

# *NAJDA UNDER CONTROL (NUC)*

**Software User Manual**

V.1.1.0

## **About this manual**

This is the user reference manual for the Najda Under Control (NUC) software. NUC is the control application that communicates via USB with the Najda board.

This manual does not contain information on the Najda hardware platform. For hardware details and installation notes, please refer to the *Najda Startup Guide*, available for download on our website **[www.waf-audio.com](http://www.waf-audio.com)**.

If you can't find in this manual the information you're looking for, please contact us via email at **[support@waf-audio.com](mailto:support@waf-audio.com)**.

Najda is an upgradeable product. Make sure you've installed the latest software and firmware combo, and that the instruction manual is relevant to the revision your running.

## A few definitions

### Setup vs. Preset

A **Setup** is a collection of DSP parameters saved in a file on the host computer. It can be (re)opened, modified and saved again. *Setup* files have the extension *nsf* (Najda Setup File) and are XML formatted so they can be viewed in a text editor.

When a *Setup* is transferred to the Najda board, it becomes locally a **Preset**. A *Preset* cannot be read from Najda back to the host computer. It cannot be modified either. In order to alter DSP parameters from a *Preset*, one must first load the corresponding *Setup* file in NUC, apply the changes then overwrite the *Preset* in Najda.

### Idle – Offline – Online – Synced

When the user launches Najda Under Control (NUC hereafter), the application is said to be in **Idle** state. The application displays only a white tab with WAF's logo.

The **Offline** state is the situation where a *Setup* file is open, but connection with the board has not been initiated. This is typical when the user creates a new *Setup* file, loads his measurements and starts thinking about the processing strategy. It's also the normal state in demo mode, i.e. when one wishes to evaluate NUC. In *Offline* state, the Monitor tab is not available and neither are any of the hardware configuration dialogues.

The **Online** state occurs when the user connects and no *Setup* file is open. NUC displays the Monitor tab and all hardware configuration dialogues are functional.

Finally, the **Synced** state takes place when connection is established and a *Setup* file has been opened. In this mode, any change to the parameters on screen is immediately transferred to the board.

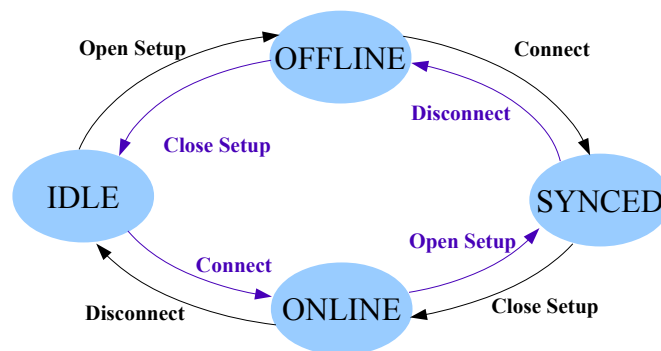


Fig. 1: Transitioning between software states

### Connected

In the whole user manual, the term **Connected** refers not only to the physical link of a USB cable between Najda board and the host computer, but also to software communication initiated with the board.

## List of keyboard shortcuts


<b>Alt-0</b>	<i>Expansion Port 0 Routing</i> Set channel routing for I <sup>2</sup> S outputs on Expansion Port 0
<b>Alt-A</b>	<i>Analogue Outs Configuration</i> Select DAC interpolation filters and set analogue output level
<b>Alt-C</b>	<i>Connect or disconnect</i> Initiate or cease communication between Najda and host computer
<b>Alt-D</b>	<i>Display Options</i> Set the type of display used. Customise sources' names.
<b>Alt-F1</b>	<i>About...</i> Display application and firmware version numbers as well as board serial number (connection required)
<b>Alt-G</b>	<i>Add Output Graph tab</i> Up to 5 Output Graph tabs may be simultaneously opened
<b>Alt-I</b>	<i>Show Input Graph tab</i> Only 1 Input Graph tab is allowed
<b>Alt-P</b>	<i>Preset Manager</i> Load or delete a Preset
<b>Alt-R</b>	<i>Remove Selected Tab</i> Close topmost tab. Closing any of the Input Processing or Channel Processing tabs will result in closing the Setup file.
<b>Alt-T</b>	<i>Bass/Treble Settings</i> Enable / disable Bass/Treble unit, adjust corner frequencies and headroom
<b>Alt-X</b>	<i>Exit</i> Quit the application
<b>Ctrl-N</b>	<i>New Setup</i> Create a new Setup file
<b>Ctrl-O</b>	<i>Open Setup</i> Open an existing Setup file
<b>Ctrl-S</b>	<i>Save Setup</i> Save changes in current Setup file
<b>F1</b>	<i>Help</i> Open the Help documentation (this file)
<b>Space Bar</b>	<i>Mute or Unmute</i> Affects only analogue outputs

## I. Connecting

In order to connect, there must be a USB cable running between the host computer and Najda board, and the latter must be switched on (i.e not in Stand By).

Select **File** → **Connect** in the menu.

In *Online* mode, the DSP keeps running the same Preset that was already running prior to connecting. The *Monitor* tab pops up and displays the Preset's sampling frequency in the status bar.



**Fig. 2:** Monitor tab in Online mode. Circled in red: DSP sampling frequency in the status bar.

In *Synced* mode (i.e. when a Setup file was open prior to connecting), the Setup file parameters are immediately loaded into the DSP. From here on, any change to the *Input Processing* and *Channel Processing* tabs will be transferred in real time so that you hear what you see.

In both modes, on connecting, outputs are muted as a safety measure. The LCD displays “*Host Connected*” and both keypad and IR receiver are disabled. Most of the commands that are usually available through keypad and IR remote controller are doubled in the Monitor's controller section.

Additionally to VU-meters, the *Monitor* tab displays important information relative to the processing status:

- The load meters show how much DSP resources are being used. It's your responsibility to ensure **the load of each core is kept below 100%**.
- Input processing clipping is displayed in the Clip Warning panel. Channel processing clipping is shown on the *ovf* top segment of each VU-meter.

## Disconnecting

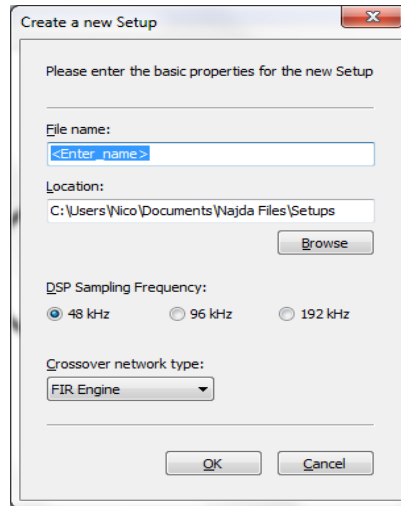
Select **File** → **Disconnect** in the menu.

Keypad and IR receiver are enabled again. If a Setup was transferred without being stored as a Preset, the LCD displays “*Setup From PC*”. This set of processing parameters is volatile: it will

remain active until you load another Preset or turn off Najda.

## II. Creating a new Setup

Select **File** → **Create New Setup** in the menu.



**Fig. 3:** 'Create a New Setup' dialogue

In this dialogue, you type the name of the Setup file that will be saved onto your computer in the selected folder. Prefer explicit file names (e.g. *3WayIirPlusSub* rather than *setup47*).

The two remaining fields set important properties of your Setup. These are the sampling frequency and the type of Setup, *they are set here once for all and cannot be altered at a later time*. Refer to section III for details on available Setup types.

As soon as the OK button is pressed, a Setup file is created on the hard drive of the host computer. This file is initially loaded with default values proper to the selected Setup type. Default values are overwritten as the user progresses in setting parameters and saves the changes.

The Setup file name appears in the window top bar. A trailing star \* means the Setup file has been modified since it was saved last time. The full path to the Setup file is displayed in the status bar.

If Najda Under Control is *Online* (i.e. if it's connected), the default Setup of selected type and sampling frequency is immediately loaded onto the board. As a safety measure, all outputs are muted. Najda is now in *Synced* mode.

## III. Setup Types

### III.1 Setup General Structure

Figure 5 shows the general processing structure common to all Setup types.

Left and right inputs go through the balance stage then across the Bass/Treble unit. Note that the latter can be bypassed (see Bass/Treble in section VI).

As a next step, the signals enter the *Input Processing section*, which is represented in the control application by the Input Processing tab.

In the Input Processing section, two additional signals, namely Sum and Diff, are generated as:


$$\begin{aligned} \text{Sum} &= (\text{Left} + \text{Right})/2 \\ \text{Diff} &= (\text{Left} - \text{Right})/2 \end{aligned}$$

Sum signal used alone is mostly useful for feeding subwoofers. Jointly, Sum and Diff signals can be used to create a 3-speaker Left-Center-Right stereo configuration.

Indeed, from the 2 expressions above, one can see that:

$$\begin{aligned} \text{Left} &= \text{Sum} + \text{Diff} \\ \text{Right} &= \text{Sum} - \text{Diff} \end{aligned}$$

This suggests the LCR configuration below.



**Fig. 4:** 3-speaker LCR configuration using the Sum and Diff Signals.  
*Note the phase inversion of the signal feeding the right speaker.*

In the Input Processing section, all 4 signals Left, Right, Sum and Diff, pass through a delay, then a gain and finally a filtering block. The nature of the filtering block (i.e. IIR or FIR) depends on the type of Setup selected.

One important feature to keep in mind is that Left and Right signals are processed by DSP core 0, meanwhile Sum and Diff signals are processed by DSP core 1. At the core level, filtering resources are allocated dynamically, hence adding filters to the path of signals Left and Right will increase the load on DSP core 0. Similarly, adding filters to the path of signals Sum and Diff will increase the load on DSP core 1.

The *Channel Processing section* consists of a routing block, in which any of the 4 pre-processed signals (Left, Right, Sum and Diff) can be routed to any channel, followed by a pre-gain, a filtering block, a post-gain, a phase inverter and finally a delay. Again, the nature of the filtering block depends on the selected Setup type.

The Channel Processing section is visualised in the Channel Processing tab.

