Introduction to the Soundskrit Mantis Headset Kit

Figure 1: Soundskrit Mantis Kit

The Soundskrit Mantis kit is designed to evaluate the performance of Soundskrit microphones and software in a headset scenario. Our most powerful headset configuration combines an omnidirectional and dipole microphone to enable several software features. There are three positions on the Mantis board with an omnidirectional microphone and an SKR0610 dipole microphone. One near the mouth as in a boom headset, one further back for a reduced boom length, and one near the ear as found in boomless, wireless headphones.

There are two software algorithms running on the Arm Cortex M7 processor for each microphone pair. The first is our adaptive beamformer which varies the polar pattern over time and frequency between omni and dipole to maximize noise rejection. This algorithm rejects a large amount of noise and sounds more natural than typical non-linear processing techniques. The second is our nonlinear DSP beamformer. Our nonlinear beamformer mutes the microphone unless speech is coming from $\pm 60°$ and further reduces non-speech noise when the microphone is unmuted. This is a very lightweight and low latency algorithm based on traditional DSP methods.

The processing all runs embedded and does not require any programs on your PC.

What's In the Box	
Mantis Board	The Mantis development board with 3 dipole microphones and 3 omnidirectional microphones.
Soundskrit PARDI audio interface	Multichannel audio interface to connect microphones over USB
Molex cable	To connect Mantis microphone board with audio interface
USB-A to USB-C Cable	A cable to connect the board to your PC.
Mounting Putty	Putty to mount the board to an existing headset for evaluation.

Table of Contents

Mantis Board Overview

The Mantis board has a Molex connector and 6 MEMS microphones. The “front” of the board is the bare side of the PCB. When using the board, face the bare side of the PCB toward the mouth. Each algorithm uses one omnidirectional microphone and one dipole microphone. This fundamental pair is located in three different positions so that you can compare the performance difference between a headset with a full-length boom, a shorter boom, and one without a boom. These fundamental pairs will be referred to as the boom pair, middle pair, and earcup pair respectively. Generally, the boom and middle pairs provide similar performance, offering the best vocal quality and highest noise rejection. The earcup pair trades off performance to embed the microphones in the earcup rather than using a boom.

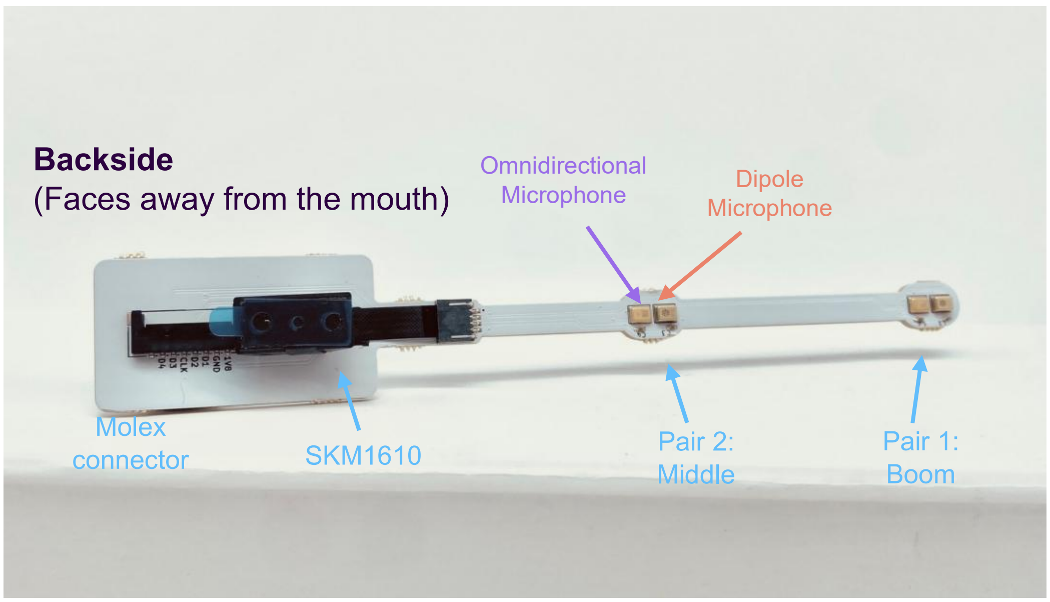

Figure 2: Mantis board microphone pairs

For this demo kit, the earcup microphone pair protrudes from the board. In the final product, the microphones can be integrated inside the headset and do not need to stick out as on this board. In a production product, like a pair of headphones, a gasket would be used to mount the microphones inside the headphones with the port holes on the same surface. For more information on integrating the microphones such that the ports are on the same surface, see our application note: [AN-300 Integration Guide for Directional MEMS Microphones](). Alternatively, Soundskrit sells ready to use modules for easy integration, more information about these modules can be found [here]().

