, b_2]})}{\frac{1}{b_2 - b_1} \sum_{k=b_1}^{b_2} S_k}$$

where

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral crest.

Spectral crest is an indication of the peakiness of the spectrum. A higher spectral crest indicates more tonality, while a lower spectral crest indicates more noise.

```
[audio,fs] = audioread('WaveGuideLoopOne-24-96-stereo-10secs.aif');
audio = sum(audio,2)/2;


noise = (2*rand(numel(audio),1)-1).*linspace(0,0.2,numel(audio));

audio = audio + noise;

crest = spectralCrest(audio,fs);

subplot(2,1,1)
t = linspace(0,size(audio,1)/fs,size(audio,1));
plot(t,audio)
ylabel('Amplitude')

subplot(2,1,2)
t = linspace(0,size(audio,1)/fs,size(crest,1));
plot(t,crest)
ylabel('Crest')
xlabel('Time (s)')
```



Spectral Flux

Spectral flux (`spectralFlux`) is a measure of the variability of the spectrum over time [12 on page 20-0]:

$$\text{flux}(t) = \left(\sum_{k=b_1}^{b_2} |s_k(t) - s_k(t-1)|^p \right)^{\frac{1}{p}}$$

where

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.

- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral flux.
- p is the norm type.


Spectral flux is popularly used in onset detection [13 on page 20-0] and audio segmentation [14 on page 20-0]. For example, the beats in the drum track correspond to high spectral flux.

```
[audio,fs] = audioread('FunkyDrums-48-stereo-25secs.mp3');  
audio = sum(audio,2)/2;
```

```
flux = spectralFlux(audio,fs);
```

```
subplot(2,1,1)  
t = linspace(0,size(audio,1)/fs,size(audio,1));  
plot(t,audio)  
ylabel('Amplitude')
```

```
subplot(2,1,2)  
t = linspace(0,size(audio,1)/fs,size(flux,1));  
plot(t,flux)  
ylabel('Flux')  
xlabel('Time (s)')
```



Spectral Slope

Spectral slope (`spectralSlope`) measures the amount of decrease of the spectrum [15 on page 20-0]:

$$\text{slope} = \frac{\sum_{k=b_1}^{b_2} (f_k - \mu_f)(s_k - \mu_s)}{\sum_{k=b_1}^{b_2} (f_k - \mu_f)^2}$$

where

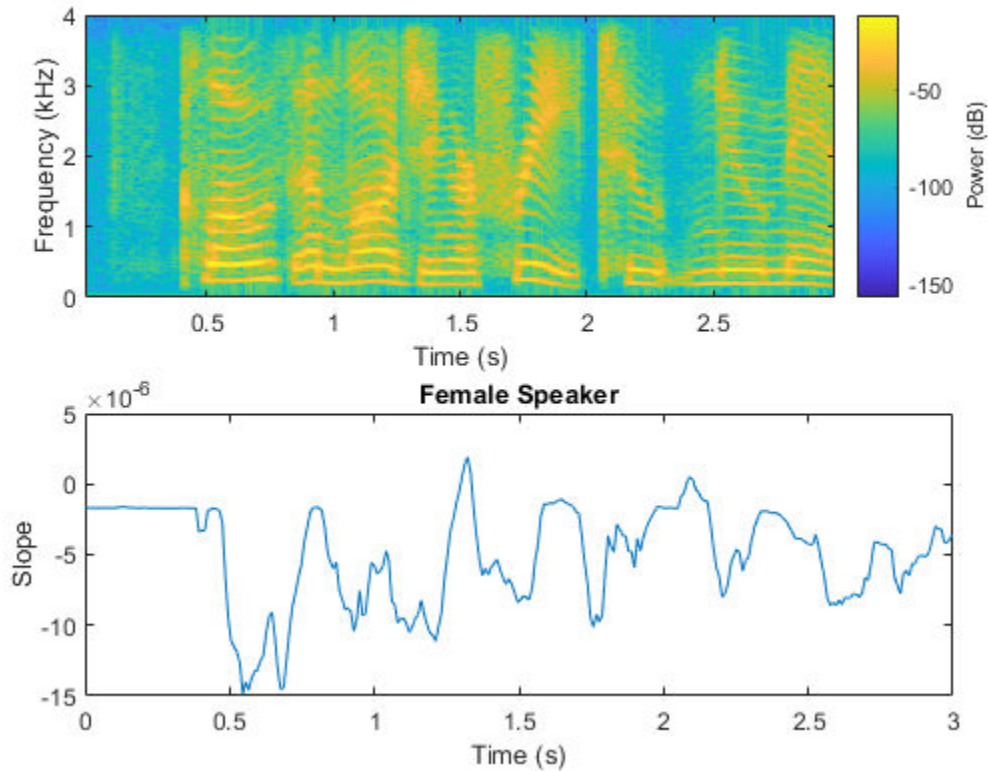
- f_k is the frequency in Hz corresponding to bin k .

