AP3 Digital Signal Processor

TECHNICAL MANUAL

VORSIS™

600 Industrial Drive, New Bern, North Carolina, USA 28562

AP3 Digital Signal Processor Technical Manual - 2nd Edition

Vorsis
600 Industrial Drive
New Bern, North Carolina 28562
tel 252-638-7000 / fax 252-637-1285

AP3 Technical Manual

Table of Contents

Chapter 1 – General Information

Chapter 2 - Controls and Functions

Chapter 3 - Vorsis AP3 GUI

General Information

Chapter Contents

General Information

Introduction

The Vorsis™ AP3 Digital Signal Processor is a highly complex piece of audio processing gear, the first product in a new line of a single-rack space audio processors. This unit occupies one 19” wide rack space (height 1 3/4”), and is 16” deep. The two-channel/stereo audio processor provides an array of analog/digital inputs and outputs, four-band parametric EQ, three-band compressor, and tunable filters along with overall AGC, limiting, and expansion. All settings can be stored and replayed as password-protected presets. The front panel display includes input, output, and gain reduction meters. Remote monitoring and operational control is possible using Ethernet-enabled GUI software.

The XLR, DB-25, DB-9 and CAT5 connectors located on the rear of chassis give you maximum installation versatility.

Energizing

Assuming the Vorsis™ AP3 Digital Signal Processor is correctly rackmounted, you may now energize it by connecting the factory supplied power cord to the rackmount unit and then plugging it into the AC mains.

Note: To de-energize the Vorsis™ AP3, unplug its AC cord from the AC mains.

I/O Connections

All audio input and output, control, Ethernet, and power supply connections are made via various connectors mounted on AP3's the rear panel. Six XLR connectors are provided for analog and digital audio input and output, and all the audio connections also appear on the DB-25 connector. Two RJ-45 connectors provide yet a third way to access the digital audio input and output. Three additional RJ-45 connectors are provided for control and Ethernet connections. The DB-9 connector provides 8 GPI inputs for the most used presets. The pinout drawings on pages 1-9 through 1-12 summarize all wiring connections.

Inputs

Both AP3 processor signal paths can be fed from three different types of inputs: microphone, analog line, and digital line.

When processing mic level (-50dBu nominal) inputs, the input is fed, either from one of the two female XLR connectors (one each for the two signal paths) or the appropriate pins on the DB-25 AUDIO I/O connector, to the internal microphone preamplifiers. The mic preamps have digitally controlled gain, up to a maximum of 70 dB, and display remarkably high performance and accuracy. Phantom power (+32VDC) is available.

The same connections can be used for line level analog audio inputs (nominal +4dBu balanced). When configured for analog line, the inputs exhibit a bridging impedance and can handle signals up to +20dBu.

A third female XLR connector, along with a paralleled RJ-45 connector, are available to handle digital line level sources; this AES digital input is also available on the DB-25 connector. The AP3 will accommodate digital inputs having sample rates of 44.1kHz or 48kHz. The sample rate of the AP3 will automatically sync up to the applied input signal's valid sample rate, allowing the AP3 to be used in an existing digital system. AES sources are by design stereo; if the AP3 is not set to be stereo the appropriate Left or Right signal within the AES stream will be applied to the signal path.

MIC / Line In—XLR-F

Pin 1 XLR1 SH – MIC1 / LINE 1 IN SH
Pin 2 XLR1 HI – MIC1 / LINE 1 IN HI
Pin 3 XLR1 LO – MIC1 / LINE 1 IN LO

Pin 1 XLR2 SH – MIC2 / LINE 2 IN SH
Pin 2 XLR2 HI – MIC2 / LINE 2 IN HI
Pin 3 XLR2 LO – MIC2 / LINE 2 IN LO

AES In—XLR-F & RJ-45

RJ-45

Pin 1 – N/C
Pin 2 – N/C
Pin 3 – AES IN HI

Pin 4 – N/C
Pin 5 – N/C
Pin 6 – AES IN LO

Pin 7 – N/C
Pin 8 – N/C

XLR-F

Pin 1 XLR SH – AES IN SH
Pin 2 XLR HI – AES IN HI
Pin 3 XLR LO – AES IN LO

Audio Inputs—DB-25

Pin 25 – MIC 1 / LINE1 IN SH
Pin 24 – MIC 1 / LINE1 IN HI
Pin 12 – MIC 1 / LINE1 IN LO

Pin 11 – MIC 2 / LINE2 IN SH
Pin 10 – MIC 2 / LINE2 IN HI
Pin 23 – MIC 2 / LINE2 IN LO

Pin 16 – AES IN SH
Pin 15 – AES IN HI
Pin 3 – AES IN LO

Outputs

The AP3's two output signals are available as individual line level analog (+4dBu, balanced) on the two male XLR ANALOG OUT connectors, and in digital AES formatted on a male XLR connector, paralleled by an RJ-45 connector for CAT5 connection. These analog and digital outputs are available also on the DB-25 connector.

Analog Out—XLR-M

Pin 1 XLR1 SH – LINE 1 OUT SH
Pin 2 XLR1 HI – LINE 1 OUT HI
Pin 3 XLR1 LO – LINE 1 OUT LO

Pin 1 XLR2 SH – LINE 2 OUT SH
Pin 2 XLR2 HI – LINE 2 OUT HI
Pin 3 XLR2 LO – LINE 2 OUT LO

AES Out—XLR-M & RJ-45

RJ-45

Pin 1 – N/C
Pin 2 – N/C
Pin 3 – AES OUT HI

Pin 4 – N/C
Pin 5 – N/C
Pin 6 – AES OUT LO

Pin 7 – N/C
Pin 8 – N/C

XLR-M

Pin 1 XLR SH – AES OUT SH
Pin 2 XLR HI – AES OUT HI
Pin 3 XLR LO – AES OUT LO

Audio Outputs—DB-25

Pin 8 – LINE 1 OUT SH
Pin 7 – LINE 1 OUT HI
Pin 20 – LINE 1 OUT LO

Pin 19 – LINE 2 OUT SH
Pin 18 – LINE 2 OUT HI
Pin 6 – LINE 2 OUT LO

Pin 2 – AES OUT SH
Pin 1 – AES OUT HI
Pin 14 – AES OUT LO

Control Ports

There are two RJ-45 connectors (one for each signal path) available for the connections of a talent microphone control panel. Functions include remote on and off, cough, and tally.

Control—RJ-45

Pin 1 – DIGITAL GROUND
Pin 2 – COUGH
Pin 3 – N/C

Pin 4 – OFF TALLY
Pin 5 – ON TALLY
Pin 6 – REMOTE OFF

Pin 7 – REMOTE ON
Pin 8 – +5V DIGITAL

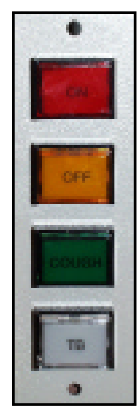

Optional ANP-RJ Announcer Panel

The optional ANP-RJ panel connects to the AP3 using a standard straight CAT5 Ethernet cable. This panel has “ON”, “OFF”, and “COUGH” switches for remote control of the microphone inputs. (Note that the “TB” button is not usable for the AP3). The “ON” and “OFF” buttons are LED illuminated. These panels can be mounted directly into a furniture surface or a turret.

To Turn the Microphone ON & OFF from a Remote Location

NOTE: The remote mic controls are meant to be used when the AP3 is processing two separate mono signals, typically microphones (see “Selecting a Signal Path” on pages 2-7 to 2-8). If the AP3 is in “stereo” mode there will be some non-intuitive interaction between the two sets of remote mic controls.

REMOTE ON — Activates the mic’s channel ON switch. Provide a momentary closure between Remote On (Pin 7) and Digital Ground (Pin 1). This will latch the mic ON. (User-supplied momentary contact switch required.)

REMOTE OFF — Activates the mic's channel OFF switch. Provide a momentary closure between Remote Off (Pin 6) and Digital Ground (Pin 1). This will latch the mic OFF. (User-supplied momentary contact switch required.)

COUGH — Temporarily Mutes the microphone. Provide a closure between Cough (Pin 2) and Digital Ground (Pin 1). This will turn the mic OFF. Note this is a non-latching mode; the mic will turn ON again as soon as the closure stops. (User-supplied momentary contact switch required.)

On and Off Tallies

ON TALLY — Lets the mic's channel ON switch control an on-air light or other "microphone on" indicator at a remote location. This control function provides a continuous closure (open collector) between On Tally (Pin 5) and Digital Ground (Pin 1) whenever the module is ON.

OFF TALLY — Lets the mic's channel OFF switch control an on-air light or other "microphone off" indicator at a remote location. This control function provides a continuous closure (open collector) between Off Tally (Pin 4) and Digital Ground (Pin 1) whenever the module is ON.

The On and Off Tallies can be used to control externally powered tally lights that require a continuous closure to function, or external tally lights (i.e., LED) can be powered from the tally output by connecting the external LED to +5V Digital (Pin 8)* and the On Tally port (Pin 5) or Off Tally port (Pin 4). In any case, current should not exceed 30 milliamps.

*We recommend a series resistor between the LED and +5V digital when you are powering the external tally from the AP3; a value of 220Ω (1/4W 5%) is suggested.

General Purpose Interface

This general purpose interface is used to remotely trigger (recall) Presets 1-8. For example, a 5VDC signal applied between Pin 1 (GPI Com) and Pin 6 (GPI 1 In), will take the AP3 preset stored in location PRE #001. The + (plus) side of the 5V signal is applied to the active selector pin (Pin 6 in the example) and the - (minus) side of the 5V signal is applied to GPI Com, Pin 1.

GPI—DB-9

Pin 1 – GPI COM
Pin 6 – GPI 1 IN
Pin 2 – GPI 2 IN
Pin 7 – GPI 3 IN
Pin 3 – GPI 4 IN
Pin 8 – GPI 5 IN
Pin 4 – GPI 6 IN
Pin 9 – GPI 7 IN
Pin 5 – GPI 8 IN

Ethernet Interface

Networked systems are connected to the AP3 via straight (pin to pin) CAT5 cable. For typical CAT5 cable pinouts see below. These connections are for communicating with the configuration computer, via a network. If you are connecting directly between the computer and the AP3 with no network in between, use a crossover cable.

Ethernet—RJ-45

Pin 1 – TXD +
Pin 2 – TXD -
Pin 3 – RXD +
Pin 4 – N/C
Pin 5 – N/C
Pin 6 – RXD -
Pin 7 – N/C
Pin 8 – N/C

Typical Ethernet Cable

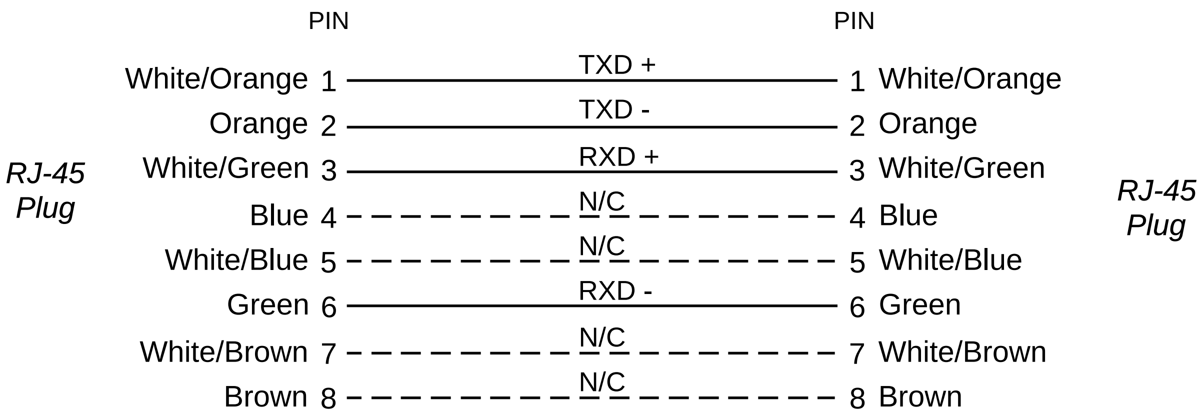

Typical Crossover Cable

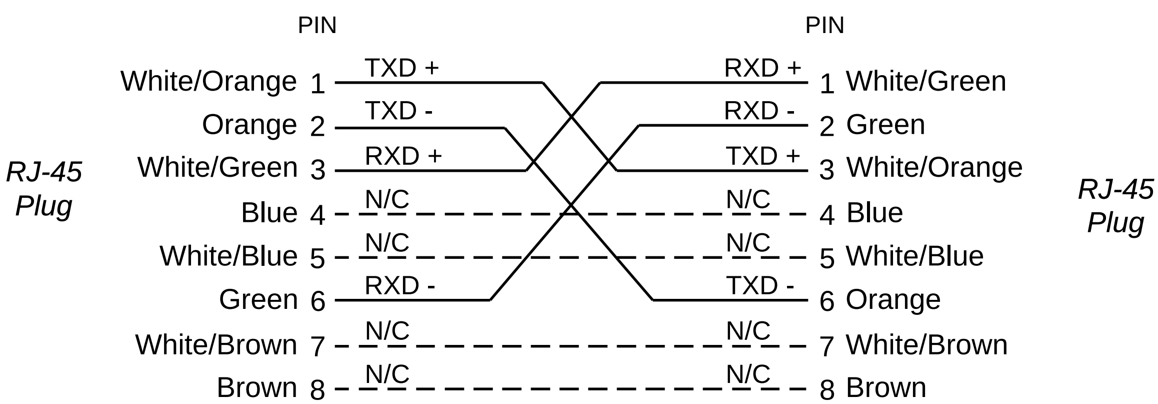

Digital Audio Connections

CABLE - All AES/EBU input and output digital audio connections are balanced and should be made using a high quality digital audio cable. Be sure to select a digital audio cable with an integral drain wire of the same wire gauge (AWG) as the twisted pair as this facilitates an easier consistent termination process. Typical AES/EBU digital audio cable has a very low characteristic capacitance per ft (pF/ft), and a nominal impedance of 110Ω. High quality digital audio cable offers better signal transmission performance versus typical analog audio cable, especially over long cable runs. Check the cable manufacturer's data sheet to be sure the cable you plan to use will work in your application.