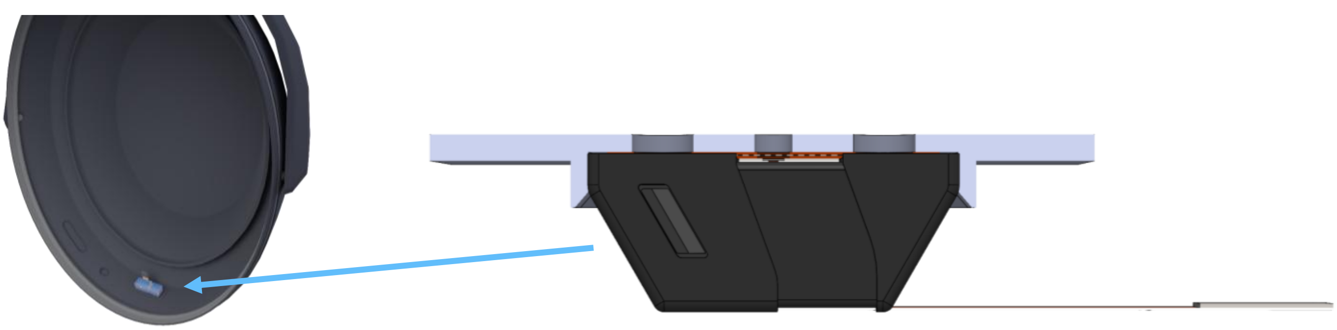

Figure 3: Example of an SKM1610 inside of a headset earcup

Audio interface overview and pinouts

The PARDI audio interface has a USB-C connector to connect it to a PC and a Molex connector to connect the Mantis board. The cable is preconnected and only needs reattachment if disconnected.

Figure 4: PARDI audio interface.

Included Beamforming

Once the Mantis board is connected, you can take a recording. The USB output will have six channels. Channels 1 and 2 are for the microphone pair at the end of the boom near the mouth, channels 3 and 4 are for the middle pair, channels 5 and 6 are for the pair that sits at the earcup. The first of each of these pairs is the adaptive beamformer and the second of each is the nonlinear beamformer. These algorithms are tuned for headset distances and will not perform well if used at a distance greater than 5-10 cm.

To learn more about combining an omnidirectional and dipole microphone to create adjustable polar patterns, check out our article: [Combining Microphone Polar Patterns].

All channels implement a high pass filter set to 40 Hz.

USB Output Channels		
Boom pair	Ch1	Adaptive beamformer
	Ch2	Nonlinear beamformer
Middle pair	Ch3	Adaptive beamformer
	Ch4	Nonlinear beamformer
Earcup pair	Ch5	Adaptive beamformer
	Ch6	Nonlinear beamformer

All polar patterns on the following page are measured from 50 cm to the microphone, since the microphone should be directional in the far field to reject ambient noise. Measurements from very close distances would look different.

The analog out jack is tied to channel 1. To output a different algorithm, put the desired processing on ch1 using tera term (see page 12 and onwards).

Channels 1, 3, 5: SimplyDSP Adaptive Beamformer

The SimplyDSP Adaptive beamformer varies the polar pattern over time and frequency between omni and dipole to maximize noise rejection. This is a non-linear beamformer which rejects a large amount of noise and sounds more natural than typical non-linear processing techniques.