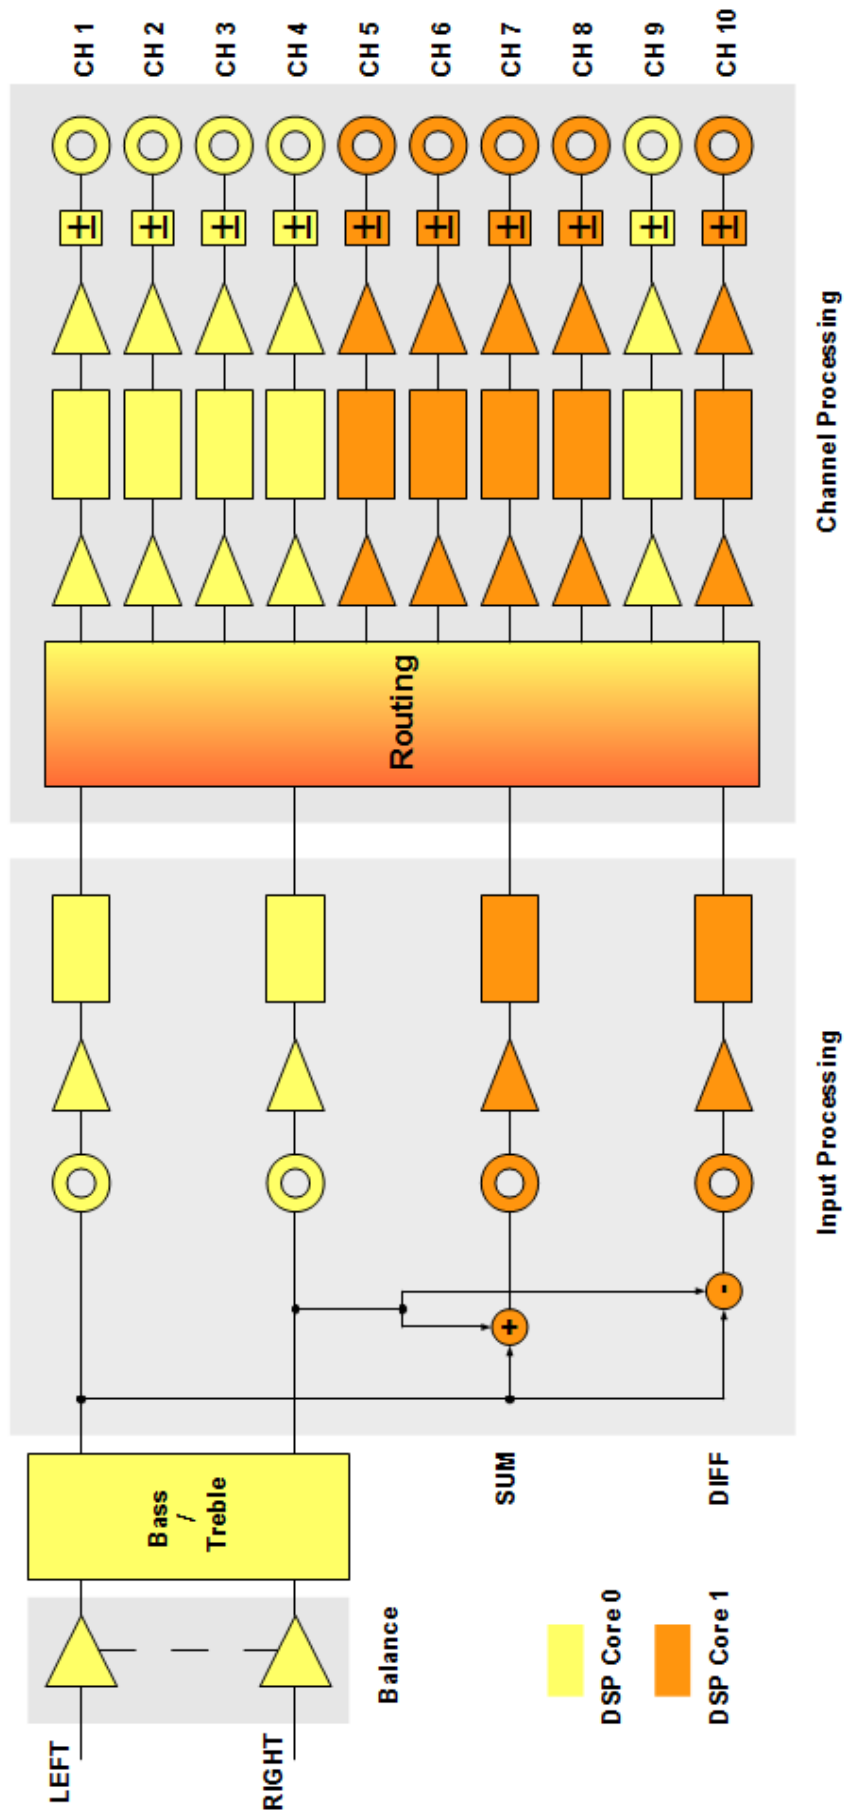


Fig. 5: General processing structure common to all setups



*Why are there 2 gains on each channel? (pre-gain vs post-gain)*

A proper gain structure is essential to preserve, as much as possible, the information contained in the original audio signal. Use the pre-gain to set signal amplitude to the limit of clipping in the subsequent filtering block, then use the post-gain to scale signals fed to your drivers.

In other words, pre-gain is solely connected to the nature of the filtering block and must be chosen as large as possible but not too large so as not to clip. Pre-gain dB values will usually be negative (i.e. attenuating the signal). Post-gain purpose is two-folded; on the one hand it's going to compensate for different pre-gain settings on different channels that have same routing, on the other hand it's going to compensate for driver unequal efficiencies. Post-gain dB value will usually be negative as well. Generally, although not always, a positive gain (in dB) is the clue of an improper gain structure.

Note that if your board includes the analogue volume chip, there's a third gain stage doubling with the post-gains, but taking place in the analogue domain. Where possible, analogue gains should be preferred to digital ones. Refer to section VII.2.

*I've set up a system with crossover filters only, i.e. no EQ, no gain. The frequency response exhibits no boost anywhere, however the signal is clipping. Why is that?*

The frequency response shows how the filters are handling steady-state signals. It doesn't show how the system is handling transients.

A steady-state signal is a sine or a collection of sines. Music and speech however contain a lot of transients.

In order to assess how the filters handle transients, we would look at other types of plots, such as for example the step response. Filters usually exhibit on these plots a certain amount of overshoot. If your listening material is normalized for 0 dB<sub>FS</sub> – as is most commercial material, then it's necessary to provide some headroom for the processing. Indeed, any amount of overshoot on a 0 dB<sub>FS</sub> transient will lead to clipping otherwise. The required headroom depends on the filters; -3 dB pre-gain is a good starting value.

*What about DSP load?*

As for the Input Processing block, Channel Processing load is split between the DSP cores. Fig. 5 shows DSP core 0 is in charge of channels 1, 2, 3, 4 and 9, meanwhile DSP core 1 processes channels 5, 6, 7, 8 and 10. Within a core, filtering resources are allocated dynamically so that additional filtering on channels 1, 2, 3, 4 and 9 will increase the load on DSP core 0 meanwhile additional filtering on channels 5, 6, 7, 8 and 10 will increase the load on DSP core 1.

If your Setup is not processing-intensive, then you don't need to care too much about DSP load as long as you make sure it's less than 100 % for each core. Indeed, what a DSP load in excess of 100% actually means is that the DSP is dropping samples (expect alias and distortion!).

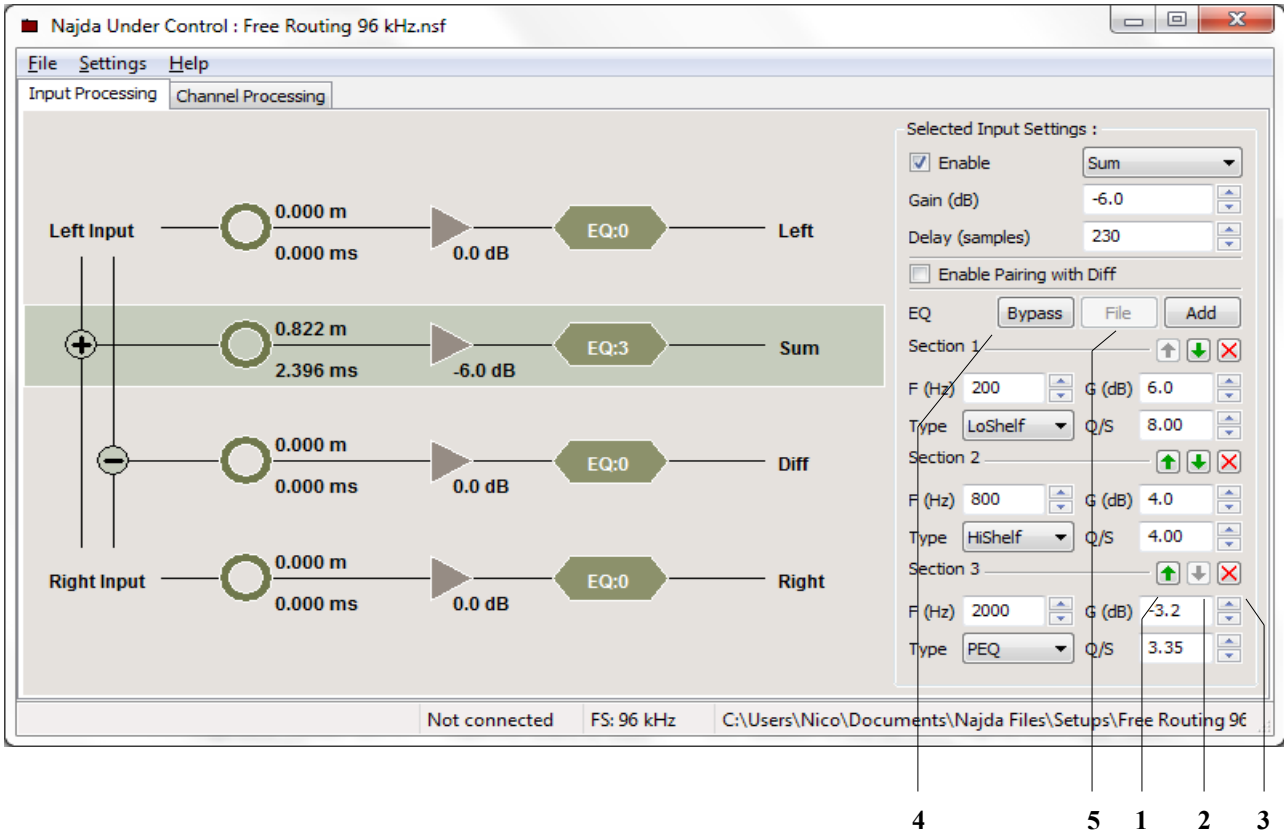
In the opposite case, if you know your Setup is going to require lots of resources (because you favour 192 kHz sampling rate, or your speakers have many ways, or you have designed long FIR filters ...), then you will want to balance the load on each core. This requires that you consider carefully how you're going to distribute the speaker channels between cores, which on its turn is going to impact the physical connections to Najda board.

Generally, your Setup is going to be quite symmetrical between left and right, so a good start is to reserve one core for the left speaker and one core for the right speaker. If you have a sub in your

Setup, put it on core 1. Indeed, core 1 has usually fewer side tasks than core 0.

Note that delay and gain settings, phase inversion and routing do not impact DSP load. In particular, a delay of N samples is identical to a delay of 0 sample in terms of DSP computational resources.

### III.2 Classic Free Routing



**Fig. 6:** *Input Processing tab in Classic Free Routing mode. Use the up (1) and down (2) arrows to reorder the parametric filters in the processing chain. To suppress a filter, click on the cross (3). The whole EQ section can be bypassed (4). Alternatively to parametric EQs, filtering can be defined by raw coefficients contained in a file (5).*

The Classic Free Routing Setup type is in essence the most common kind of processing strategy that's made available on marketed DSP cross-over units. This section will describe the peculiarities of the commands in the control application and will provide an overview of the DSP implementation.

The figure above shows the Input Processing tab of the IIR (aka Classic Free Routing) mode, where the SUM processing path has been highlighted.

In order to highlight (i.e. select) a signal, click with the mouse on the block diagram at the appropriate height, or select the input in the top right combo box. The remaining components in the right panel are immediately filled with the current user settings.

In the example above, there's a 230-sample delay, which corresponds to a 2.396 ms delay at 96 kHz sampling frequency (FS field in status bar). One can change the delay by entering a new value in the right panel. An alternate way is to click on the ring in the block diagram and use the up and down keyboard keys to alter the setting. Generally, clicking on a specific block will give software

focus to the associated component in the right panel so that you can immediately alter the parameter value.

The pre-gain here is set to -6 dB. In order to change the pre-gain setting, type in a new value in the right panel, or click with the mouse on the gain symbol (the triangle) and use the up and down keys.

The same processing path shows that 3 biquads are in use. A glance at the right panel reveals that these filters are parametric EQs defined each by a set of data.