CONNECTORS - The AES/EBU connections are made with the supplied DB-25 male mating connector, or with XLR or RJ-45 connectors. The DB-25 crimp style connector will accept wire gauge 22 - 28AWG. Please refer to the RJ-45 and XLR mating connector manufacturer's recommendations for termination instructions.

Unbalanced Analog Connections

ANALOG INPUTS — Wire to the processor input end with typical shielded, two conductor cable (like Belden 9451), just as if you were connecting a balanced source. At the unbalanced source machine's output, connect the + output to the HI input wire and connect the source machine GND wire to LO. Connect the shield at the AP3 end only.

Note: Unbalanced analog sources typically have -10dBv (316mV RMS) signal levels and will not match the processor nominal operating level of +4dBu (1.23V RMS). We highly recommended that you first externally balance any unbalanced sources you plan on connecting to the Vorsis™ AP3. Many third party "match boxes" are commercially available for this.

ANALOG OUTPUTS — The AP3 analog outputs use a balanced output circuit which behaves exactly like the secondary of a high-quality transformer, with no center tap—this output is both balanced and floating. For unbalanced operation, either the HI *or* LO side of the analog output must be strapped to ground of the unbalanced input, with the output taken from the other side. (Normally you would strap LO to ground, and use HI to feed your unbalanced equipment input.) Leave the SH floating at one end.

Unbalanced Digital Connections (SPDIF)

SPDIF INPUTS — The SPDIF (Sony/Phillips Digital Interface) or "consumer" digital audio interface is a two wire unbalanced signal typically on a single RCA style connector. Note that the SPDIF signal level of approximately 500mV and 75 ohm impedance does not correctly match the processor AES inputs. We highly recommend using a "balun" or format converter when interfacing "consumer" grade source devices to the processor.

In cases where a consumer grade device must be interfaced and the appropriate matching device is not available, try wiring the SPDIF center conductor (HOT) to the HI input pin and SPDIF shell (ground) to the LO input. Connect SH at the processor end only.

SPDIF OUTPUTS — The Vorsis™ AP3 digital outputs are fixed, professional, AES-3 formatted outputs. SPDIF consumer format is not supported. Use an external format converter to connect the digital outputs to consumer gear.

Mic/Line In - XLR-F

XLR1-F

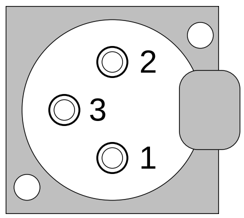

PIN 1 XLR SH - MIC 1 / LINE 1 IN S
PIN 2 XLR HI - MIC 1 / LINE 1 IN HI
PIN 3 XLR LO - MIC 1 / LINE 1 IN L

XLR2-F

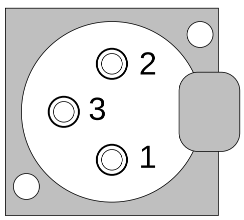

PIN 1 XLR SH - MIC 2 / LINE 2 IN S
PIN 2 XLR HI - MIC 2 / LINE 2 IN HI
PIN 3 XLR LO - MIC 2 / LINE 2 IN L

AES In - XLR-F & RJ-45

XLR-F

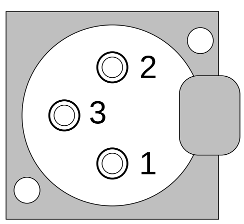

PIN 1 XLR SH - AES IN SH
PIN 2 XLR HI - AES IN HI
PIN 3 XLR LO - AES IN LO

RJ-45

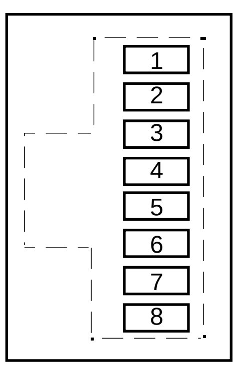

N/C
N/C
AES IN HI
N/C
N/C
AES IN LO
N/C
N/C

Audio Input / Output

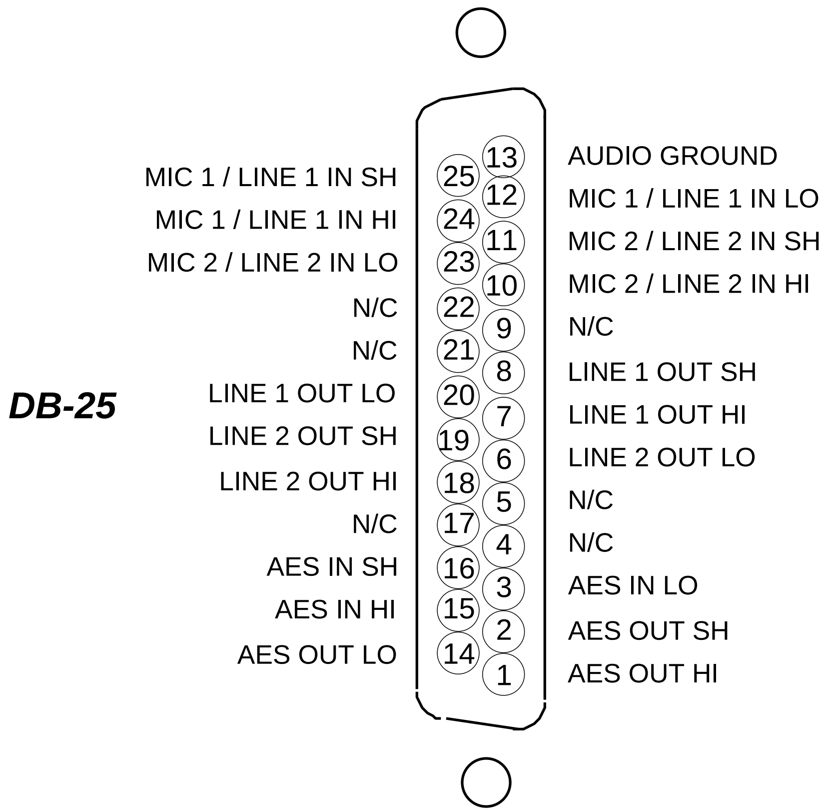

Analog Out - XLR-M

XLR1-M

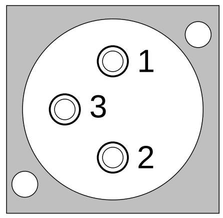

PIN 1 XLR SH - LINE 1 OUT SH
PIN 2 XLR HI - LINE 1 OUT HI
PIN 3 XLR LO - LINE 1 OUT LO

XLR2-M

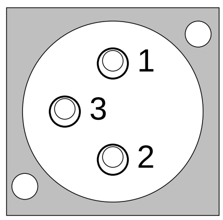

PIN 1 XLR SH - LINE 2 OUT SH
PIN 2 XLR HI - LINE 2 OUT HI
PIN 3 XLR LO - LINE 2 OUT LO

AES Out - XLR-M & RJ-45

XLR-M

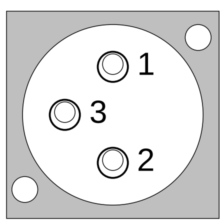

PIN 1 XLR SH - AES OUT SH
PIN 2 XLR HI - AES OUT HI
PIN 3 XLR LO - AES OUT LO

RJ-45

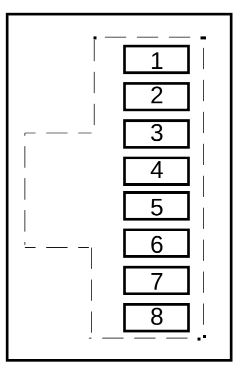

N/C
N/C
AES OUT HI
N/C
N/C
AES OUT LO
N/C
N/C

Control - RJ-45

MIC 1 RJ-45 #1

1	DIGITAL GROUND
2	COUGH
3	N/C
4	OFF TALLY
5	ON TALLY
6	REMOTE OFF
7	REMOTE ON
8	+5V DIGITAL

MIC 2 RJ-45 #2

1	DIGITAL GROUND
2	COUGH
3	N/C
4	OFF TALLY
5	ON TALLY
6	REMOTE OFF
7	REMOTE ON
8	+5V DIGITAL

Ethernet - RJ-45

ETH RJ-45

1	TXD +
2	TXD -
3	RXD +
4	N/C
5	N/C
6	RXD -
7	N/C
8	N/C

GPI

GPI DB-9

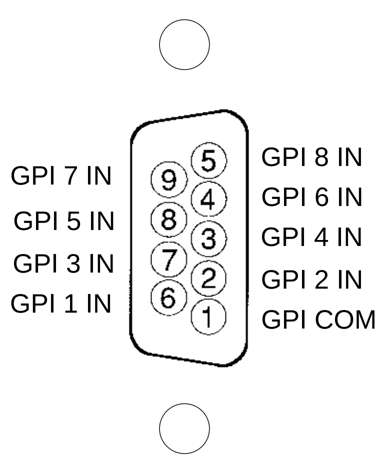

AP3 Controls and Functions

Chapter Contents

AP3 Controls and Functions

As internally complex as is the AP3, and as small as is the front-panel, it is obviously not possible to provide a control-per-function, as desirable as that may be. However, considerable effort has gone into making the front-panel controls logical and usable.

A wide variety of both frequency-response shaping and dynamics tools (on the right-hand side of the front panel) are available for selection and adjustment in each AP3 signal path via the switches on the right-hand side of the front panel. Despite the considerable overlap in functionality, these will be grouped into Signal Conditioning, Equalization, and Dynamics. This also approximately represents (exceptions will be noted) the order in which these tools are applied to the signal paths.

Most of the functions described in this chapter are covered in greater detail in Chapter 3.

Selecting a Signal Path

The front panel Channel 1 and 2 buttons determine which of the two AP3 signal paths is being operated upon - the appropriate path's button will be lit; pressing the other button will cause it to be lit instead, and command and display now act on that path. The AP3 can be made "stereo" (i.e. both paths being controlled together) by simultaneously holding down both Channel 1 and 2 buttons for two seconds; both buttons will light signifying both are under command, and henceforth both paths track together.

Both Channel 1 and Channel 2 will adopt the settings of ***the path which was last selected*** for control at the time of making the AP3 stereo.

An important exception to this adoption of a single set of parameters for stereo is Input Gain settings; if there is a difference between the Channel 1 and Channel 2 paths' input gains (such as could easily occur with two microphones, even as a stereo pair) this offset will be maintained between the two input gains as the now "stereo" input gain is adjusted. This offset will be ignored at gain extremes (i.e. both paths will ultimately reach max or min gain) but the offset will be retained and reapplied when the gain returns to saner values.

Additionally, the dynamics signal processing will become married, such that independent Channel 1 and Channel 2 gain variation and resulting stereo image shifting cannot occur.

Reverting to two independent paths again is achieved by holding down either the Channel 1 or 2 button for two seconds. Both path's settings adopt, until subsequently altered, the settings attained during the "stereo" usage.