- μ_f is the mean frequency.
- s_k is the spectral value at bin k . The magnitude spectrum is commonly used.
- μ_s is the mean spectral value.
- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral slope.

Spectral slope has been used extensively in speech analysis, particularly in modeling speaker stress [19 on page 20-0]. The slope is directly related to the resonant characteristics of the vocal folds and has also been applied to speaker identification [21 on page 20-0]. Spectral slope is a socially important aspect of timbre. Spectral slope discrimination has been shown to occur in early childhood development [20 on page 20-0]. Spectral slope is most pronounced when the energy in the lower formants is much greater than the energy in the higher formants.

```
[female, femaleFs] = audioread('FemaleSpeech-16-8-mono-3secs.wav');  
female = female./max(female);
```

```
femaleSlope = spectralSlope(female, femaleFs);  
t = linspace(0, size(female, 1)/femaleFs, size(femaleSlope, 1));  
subplot(2, 1, 1)  
spectrogram(female, round(femaleFs*0.05), round(femaleFs*0.04), round(femaleFs*0.05), femaleSlope, 'y');  
  
subplot(2, 1, 2)  
plot(t, femaleSlope)  
title('Female Speaker')  
ylabel('Slope')  
xlabel('Time (s)')
```

Spectral Decrease

Spectral decrease (`spectralDecrease`) represents the amount of decrease of the spectrum, while emphasizing the slopes of the lower frequencies [1 on page 20-0]:

$$\text{decrease} = \frac{\sum_{k=b_1+1}^{b_2} \frac{s_k - s_{b_1}}{k-1}}{\sum_{k=b_1+1}^{b_2} s_k}$$

where

- s_k is the spectral value at bin k . The magnitude spectrum is commonly used.

- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral decrease.

Spectral decrease is used less frequently than spectral slope in the speech literature, but it is commonly used, along with slope, in the analysis of music. In particular, spectral decrease has been shown to perform well as a feature in instrument recognition [22 on page 20-0].


```
[guitar,guitarFs] = audioread('RockGuitar-16-44p1-stereo-72secs.wav');
guitar = mean(guitar,2);
[drums,drumsFs] = audioread('RockDrums-44p1-stereo-11secs.mp3');
drums = mean(drums,2);
```

```
guitarDecrease = spectralDecrease(guitar,guitarFs);
drumsDecrease = spectralDecrease(drums,drumsFs);
```

```
t1 = linspace(0,size(guitar,1)/guitarFs,size(guitarDecrease,1));
t2 = linspace(0,size(drums,1)/drumsFs,size(drumsDecrease,1));
```

```
subplot(2,1,1)
plot(t1,guitarDecrease)
title('Guitar')
ylabel('Decrease')
axis([0 10 -0.3 0.3])
```

```
subplot(2,1,2)
plot(t2,drumsDecrease)
title('Drums')
ylabel('Decrease')
xlabel('Time (s)')
axis([0 10 -0.3 0.3])
```



Spectral Rolloff Point

The spectral rolloff point (`spectralRolloffPoint`) measures the bandwidth of the audio signal by determining the frequency bin under which a given percentage of the total energy exists [12 on page 20-0]:

$$\text{Rolloff Point} = i \quad \text{such that} \quad \sum_{k=b_1}^i |s_k| = \kappa \sum_{k=b_1}^{b_2} s_k$$

where

- s_k is the spectral value at bin k . The magnitude spectrum and power spectrum are both commonly used.