Note that signal paths that are not used can be disabled in order to save processing resources. Also, input channels can be paired so that the changes made are reflected immediately on the paired channel too.

### III.2.1 Built-in EQ types


The first filter in the processing chain is a low-pass shelving filter with 200 Hz characteristic frequency, 6 dB gain and a slope factor  $S$  of 8.

The second filter is a high-pass shelving filter with 800 Hz characteristic frequency, 4 dB gain and a slope factor  $S$  of 4.

The characteristic frequency is defined as the filter mid-excursion point.

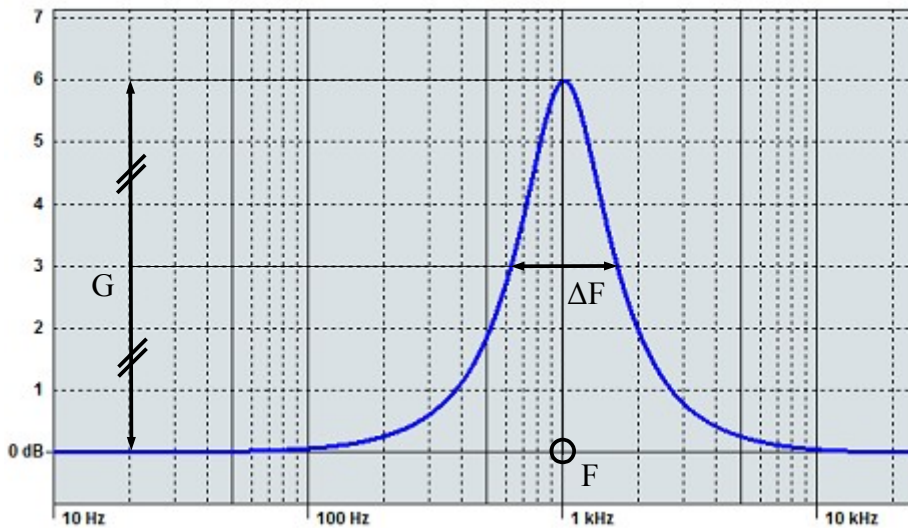
The slope factor is limited to the range 3 to 12, with the latter being the steepest slope achievable.

The slope of the filter, in dB/oct, is proportional to the  $S$  factor.



**Fig. 7:** Illustration of the shelving filter parameters.  
*This is a low-pass at 1 kHz, 6 dB gain and a slope factor of 8.*

The third and last filter in figure 6 is a peaking filter set at 2000 Hz, with a negative gain of -3.2 dB and a quality factor  $Q$  of 3.55. The higher the  $Q$ , the narrower the filter. See figure 8.

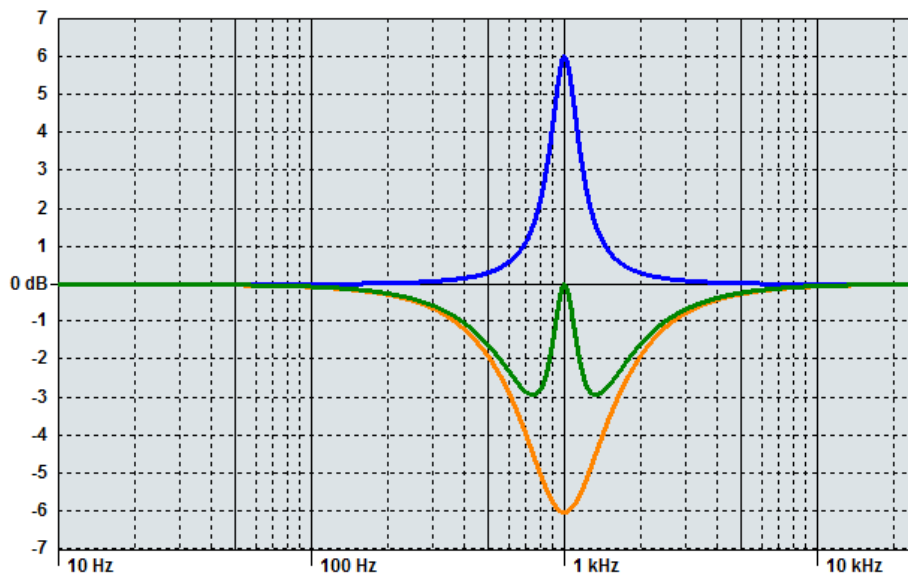


**Fig. 8:** Illustration of the peaking filter parameters.  
 This one is centred about 1 kHz, 6 dB gain and a quality factor  $Q = F/\Delta F$  of 1.

### III.2.2 EQ sections ordering

As illustrated in fig. 6, EQ sections can be reordered. Although EQ sections ordering doesn't impact what you see on the frequency response plot, it does impact the amount of headroom you need to allocate and, hence, the sound quality of your filtering strategy.

Let's see now a simple example of headroom to channel ordering dependency. Fig. 9 below shows two peaking filters centred about 1 kHz. The blue line is 6 dB gain with  $Q = 3$ , and the orange line is -6 dB gain with  $Q = 1$ .



**Fig. 9:** In blue, PEQ 1 kHz 6 dB  $Q=3$ . In orange, PEQ 1 kHz -6 dB  $Q=1$ . In green, their combination.

The cascaded combination of the two filters is shown in green. It's the same regardless of sections ordering – i.e. regardless of whether the signal is first processed by the 'blue' section then by the 'orange', or the way around.

If you choose to filter through the blue section first, then you need to allocate at least 6 dB of headroom. It doesn't matter that the boosting filter is followed by a compensating dip: in absence of the specified headroom, the 'blue' filter will clip in the first place and the *ovf* segment of that channel will light on in the Monitor. Note that by allocating 6 dB of headroom, you're shaving off one bit of resolution from the 24 original resolution bits of your signal.

Consider now the reversed ordering. The orange filter doesn't require any headroom other than transient overshoot buffer (see transients handling on p.7). Also, the orange filter attenuates that same band that the blue filter is boosting. As a consequence, this section ordering doesn't call for specific headroom allocation and should be preferred to the former.

The case above is synthetic and trivial. Real cases are likely to involve more sections and require a more careful look as the number of sections increases (the number of possible ordered sequences grows as the factorial of the number of sections!). The aim of section ordering is always to minimize headroom while prevent clipping.

EQ section ordering is only available with built-in EQs, it's not available with custom biquads. Indeed, with custom biquads, sections are ordered as per the supplied file.


### III.2.3 Custom biquads

Each filtering block consists of N cascaded second-order IIR sections, also known as biquads. For those inclined to maths, the Z-transform of such a complete block is given by ( $N \leq 9$ ):

$$H(z) = \prod_{n=0}^{N-1} \left( \frac{b_{0,n} + b_{1,n} \cdot z^{-1} + b_{2,n} \cdot z^{-2}}{1 + a_{1,n} \cdot z^{-1} + a_{2,n} \cdot z^{-2}} \right)$$

where  $b_{0,n}$ ,  $b_{1,n}$ ,  $b_{2,n}$ ,  $a_{1,n}$  and  $a_{2,n}$  are the coefficients of section  $n$ .

Each biquad is implemented in Direct Form 1 with all coefficients being 48-bit. This is illustrated in the figure below.



**Fig. 10:** *Najda's biquad implementation. Note the – sign in front of  $a_1$  and  $a_2$  coefficients. Incoming audio is 24-bit, outgoing audio is rounded back to 24-bit but the whole the internal path is 48-bit.*

Alternatively to defining parametric filters, one can load raw biquad coefficients contained in a text file. This feature, known as *custom biquad*, allows implementing types of filters other than the ones readily available, such as Linkwitz transformers for example. We'll describe now the required file format.

A maximum of 15 biquads may be loaded in each EQ section. When opening a file, Najda Under Control expects to find the coefficients in a column in the following sequence:  $b_0$ ,  $b_1$ ,  $b_2$ ,  $a_1$ ,  $a_2$ . Note in particular that  $a_0$  is implicitly unity and is therefore omitted. 5 coefficients fully describe one biquad, thus the file can contain up to 75 coefficients. A file containing more than 75 coefficients will be rejected, as will be a file containing any number of coefficients which is not a multiple of 5. See Appendix A for examples of compatible files.

Coefficients are checked for stability: if the file contains one or several unstable biquads, then the whole file is rejected. Also, the global gain of each biquad cannot exceed 24 dB – so the file will also be rejected if any coefficient is greater than  $2 \cdot 10^{24/20} \approx 31.69$ .

After we tested several third-party applications generating biquad coefficients, we found out that they use a different sign convention. Indeed, all these applications compute coefficients for second-order sections defined by (note the '-' signs in the denominator):

$$H_{biquad}(z) = \frac{b_0 + b_1 \cdot z^{-1} + b_2 \cdot z^{-2}}{1 - a_1 \cdot z^{-1} - a_2 \cdot z^{-2}}$$

In order to allow the loading of files generated by these applications, NUC will invert the signs of all  $a_1$  and  $a_2$  coefficients. This is important to keep in mind only if you intend to write your own routines for coefficients computation (you don't need to care otherwise).

Parametric EQs and custom biquads are mutually exclusive: this means that you can't have both in the same EQ block. The bypass switch is effective also with custom biquads, so it's safer to reserve these for filters others than crossovers.

It's often useful to inspect visually the files generated by third party applications. Indeed, if you spot a section defined by  $b_0 = 1$ ,  $b_1 = 0$ ,  $b_2 = 0$ ,  $a_1 = 0$ ,  $a_2 = 0$  – just delete that section because it's a pass-through that uses DSP resources without any effect. Similarly, if you spot a section with following coefficients  $b_0 = \beta$ ,  $b_1 = 0$ ,  $b_2 = 0$ ,  $a_1 = 0$ ,  $a_2 = 0$  where  $\beta$  is any number, then this section is a simple gain of value ( $20 \cdot \log_{10} \beta$ ) in dB: see if you can suppress that section and report the corresponding gain in one of the pre- or post-gains.

### III.2.4 More controls (Channel Pairing and Channel Enabling)

Channels may be paired in order to speed up the definition of a setup when there's a certain amount of symmetry between the left and right speakers.

On the Input Tab, the left and right channels can be paired, as can the sum and diff channels.

On the Output Tab, the default pairing is as follows: 1 ↔ 5, 2 ↔ 6, 3 ↔ 7, 4 ↔ 8, 9 ↔ 10.

This is the recommended pairing strategy as it helps maintaining a balanced loading of the DSPs (refer to Fig.5).

An alternate pairing strategy has however been made available where the paired channels are laid in a successive order, i.e 1 ↔ 2, 3 ↔ 4, 5 ↔ 6, 7 ↔ 8, 9 ↔ 10. Switching between pairing strategies is achieved by clicking on the pairing type button (see fig. 11).

Pairing two channels doesn't necessarily result in the channels being identically parametrised.

Indeed, you can first pair the channels and apply the common settings, then unpair and apply individual settings to each channel.

Each processing path also features an Enable switch. It's recommended to disable all the channels

that are not going to be used at all in the final form of the setup. By disabling unused channels, you lower the load on the DSPs. By default when creating a new setup, all channels are enabled.