Basic Setup

Before you can start tweaking your signal from the front panel there are a few basic functions you must access and configure for your signal chain; selecting a source, setting mode, and input/output level control.

Source - Press the INPUT switch to select the following source type - Mic, Line, or AES (digital). Also set M/S - L/R (sum & diff or stereo) to L/R for “normal operation”.

TIP: Some function settings are “paged”. See page 2-7 for a description of the paging controls, and see the chart in page 2-13 for menu page location of various functions.

Output MODE - For Stereo operation set M/S-L/R to L/R. For MONO operation set M/S - L/R to M/S and use the L channel output.

GAIN - Press the GAIN switch to access the input/output gain controls.

AES Clock Error Messages - will flash obnoxiously if the Sample Rate is set to AUTO and no AES input is connected or if your AES source and the AP3 don’t agree on a sample rate. See the AES Error Messages section (page 2-11) to decipher these.

Signal Conditioning

Generally these functions remove undesired signals and artifacts from the source material.

High-Pass Filter

A fast rolloff-rate high-pass filter, adjustable between 20Hz and 1kHz, serves to aid removal of low-frequency signals and noises, such as air-conditioning rumble, table thumps, or in the case of a TELCO source, line-borne noises below the conducted frequency range (typically 300Hz).

Notch

The notch has been optimized in terms of width and depth to do the most damage to an undesired signal (such as supply-line hum, a TV/ computer line-timebase while, or an annoying instrument “ring”) whine doing the least damage to the desired audio; it is at once broad enough to be reasonably easy to use in “finding” the offending signal, yet narrow enough not to impart excessive coloration to adjacent frequencies. It is tunable over the entire audio range of 20Hz – 20kHz.

Low-Pass Filter

A fast rolloff-rate low-pass filter, adjustable between 20kHz and 1kHz, aids in the removal of distracting high-frequency noises, or to deliberately restrict the program material’s bandwidth.

De-Esser

The de-esser is a dynamic filter which reduces the signal energy within a chosen band of frequencies when it has been sensed to exceed a desired level. This is a valuable tool in reducing prominent sibilance (pronounced “SSS” noises) on some voices, particularly if poorly recorded or subject to poor or overly close microphone technique. It is also often used to help control the “honky” effect on phone lines.

Expander

Sometimes called “downwards expansion”, the expander automatically reduces the signal level further once it has fallen below a chosen threshold. This is commonly used to “gently” or partially turn off a source if no valid signal is present. Examples are the suppression of room noises when no-one is talking into a microphone, reduction of prominent breath noises in speech, or the muting of phone-line noise if used on a TELCO circuit. Such uses become more necessary if the signal is later subjected to high degrees of compression, either within the AP3 and/or by subsequent airchain processing - compression tends to make room/breath/ TELCO noises all the more objectionable by raising their level. Among more subtle usages can be to reduce an environment’s apparent reverberation time, and its effects, by accelerating the decay of the reverb “tail”; similarly with an instrument which has excessively long “ringing”. The AP3’s expanders have sufficient adjustable range to perform these subtleties, yet at the other extreme act as a precipitous gate.

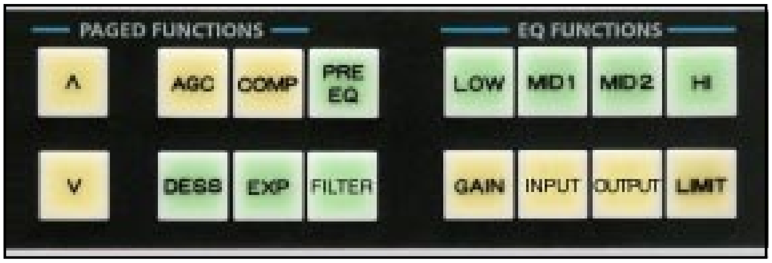

Equalization

Parametric Equalizer

A powerful four band parametric equalizer, each band adjustable in center-frequency, bandwidth and degree of lift or cut is available in each signal path. It is selectable to be prior to or after the compressor (“pre” or “post”), each position having benefits depending on the application and program material.

Dynamics

The Dynamics processing within the AP3 is designed to maintain an optimum average level range (AGC, Automatic Gain Control), allow the increase of signal “density” or apparent “loudness” (Compressor), and prevent the output signal from exceeding a critical pre-determined level (Limiting).

Multiband Compressor - “Production”

The compressor within the AP3 is a multi-band processor. It divides the signal spectrum up into three paths (nominally “LF”, low frequencies, “MF”, mid frequencies, and “HF”, high frequencies) each of which has its own individual compressor, the outputs of which are then recombined. The carefully matched and contrived dividing filters are adjustable in frequency to suit differing program types, voices, or sources. Each of the three band’s compressors has individually adjustable parameters, making possible a wide range of signal control and dynamic equalization. The compressor’s intent in the AP3 is to provide a high-definition and well-packaged signal (free of all the cross-modulation and “pumping” characteristic of single-band compressors) which is better able to transcend later aggressive processing without being hacked into unrecognizability by it.

Multiband Limiter - “Airchain”

This is an alternate three-band limiter to the above three-band compressor, which is intended to easily produce the highly aggressive processing often required in a radio and TV sound “airchain”. In short, it is designed to get LOUD, as opposed to the compressor which is designed to be transparent.

AGC

Multi-banded and operating as an integral part of each band of the multiband compressor, the AGC (Automatic Gain Control) observes the long-term signal energy and adjusts the signal levels accordingly to maintain them within an optimal range. Typically AGC would compensate for someone moving somewhat off-microphone, a reader steadily becoming tired, a DJ who can't read meters, or other such examples of relatively long-term applied signal level variation. A common set of attack and release time-constant values is applied across the three discrete AGC elements of the multi-band processor.

Limiter

Last in the signal processing chain is a precision peak limiter, which is intended to contain the AP3's output within desired level constraints, most particularly to avoid clipping within or overdriving signal destinations.

Metering

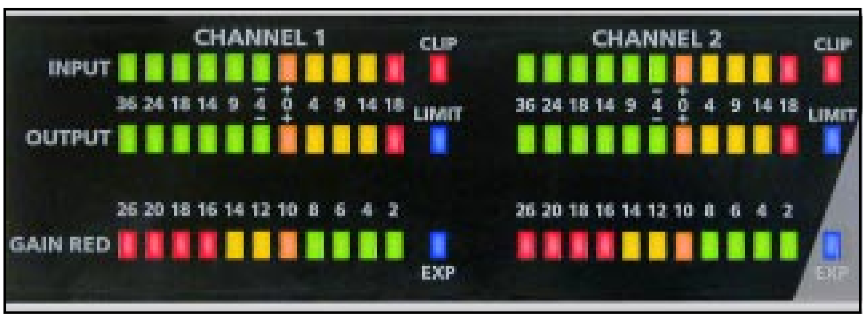

On the left-hand side of the front panel, each signal path has bargraph metering for (top) input signal level (adjustable by the input gain controls), output signal level (middle), and, as a default, composite gain-reduction (bottom). This latter indicates the gain-reduction being applied to the signal by any of the applicable processes: expansion, AGC, compression, or limiting.

Slightly distanced at the right-hand end of each of these bargraphs is a separate LED indicator, to show respectively for the input, output, and gain-reduction bargraphs input clipping (overdrive), output limiter activity, and expander activity. These last two will show if any indicated gain-reduction is due to these functions, regardless of control panel settings.

As mentioned, the gain-reduction bargraph reads by default a composite of all effecting signal-path functions: expansion, AGC, multi-band compressor, and output peak limiter. If, however, one of those functions is selected onto the Main Control Group, then the gain-reduction meter reacts solely to that function; in other words, if the output peak limiter ("Limit") is selected, the bargraph just indicates the gain-reduction due to the limiter. Similarly with the De-Esser, AGC, and Expander.

The multi-band compressor is a slightly different case in point; when selected ("Comp") the bargraph indicates a composite of the three bands' compressions. Overpressing one of the three encoders causes the bargraph to indicate just the gain-reduction for that band, for as long as the encoder is pressed. This permits rapid access to and visualization of the behavior of the multi-band compressor.

Main Control Group

The main method of effecting changes to any of the functions is a Main Control Group of three knobs, driving "endless" digital encoders, their associated 8-character alpha-numeric displays which show the values of the parameters under control at that point, and an "IN" (On/Off) button.

The signal processing functions upon which this Main Control Group is operating are determined by a set of buttons on the right-hand side of the front panel. Simplistically, pressing one of these buttons instantly makes the Main Control Group indicate and display parameters relevant to the button pressed; e.g. pressing LIMIT will bring the limiter's parameters into the MCG for inspection or adjustment.

Many signal processing functions have more than three elements to adjust; this is indicated by the illumination of a pair of paging buttons (marked with up and down arrows) adjacent to the MCG which allow jumping between groups of adjustable elements. For example, the De-Esser has Threshold, Attack, and Release times adjustable on one "page"; pressing one of the up/down buttons will instead display another "page" of De-Esser elements, such as Frequency, Bandwidth, and In/Out. The up/down buttons will cycle through available pages, although in different directions.

TIP: This section describes the menu paging controls.

A variation on the above theme (of paging in different elements of a selected signal processing function) occurs when the multi-band compressor is selected. In this case, the like element of each of the three bands is displayed simultaneously; e.g. the three bands Thresholds will be up on the MCG at once. Paging with the up/down buttons will cycle through each of the variable parameters of the three compressors. This greatly facilitates the adjustment of the triband processor; one is likely to be operating on simultaneous adjustment of, say, output level trims across the bands, which having to jump from band-to-band would make highly tedious.

Presets

A Preset is a snapshot of all the AP3's settings and configuration, with the intention of being able to recall the preset and restore the unit to exactly the same condition at a future juncture. The entire system (i.e. both signal paths, not one signal path at a time) is saved as a preset.

The AP3 is shipped with a number of likely useful presets included which are ready to use as-is, or as departure points for customization. All presets other than these will initially be "soft" and capable of being modified and

over-written. 180 such presets are available. AP3 presets are “sticky” as opposed to “tunable”. What this means is that if an AP3 preset is invoked and adjustments are made, whether by hardware front-panel or by the GUI, those changes do not become part of the preset unless deliberately saved back into it.

A DB-9 connector on the rear panel allows preset numbers 1 to 8 to be invoked, by application of a 5VDC signal between its associated DB-9 connector pin and the GPI COM pin. This is a simple and direct way of allowing AP3 configurations to be changed in accord with changing users, presenter or station day-parts.

It is envisaged that most configuration of presets will be performed using the AP3’s corresponding GUI PC application, but it is nevertheless still entirely possible to create and manipulate presets from the hardware front-panel.

Immediately to the left of the Main Control Group of encoders and switches are three buttons for preset manipulation:

Presets Button

In the same manner as the function attention buttons on the right hand side of the panel, pressing this button calls preset management into the Main Control Group, and illuminates. The left-most encoder when rotated scrolls through the saved presets, the reference number of which is displayed in the above alpha-numeric display. Any text / label associated with the preset is displayed in the central display, and any attributes such as “locked/unlocked” (write-protection) will show in the right-most display.

Take Button

Upon being pressed and held for two seconds (it illuminates then darkens in this period), the preset showing in the displays will be immediately impressed into the AP3 which will become active with the contained settings and configuration.

Save Button

Pressing and holding this button for in excess of two seconds results in the existing and operating AP3 settings and configuration being recorded in the next available empty preset slot.

Modify (Take and Save simultaneously pressed)

Pressing both "Take" and "Save" buttons together for two seconds saves the current settings back into the preset showing in the displays.

These allow all operationally useful preset actions to be achieved easily from the front panel, including preset order rearrangement. There are, however, a number of preset management functions which are only done using the associated GUI:

• Preset order rearrangement while the unit is "live" and in use, and when changing the actual settings of the AP3 would be otherwise unwelcome.

• Labeling. We refuse to inflict the simply awful dial-up-a-letter-at-a-time Chinese-water-torture thing on anybody.

• Attributes, such as write-protection of created "gospel" presets.

• Transferring of presets from one AP3 to another.

System Features and Settings

"System" is a grab-bag of once-and-for-all set-up functions, or occasionally required features. This sub-menu is accessed by over-pressing (holding down for longer than two seconds) the "Presets" selection button. The "Presets" button slowly flashes while "System" is accessed.

The system element or feature is named in the left-most alpha-numeric display, while its present settings are displayed in the middle and right-most display. The elements may be cycled through using the left-most encoder.

Password

The password sub-system allows the user to enter a password in order to lock or unlock the front panel, and also to establish a new or change an existing password.