- b_1 and b_2 are the band edges, in bins, over which to calculate the spectral rolloff point.
- κ is the specified energy threshold, usually 95% or 85%.

i is converted to Hz before it is returned by `spectralRolloffPoint`.

The spectral rolloff point has been used to distinguish between voiced and unvoiced speech, speech/music discrimination [12 on page 20-0], music genre classification [16 on page 20-0], acoustic scene recognition [17 on page 20-0], and music mood classification [18 on page 20-0]. For example, observe the different mean and variance of the rolloff point for speech, rock guitar, acoustic guitar, and an acoustic scene.

```
dur = 5; % Clip out 5 seconds from each file.

[speech,fs1] = audioread('SpeechDFT-16-8-mono-5secs.wav');
speech = speech(1:min(end,fs1*dur));

[electricGuitar,fs2] = audioread('RockGuitar-16-44p1-stereo-72secs.wav');
electricGuitar = mean(electricGuitar,2); % Convert to mono for comparison.
electricGuitar = electricGuitar(1:fs2*dur);


[acousticGuitar,fs3] = audioread('SoftGuitar-44p1_mono-10mins.ogg');
acousticGuitar = acousticGuitar(1:fs3*dur);

[acousticScene,fs4] = audioread('MainStreetOne-24-96-stereo-63secs.wav');
acousticScene = mean(acousticScene,2); % Convert to mono for comparison.
acousticScene = acousticScene(1:fs4*dur);

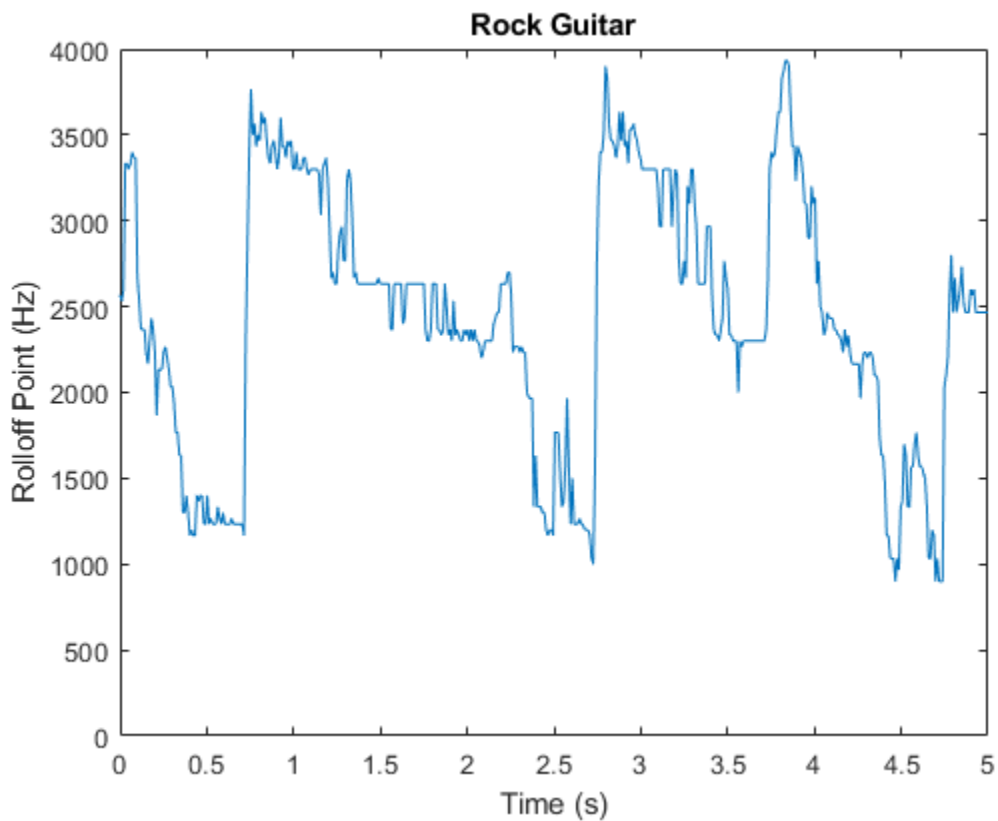
r1 = spectralRolloffPoint(speech,fs1);
r2 = spectralRolloffPoint(electricGuitar,fs2);
r3 = spectralRolloffPoint(acousticGuitar,fs3);
r4 = spectralRolloffPoint(acousticScene,fs4);

t1 = linspace(0,size(speech,1)/fs1,size(r1,1));
t2 = linspace(0,size(electricGuitar,1)/fs2,size(r2,1));
t3 = linspace(0,size(acousticGuitar,1)/fs3,size(r3,1));
t4 = linspace(0,size(acousticScene,1)/fs4,size(r4,1));

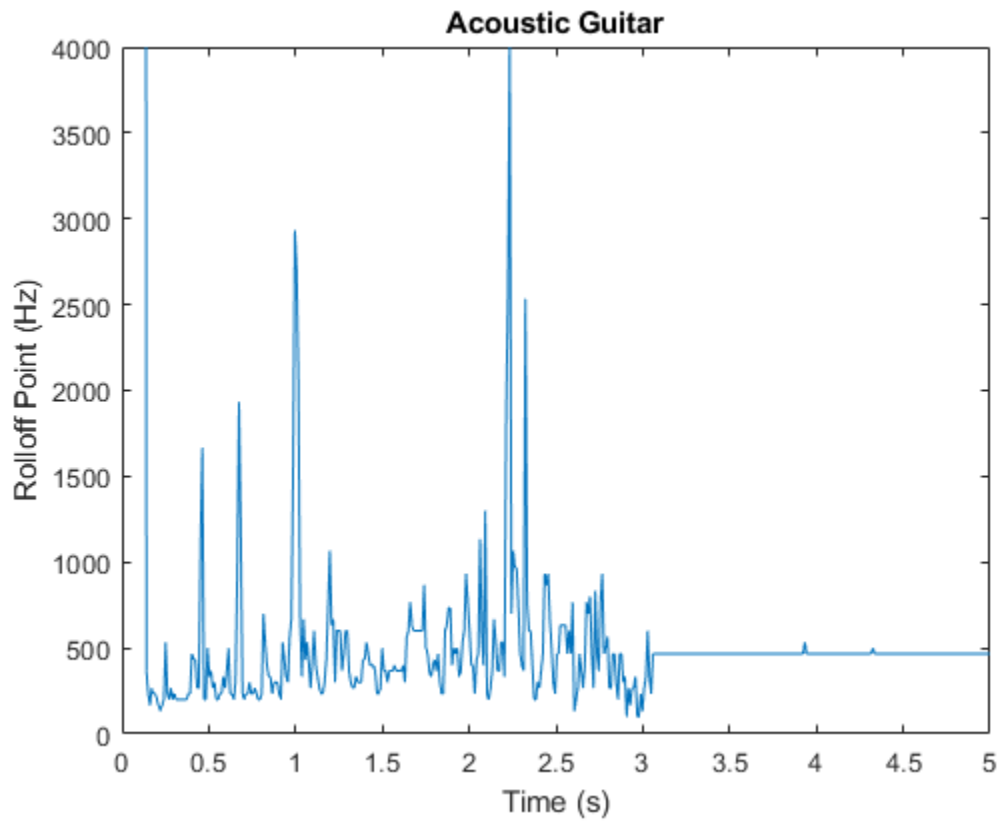
figure
plot(t1,r1)
title('Speech')
ylabel('Rolloff Point (Hz)')
xlabel('Time (s)')
axis([0 5 0 4000])
```