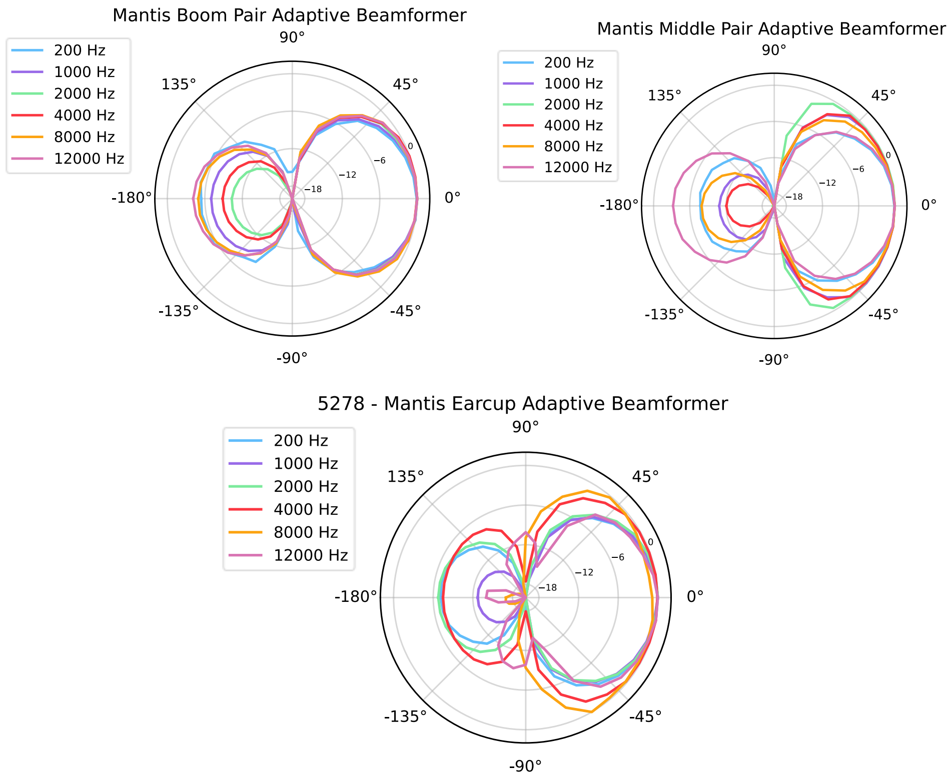

Figure 5: Adaptive beamformer polar patterns at each position

Channel 2, 4, 6: SimplyDSP nonlinear beamformer

Our nonlinear beamformer uses phase information provided by the omnidirectional-dipole pair to mute the microphone unless speech is coming from ±60°. This “activation zone” is pictured below. When the microphone is unmuted, the cardioid beamformer is used and this phase information is used to further reduce non-speech noise.

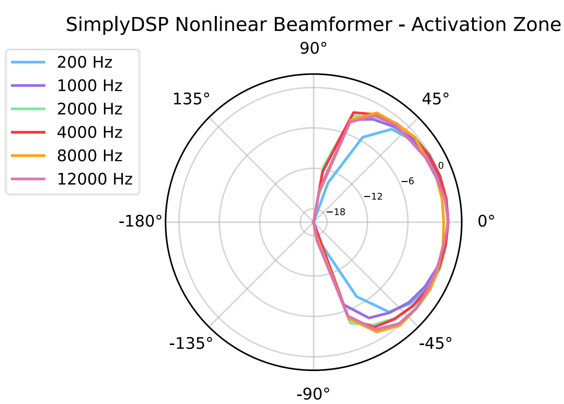

Figure 6: Nonlinear beamformer activation zone

How to Use the Mantis Board

There are two methods to mount the Mantis board to a headset and take recordings. First, for quick demos you can use the included putty to mount the board to the left side of your headset with the cable hanging down. Place the putty on the space of the board as shown in the image below. Adjust the boom to be positioned near the corner of your mouth.

Figure 7: Putty on the Mantis board

Figure 8: Mantis board with the cable facing down

Second, to improve stability for longer term use, mount the board to the right side of your headset with the USB connector facing up. Route the cable over the top of the headset and secure it to the headband with a Velcro cable tie or tape.

Figure 9: Mantis board with cable routed over the headband

The frequency response of the Mantis is tuned for positions as shown in the two photos above. Mounting the board differently, e.g. above the ear or much lower, might lead to a sub-optimal sound.

Recording Audio with the Mantis

Important: To evaluate the demokit using Soundskrit processing only, disable any operating-system audio enhancements (noise suppression, AGC/leveling, spatial sound, etc.) on the Soundskrit Mantis audio device before recording.

Windows 11/10 (recommended):

- Open **Settings > System > Sound**
- Under **Input**, select the Soundskrit demokit microphone device to open its **Properties**
- Set **Audio enhancements** to **Off** (and disable **Spatial sound** if present)

To record audio with the Mantis board, we recommend installing [Audacity](). Audacity is a trusted, free to use, multiplatform suite of tools for recording and working with audio files.

Once you have installed Audacity, configure the software for use with the Soundskrit PARDI board. Configure the settings as listed below:

Audio Host	Windows WASAPI
Input[1]	*Microphone (Soundskrit Mantis)* – Ensure the non-loopback version of the driver is selected
Output	Your listening device
Channels	6 Recording Channels

[1]WASAPI has two versions of each input option, the regular and the loopback, to record from all six channels the non-loopback mode must be selected. If this is configured correctly, there will be six recording channels listed, while the loopback will only allow two. Non-loopback is typically the second of the two versions.