Through the "Presets" selection (which as noted above is still accessible however "locked" the box may be) the default "System" selection is for "Password".

(Should the AP3 already be "unlocked", i.e. operating without password protection, or having been opened using a password already, this first selection will instead be "New Pass"; see below.)

Password selection (actually a "pass-number"), consists of two numbers between 0 and 255 entered on the mid and right-most encoders in true safe-cracker fashion. A selection of "000" "000" represents no passwording. The middle encoder is turned until the indicated number is correct, then the knob is "dobbied" and selection is repeated with the right-most encoder. Once the required password has been dialed up, if it is accepted, the number disappears and "UNLOCKED" appears on the middle display; should - heaven forbid - the password be wrong, "WRONG PASSWORD" will display embarrassingly for three seconds before disappearing, prompting for another attempt, scoundrel.

Tip: DOBBY (pronounced dah-bee) - means to quickly press and release an encoder knob.

New Password

Unless the AP3 is fully unlocked (by using the password, if in use) this feature will be inaccessible.

“New Pass” will appear in the left-most display, and the two 0-255 numbers are dialed in using the center and right-most encoders, “dobbying” to set the number. “REPEAT” will appear in the mid display for three seconds before blanking; the new password number is entered again in the same fashion. If the two numbers agree, “PASSWORD ACTIVE” will appear in the mid display; otherwise “MISMATCH” will indicate irritatingly for three seconds prior to the display blanking as it makes ready for another attempt.

Access

This defines the amount of control an unattended AP3’s control surface is allowed when locked with a password. Dialing through with the center encoder brings up the following options:

“All” — All functions and features of the AP3’s control surface are available.

“Presets” “Gains” — The input gain controls remain active and preset selection is accessible.

“Presets” — Preset selection available only.

“None” — No access, other than to “Password”.

“Inputs” “Follow” “Presets” / “Inputs” “Stay” “Same”

One or the other of these will be displayed across the three alphanumerics. They may be cycled between by rotating the center encoder.

A serious operational concern is that a freshly selected preset may or may not have appropriate settings for the input selection and gains contained within it, and similarly that once an AP3 has good input and gain settings, just differing processing is all that is required of a preset. Contrarily, sometimes everything within the preset, input selection, gain and all, is necessary.

“Inputs Follows Presets” (default) when selected for an AP3 allows the device to accept all preset information, while in “Inputs Stay Same” the AP3 ignores input selection and gain settings within presets.

It should be noted that, regardless of this setting, any presets saved in the AP3 do contain the existent input settings.

Sample Rate

When using the AES digital input, the AP3 conveniently automatically follows the incoming sample rate, providing it is a valid value.

For analog inputs, though, there are options. One of four displays will show in the “System” sub-menu:

- **"SampRate" "Analog" "48kHz"**

  The AP3's sample rate when using an analog input will be an internally generated 48kHz.

- **"SampRate" "Analog" "44.1kHz"**

  Internally generated 44.1kHz.

- **"SampRate" "Analog" "Auto/48"**

  The clock will be internally generated 48kHz, *unless* a valid AES stream is being presented to the AES input; in that case the AP3 will follow that clock.

- **"SampRate" "Analog" "Auto/44"**

  Internal 44.1kHz, unless a valid external AES clock is applied.

AES CLOCK ERROR
If AUTO/48 or AUTO/44.1 is selected and no digital input is present, the displays will flash "AES CLOCK ERROR"

Since the item being varied is the last display, the right-most encoder cycles through the options.

In the event that the AES digital input is being used but the data stream is sensed to be invalid, the AP3 will mute, and "AES" "Clock" "Error" will flash infuriatingly across the three displays until such time as either (a) a valid stream is restored or (b) any button on the control surface is pressed. From then on, the error message will flash for three seconds every fifteen seconds, unless activity is sensed on the front panel.

If an "Auto" option (above) is selected for analog input clock source, and the AES source is invalid, the same "AES" "Clock" "Error" will appear; however, the AP3 will not mute, since it will be operating off its own internal clock instead.

Sample Rate Error Messages

"AES" "CLOCK" "ERROR" — displays if AUTO is selected and no AES clock is present, or if the signal is INVALID.

"SAMPLE" "RATE" "MISMATCH" — displays if the sample rate connected to the AES input doesn't match the selected sample rate (AUTO/48 or AUTO/44.1).

IP Address

An AP3 needs telling who it is, from a networking point of view, so that its associated GUI knows how to find it, and what to call it. The device's IP address is entered in a similar "safe-cracking" manner to the password (above):

An IP address consists of four numbers between 0 and 255 (Vorsis will supply the necessary intervening "dots" free of charge); they are shown across the center and rightmost displays. These four numbers are set up one at a time, in left-to-right sequence, using the center control encoder;

rotating the center encoder will cause the first number to "dither"; once the desired number is reached, the center knob is dobbied, whereupon the second number will dither - this should be dialed-in and dobbied, too. Likewise for the third and fourth numbers.

Tip: DOBBY (pronounced dah-bee) - means to quickly press and release an encoder knob.

Entry by dobbying of the last (fourth) number makes the overall displayed number the IP address by which that AP3 is known to the world.

Subnet / Gateway

Entered in exactly the same manner as "IPAddres"; these are sometimes necessary TCP/IP routing codes. "Gateway" is presently unused.

Changes made to IP address, subnet, and gateway take effect after a subsequent power-cycle.

AP3 FRONT PANEL CONTROL FUNCTIONS

TABLE 1

SWITCH	PAGE	MODE	PARAMETERS		
			1	2	3
GAIN	1	N/A	INPUT GAIN	OFFSET (BALANCE)	OUTPUT GAIN
INPUT	1	N/A	M/S or L/R	SYM IN / OUT	PRE-EMPH
	2	N/A	SOURCE	Ø 1 (L)	Ø 2 (R)
OUTPUT	1	N/A	EMPHASIS	FILTER	M/S - L/R
LIMIT	1	N/A	THRESH	ATTACK	RELEASE
	2	N/A	CLIP	ST LINK	IN/OUT
	3	N/A	POSMOD	BACKOFF	
AGC	1	N/A	ATTACK	RELEASE	IN/OUT
COMP			PROCMODE (PRODUCTION or AIRCHAIN)		
	1	PRODUCTION	XRF 1	XRF 2	IN/OUT
	2	PRODUCTION	TRIM (LO)	TRIM (MID)	TRIM (HI)
	3	PRODUCTION	RATIO	RATIO	RATIO
	4	PRODUCTION	RELEASE	RELEASE	RELEASE
	5	PRODUCTION	ATTACK	ATTACK	ATTACK
	6	PRODUCTION	THRESH	THRESH	THRESH
	7	PRODUCTION	M DRVE	MAKE UP	
	1	AIRCHAIN	XRF 1	XRF 2	IN/OUT
	2	AIRCHAIN	TRIM (LO)	TRIM (MID)	TRIM (HI)
	3	AIRCHAIN	CLIPTYPE	CLIPTYPE	CLIPTYPE
	4	AIRCHAIN	CLPBKOFF	CLPBKOFF	CLPBKOFF
	5	AIRCHAIN	RELEASE	RELEASE	RELEASE
	6	AIRCHAIN	ATTACK	ATTACK	ATTACK
	7	AIRCHAIN	DRIVE	DRIVE	DRIVE
	8	AIRCHAIN	M DRIVE	OUTLEVEL	FREEZE
PRE EQ	N/A	N/A	PRE/POST	N/A	IN/OUT
DESS	1	N/A	THRESH	ATTACK	RELEASE
	2	N/A	FREQ	BAND WID	IN/OUT
EXP	1	N/A	THRESH	RATIO	DEPTH
	2	N/A	N/A	N/A	IN/OUT
FILTER	1	N/A	HPF FRQ	NTCH FRQ	LPF FRQ
	2	N/A	IN/OUT	IN/OUT	IN/OUT
LOW	1	N/A	FREQ	BAND WID	LEVEL
	2	N/A	N/A	N/A	IN/OUT
MID1	1	N/A	FREQ	BAND WID	LEVEL
	2	N/A	N/A	N/A	IN/OUT
MID2	1	N/A	FREQ	BAND WID	LEVEL
	2	N/A	N/A	N/A	IN/OUT
HI	1	N/A	FREQ	BAND WID	LEVEL
	2	N/A	N/A	N/A	IN/OUT

TABLE 2

SWITCH	FUNCTION
PRESETS	Press once to access Preset List, Scroll List.
TAKE	Press and Hold to take Preset.
PRESETS	Press and Hold to access System Menu.
SAVE	Stores New Preset.
TAKE + SAVE	Overwrites currently displayed Preset.

TABLE 3

SYSTEM MENU	
Press and hold PRESETS switch to enter System Menu. Scroll soft encoder 1 to view and change System settings.	
Password	Enter Password [ØØØ-999] [ØØØ-999]
Lock	Locked - Unlocked
Access	Presets - Presets + Gains - ALL
Inputs Follow Presets or Inputs Stay Same	
Sample Rate	AUTO/48K, AUTO 44.1K, 48K, 44.1K
IP Address	User defined - 192.168.1.191 default
Subnet	User defined - 255.255.255.0 default
Gateway	Not supported - 255.255.255.255 default
MAC	Not editable

Vorsis AP3 GUI

Chapter Contents

Vorsis AP3 GUI

Getting Started

A PC application program to control one or many AP3s comes with the unit. It is a Windows™ Graphical User Interface (GUI) program, intended to be straightforward in use, controlling and displaying the powerful features of the AP3 to best effect. It is supplied as a self-installing program which may be executed from any directory on the computer, the resulting files being installed under "Program Files\Vorsis\AP3"; a "Vorsis AP3" icon will appear on the desktop, double-clicking on which will start the GUI program. A registration key is required from Vorsis to operate the program longer than a week.

NOTE: A basic familiarity with Windows navigation techniques is assumed. If terms such as "left-click", "right-click", "double-click", and "drag-and-drop" seem alien, operating the AP3 GUI could be problematic.

Connecting the AP3 and the GUI

The AP3 and PC running the AP3's GUI program can be straightforwardly connected together over a standard Ethernet Local Area Network (LAN).

The AP3 should be installed, powered-up, and verified to be operating normally; a CAT5 cable connected to the 100baseT LAN should be inserted into the appropriate socket on the rear. Likewise, the GUI should be installed on the desired PC which is verified to be working correctly with the LAN.

It is also possible to work without a LAN by connecting the AP3 and the PC Ethernet ports together using a crossover cable.

Constraints for the system are not extraordinary, in that that the PC should be preferably at least 1GHz in speed, and that the screen be at least 1024 x 768 pixels in size. The LAN should be 100baseT (100MHz) capable. It is indeed possible for the GUI to be run on a portable computer connected to the LAN by wireless (802.11b as a minimum) as long as the AP3 itself "sees" a 100baseT network. The dodgier the radio link, though, the more may be the impact on smoothness of the GUI's real-time graphics.

The AP3 does not accept DHCP (automatically assigned network addressing) and needs to be told its own network-unique address; this can be entered into the AP3 from its front panel ("System" / "IPAddress" - see Chapter 2). Most DHCP networks assign addresses starting at the bottom of the group 192.168.000.xxx; choosing an address high in that group, say 192.168.000.200 will likely keep it out of the way of the busiest DHCP addressing. However, it is always wise to check with your local friendly network administrator. If the network warrants it, the "Subnet" and "Gateway" values may need to be differed, but the suggested values of 255.255.255.000 (Subnet) and 255.255.255.255 (Gateway) will suffice for all but complex situations.

For security, the AP3 will not act as a "web device" or support open connection to the Internet. However, connection between a GUI and AP3s through the Internet between remote locations is entirely possible by arranging for the local LANs upon which they are residing (say "Studios" and "Transmitter_Site") to be connected via a "Virtual Private Network", or VPN; this acts as a "tunnel" through the Internet allowing the participants to behave as if they are all on a single big LAN. Setting up a VPN – although not difficult - is beyond the scope of this manual owing to the uniqueness of each situation.

Once the AP3 has a valid address, the GUI can find it on the network. Double-click on "Devices" on the GUI; click "Add"; click "Browse" - most times the AP3's IP address will promptly show in the appropriate field; enter a name for the AP3 if needed (say "Production") and click "OK". Alternatively, directly enter the new AP3's IP address and its pet name, then "OK". The GUI should immediately connect to the AP3, and "Online" will display. Once the GUI has been made aware of the AP3's existence in this manner, it will always appear in the list of "AP3 Devices", and be instantly accessible.