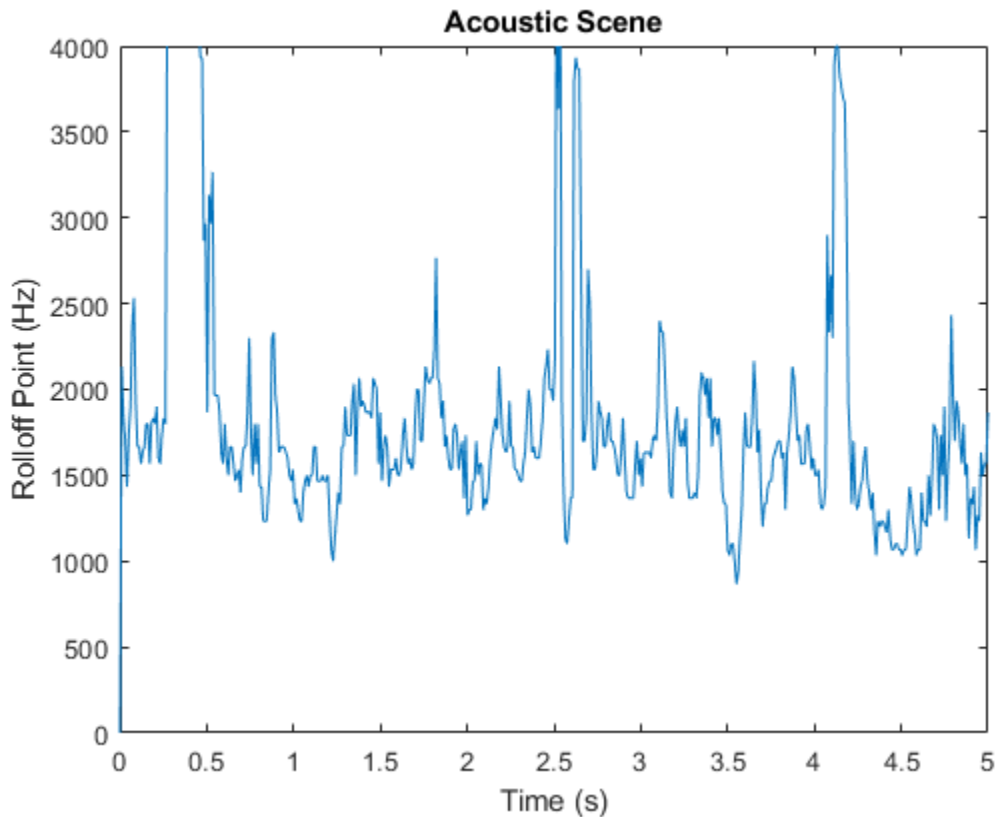
```
figure
plot(t2,r2)
title('Rock Guitar')
ylabel('Rolloff Point (Hz)')
xlabel('Time (s)')
axis([0 5 0 4000])
```



```
figure
plot(t3,r3)
title('Acoustic Guitar')
ylabel('Rolloff Point (Hz)')
xlabel('Time (s)')
axis([0 5 0 4000])
```



```
figure
plot(t4,r4)
title('Acoustic Scene')
ylabel('Rolloff Point (Hz)')
xlabel('Time (s)')
axis([0 5 0 4000])
```



References

- [1] Peeters, G. "A Large Set of Audio Features for Sound Description (Similarity and Classification) in the CUIDADO Project." Technical Report; IRCAM: Paris, France, 2004.
- [2] Grey, John M., and John W. Gordon. "Perceptual Effects of Spectral Modifications on Musical Timbres." *The Journal of the Acoustical Society of America*. Vol. 63, Issue 5, 1978, pp. 1493-1500.
- [3] Raimy, Eric, and Charles E. Cairns. *The Segment in Phonetics and Phonology*. Hoboken, NJ: John Wiley & Sons Inc., 2015.
- [4] Jongman, Allard, et al. "Acoustic Characteristics of English Fricatives." *The Journal of the Acoustical Society of America*. Vol. 108, Issue 3, 2000, pp. 1252-1263.