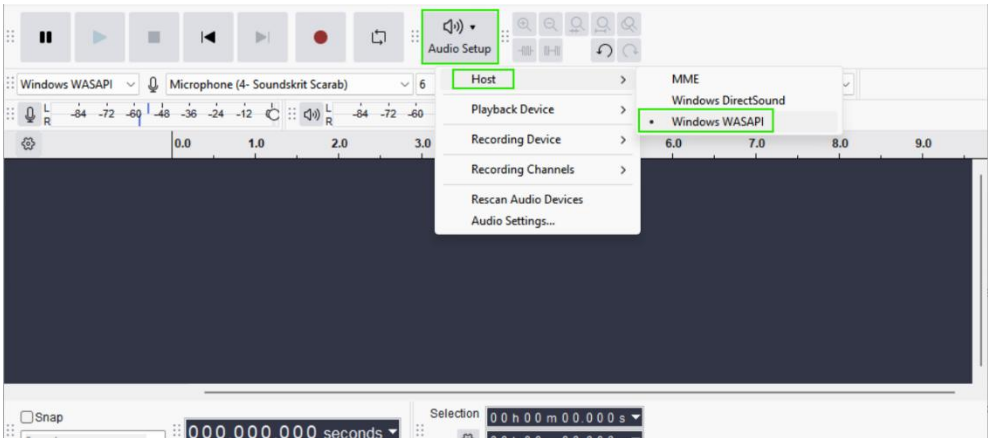

Figure 10: Audacity host selection

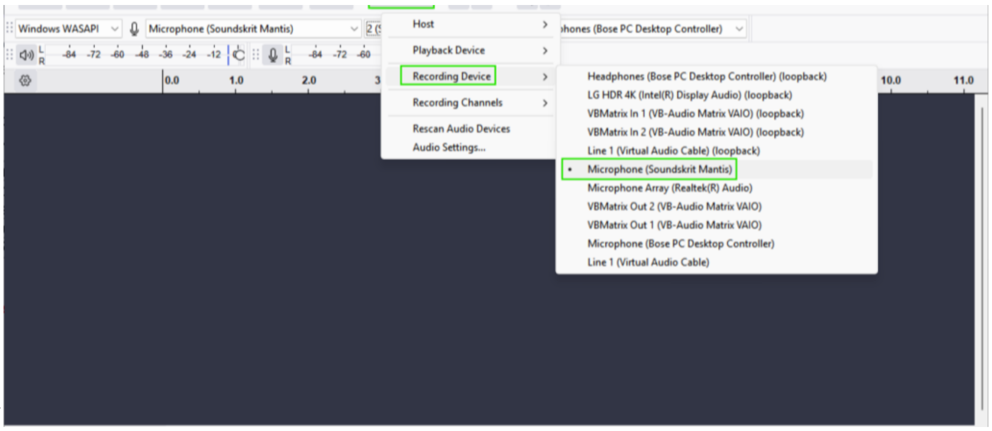

Figure 11: Audacity device selection

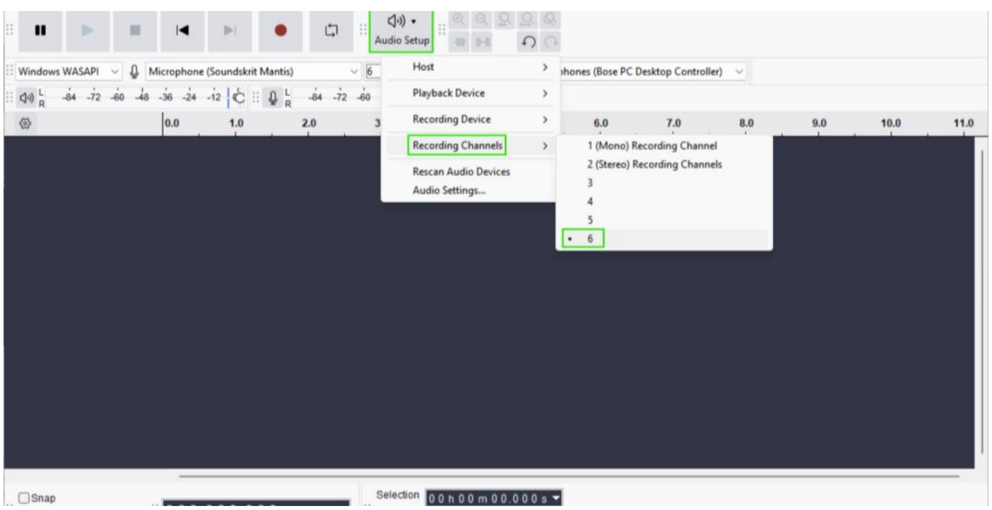

Figure 12: Audacity channel count selection

When you take a recording, all 6 channels will be recorded simultaneously. The microphone gain might be initially very low, this is on purpose the default to avoid signal clipping. You can either increase the recording gain in Windows from the default 54% to a level better suited to your recording situation or amplify the signals in Audacity. To do so, select all tracks by pressing CTRL+A and apply the same gain value to them by selecting in the top menu Effect → Amplify. In the pop up, select a gain value below the clipping threshold, e.g. by setting the “New Peak Amplitude” to -3 dB.