TIP: If "Trying" is continuosly displayed in the section of the title bar, then try:
• go Offline then Online, or
• cycle power on the AP3

Using the GUI

The GUI screen may be maneuvered around the computer's screen by left-clicking-and-holding on the "Vorsis" logo, and dragging the GUI to the required position. In true "Windows" fashion, the GUI may be minimized (taken off the screen) and closed by way of the familiar controls at the extreme top right.

Description of the features and functions of the GUI will follow different discrete regions of the screen; it will begin with the most eye-catching region containing the Dynamic Displays, then the Control Area above it, followed by the Side Bar buttons, and lastly the Title Bar.

Dynamic Displays Region

Frequency-Domain Graph

Visually, the most important element on the main screen, dominating the lower half of it, is a large graphical display; the standard audio frequency of 20Hz – 20kHz is ranged in logarithmic form across this graph with gradations at salient frequencies. The vertical axis scale changes according to context, although the overall vertical range remains at 30dB, with gradations every 3dB.

A number of real-time dynamically varying spectra and other frequency-response-versus-time effects may be displayed simultaneously. A row of large buttons under the graph determine which are displayed, and also determine the presence or otherwise of equalizer frequency response, filter responses, de-esser sidechain filter response, etc. (The de-esser's dynamic frequency response shows here, too, as a dynamic "upside-down shark-fin" when in operation.):

INPUT — A Fast-Fourier Transform (FFT) based real-time spectral analysis of the selected (Left/Right, A/B) input signal.

OUTPUT — This is a real-time spectral analysis of the selected output signal. (The scaling of either the input or output display may be altered by clicking on and vertically sliding the right-hand "dB" scale, to best position the curve.)

COMPRESSOR — A curve representing the instantaneous gain-reduction by frequency of the three-band compressor / multiband limiter is drawn.

AGC — Displays the gain-reduction due to AGC action across the spectrum; since the AGC is three-band, a curve joining the LF, MF and HF gain reductions is drawn.

BOTH — This displays a curve of the composite (instantaneously greatest) gain reductions spectrally from both the AGC and the compressor.

VIEWS — ”Views” cycles through a short set of display options:

- “Live Audio” - the above selection of real-time information,
- “EQ/Filters” – the frequency responses of applied processing,
- “Both” - dynamic information and response curves together.

Bargraph Metering

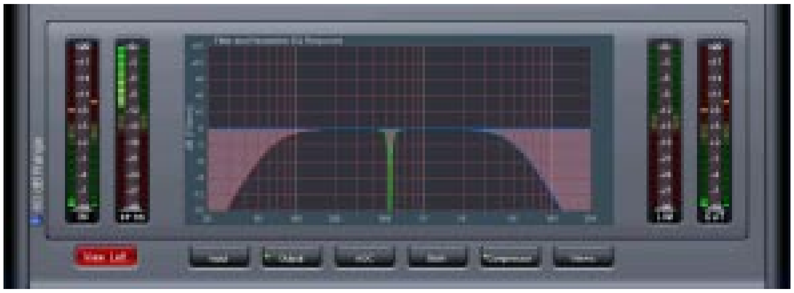

Adjacent to the graph, both to the left and right, and sharing its 30dB vertical range (0-30dB of gain reduction, 30dB signal level meter range), are a number of bargraph level indicators. They indicate signal levels and gain reductions of elements in the signal path; those ahead of the multiband compressor (input side) are to the left of the main graph; those after to the right:

Left of Graph:

- left input signal level
- right input signal level
- “XP” expander gain reduction
- “DS” de-esser gain reduction

Right of Graph:

- left limiter gain reduction
- right limiter gain reduction
- left output signal level
- right output signal level

The level meters are dual-indicating peak-over-average, with peaks riding as a lone “dot” over a solid bar-graphed average. An “Over” indication is at the top (but slightly distanced from) the input level bargraphs.

To the left of the screen is a small button labeled “60 dB Range”, allied to the input signal level bargraph; when clicked, it toggles between normal (30dB, un-ticked) range and 60dB (ticked). This latter is useful for “finding” signals during initial set-up, or when the signal source is of high dynamic range, such as live classical.

Signal Path Selection

Beneath the input metering bargraphs is a large button, which toggles between “View Left” (colored red – port) and “View Right” (green – starboard) if the AP3 is set up for stereo, or “View A” and “View B” if operating as two discrete channels. This selects which of the two signal channels within the AP3 is being operated upon / displayed.

TIP: “A” and “B” correspond to the “Channel 1” and “Channel 2” labeling on the front of the AP3 unit.

Control Area Region

This is the large "enclosed" area directly above the graph and meters.

All the variable knobs and faders are within this area. To operate them, either:

• place the cursor above the desired control, and use the mouse's scroll wheel to "increase" or "decrease" the control's position, or,

• click on the control, and slide the mouse to move the control - up or right increases the control's indicated position, down or left decreases it. In either case, the resolution of the control variation can be increased - made finer - by holding down the keyboard's control key ("Ctrl") while adjusting the control as above.

Left-to-right across the upper portion of this area are a series of buttons corresponding to individual or logically grouped processing elements in the signal flow, in input-to-output order. A small indicator at the left end of each button indicates if any signal processing within that block is currently active. Left-clicking a button gains access to the controls for that portion of the signal-processing.

NOTE: The order of EQ and compression can reversed on the AP3 by the condition of "Pre-EQ" selection. This is achieved on the GUI by clicking upon the "SWAP" button linking the EQ and compressor blocks visible in either of these functions' control screens.

Input

Representing input source selection, input gain control, balance/offset, symmetry decorrelator, absolute and relative L/R phase, M/S capability, and pre-emphasis selection.

Symmetry

Specifically intended for use with voice, for greatest benefit it needs to occur before any dynamics processing has taken place. By decorrelating the fundamental and harmonic elements of a voice (far less painful than it seems) the relatively high asymmetry ("positive" going signal peaks vs. "negative" going peaks) of some 8 to 10dB of a normal voice can be substantially reduced, with little audible detriment. This, by in effect reducing the "peak to average" ratio of the voice signal, can allow a significant increase in apparent voice loudness in a dynamically controlled path.

Phase

Phase reversal of the signal path, in case there is an inadvertent reversal elsewhere in the signal chain of which the AP3 is a part.

Filter

High-pass, low-pass filters, and notch.

High-Pass Filter

A fast rolloff-rate high-pass filter, adjustable between 20Hz and 1kHz.

Notch

The notch is tunable over the entire audio range of 20Hz – 20kHz.

Low-Pass Filter

A fast rolloff-rate low-pass filter, adjustable between 20kHz and 1kHz.

Expander

An expander such as the one in the AP3 (often called a "downward expander") is a useful tool for reducing unwanted background noises. These could be variously air-conditioning rumble or noise, background conversation, phone-line noises, recording hiss, etc. It is also useful for reducing the inevitable general increase in background noise of some recorded material when subject to heavy compression. A common usage in live sound is to effectively turn a microphone off when not being talked/ sung into, so as to reduce corruption of a mix or reduce the chances of feedback with an unwanted open microphone.

The expander is slightly counter-intuitive when first encountered, in that unlike nearly any other processing element it is active - i.e. working, attenuating away the input signal - when the input signal is at its quietest, at or below the threshold. If the expander is on, there will be gain reduction when no signal is present. The gain-reduction reduces as the threshold is approached, and there is none above the threshold.

The controls are:

THRESHOLD — Below which the automatic attenuation starts to take effect.

RATIO — Being the proportion of how many dB the input signal is attenuated for every dB it drops below the threshold. 1:3 indicates 18dB loss for 6dB drop in input signal level.

DEPTH — The maximum amount the expander is permitted to reduce the input signal level.

OPEN — The time-constant of the rate at which the expander un-attenuates, or opens; sometimes called "attack".

HANG — An adjustable period of time the expander remains open without attenuating, before starting to close. Handy to keep the expander open during, say, speech inter-syllables or other short pauses, without having to resort to excessively long...

CLOSE — ... close times, being the rate at which the expander attenuates away the input signal once below the threshold.

An input/output plot, a graphical representation of the relationships between threshold, ratio, and depth, is on the Expander Control screen of the GUI; it is a handy visual aid.

Almost always, the trick is to set the threshold of the expander - below which it starts to attenuate away the input signal - high enough to capture the noise, but not too high as to snatch at the lower levels of the desired parts of the program material. That can sound really irritating.

Sometimes the gain reduction is required to be subtle so as not to draw attention to the fact that the expander is in operation; under these conditions shallow expansion ratios, such as 1:1.5 or 1:2 are preferred, as are restricted depth - 6dB, or 10dB, is plenty and makes a substantial subjective improvement to the noise.

These, too, are the kind of settings used for another application of an expander: effectively shortening an excessively long room reverberation time, or an instrument's ring-out that is overly persistent. In these cases the threshold is set somewhat higher, well up into the desired audio levels - in this way the attenuation becomes part of the overall sound, but the gentle ratio prevents a sense of anything "odd" happening. Again, relatively shallow depths of 12dB or so are plenty to achieve the desired effect.

More aggressive expansion, or "gating", is accomplished with steep ratios (1:3, 1:5) and with shorter open and close times than for "unobtrusive". It is still best not to go overboard with depth - even just 14dB, 20dB tops, is enough to make a signal "disappear" in the context of a mix; the whole gating sound, especially surprisingly its opening, is less obvious with shallower depth. Sometimes the "Surprise!" element is required, though, for effect.

De-Esser

The de-esser is a dynamic equalization section, designed to restrict the output level of a chosen frequency range within the spectrum. The classic usage is to hold down the sometimes prominent and objectionable sibilant “hissy” noises in speech, particularly if poorly recorded or subject to poor or overly close microphone technique. It is also often used to help control the “honky” effect on phone lines; mercifully, with digital phone systems this is less of an issue than it once was, but it still crops up. The controls are:

THRESH — Threshold is the signal level (within the de-esser’s controlled bandwidth) at which the de-esser holds the output level, again, within the de-esser’s bandwidth. In other words, a signal within the de-esser’s band will not be allowed to exceed this threshold level.

ATTACK — The time-constant controlling how quickly the de-esser responds to a signal exceeding the threshold.

RELEASE — The time-constant controlling how long the de-esser takes to recover from a signal that has exceeded the threshold.

FREQUENCY — This is the center-frequency of a filter which determines the band of frequencies over which the de-esser operates. It is operated in conjunction with . . .

BANDWIDTH — ... which determines how wide on either side of the center-frequency the de-esser is sensitive.

The de-esser screen on the GUI brings up a graphical representation on the big frequency-domain graph of the de-esser’s filter, as controlled by “Frequency” and “Bandwidth”; the 3dB-bandwidth of the filter is arranged to be at the center-line of the graph for easy visualization. Particularly when the input signal spectral analysis is being displayed, it is often easy to visually “tune in” the de-esser to the visual clues of the objectionable sound on the analyzer.

Once the threshold within the de-esser’s band has been broached, an “inverted shark’s fin” reaches down from the top of the frequency-domain graph, indicating the width in frequency and depth of attenuation due to the de-esser.

Regardless of the visual clues available thanks to the GUI, the most direct way of tuning in the de-esser is the old-fashioned way - by ear. Make the “Bandwidth” small, narrow; crank down the threshold (i.e. make it more sensitive), then carefully tune the “Frequency” until the undesired artifact is getting audibly affected. Careful (!) adjustment of “Bandwidth” and “Threshold” is then undertaken to best control the artifact. It is nearly always a compromise - as with any “garbage disposal” type processing - between reducing the undesired artifact, and mitigating its effect on the rest of the sound.

A less conventional use for the de-esser is as a dynamic EQ. Using it to shape the frequency response at higher signal levels (leaving it alone at lower levels) can be far less sonically obvious than digging a big hole with conventional EQ. An example would be using the de-esser to control the mid and mid-high frequencies of a music bed over which is to be laid a voice-over; the bed can be kept loud and energetic - and still natural-sounding - while not interfering with the important voice frequencies.

It should be noted that the de-esser is quite sensitive to changes in input level; ensure that the input signal is at a reasonable level, and stays that way, for best and consistent effect.

Emphasis and De-Emphasis

Emphasis is available at the AP3’s input, and both pre-emphasis and de-emphasis at the output.

Emphasis (also known as pre-emphasis, particularly in circumstances where it is operating in conjunction with a complementary de-emphasis) is the imposition on the frequency response of a system a fixed and known frequency-dependent gain; the most usual implementation – and that used here - is an emulation of the fixed 75uS (Americas) or 50uS (elsewhere) high-frequency boost used in FM broadcast transmission. These, roughly speaking, have a turnover frequency (at the 3dB gain point) of 2kHz and 3kHz respectively, with a 6dB-per-octave gain increase above.