- [5] S. Zhang, Y. Guo, and Q. Zhang, "Robust Voice Activity Detection Feature Design Based on Spectral Kurtosis." *First International Workshop on Education Technology and Computer Science*, 2009, pp. 269-272.
- [6] Misra, H., S. Ikbal, H. Bourlard, and H. Hermansky. "Spectral Entropy Based Feature for Robust ASR." *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- [7] A. Pikrakis, T. Giannakopoulos, and S. Theodoridis. "A Computationally Efficient Speech/Music Discriminator for Radio Recordings." *International Conference on Music Information Retrieval and Related Activities*, 2006.
- [8] Pikrakis, A., et al. "A Speech/Music Discriminator of Radio Recordings Based on Dynamic Programming and Bayesian Networks." *IEEE Transactions on Multimedia*. Vol. 10, Issue 5, 2008, pp. 846-857.
- [9] Johnston, J.d. "Transform Coding of Audio Signals Using Perceptual Noise Criteria." *IEEE Journal on Selected Areas in Communications*. Vol. 6, Issue 2, 1988, pp. 314-323.
- [10] Lehner, Bernhard, et al. "On the Reduction of False Positives in Singing Voice Detection." *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014.
- [11] Y. Petetin, C. Laroche and A. Mayoue, "Deep Neural Networks for Audio Scene Recognition," *2015 23rd European Signal Processing Conference (EUSIPCO)*, 2015.
- [12] Scheirer, E., and M. Slaney. "Construction and Evaluation of a Robust Multifeature Speech/Music Discriminator." *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1997.
- [13] S. Dixon, "Onset Detection Revisited." *International Conference on Digital Audio Effects*. Vol. 120, 2006, pp. 133-137.
- [14] Tzanetakis, G., and P. Cook. "Multifeature Audio Segmentation for Browsing and Annotation." *Proceedings of the 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 1999.
- [15] Lerch, Alexander. *An Introduction to Audio Content Analysis Applications in Signal Processing and Music Informatics*. Piscataway, NJ: IEEE Press, 2012.
- [16] Li, Tao, and M. Ogihara. "Music Genre Classification with Taxonomy." *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2005.

- [17] Eronen, A.j., Vt. Peltonen, J.t. Tuomi, A.p. Klapuri, S. Fagerlund, T. Sorsa, G. Lorho, and J. Huopaniemi. "Audio-Based Context Recognition." *IEEE Transactions on Audio, Speech and Language Processing*. Vol. 14, Issue 1, 2006, pp. 321-329.
- [18] Ren, Jia-Min, Ming-Ju Wu, and Jyh-Shing Roger Jang. "Automatic Music Mood Classification Based on Timbre and Modulation Features." *IEEE Transactions on Affective Computing*. Vol. 6, Issue 3, 2015, pp. 236-246.
- [19] Hansen, John H. L., and Sanjay Patil. "Speech Under Stress: Analysis, Modeling and Recognition." *Lecture Notes in Computer Science*. Vol. 4343, 2007, pp. 108-137.
- [20] Tsang, Christine D., and Laurel J. Trainor. "Spectral Slope Discrimination in Infancy: Sensitivity to Socially Important Timbres." *Infant Behavior and Development*. Vol. 25, Issue 2, 2002, pp. 183-194.
- [21] Murthy, H.a., F. Beaufays, L.p. Heck, and M. Weintraub. "Robust Text-Independent Speaker Identification over Telephone Channels." *IEEE Transactions on Speech and Audio Processing*. Vol. 7, Issue 5, 1999, pp. 554-568.
- [22] Essid, S., G. Richard, and B. David. "Instrument Recognition in Polyphonic Music Based on Automatic Taxonomies." *IEEE Transactions on Audio, Speech and Language Processing*. Vol 14, Issue 1, 2006, pp. 68-80.