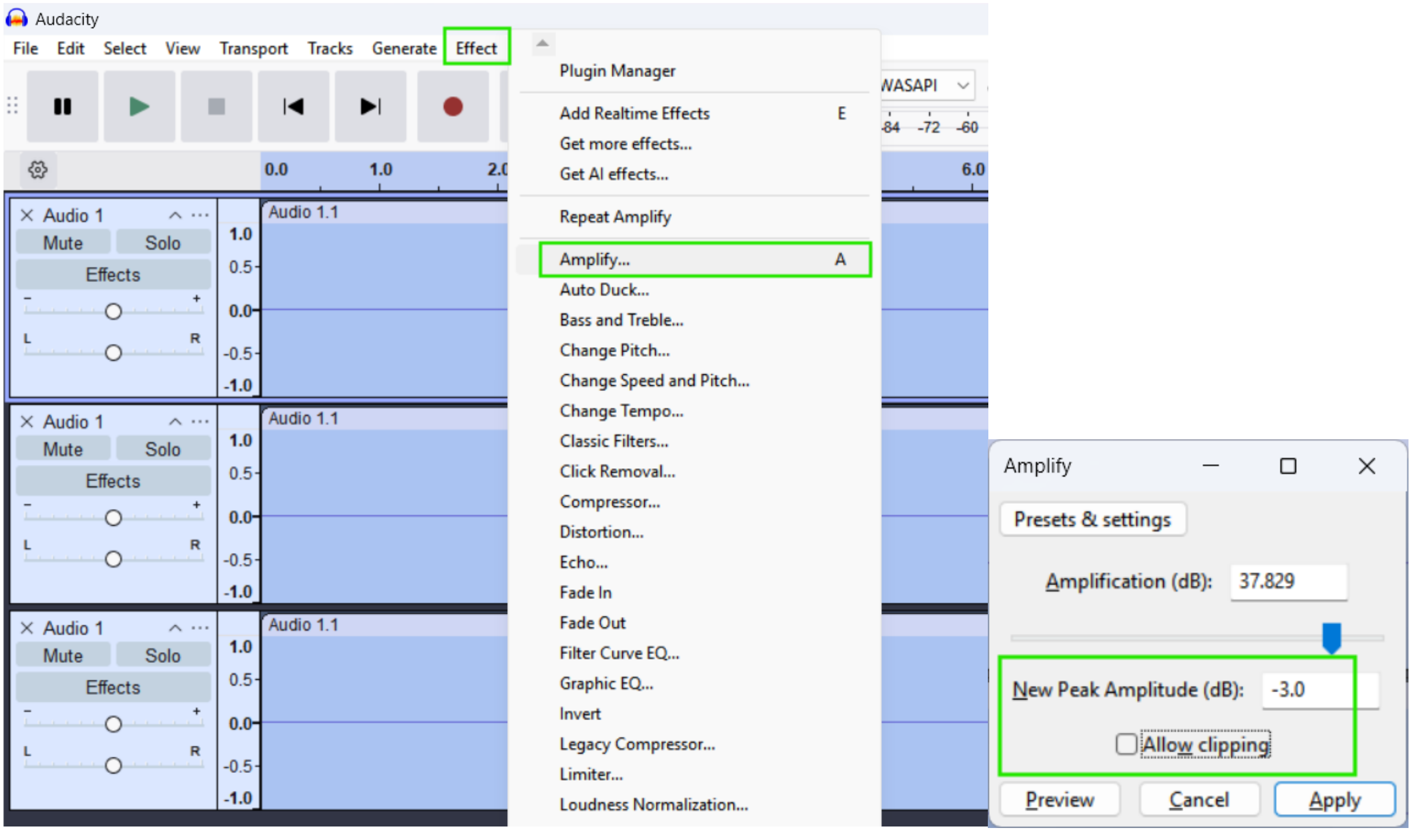

Figure 13: Applying gain to all recorded tracks.