The reasoning is that since most program material has decreasing energy with progressively higher frequency above the mid-range, higher frequency signals may be boosted accordingly and still not breach headroom or deviation limits. Why? So that a complementary de-emphasis filter at the far end of the transmission chain (receiver) can reduce them equally, so as to render a flat frequency response again. Why? So that high-frequency noise inherent to and introduced into the transmission chain may be reduced by the rolloff of the de-emphasis filter, too, dramatically improving the noise performance of the path.

Since the AP3 has great utility in both FM and AM broadcast paths, a highly comprehensive set of pre- and de-emphasis controls is made available. In approximate order of complexity, typical applications are listed:

Transmitter pre-emphasis

The usually built-in pre-emphasis in FM stereo encoders and exciters may be replaced by pre-emphasis in the AP3. The advantages are that it is then possible to monitor pre-transmitter exactly what the overall processed signal is doing; the STL (Studio-to-Transmitter Link, often a radio itself) may be included within the pre-emphasized section of the path, and avoiding (by its disablement) a separate pre-/de-emphasis sub-path within that link; the excellent zero-overshoot output peak limiter in the AP3 captures any overshoots that may spring from the pre-emphasis.

Using the pre-emphasis from the "Output" section allows the modified signal to be bounded by the output peak limiter, as described. Alternatively, pre-emphasizing at the "Input" allows the full gamut of processing to be operational on the pre-emphasized signal; this can be advantageous, in that the multiband compressor/limiter in particular can offer a greater degree of control over the now-contoured signal, and leave less to the output limiter to catch; this has sonic benefits. This also makes the HF band of the multiband processor "work harder", which may at times be of benefit if "LOUD" is the desired end result, but at the very least will require deliberate attention and scrutiny - "too much of a good thing" department. The triband processor certainly will have to be set up differently for pre-emphasis; it will otherwise do a good job of neutralizing the boosted high frequencies, with the net result it will almost be as if pre-emphasis were never even applied!

When adjusting processing pre-emphasized at the input, it is "vital" that monitoring is done with complementary de-emphasis in place, either within the AP3, or through end-to-end off-air monitoring. "Raw" pre-emphasis sounds very, very bright and indeed unnatural, and it is very easy to inadvertently adjust things such that the overall chain sounds dull, since an innate tendency is to make a path sound "flat", or at least more familiar and natural.

Transmission path emulation.

Particularly when an AP3 is being retrofitted into an existing airchain path, or one which maybe does not allow the flexibility to early-pre-emphasize, as above, the AP3 may be used to internally capture, control and so neutralize any downstream "surprises" from pre-emphases in STLs, exciters, or encoders.

This is achieved by pre-emphasizing at the "Input", and complementarily de-emphasizing at the "Output". In this way, the processing "sees" a pre-emphasized signal (such as would exist in an STL and/or sender) and can control it appropriately, catching things and handling them perhaps more gracefully than the more rudimentary limiters downstream might. Additionally, a fully processed yet "flat" signal is available for monitoring; DJs like to hear the processed signal, rather than the usually much less exciting sounding console output.

So, in short: The pre-emphasis is to emulate that required of a typical transmission scheme. By bringing it into the AP3, superior control over the overall signal path, and superior airchain processing, may be achieved. The de-emphasis in the box allows for normalization of a pre-emphasized signal within the box either for monitoring, or where a "flat" signal is required for onward passage.

M/S

This is a technique for stereo used in many places for different reasons, and is often known as "Sum / Difference" processing.

An "M/S" switch, as implemented in the AP3, will turn a conventional left/right stereo pair of signals into a sum-and-difference pair; the "left" path then becomes a mono sum of the applied left and right signals, the "right" path now carrying the difference (left minus right) of the applied pair. Conversely and conveniently, the same "M/S" processing turns an applied sum-and-difference pair into their constituent left and right signals. Neat, eh?

The AP3 has "M/S" switches at both "Input" and "Output". This allows a wide range of additional processing possibilities:

M/S microphones

The name "M/S" in fact comes from "Mid / Side" - a stereo microphone technique where instead of a left and right sensing pair of microphones, there is a single forward-sensing (or omni) microphone, and a second figure-of-eight sensing microphone aimed sideways. Particularly in classical and other acoustic-space-rich style recordings, this technique offers latitude for easy remote-pattern and stereo image manipulation, and far better sound-stage positional accuracy than possible with a conventional pair of cardioid or super-cardioid microphones.

This sum/difference pair may be converted to conventional stereo by either the "M/S" switch at the input or output of the AP3; the latter offers the possibility of differential processing (meaning different between sum and difference paths) way beyond the usual 'width' control afforded by merely adjusting their relative gains.

Mono Airchain Processing (e.g. for AM)

The "M/S" switch will present a mono sum of the applied left and right signals from a stereo production or air-suite into the "Left", or "A" path of the AP3. The "Right" or "B" path would be ignored in this case.

FM Airchain Processing

Particularly in conjunction with the earlier described internal pre-emphasis, the "sum" and "difference" signals required for stereo multiplex encoding may be created within the box; this, again like local pre-emphasis, allows greater control over the ultimately broadcast signal.

Sum / Difference Processing for fun

With both the "Input" and "Output" "M/S" switches set, a stereo source will emerge from the box in stereo, yet the whole gamut of intervening processing will operate on their sum and differences. If the box is kept in "Stereo" mode, with identical processing occurring on the two paths, the net effect will be substantially the same as if the "M/S"ing hadn't been done (dynamics tend to act slightly differently when driven 6dB harder, as will occur on a substantially "mono" source, but otherwise it's pretty much a wash). If, however, once baseline configuration of the two paths has been done in "Stereo" mode, the AP3 is then made to operate in "Two Monos" mode, things get interesting.

Width

As alluded to under "M/S Microphones", changing the difference path gain in relation to the sum affects the apparent stereo "width"; reducing the difference to zero leaves solely the sum, i.e. mono; if the path gains are equal normal stereo occurs, but if the difference gain exceeds the sum, then a "super-wide" image is possible. It is quite extraordinary to "hear" sonic images come from outside the physical separation of a pair of speakers. This technique is called "shuffling", and like all good parlour-tricks is best used in moderation - it is too easy to create "hole in the middle", for the stereo effect to disappear and with two mangled images huddling around the speakers. Bear in mind that most good recording guys know this trick already and use it to some degree, usually on reverbs; an inadvertent excess added by your dainty hand can sound horrible.

Mono Low Frequencies

For disparate reasons, it has long been common for the heavy energy containing low-frequency elements of a mix to be panned central on recordings. Thankfully setting aside the constraints of disc-cutting(!), the main reason is that it is possible to get up to 6dB more level when using both channels mono-like instead of just one, as would be the case of something panned hard left or right. (More level = loud = good, Igor!) However, some program material isn't like that, either because it's old ("stereofied" fifties and sixties singles, Beatles in particular) and having instruments panned hard; the producers didn't get the message; were trying to be artsy; or it just naturally isn't like that (jazz, classical). In the kind of world the AP3 inhabits, centering the LF energy has many benefits: in FM the difference subcarrier is excited less, which translates into less "scratchiness" on marginal signals; the LF can be made louder (that two channels vs, one thing again); perceptual encoders can create more efficient data compression with less difference energy to "aim".

By high-pass filtering away the "difference" LF energy, one is just left with that in the "sum" path, i.e. mono. Compression / limiting downstream will allow any additional gain to be realized.

Naturally, this trick is not limited to just LF; particularly in conjunction with "Width" (above), filtered, EQ'd and dynamically modified "difference" paths can lead to some interesting, some wild effects.

But a word of warning. If the station's music library, shall we say, owes much to the glory days of Napster, or otherwise contains aggressively perceptually encoded material, one should be highly circumspect of doing much in the way of "sum"/"difference" processing. Some perceptual encoders play games with, indeed take great liberties with, 'difference' information during encoding, using the absolute bare minimum to convince the ear of an image's stereo positioning. By and large it works very well in context, but not when the difference information is isolated or exaggerated - like by the very techniques just described above. Then, there can be excruciatingly nasty-sounding images at best, space-alien "beepity-boop" noises at worst. (Go on, try it.)

Output Limiter

Understanding the output limiter is also key to grasping the Multiband Limiter's operation, detailed later.

The AP3's output limiter is designed to strictly control the unit's maximum output level such that overdrive or clipping the input of a downstream device is prevented. Typical concerns are the introduction of distortion products in perceptual encoding systems (resulting in unnecessarily high processing/bandwidth, inefficient encoding, and "funny noises"), FM sender over-deviation (naughty), and AM sender over-modulation (calamitous). A set of recommended parameters to achieve this performance are given; in addition, though, the limiter may be used with more relaxed settings for dynamic control not requiring these stringent constraints.

There are a number of controls on the AP3's output limiter which are unusual or seemingly unconventional. But first, the easy ones:

THRESH—The signal level which output level is desired not to exceed.

ATTACK—The period of time an applied signal must instantaneously exceed the Threshold in order to reduce the output signal down to the threshold. 0.5mS is recommended for zero-overshoot limiting.

RELEASE—The time the limiter's gain takes to recover from a typical degree of limiting once the signal no longer exceeds the threshold. 50mS is recommended for transparent zero-overshoot limiting.

GAIN—This is used to adjust the output level of the limiter; a typical use is if it is required to have the limiter "brick wall" at an effective threshold of +20dB (which equates to 0dB full-scale digital). The limiter's threshold would be set at say +8dB (allowing 12dB "headroom" for the limiter to operate without constriction), and then +12dB of gain to be added using the "Gain" control.

CLIP ENABLE—When enabled a zero-attack-time limiter at the set threshold level is activated, beyond which no overshoot is then possible. It is used to contain transients shorter than the limiter's attack-time and other such anomalies that escape the limiter. Used as such, it is sonically invisible.

Warning: "Clip" is intended to be used in conjunction with the limiter; serious distortion can result otherwise. Alone it can be anything other than sonically invisible.

CLIP BACKOFF—The indicated threshold on the output limiter represents the peak output level; it is actually the level of the clipper. Ordinarily, if the limiter is running with it's default values for "look-ahead zero-overshoot" (0.5mS Attack, 50mS Release) the limiter will catch pretty much everything leaving the clipper to take care only of the occasional sub-attack-time transient and such. If the attack time is lengthened, however, longer transients will escape through it; at the longer attack times (say 50mS or longer) the limiter ceases to be a peak limiter and tends to become more of an averaging detector. (This is readily observed with nearly any peak limiter - the peak output increases by 6 - 10dB, or even more depending on source material, when the attack time is increasingly relaxed.)

Under this condition - long attack time, for a less "tight" sound - peak signals many dB higher than the threshold will be allowed through the limiter. If "Clip" is enabled, distortion as these are peremptorily removed can occur. The "Backoff" control allows the limiter's threshold to be reduced in relation to the clip threshold (set by "Threshold"); this affords the running of looser limiter attack times whilst retaining the absolute maximum output limit set by the Clip.

In general, it is set by ear; like many other settings, the degree of clipping which is tolerable is truly subjective; no amount of blinky-lights or cute graphics can help with this decision. Naturally, looser limiter attack times often sound better, and the clipper helps maintain peak level control; the result though is less 'loud' in comparison to the limiter set up for zero-overshoot. A judgment call.

The "Backoff" control, which adjusts the limiter's threshold to a value lower than that of the clip level (set by 'Threshold'), allows one to run looser attack times (often less "constrained" sounding) without the resulting excess level clipping, or clipping too regularly.

POSITIVE MODULATION—This is the percentage by which the positive-going output level can exceed the negative-going, represented by the output limiter's "Threshold". The FCC presently allow 125% for AM senders in the US (some other countries being more liberal). This affords a noticeable increase in on-air loudness.

The positive-going Limiter and Clip threshold is increased by the indicated amount over the negative-going threshold, which remains fixed at the value set by "Threshold".

For all normal purposes, this control should be set to 100%.

(Should be ignored for all applications other than AM broadcast, except for the bizarre experiments of obscure PhDs somewhere in academia.)

Output Limiter: In depth

To achieve guaranteed lack of overshoot, advanced look-ahead and up-sampling algorithms are employed, to respectively catch most overshoots before they occur, and mitigate untoward "aliasing" effects ("grittiness") that plague lesser digital limiters. When set up correctly, the limiter will not allow excess output at all, and do so gracefully and sonically transparently.