### III.2.5 Classic Free Routing specifications

Processing block	Implementation
Gains	-24 dB to 12 dB
Input Processing delays	All up to 8191 samples except DIFF 4095 samples
Output Processing delays	All up to 4095 samples
Crossover filters	Butterworth 6, 12, 18, 24, 36, 48 dB/oct Bessel 12, 18, 24, 36, 48 dB/oct Linkwitz-Riley 12, 24, 36, 48 dB/oct
Max number of second-order sections per EQ block	15
EQ	Peaking Low shelf High shelf Custom
Coefficient bit width	48-bit
Phase invert	+/-

**Tab. 1:** Summary of the Classic Free Routing processing blocks. Bessel filters are phase-matched.

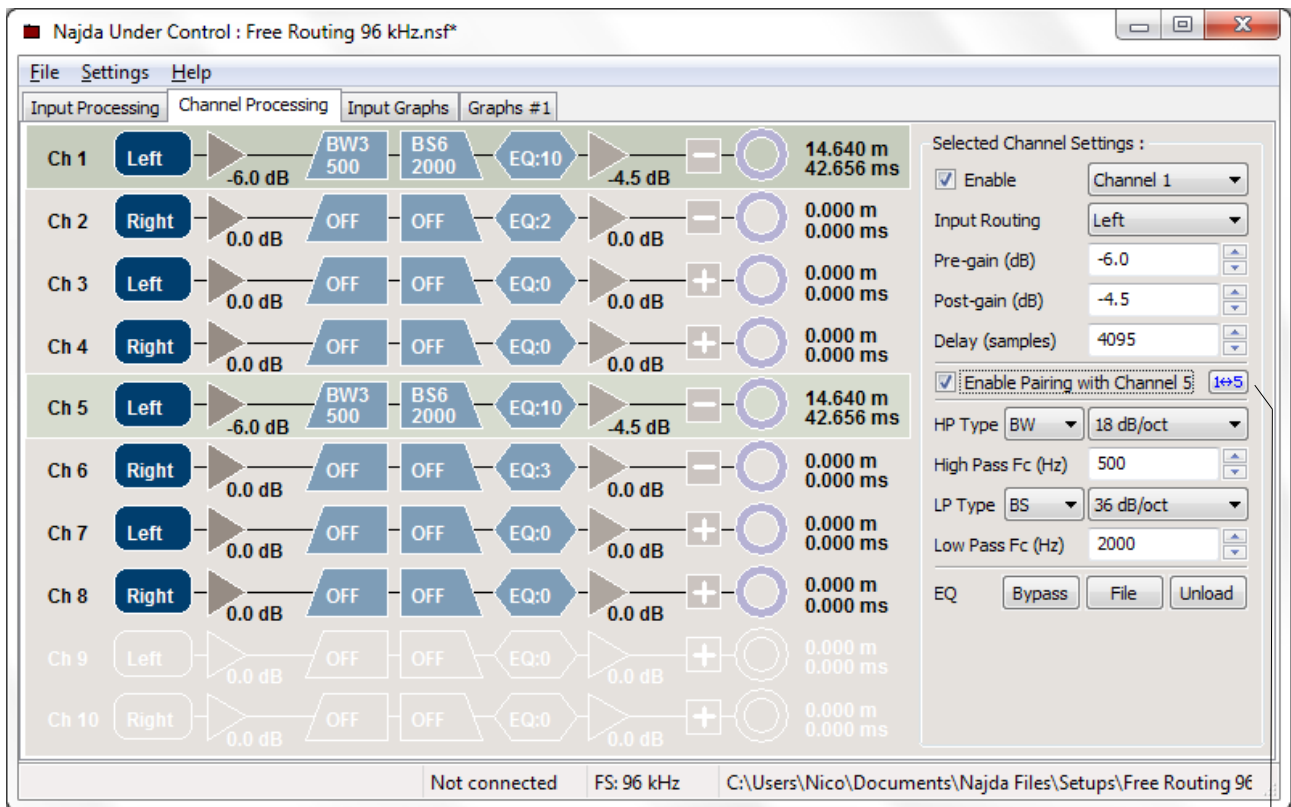
### III.2.6 Default Classic Free Routing Setup

Upon creating a new Classic Free Routing Setup, NUC will create a setup file with default parameters. The default Classic Free Routing setup file implements a *pass-through*, routing left input to channels 1, 3, 5, 7, 9 and right input to channels 2, 4, 6, 8, 10.

In particular:

- All delays set to 0 sample.
- All gains (pre- and post-) set to unity (0 dB).
- All crossover filters OFF (flat unity response).
- All filtering blocks contain no filter, they just let the signal pass through unaltered.
- All polarities positive.

When you create a new Setup in *Synced* mode, NUC will immediately transfer all default parameters and mute outputs as a safety measure. Before you unmute, you must make sure your setup is not likely to damage your speakers. A good policy is to keep the volume low at design time – or design the setup in *Offline* mode.



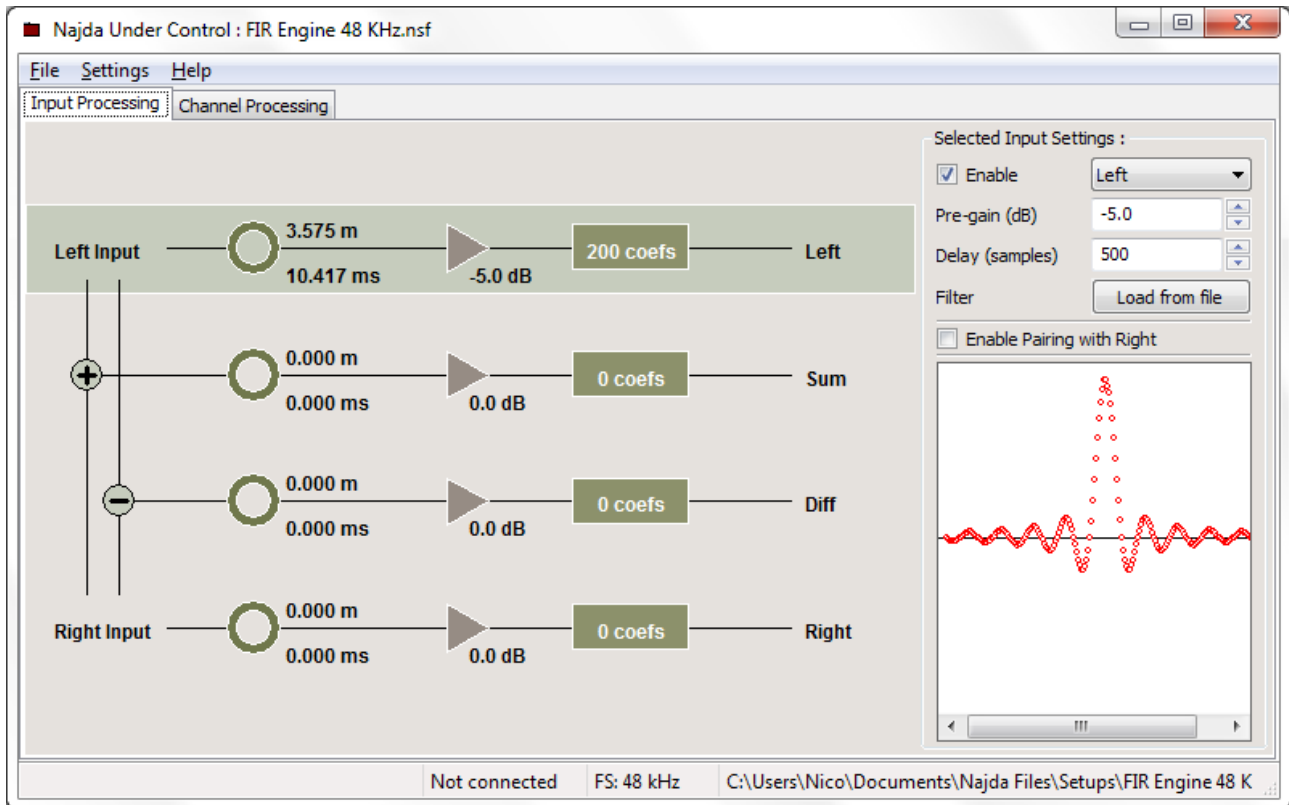
Pairing type switch

**Fig. 11:** View of the Output Processing tab, with channel 1 highlighted.  
 Note that 10 custom biquads are in use and that phase is inverted and that channel 1 is paired to channel 5.  
 Also, Channels 9 and 10 have been disabled.



### III.3 FIR Engine

The FIR Engine mode is identical to the Classic Free Routing mode (section III.2) except that all filtering blocks are of FIR type. In the screenshot below, 200 coefficients have been loaded in the left input processing path. The same processing path also carries a 500-sample delay and a -5 dB gain.



**Fig. 12:** *Input Processing tab in FIR Engine mode.*

#### III.3.1 FIR Filtering Model

The FIR Engine mode implements a direct convolution in each filtering block, defined as:


$$H(z) = \sum_{n=0}^{N-1} h[n] \cdot z^{-n}$$

where  $H$  is the Z-transform and  $h$  is the impulse response of the filter of length  $N$ .

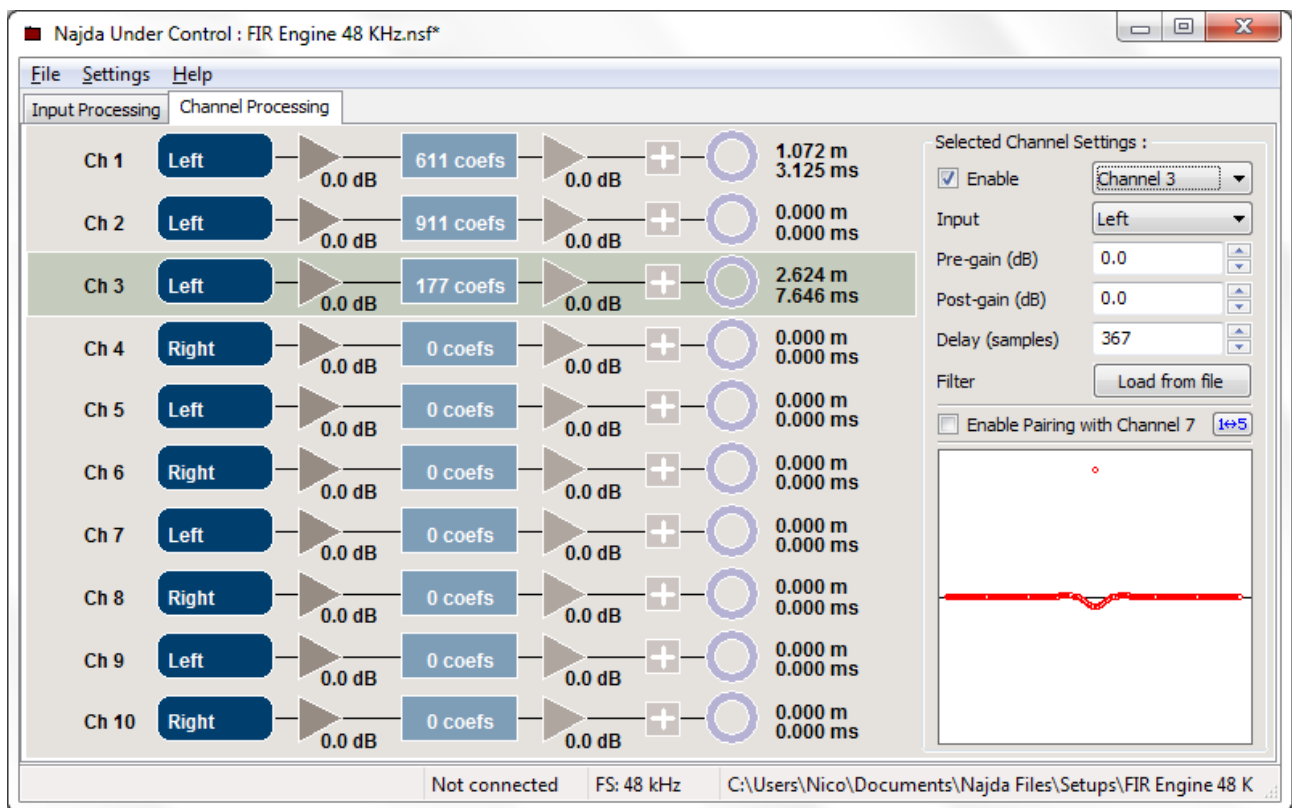
Najda implements the convolution in its direct form, depicted in figure 13. Raw FIR coefficients must be supplied in a text file – refer to Appendix B for examples of such files. All coefficients must be in the interval  $] -1, 1 ]$ , i.e. 1 is allowed but -1 is not. See table 2 for the number of coefficients allowed; files containing more coefficients will be rejected.