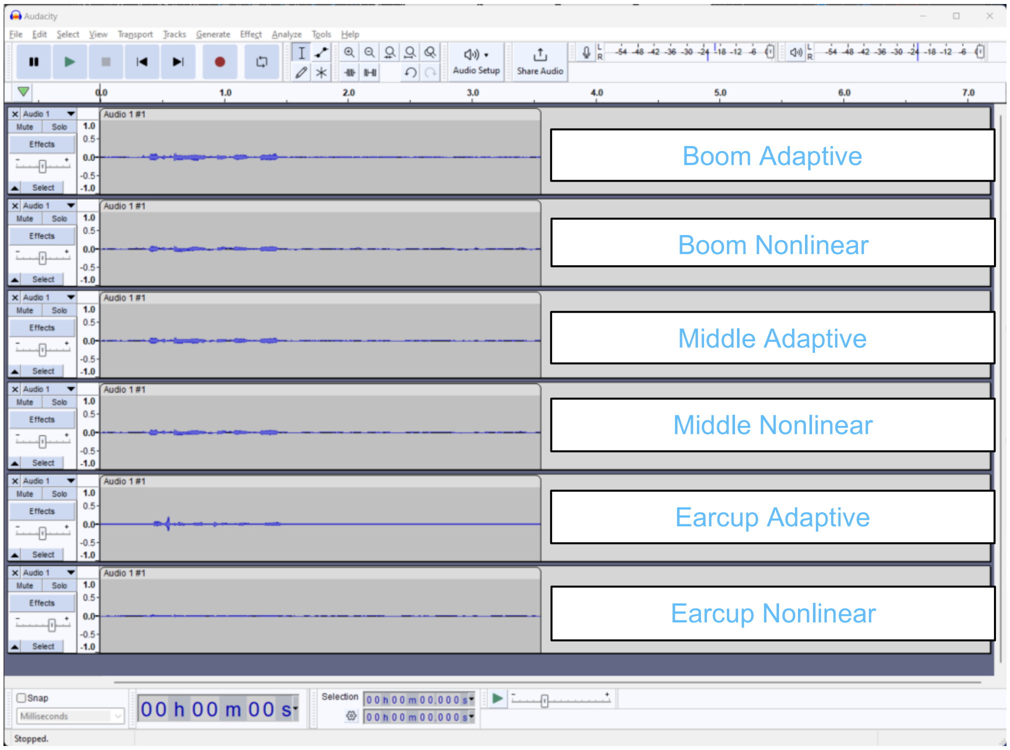

Figure 14: Output channels

Recommended Demos

Below are instructions on how to demonstrate each of the processing features available in this kit.

First, play ambient noise from a speaker a couple meters away from the Mantis board. The speaker should be enough distance away to simulate diffuse noise such as from a crowd or general background noise. Try talking into the board.

Figure 15: Mantis board positioned near the corner of the mouth

Both algorithms will reduce the noise, the adaptive beamformer will vary the polar pattern over time and frequency to maximize noise rejection, providing strong noise reduction while sounding more natural than typical non-linear processing. The nonlinear beam will have stronger noise reduction and completely mute when the user is not talking, but there will be some distortion in the voice.

Next, to test the nonlinear beamformer, have the Mantis board facing your mouth. Talk into the microphone. The voice should come through full and clear. Have a colleague speak into the other side of the board, their voice will be completely muted. Then, have them speak at the same time as you near the board. Despite the competing speech, your voice will come through clearly and the microphones will be completely muted when you are not speaking.

Use with Video Conferencing

To test out the Mantis with video conferencing, the output must be configured such that there is only an output on channel 1 and all other channels are muted. To configure the firmware, you will need to install an SSH terminal emulator such as Tera Term. First, install Tera Term and in the new connection menu, connect to the Mantis board. The Mantis board will be listed as a serial device named COM**X**: USB Serial Device (COM**X**)

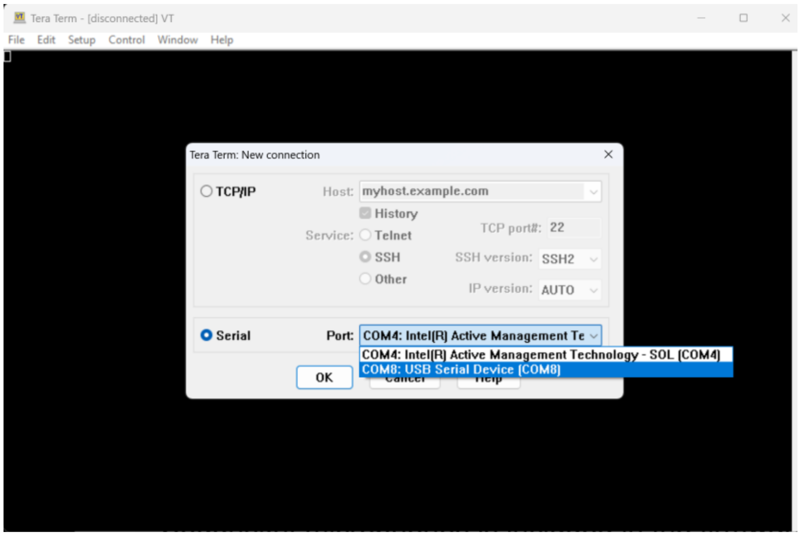

Figure 16: Tera Term new connection

Once connected, go to Setup → Terminal, and configure the new-line settings as LF for both receive and transmit and enable local echo.

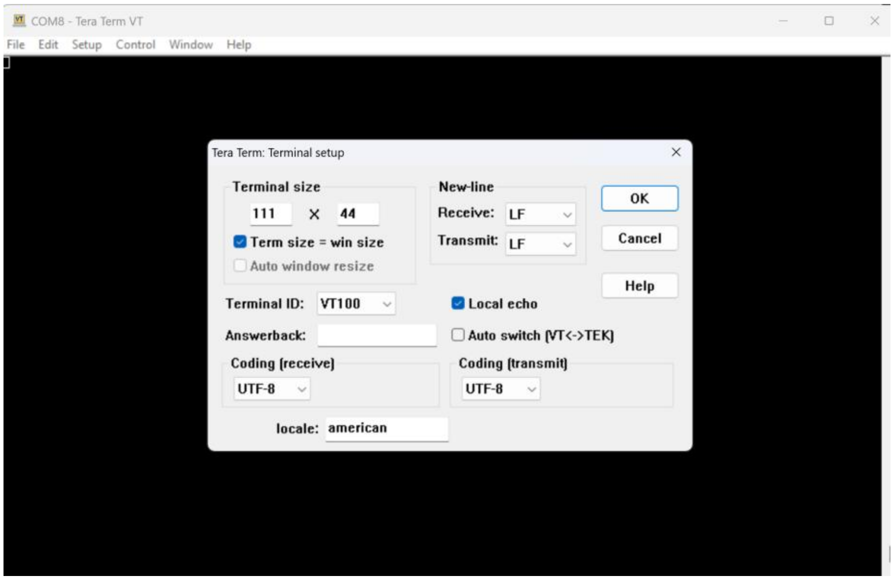

Figure 17: Tera Term terminal setup

Now you can interface with the Mantis board firmware. To begin, reset the audio configuration using the command ***reset***. To view the current audio configuration at any point, send the command: ***get audio***.