"Look-ahead" refers to the technique of fractionally delaying the input signal such that the limiter's detecting sidechain has a chance to react, generating an appropriate control signal at the exact time it is required to control the passing audio signal. A short delay is used so as not to incur problematic additional latency, and also to minimize "escapees". These refer to impulses and transients which are shorter than the attack time of the limiter, and which, obviously, cannot then be accurately captured and controlled regardless of the look-ahead technique.

The longer the attack time, the more energy escapes the limiter; at (not very great) extremes, a clipper - by definition zero-attack time - allows no escapees, while a 100mS attack time is essentially an averaging detector, allowing sometimes wildly high instantaneous peaks through.

To this end the AP3's look-ahead limiter with attack times shorter than 1mS dialled in (0.5mS is the specified and recommended value) captures almost all transients found in program material; the downstream zero-attack-time limiter - "clipper" - takes care of the rest. A release time of 50mS (recommended) on the limiter will catch and control signal peaks quite transparently; by itself it can add a considerable amount of "LOUD", but at the cost of the well-known sonic detriments that come from excessive broadband limiting – and which are the reasons for multiband processing in the first place! It is recommended that the output limiter just be used to control signal peaks just a few dB deep - to this degree it is quite transparent. Deep processing should be left to either of the triband processors.

AP3 Features for AM Broadcast

The AP3 is an ideal AM broadcast airchain processor given the inherent multiband and broadband capabilities and a number of key AM-specific features:

Brick-wall Bandwidth Filters.

These are accessible on "Output". Three bandwidths are available:

4.5kHz—common where 9kHz channelization is in use (outside the Americas).

5kHz—required should an AM station be simultaneously broadcasting IBOC, and also recently proposed as at least a "night-time" bandwidth standard within the US, aimed at reducing adjacent-channel chatter.

10kHz—the US NRSC AM bandwidth standard.

Emphasis

It is common to add a rising gain-vs.-frequency emphasis to combat the inevitable high-frequency rolloff of the IF filters in AM receivers. The NRSC standard, for use in conjunction with the above 10kHz brick-wall filter, is for use of a modified 75uS pre-emphasis, conveniently available on your local friendly AP3's "Input" or "Output". The "modification" consists of a gentle rolloff of the pre-emphasis slope approaching 10kHz, and with the AP3 may be nicely emulated by applying the low-pass filter ("Filters") set to a frequency of approximately 12kHz.

Truth be told, the niceties of matching exactly such a paper specification get substantially blown away when complex processing is used, but arranging for a non-excessive rising trend in response either by pre-emphasis, multi-band processing, or EQ to counter receiver IF "droop" is a good idea; the brick-wall filter will keep sideband energy happily within the mandated spectral mask.

Asymmetric Output Peak Limiter

The look-ahead output limiter in conjunction with its zero-attack-time "clipper" ensure zero overshoot, and allow one to get very close to sender linear modulation extremes with confidence. The limiter's "Threshold" represents the negative-going peak level, whilst the positive-going threshold is separately controllable to afford up to 150% positive modulation. (The current US limit is 125%.).

Bear in mind that if set at say 130%, that is 130% of whatever depth the negative-going level is set to in the real-world on the actual sender; if that happens to be 90%, then the "actual" positive modulation will be 90% of 130%, or 117%. For a typical real-world depth of 95% negative modulation, the "Positive Modulation" control would need to be set at 131% to achieve actual 125% sender positive modulation. Confused yet?

Pos. Mod. Dial % = (100 / actual neg. mod.%) * desired Pos. Mod. %

The closer to 100% negative modulation one has the nerve to achieve, the closer the displayed and actual positive modulation numbers will agree.

Mono

In the event that there is no mono program source available from the console / program suite, the left and right signals may be applied to the appropriate inputs on the AP3; the "Input" "M/S" switch will sum the left and right sources to mono in the "Left" or "A" path of the unit, which can be set up for mono processing, the "Right" or "B" path being ignored.

Parametric Equalizer

This powerful four band parametric equalizer is selectable to be prior to or after the compressor ("pre" or "post").

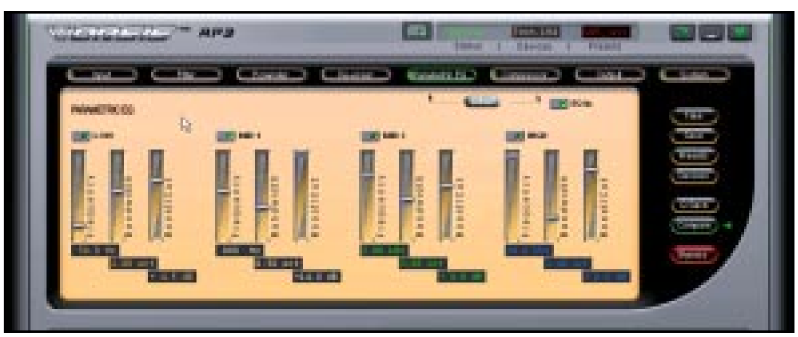

Multiband Processing

The three-band dynamics section is the keystone of the AP3. Two different styles of multiband processing are available in the AP3 II, entitled loosely "Production" and "Airchain". The former uses a very high-quality soft-knee compressor for each of the three bands, and is intended to achieve excellent transparent control over "raw" program sources, whether musical or spoken voice. The emphasis is on using the strengths of multiband processing to render a surprising amount of dynamic control "invisibly" and with a minimum of audible artifacts.

"Airchain", on the other hand, is geared to aggressive attention-grabbing processing, as for better or worse is a feature of the broadcasting landscape. Its reason for existence is to be loud. There are three frequency bands comprising "Climps" - extremely clean zero-overshoot limiters each followed by zero-attack-time "clippers" - which can dramatically reduce the peak-to-average ratios of the audio within their passband, and hence substantially increase their apparent loudness. A choice of "clip" styles allows the use to be better tailored to differing types of program material.

The choice between these two styles of multiband processor is made by check-boxes on the 'Multiband' screen of the GUI, and as a "ProcMode" selection under "Multi" on the unit itself. A description of each follows.

Multi-Band AGC/Compressor - "Production"

This affords both highly transparent compression free of the artifacts endemic to normal compressors, and yet at the other extreme wild dynamic equalization effects.

It was felt an unnecessary disservice to "dumb down" the control of this compressor; all the "handles" allowing full control over it are made available. Yes, at first glance the wealth of controls can seem overwhelming; cognizant of this, pre-cooked presets and the addition of a couple of controls have made the multi-band compressor simpler to use, with little need to delve into deeper mechanics unless really desired.

The basic operational model for the compressor is that of the classic UREI 1176 or DBX 160 units of yore; the compressor "core" is enclosed within an input level, or "Drive" control, and an output level, or "Makeup" control. Simply, the depth of compression is adjusted by driving the compressor harder or more gently using the "Drive" control, while the output level is adjusted accordingly with the "Makeup" control. These controls are at the beginning and end of the top row of controls on the GUI's compressor screen, so are easy to find.

The wide variety of presets supplied with the AP3 have compressor settings ranging everywhere from slight to brutal; it is recommended that some time is spent calling in different presets and playing with the “Drive” and “Makeup” controls on each to get a feel for their effect on program material. From this will eventually spring the desire to explore the compressor’s entrails.

Why three bands? Even the best of broadband compressors (including the ones that play games with their sidechains to mitigate the effect) suffer from auto-modulation; this most usually manifests itself as heavy low-frequency information (kick drum, bass) audibly controlling and pumping the level of higher frequencies. Splitting away the bass region into a separate compressor from the rest of the spectrum gets rid of the worst of this effect (two-band unit) but similar, if lesser, issues between mid and high frequencies then become plain. The advantages of a second split between mids and highs (three-band) primarily stem from the highly tempting ability to treat the mid frequencies (by and large the intelligence-bearing frequency range) differently from high frequencies. Of course this line of thinking can and does extend to five, six or more band processors; three was felt to convey most benefit, the most “band for the buck”, for the Vorsis AP3’s intended applications.

Multi-band processors have the reputation - born of being able to apply far deeper and more aggressive compression than broadband compressors for the same amount of sonic “damage” - of being LOUD. Yes, indeed, as a bus compressor the AP3 can be persuaded to make a VU or PPM meter stand still, but the unit’s greatest strength is its ability to provide firm and positive dynamic control over originating source material quite transparently.

The heart of the three-band compressor is a highly regarded full-blown soft-knee compressor very similar to the design used as a broadband compressor in Wheatstone Corporation’s high-end TV and radio broadcast consoles. As such, the design’s ability to get deep compression and level protection without “funny noises” was paramount. A good start for a multi-band processor.

In the AP3’s compressor, the input signal is fed through a crossover, producing nominally High Frequency, Mid-Frequency and Low Frequency bands; the high-accuracy phase-accurate crossover frequency points are determined by the “HI X” and “LO X” controls; these are vertically disposed on the right-hand side of the GUI’s compressor screen. These crossovers are fairly gentle, so as to minimize odd-sounding spectral transitions between the frequency bands.

These three bands are then fed into three identical compressors, signified blue for HF, yellow for MF, and red for LF on the GUI. These colors are also used on the dynamic graph, indicating gain-reduction by

frequency. Each of these compressors has the normal controls to be expected of a good compressor. In order:

THRESH — Threshold, the level (within the frequency band) above which gain reduction is applied.

RATIO — The proportion by which a signal exceeding the threshold is reduced in level. "3:1" means that a change in signal level above the threshold by 15dB will be reduced to only a 5dB change in level.

ATTACK — How quickly the compressor responds to a signal exceeding the threshold.

RELEASE — The speed with which the compressor recovers as the exciting input signal reduces or disappears.

These all act just as would be expected on an ordinary compressor. The nature of a multi-band compressor means that slightly more complex approaches to inputs and outputs have to be employed, paradoxically to make the unit simpler to operate. In particular, the relationships of "Drive" and the three "Thresholds", and again "Makeup" and the three "Trims" will be explored.

On the GUI's compressor screen it is no accident that the three bands' "Threshold" controls are in the same vertical line as "Drive", and that the three bands' "Trim" controls are in line under "Makeup".

In the AP3, the three traditional "Threshold" controls may be regarded as simply offset trims per band of the overall "Drive" control. In other words, it is the **differences** between the three threshold values that are far more salient than their absolute values, which get subsumed by the overall "Drive" control. For example, it is common (since there is generally less energy at higher frequencies) to want to reduce the HF band's threshold (i.e. make it compress sooner) in relation to the lower bands; in that way all the bands are more likely to start compressing at the same time with applied program material. Again, it is the three thresholds' relative differences that are important, not their absolute indicated values.

Similarly, the three bands' "Trim" controls may be regarded as offsets to the overall "Makeup" control; it is usual to need to adjust a band's output in relation to the other two for best tonal balance, or lunatic effect.

The power of these approaches is that overall amounts of compression and output level may be adjusted without worrying about disrupting the hard-won internal balances of the three compressor bands settings.

Transparent Compression

This is for applications where compression is necessary, but all effort needs to be made to make it transparent, that there be little clue that any compression is taking place. Multi-band compressors can be pretty good at this, but nevertheless require careful application.

First consideration is that the sound, the spectral balance, should not seem to change regardless of the depth of compression, or in comparison to out of compression. To this end, the thresholds of the compressors should be adjusted such that with the expected type of program material (Speech? Classical? Jazz?) the compressors should all start to act simultaneously. Differing source material types will have differing overall spectral balances, and need adjusted thresholds and crossover points accordingly to achieve this hoped-for simultaneity.

Compressor ratios should be gentle, 1.5:1 or 2:1 outside, and the same across the three bands; attack and release times should be fairly relaxed - mid-scale on the controls is not a bad start.

The band output "Trim" controls will need to be carefully adjusted such that a natural spectral balance when in compression is not dissimilar to that in the unprocessed source material. It won't take much, a few dB here or there across the three.

For pre-recorded music, a subtle "sweetening" effect may be introduced by reducing the mid "Trim" just a few dB, and adjusting the crossover frequencies to achieve a gentle dynamic "smiley" curve.

With nearly everything except modern music, the center of energy is in the mid-range: the crossovers are likely best disposed to bracket this energy region, rather than divide it. This is true with speech especially, where the 200Hz to 4kHz region carries the brunt of intelligence. Disrupting this critical spectral area with a crossover (implying differing processing either side of it) is asking for an unnatural effect.

Compressing for effect

No.1 - ignore all the guidelines under "transparency".