We recommend creating a pair of simple FIR files, which will be useful in various situations.

The first file is simply an empty file and could be named *fir0.txt*. It implements an open-circuit and can be used to unload coefficients you wish to discard.  
 The second file contains only '1' and can be named *fir1.txt*. It implements a pass-through. See Appendix B for more details.



**Fig. 13:** Najda's implementation of the convolution. The accumulator is 56-bit wide (8:48).



**Fig. 14:** View of the Output Processing tab in FIR Engine mode, implementing here a 3-way crossover on channels 1, 2 and 3. Note the routing section on the left.

### III.3.2 FIR Engine specifications

Processing block	Implementation
Gains	-24 dB to 12 dB
Input Processing delays	All up to 8191 samples except DIFF 4095 samples
Output Processing delays	All up to 4095 samples
FIR length	Max 1023 coefficients on Left, Right, Sum, Ch. 1, 2, 3, 5, 6, 7. Max 511 coefficients on Diff, Ch. 4, 8, 9, 10.
Coefficient bit depth	24-bit
Phase invert	+/-

**Tab. 2:** Summary of the FIR Engine mode processing blocks.

### III.3.3 Default FIR Engine Setup

Upon creating a new FIR Engine Setup, NUC will create a setup file with default parameters. The default FIR Engine file implements a *block-all*, and routes left input to channels 1, 3, 5, 7, 9 and right input to channels 2, 4, 6, 8, 10.

In particular:

- All delays set to 0 sample.
- All gains (pre- and post-) set to unity (0 dB).
- All filtering blocks contain no coefficient, hence they block the signal and act as open circuits.
- All polarities positive.

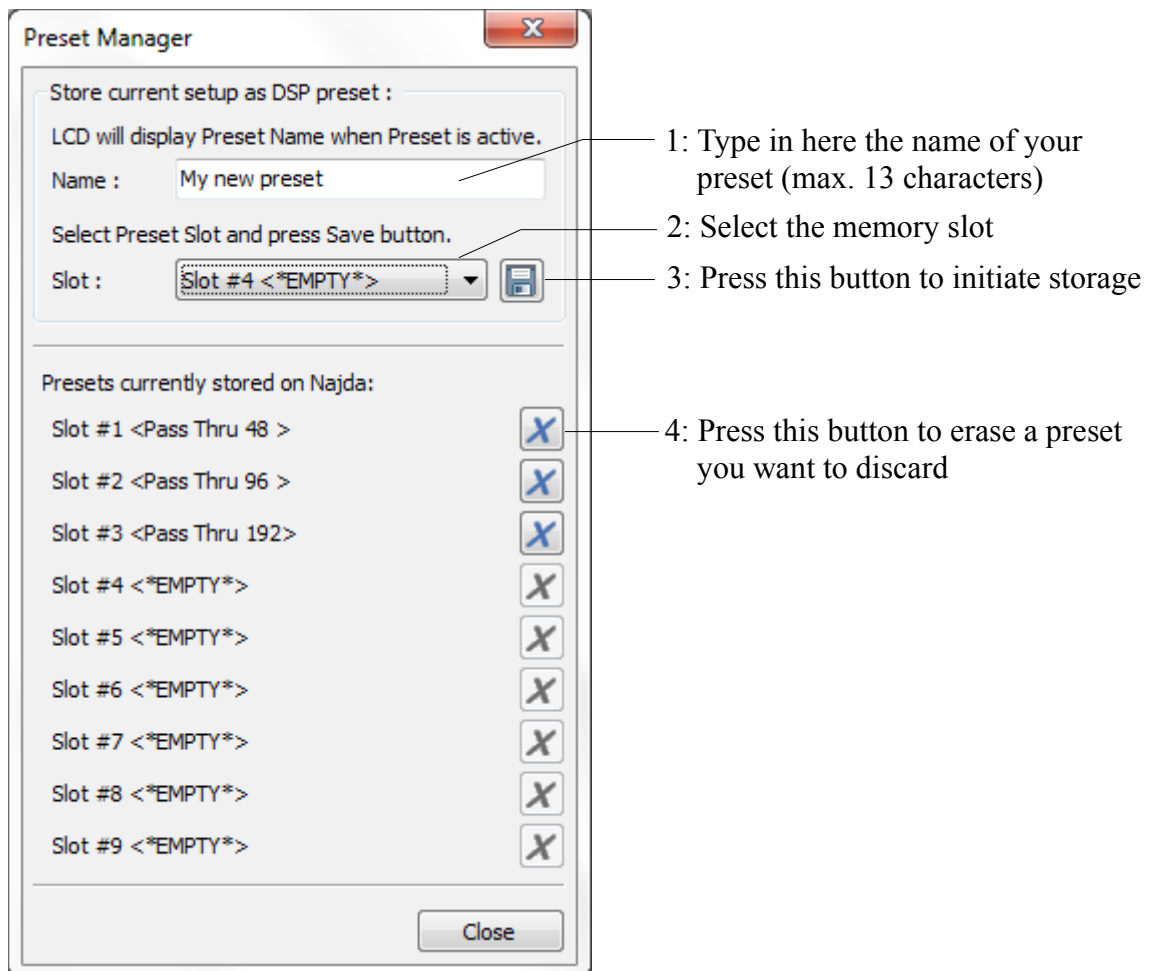
In order to create a pass-through FIR, you must load the unity filter (see file *fir1.txt* described in section III.3.1) into each filtering block.

## IV. Preset Management

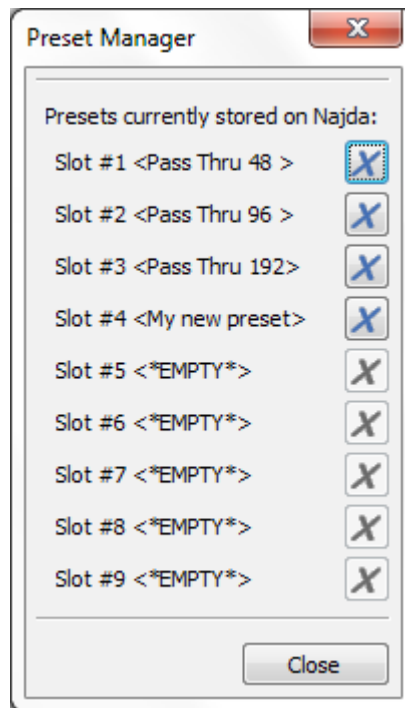
Up to 9 presets can be stored into Najda onboard memory. A limitation applies to FIR Engine presets, which can potentially involve large amounts of data; FIR presets can only be stored into memory slots 1 to 3. Classic Free Routing (IIR) presets can be stored into any slot.

In order to store or delete a preset, it should be fairly obvious that Najda must be connected. As a matter of fact, the menu item (**File** → **Preset Manager**) is greyed out and the corresponding dialogue is not available until connection is established.

The Preset Manager dialogue in Synced mode is shown in figure 15 below. This dialogue displays the name of the preset stored in each of the 9 memory slots and lets you transfer the current setup into the preset slot of your choice. It also allows you erasing a preset you'd like to discard. Note the transferred setup corresponds to the most current configuration (i.e. it includes the changes you wouldn't have saved in the setup file). The name typed in this dialogue will appear on the LCD display when the preset is active.



**Fig. 15:** *Preset Manager dialogue in Synced mode*



**Fig. 16:** *Preset Manager dialogue in Online mode*

The appearance of the Preset Manager dialogue is slightly different in Online mode, as pictured in figured 16. In Online mode, you cannot store a preset because you haven't opened a setup file. The Preset Manager dialogue still displays the list of presets and lets you erase the ones you'd want to discard.

As a final note, there must be at least one preset stored into Najda's memory otherwise the unit wouldn't be able to boot. This explains why the Preset Manager dialogue will only allow you erasing all but one preset.

## V. Graph Utility

NUC includes a flexible graph utility that displays the expected processing outcome. One single Graph tab may be opened in relation to the Input Processing block, and up to five Graph tabs are allowed for the Channel Processing block. Having one separate Output Graph tab per speaker box is a healthy approach.

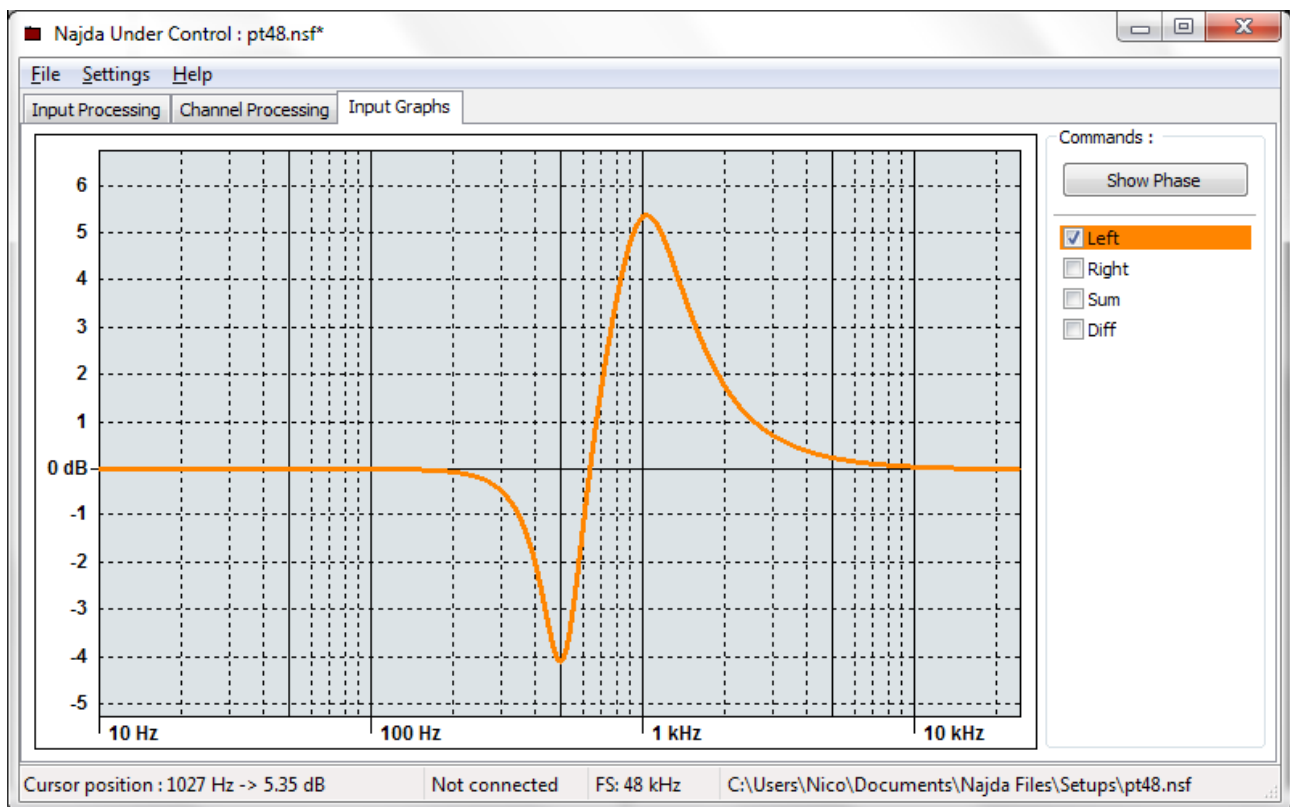
### V.1 Input Processing Graph

To display the Input Graph tab, select in the menu **File** → **Show Input Graph Tab** (or press Alt-i). This menu item is available only if you have opened or created a setup file.