Figure 18: Current audio configuration

Here you can see the current configuration of the audio and how they are connected to the USB output channels. For example, USB output channel 1 uses the microphones on the boom pair, input pair 1, applies the adaptive beamformer, labeled as "adaptive", and applies the highpass filter at 40 Hz.

For the nonlinear beamformers, the highpass filter is set to 40 Hz, consistent with all other channels.

Next, mute all output channels using the command: ***mute a***

Figure 19: Muting the output channels

Next, we will set usb-ch-1 to be the output configuration you want to use over video conferencing.

The command to set the USB output is ***link <input pair> <algorithm type> <output>*** where <input pair> is one of the three microphone pairs 1-3 and <output> USB channel1-6. <algorithm type> is the type of processing which should be applied, type "get algorithms" to display all available algorithms.

The currently available algorithms are the following:

"omni-raw" → Unprocessed omnidirectional microphone
"dipole-raw" → Unprocessed dipole microphone
"omni-eq" → Equalized omnidirectional microphone

"dipole-eq" Equalized dipole microphone
"hd-beam"   A high directivity beam pattern using the omni-dipole mic pair
"adaptive" → An adaptive beam pattern that varies over time and frequency to maximize noise rejection
"nonlinear" → A nonlinear narrow beam pattern using the omni-dipole mic pair

For a voice call, use the link command to route the desired algorithm to USB output channel 1. For example: link 1 adaptive 1

```
link 1 nonlinear 1
get audio

====================================== AUDIO FLOW ==========================================
------------------------------------------------------------------------------------------------
      INPUT       :          ALGORITHM          : AI DENOISER :   GAIN   :  HIGHPASS  :  USB OUTPUT
------------------------------------------------------------------------------------------------
 omni 1     ----->:
                  :-> nonlinear            --> off       --> 0.0dB  -->   80 Hz  -->  usb-ch-1
 dipole 1   ----->:
      xxxxx          mute                                                             usb-ch-2
      xxxxx          mute                                                             usb-ch-3
      xxxxx          mute                                                             usb-ch-4
      xxxxx          mute                                                             usb-ch-5
      xxxxx          mute                                                             usb-ch-6
```

Figure 20: Example: routing the long boom nonlinear beam to USB channel 1.

With these commands, we can set the output channel to be either of the beamforming algorithms or as the equalized single omnidirectional or dipole microphones. For comparative purposes, we always recommend equalized outputs for the most direct comparison.

For example, to use the long boom adaptive beamformer on channel one for video conferencing, use the commands below:

reset
mute a
link 1 adaptive 1
get audio

Additional Support

For further information on Soundskrit's products, visit our website at http://www.soundskrit.ca where you can find more application notes, datasheets, and purchasing information. If you have any questions or need technical support, please reach out to applications@soundskrit.ca.

Revision History

Revision Label	Revision Date	Sections Revised
-	March 2024	Initial release
Rev A	February 2025	Changes related to PARDI V4 audio interface, Updated nonlinear plot, Updated figure 3 with SKM1610
Rev B	March 2025	Minor updates, plot updates
Rev C	February 2026	Updated channels 1, 3, 5 from high-directivity linear beamformer to adaptive beamformer. Updated highpass filter from 80 Hz to

		40 Hz. Removed set highpass command (no longer supported in firmware). Added adaptive algorithm to algorithm list.

Soundskrit developed the first high-performance directional MEMS microphone on the market, leveraging years of research in bio-inspired MEMS based on how spiders and other insects in nature hear. In combination with Soundskrit's in-house audio processing algorithms, directional microphones can be used to capture and isolate any sound in an environment with a fraction of the size, power, and computation of traditional omnidirectional-based microphone arrays.

Soundskrit was founded in 2019 and is headquartered in Montreal, Quebec with an R&D facility in Ann Arbor, Michigan.

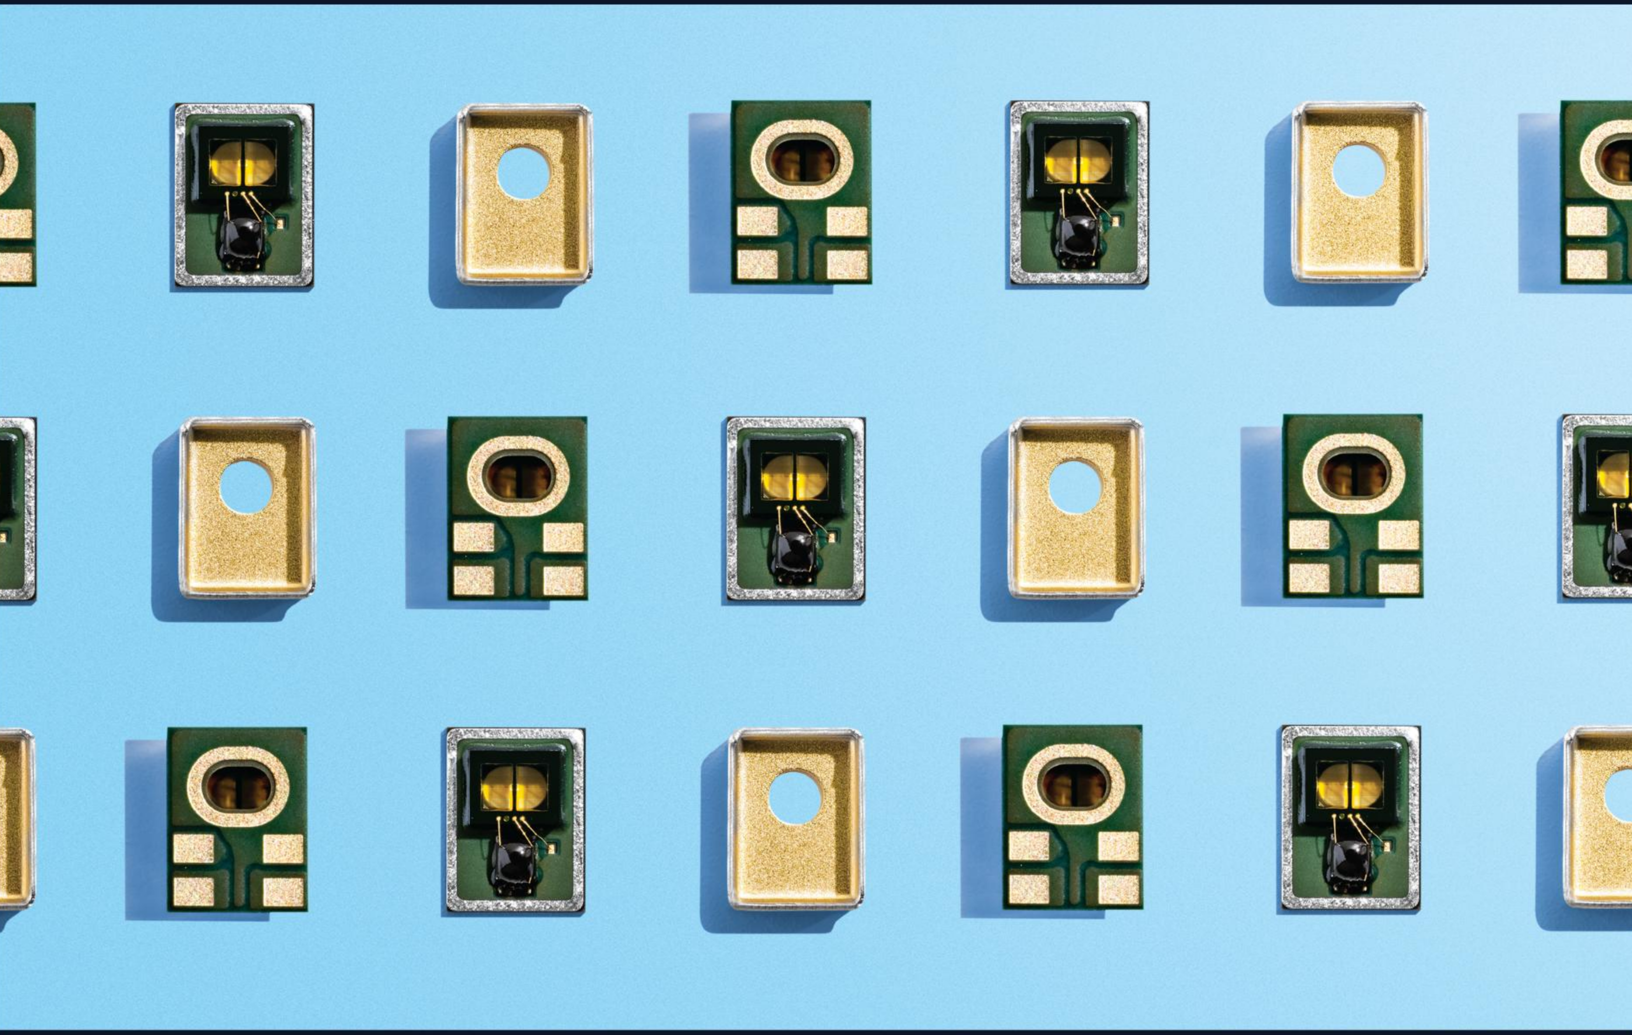

soundskrit
Hear the impossible.