Select in the right panel what processing paths you wish to view (Left, Right, Sum and/or Diff).


By default, this graph shows the magnitude of the processing in the range +/- 24 dB, with the 0 dB line in the middle of the window.

With the mouse wheel (or the scrolling feature of the mouse pad), you can alter the range of magnitudes. You can also drag the 0 dB line up or down (left click on the graph and move the mouse up or down). By double-clicking on the graph, you're restoring the initial graph view (+/- 24 dB).



**Fig. 17:** *Input Graph tab showing magnitude response of filters applied to the Left input signal.*

The button in the top right corner allows toggling between magnitude and phase views (fig. 18). There's no zooming nor dragging feature in the phase view.



**Fig. 18:** Phase view of the above filters.

Note that while hovering over the graph, the mouse coordinates are displayed in the status bar.

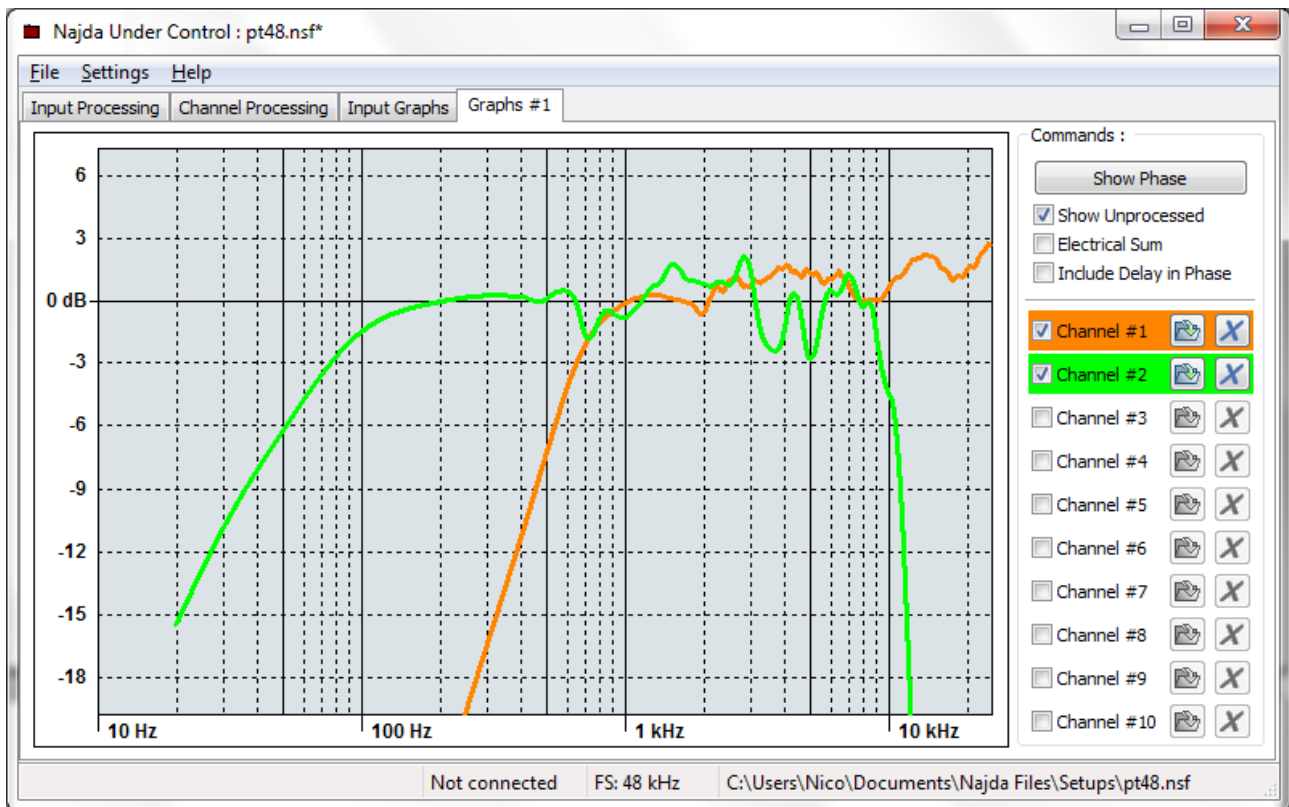
## V.2 Output Processing Graph

The Output Graph utility is functionally similar to the former Input Graph. However, it has a few extra features described hereafter. To open a new output graph tab, select **File** → **Add Output Graph tab** in the menu (or press Alt-g).

### V.2.1 Loading of frequency responses

In the screenshot below, we have set a graph to display the individual responses of the drivers forming a 2-way bookshelf speaker. Channel 1 is dedicated to the tweeter (orange) and channel 2 holds the woofer's signal (green).

To load a frequency response file, click on the button near the channel index number. To unload a file, click on the button with a cross in it. The expected file format is CSV-like, with the first column being the frequency, the second column being the measured magnitude in dB at the former frequency and the optional third column being the measured phase in degrees. See Appendix C for examples of compatible files.



**Fig. 19:** Output graph loaded with the frequency responses of the drivers for a 2-way project.

In order to accommodate measurements by users who don't own a SPL calibration reference, the output graph utility embeds a routine for automatically aligning level-wise the plots of several drivers. This routine works the following way:

- If only one single frequency response is loaded, the output graph will offset the curve so that it crosses the 0 dB line at 1 kHz.
- If successive frequency responses are loaded, the graph utility will pick as the reference the response that exhibits the largest magnitude at 1 kHz. The remaining curves will be offset so that the relative magnitude difference is maintained.

Here's the recommended approach for measuring all the drivers of one speaker in their enclosure:

- Position your measuring microphone (at 1 m in the axis, off-axis or at listening position...).
- Select a suitable level for the test signal. It has to be high enough for the measurement purpose, but not too high so you won't damage any of the drivers.
- Measure each driver successively *without moving the microphone* and *without altering the test signal level*.

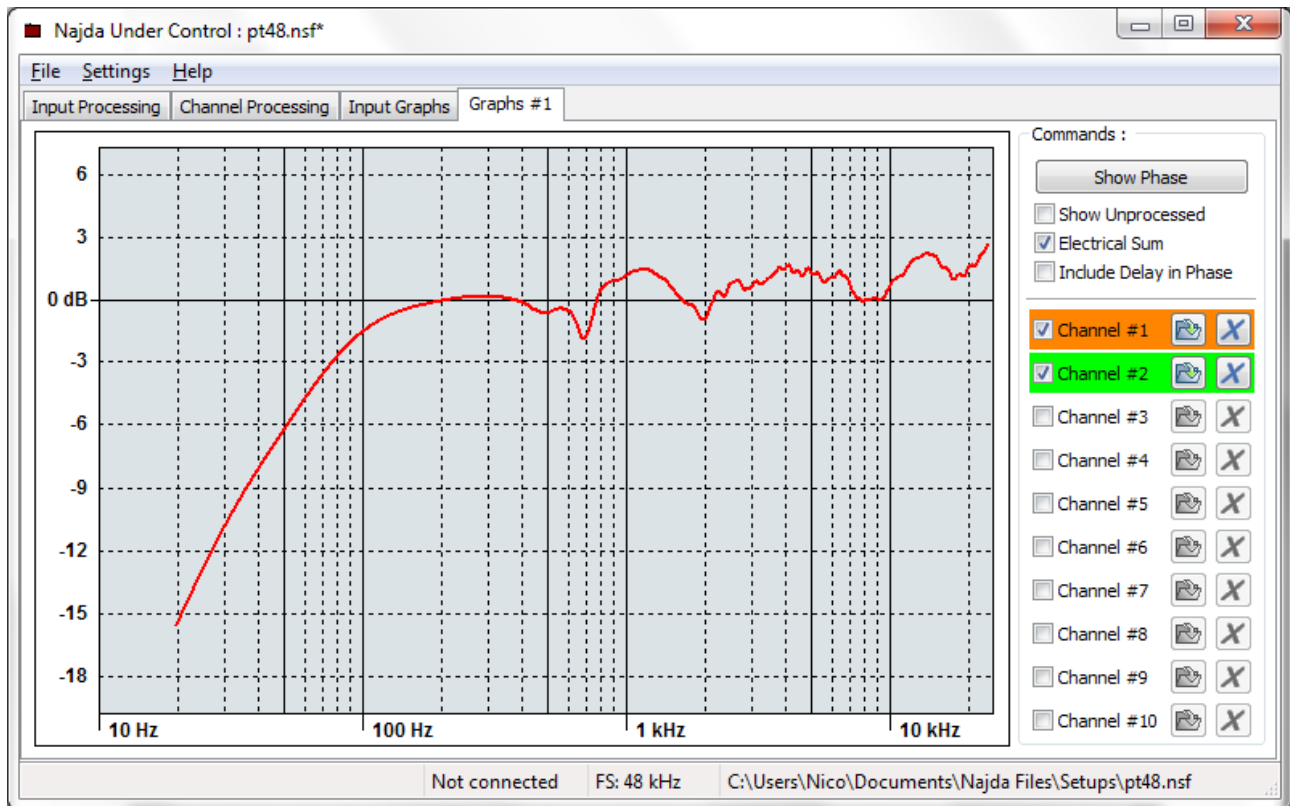
### V.2.2 Processed/unprocessed views

The graph utility displays the global outcome of the processing strategy, i.e. the combined effect of the filters, pre- and post-gains, phase reversals and, optionally delays (see V.2.3 Electrical Sum). In order to quickly assess how the setup is actually performing, you can toggle between the processed and unprocessed views using the *Show Unprocessed* checkbox.



## V.2.3 Electrical sum

The output graph also includes a summing function that allows viewing how all the drivers on that tab will behave jointly. In the screenshot below, we have set a LR crossover at 700 Hz to the pair of drivers in fig. 19.



**Fig. 20:** Electrical sum of the 2 drivers in fig. 19 with a LR network at 700 Hz.

Phase reversal will not impact the magnitude responses of individual drivers. It will however dramatically impact the magnitude summed response. Similarly, the delays will not impact the individual drivers' responses, but it'll have a dramatic effect on the sum.

We have elected to display optionally the delay effect because in some circumstances it's relevant to crossover network design, and sometimes not.

If your network design is IIR based, then you're most probably going to use delays for aligning temporally drivers that are offset acoustically because they're mounted on the same plane panel. In this case you should leave the delay out of the sum.

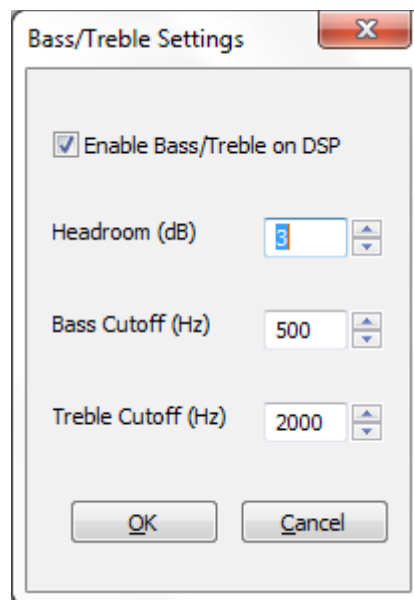
On the other hand, if your design is based on FIR filters, then you will also use delays to compensate for the inherent delay difference exhibited by filters of different lengths. In this one case, you should at some point include the delay in the sum in order to make sure the network is behaving according to your expectations.

## VI. Bass/Treble Settings Dialogue

Najda includes a Bass/Treble unit (see figure 5) that acts like a tone control independent of the processing parameters. These filters are 1<sup>st</sup> order shelving filters implemented according to the Mitra topology; they allow boosting or attenuating bass and treble content in dB steps by up to +/- 6 dB and offer some configuration options via the Bass/Treble Settings dialogue.

First thing first: this unit can be disabled. This is beneficial if you don't intend to use any tone correction as you save the corresponding DSP computational resources. In order to enable or disable the Bass/Treble unit, tick or untick the checkbox on top of the dialogue (fig. 16).

Additionally, the corner frequencies of the shelving filters can also be customized. Corner frequencies are defined as the mid-excursion point in the transition band (see fig. 7). Finally, as for any filter, the Bass/Treble unit requires some headroom in order to prevent clipping. The amount of required headroom depends on the amount of boost you intend to apply, so the headroom is also configurable via the Bass/Treble Settings dialogue.



**Fig. 21:** *Bass/Treble Settings dialogue*

Bass/Treble settings are independent of the processing parameters. They apply to all the presets stored in Najda and must be seen as a global parameter subset.

## VII. Hardware Configuration Dialogues

This last paragraph describes the remaining hardware configuration dialogues. As for the Bass/Treble unit (see section VI), these settings act globally and are not tied to any preset.

### VII.1 Expansion Port 0 Routing

Expansion Port 0 comprises 3 I<sup>2</sup>S lines that can convey 6 audio channels and are useful if you wish to experiment with external DACs. Any of the 10 processing channels can be freely routed to these lines.

In the example screenshot below, channels 9 and 10 are routed to the I2S\_EXT0\_LINE1 pin of the expansion port. One use of such a setting could be an external DAC section with headphone amp. Please refer to the *Startup Guide* for Expansion Port 0 pinout.

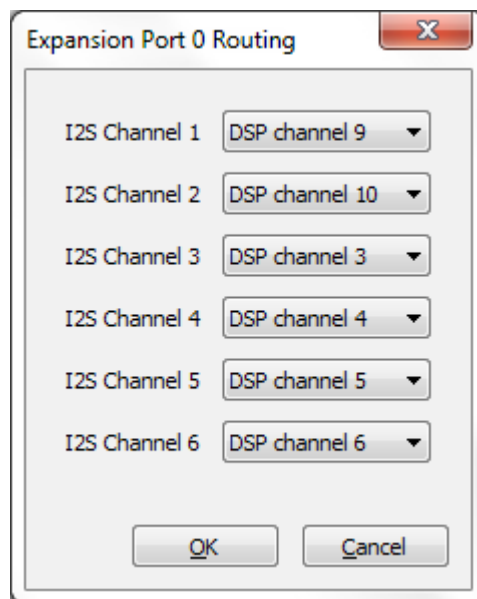


Fig. 22: Expansion Port 0 Routing dialogue

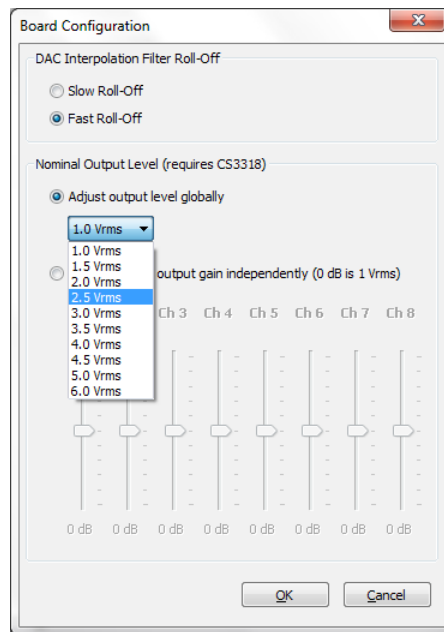
### VII.2 Analogue Outputs Configuration

This dialogue allows selecting the onboard DACs interpolation filters and adjusting the output nominal level.

Two sets of interpolation filter coefficients are made available with the DACs in the CS42428 chip, they differ in the steepness of the transition band. For details on these filters, please refer to the Cirrus Logic CS42428 datasheet.

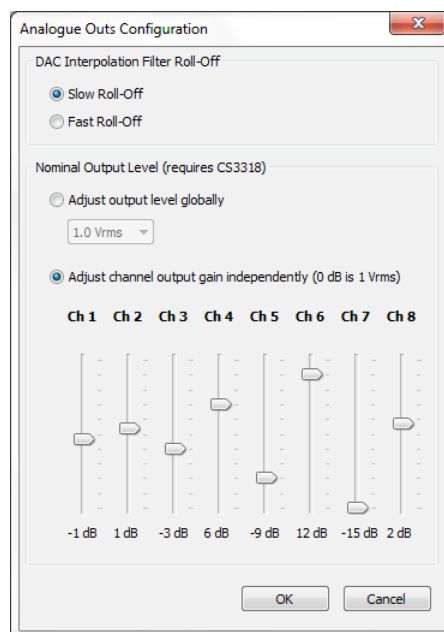
In the same dialogue, you can adjust the nominal output level, either globally for all channels or independently channel-wise.

You'll want to adjust the output level globally when you wish to adapt Najda's output to the analogue level expected by your power amps. Please check your amps documentation for their nominal input level. Global output level can be set anywhere between 1 V<sub>rms</sub> and 6 V<sub>rms</sub>.



**Fig. 23:** *Analogue Outputs Configuration dialogue for global nominal output level setting*

Alternatively, you'll want to set output levels independently when you wish to substitute digital post-gains with analogue gains – and therefore compensate in the analogue domain for drivers' efficiency discrepancies. Output level can be adjusted between -15 dB and +15 dB (0 dB is 1 V<sub>rms</sub>).



**Fig. 24:** *Analogue Outputs Configuration dialogue for independent channel output level setting*

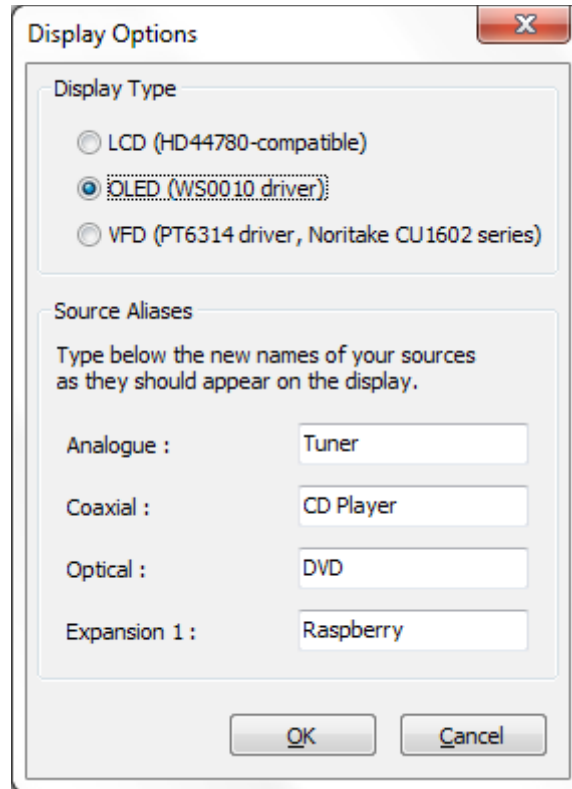
### VII.3 Display Options

The purpose of this last dialogue is two-folded.

Since the Najda supports 3 types of displays, namely LCD HD44780 or compatible, OLED

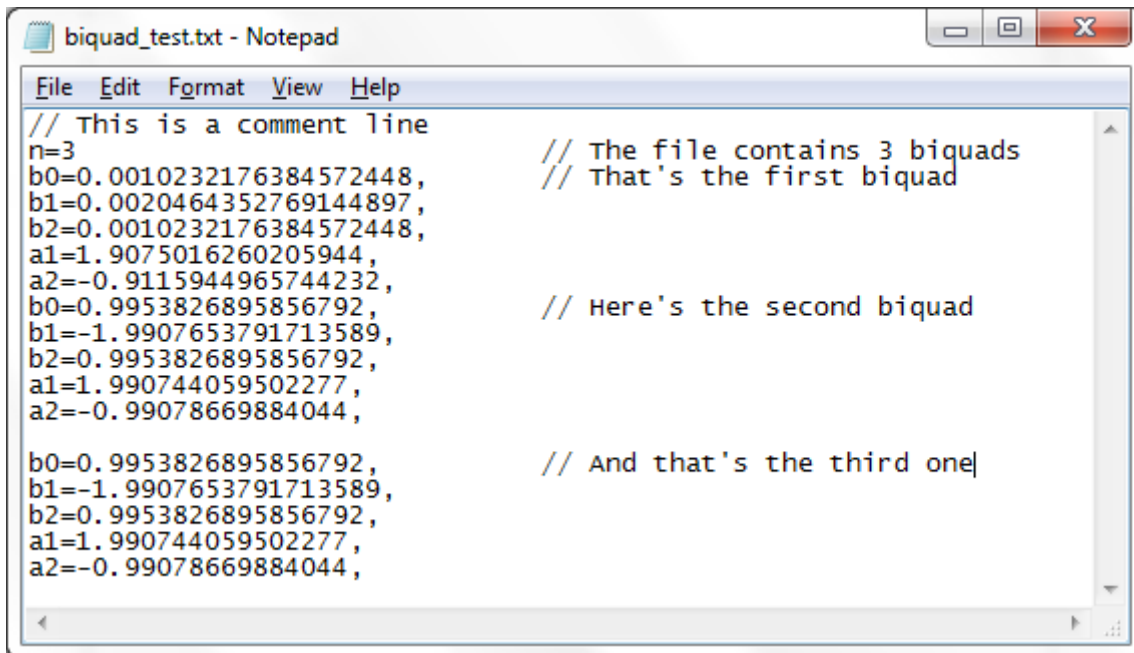
WS0010 and VFD PT6314, it's required to let the board know what display has been installed. Declaring the wrong type will result in display artefacts and inappropriate menu options. Setting of the display type should happen only once when you connect for the first time. It can also be required after you have installed a new display or after updating the firmware.

This dialogue allows also customising the names of the sources as they appear on the display.



**Fig. 25:** *Display Options dialogue.*

## Appendix A: Custom Biquads File Examples



```
biquad_test.txt - Notepad
File Edit Format View Help
// This is a comment line
n=3                                // The file contains 3 biquads
b0=0.0010232176384572448,         // That's the first biquad
b1=0.0020464352769144897,
b2=0.0010232176384572448,
a1=1.9075016260205944,
a2=-0.9115944965744232,
b0=0.9953826895856792,           // Here's the second biquad
b1=-1.9907653791713589,
b2=0.9953826895856792,
a1=1.990744059502277,
a2=-0.99078669884044,

b0=0.9953826895856792,           // And that's the third one
b1=-1.9907653791713589,
b2=0.9953826895856792,
a1=1.990744059502277,
a2=-0.99078669884044,
```

The example text file above contains definitions for 3 biquads.

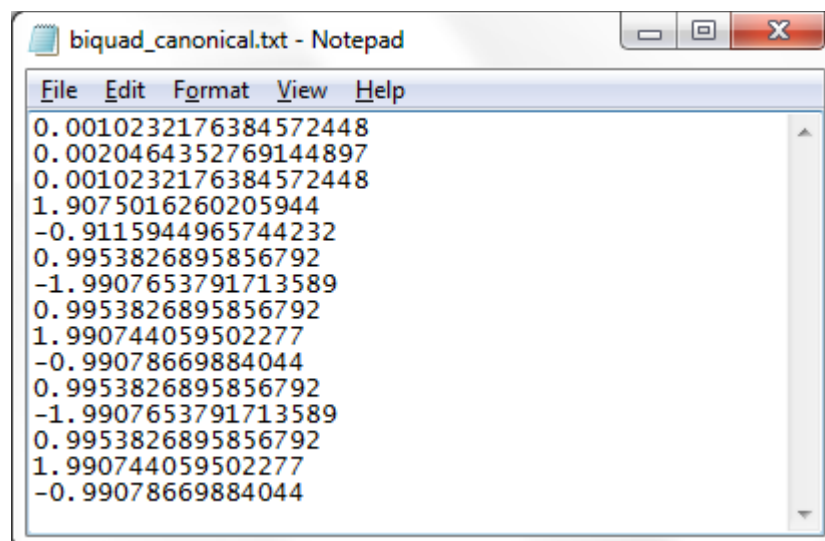
The top line is just a comment, ignored by NUC.

The 2<sup>nd</sup> line is also ignored: NUC collects coefficients and deduces the number of biquads.

There's a break before the 3<sup>rd</sup> biquad definitions: line breaks are fine too.

The coef identifiers (*b0*, *b1* etc) and the coma also are optional.

The coefficients' order is important: NUC expects to find the following sequence:  $b_0$ ,  $b_1$ ,  $b_2$ ,  $a_1$ ,  $a_2$ .

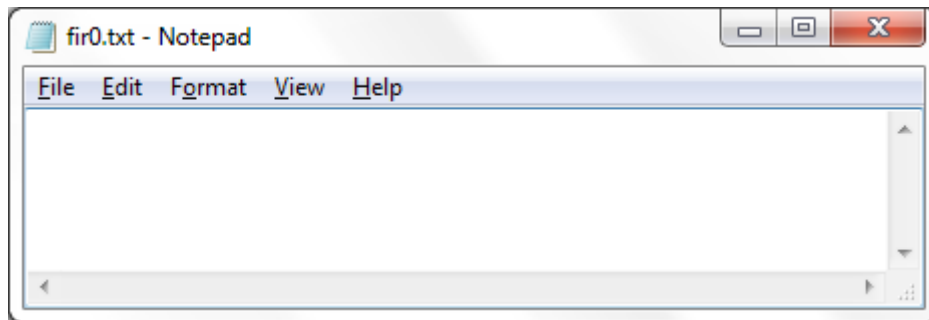


```
biquad_canonical.txt - Notepad
File Edit Format View Help
0.0010232176384572448
0.0020464352769144897
0.0010232176384572448
1.9075016260205944
-0.9115944965744232
0.9953826895856792
-1.9907653791713589
0.9953826895856792
1.990744059502277
-0.99078669884044
0.9953826895856792
-1.9907653791713589
0.9953826895856792
1.990744059502277
-0.99078669884044
```

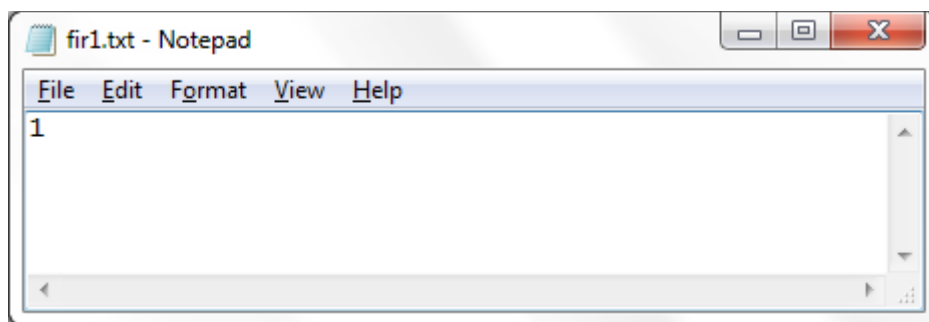
The text file above is the stripped down version of the former. The 2 files, as seen by NUC, are functionally equivalent.

## Appendix B: FIR Coefficients File Examples

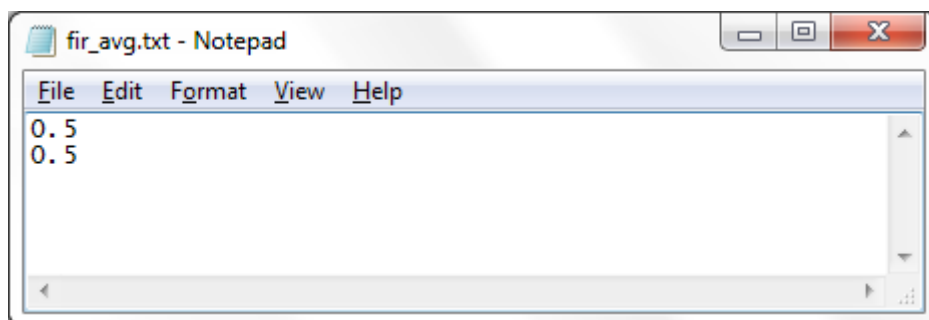
The simplest FIR file is the open-circuit: it's just an empty file (*fir0.txt*).



The pass-through FIR file contains one single coefficient (*fir1.txt*).



Below is the moving average filter: the output is equal to the average of the previous and current inputs (*fir\_avg.txt*).

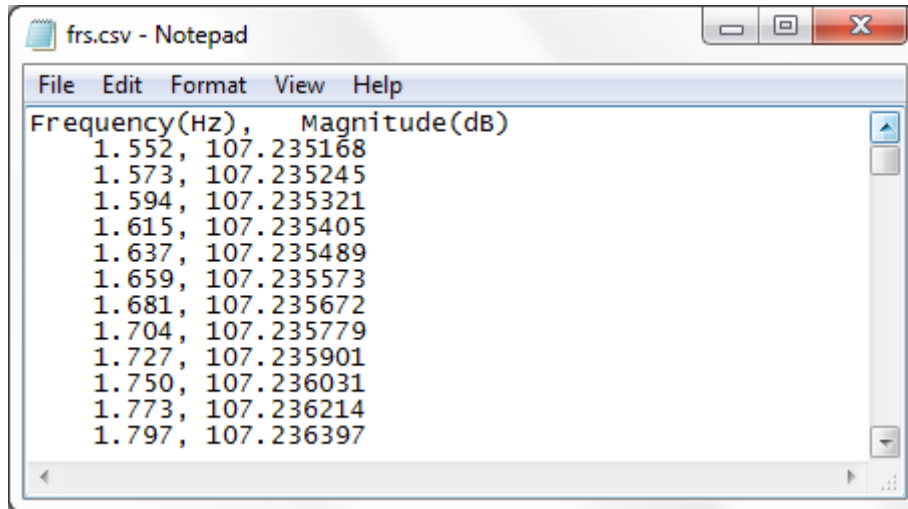


## Appendix C: Examples of Compatible Frequency Response Files

Below is the start of a compatible frequency response file.

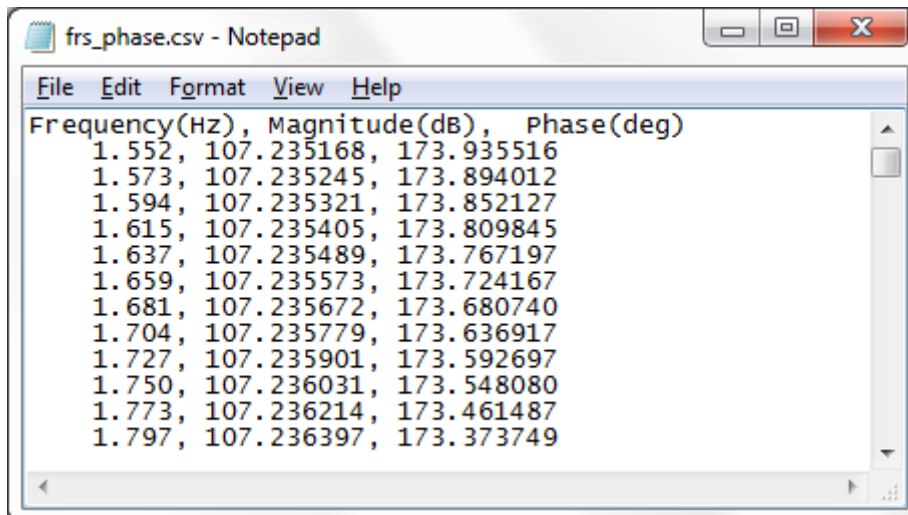
Note the first line (*Frequency (Hz), Magnitude (dB)*) is optional and can be omitted.

This file contains magnitude information – but no phase information. NUC will assume phase is null at all frequencies.



```
frs.csv - Notepad
File Edit Format View Help
Frequency(Hz), Magnitude(dB)
1.552, 107.235168
1.573, 107.235245
1.594, 107.235321
1.615, 107.235405
1.637, 107.235489
1.659, 107.235573
1.681, 107.235672
1.704, 107.235779
1.727, 107.235901
1.750, 107.236031
1.773, 107.236214
1.797, 107.236397
```

The next file is quite similar to the former – except that it contains a third column with phase information. Again, the first line is optional: NUC sees it as a comment line.



```
frs_phase.csv - Notepad
File Edit Format View Help
Frequency(Hz), Magnitude(dB), Phase(deg)
1.552, 107.235168, 173.935516
1.573, 107.235245, 173.894012
1.594, 107.235321, 173.852127
1.615, 107.235405, 173.809845
1.637, 107.235489, 173.767197
1.659, 107.235573, 173.724167
1.681, 107.235672, 173.680740
1.704, 107.235779, 173.636917
1.727, 107.235901, 173.592697
1.750, 107.236031, 173.548080
1.773, 107.236214, 173.461487
1.797, 107.236397, 173.373749
```



## Table of contents

About this manual .....	2
A few definitions .....	3
List of keyboard shortcuts .....	4
I. Connecting .....	5
II. Creating a new setup .....	6
III. Setup types .....	6
III.1 Setup General Structure .....	6
III.2 Classic Free Routing .....	10
III.2.1 Built-in EQ types .....	11
III.2.2 EQ sections ordering .....	12
III.2.3 Custom biquads .....	13
III.2.4 More Controls (Channel Pairing and Channel Enabling) ..	14
III.2.5 Classic Free Routing specifications .....	15
III.2.6 Classic Free Routing default setup .....	15
III.3 FIR Engine .....	17
III.3.1 FIR filtering model .....	17
III.3.2 FIR Engine specifications .....	19
III.3.3 FIR Engine default setup .....	19
IV. Preset Management .....	20
V. Graph Utility .....	22
V.1 Input Processing Graph .....	22
V.2 Output Processing Graph .....	23
V.2.1 Loading of frequency response .....	23
V.2.2 Processed vs. unprocessed views .....	24
V.2.3 Electrical Sum .....	25
VI. Bass/Treble Settings dialogue .....	26
VII. Hardware configuration dialogues .....	27
VII.1 Expansion Port 0 Routing dialogue .....	27
VII.2 Board Configuration dialogue .....	27
VII.3 Display Options .....	28
Appendix A: Custom biquads file examples .....	30
Appendix B: FIR coefficients file examples .....	31
Appendix C: Example of compatible frequency response file .....	32