A good example of compressing for effect can be the transformation in spectral balance and indeed overall sound between when the source signal is below threshold, and when the source is loud enough to be well into compression. Such is the "Wolfman Jack" effect. (For those unaware, but to hopelessly understate his history, this character was a disk-jockey between the 50's and 80's with phenomenal command of microphone technique. When speaking quietly, his voice was deep, rich, gravelly and envelopingly warm, but as he started to wind up eventually to screaming, it transformed into a thin hard shriek.) Such is an effect, for want of the native human talent, for which a three-band compressor was made.

An approach is to arrange for the three bands to have descending amounts of compression; the LF band would be highly compressed at a high ratio, with the LF crossover set to encompass the lower speech fundamentals. The mid threshold would be set higher and be at a middling ratio, and the HF threshold set to trigger a touch higher yet and with a gentle slope, with the HF crossover low enough to capture a voice's higher elements and fricatives.

The effect would be that the highly compressed LF range would predominate at low levels, but would be held down relatively below the MF and eventually the HF bands as the source level increases. At high levels, the HF band would predominate.

In general, when compressing for effect (any effect), there are a couple of items worth bearing in mind:

• Listen intently as the crossover frequencies are adjusted; moving them even just an octave can have huge effect. Crossing over "LO XO" below 100Hz can lend a lovely deep "Deutche HiFi" feel to the bass, but leaves the low-mids at the mercy of the mid compressor (or vice-versa). Likewise the "HF XO"; it can be a trial finding the right crossover point. Placing it high (8kHz, 10kHz) can lend a real sweetness to the extreme highs, but at the expense of potential loudness to be gained by sharing the mid-frequency compressor's load.

• Be aware of over-compressing the high-frequency band. This can lead to the "Gabriel Effect", named after Peter Gabriel's song "In Your Eyes", in which a triangle is tastefully laid way back in the mix. If animals flee and eyes water, it is a possible indication of having overcooked the highs somewhat.

• Steep compression slopes, say 3:1, 5:1 or even higher, and short attack and release times, can overall lend more "urgency" to the sound, and tend to make it sound louder than its actual level would imply. Not too short on the attack time, though - a feature of the compressor design used is that it can act as either a peak-sensing detector, for use in peak limiting, or as an average-sensing detector for compression. Attack times shorter than say 10mS are by default starting to make the compressor act more as a peak sensor. The end-effect is that the output actually starts to get "quieter" as the attack time is reduced below that point, since the peaks rather than the average energy are being captured; the peak-to-average ratio can be anywhere from 6 to 12dB, depending on the source. Although peak-sensing can sometimes be interesting, it is not necessarily best for Loud. If Loud is the goal, use of the Multiband Limiter option ("Airchain") is the route.

AGC

The AGC is an additional sidechain operating in conjunction with each of the three compressors; generally the AGC has much slower integration times for attack and release, so that it responds more to the medium or long-term energy of the program material in a given band, rather than more typically energetic actions of the compressor. Its time-constants are controlled by “AGC Attack” and “AGC Release”, in the middle of the top row of controls on the GUI’s compressor screen. A single set of time-constants for all three bands was felt more than adequate, and has proven so. This single set of parameters does not imply that the three bands’ AGCs are coupled in any way; all three operate independently.

The AGC may be used alone (by enabling it and turning off the compressor) for gentle control of inconsistent input signals. It can also be used in conjunction with the compressor; in this case, the AGC action rides ‘underneath’ the compressor; depending on the time-constant settings and character of the program material, the AGC can follow the average long term level changes and create a gain-reduction “bed” some 12dB to 6dB under the compression gain-reduction peaks. The AGC rides the general level, the compressor processing for effect takes place on a consistently controlled signal.

On sudden application of an input signal, the faster compressor’s attack captures the onslaught, with the AGC eventually catching up. On release, the effect is identical to the much-vaunted “two-slope release” of classic compressor units such as the Audio and Design F760xrs and Joemeek SC2. On departure or reduction of the input signal the usually faster compressor release predominates until its gain-reduction contribution falls below that of the AGC, whose much slower release rate takes over. A big advantage of the compounded processes is that on normal program material, the compressor does not have to “move as far” to capture signal peaks, so reducing the “snatching” which can occur at deep compression onset.

For most purposes the combination AGC/compression is the most transparent; if the intention is wild effect, deriving all the gain-reduction from the compressor alone is probably better.

Multiband Limiter - “Airchain”

No bones about it, and no pretence: this processing section’s sole reason for being is to make things LOUD. As with the three-band compressor, it was felt that ultimately the user is better served by having all the “handles” on the processing available. There are a different set of controls for this Multiband Limiter, although the two truly important ones operationally remain similar to the “Production” style processor: “**Drive**” which adjusts overall depth of limiting and gain-reduction, and “**Output Level**” which sets the overall output from the multiband stage.

These two, assuming that initial operation is going to be from an existing preset, are the only ones of concern, at least until confidence has risen and the control screen no longer seems as intimidating as Concorde’s flight deck. The specific controls for the three bands are ranged horizontally across the GUI screen:

DRIVE — Yes, Virginia, there are four controls called “Drive” on the GUI. The top one is the important one - it is the master drive for the whole three-band section. The three below it are for the high-frequency, mid-frequency, and low-frequency processing bands. These allow the relative amounts of signal drive into each band to be varied; the range is +8dB to -20dB, more than enough to (over) emphasize a band, or effectively remove it. In practice these should rarely be more than a very few dB adrift of the nominal **0dB**.

ATTACK — The period of time an applied signal within the band must instantaneously exceed the threshold in order to reduce the output signal down to the threshold. **0.5mS** is recommended for zero-overshoot limiting. Free-form range is 0.2mS to 1Sec.

RELEASE — The time the band’s limiter’s gain takes to recover from a typical degree of limiting once the signal no longer exceeds the threshold. **50mS** is recommended for transparent zero-overshoot limiting. Overall range is 33mS to 1Sec.

BACKOFF (CLIP BACKOFF) — The discussion on “Backoff” as applied to the Output Limiter (on Page 31) is hereby set as homework. Review is strongly recommended since it is directly appropriate to the limiters used in this section. For the recommended zero-overshoot / zero-attack settings, “Backoff” should be set at **0dB**.

In short, “Backoff” is the amount that the limiter’s threshold is reduced below it’s downstream zero-attack-time “clipper” threshold. It becomes relevant if the attack time is made greater than approximately 2mS. Slower attack times allow greater peaks to ‘escape’ through a limiter; reducing its threshold with “Backoff” can prevent too much signal being clipped, which could potentially sound unpleasant. The range is 0dB to -10dB.

CLIP “SOFT” / “HARD” — The zero-attack-time “clipper”, is downstream of the band’s limiter. “Hard” clipping is, in effect, saturation. It does not affect the signal at all unless its threshold is breached. “Soft” clipping starts bending “gracefully” well before the saturation level, reducing the peak-to-average and hence increasing the apparent loudness of many program types. And also making a mushy mess of some other program types, which is why there’s a choice between the two styles.

TRIM — There is a ‘Trim’ control for each of the HF, MF and LF bands. They control the relative amount of each band contributing to the overall section’s output; the range is +/-18dB, plenty enough to emphasize a band, or effectively remove it. In practice these should rarely be more than a very few dB adrift of the nominal **0dB**.

The “Trims” are probably the most powerful controls in determining the ‘sound’ of the Multiband Limiter; they can be thought of as Big Knobs on a Gigantic three-band EQ. Paraphrasing, in general these need to be handled with care, and small adjustments make a big sonic difference.

There are a number of controls common to the three bands, or affecting the overall Multiband Limiter section: “AGC Attack”, “AGC Release”, and “AGC In’”:

These control the operation and time constants for a three-band AGC system, which when selected underlies the operation of the three bands’ limiters. The time constants are common to all three bands, but this is not to imply they are coupled; the three band’s AGCs operate independently. The time constants are usually set to be very significantly slower than those of the limiters; the intention is to provide an average-energy gain-riding, along with the instantaneous level control of the limiters.

In combination, the triband AGC and the triband “Climps” achieve enormously deep amounts of gain-reduction far fewer ‘breathing’ artifacts and such, as would occur with a corresponding depth of solely limiting. It may be thought of, and works in fact, as a few dB of instantaneous limiting riding atop the AGC’s background gain-reduction.

This way of operating is highly suited to airchain processing; in addition to the requirement for the processor to do its “loud magic”, compensation for sometimes wickedly askew signal levels from the studio complex has to occur; these can often be 20dB adrift moment-to-moment, disc-to-disc. The AGC does the auto-levelling, the limiters do the “loud”. Obviously, in order to accommodate very low occasionally expected source signals, for the normal or higher levels the airchain processor will already be very deep into gain reduction, far more than the limiters need to work effectively. In days of old, and indeed still today, it was common to use a 'pre-processor', to ensure that a consistent level was sent to the limiters so that they could operate in their "sweet spot", above and below which their sound suffered for a number of reasons. The triband limiters in the AP3 have no such limitations and will happily work 30+dB deep, but that can sound breathtaking when they "come up for air" in a program gap; AGC and "Freeze" (below) solve that issue. The integrated AGCs in the AP3 dispense with the need for pre-processing entirely.

What should be borne in mind here is that "very" deep gain reductions will be the norm, are expected and encouraged, using the Multiband Limiter; they are the combination of that incurred normally from a pre-processor / AGC and that of the main airchain processor. Be not of faint heart when the gain-reduction meter zooms well beyond 20dB.

"Freeze" — An adjunct to or alternative to the expander, for use specifically in the Multiband Limiter. If the freezer senses the input signal's instantaneous level to the section dropping below the threshold set by the "Freeze" control for more than 1/4 second, it applies a *very long* time-constant to the AGC (approx 30 seconds). This prevents the gain suddenly rising precipitous amounts, causing "breathing" and such noises, and also means the AGC and limiters' sidechains have much less "catching up" to do when the signal reappears.

A typical threshold is -40dB; this is below the level of any meaningful signal, and above room noises and well above system noises.

Loud! - The life and times of the Airchain Multiband Limiter

No pretences here about transparency or naturalness; the name of the game is Loud. This is where the Multiband Limiter (Airchain) option excels, Aggressive settings, "Drive" high, and short attack and release times.

MIDS — The ear is most sensitive at mid frequencies - aim to compress hard in this region, more so in fact than elsewhere. The less compressed "space" in the adjoining bands can add significantly to an illusion of loudness. Consider EQ'ing in this frequency region, too. EQ'ing pre-compressor has the benefit that an amount of EQ which would be impossible otherwise (through eating up headroom and just plain sounding horrid) just happily gets squashed by the compressor, adding a density and power rather than a screeching sense of over-EQ'ing.

BASS — Bass can be recovered, indeed exalted. After the three-band limiting, there is often "room" to bring up the level of the bass, which *definitely* lends a sense of increased weight and loudness to the whole; it is probably this region which benefits the most from such processing. Subtle EQ'ing post-limiting - rather than simply turning up the trim - to bring up those specifically required elements of the sound, can add worthwhile enhancement.

HIGHS — As with all rules which are merely to be broken, it is a good idea not to squash the highs too much. It generally gives one a headache. Fortunately the nature of energy distribution by frequency in music is such that there is less signal level in this top band than the other two, it doesn't drive its limiter as hard and remains fairly "natural". Since as a consequence less gain-reduction is applied to it, the high-frequency "Trim" control needs usually to be reduced a few dB in relation to the LF and MF in order for the output spectral balance to be somewhat sane.

Creating a "sound"

Talking of spectral balance, far more than anything else, the "sound" of the processing (that with which you are trying to impress the Program Director) is controlled by the three "Trim" controls. As earlier alluded, these act as three knobs on a giant three-band EQ. Although there is very wide latitude here, an effort should be made to keep the average of the trims' settings around 0dB, with overall lift-and-cut being done by the "Output Level" control. By the time the signal has passed through the multiband limiter, each band's level is well-and-truly fixed in relation to the other two, and this relationship is at the mercy of the "Trims". Nearly anything can change spectrally ahead of the multiband (including EQ - subtle EQ in particular) and its effect will as good as disappear, with the three-band's relationship overwhelmingly impressed upon it. The consistency of this sound is one of a multiband processor's greatest strengths. Considerable thought and listening to (unprocessed) material to find a sound one would like to emulate with the box is well worthwhile; it is always easier working towards a known sonic goal. Post EQ'ing to provide detail shading to the spectrum, having used the "Trims" as broad-brushes, is a good approach.

Deep Gain-reduction

One should not be shy with depth of limiting in this section; a little is absolutely not enough. Gain-reduction depths (AGC plus Limiting) should be of amounts to make genteel ladies swoon and fly totally against one's better nature, previous training, and instinct. There is a master "Drive" control there: Use it. There are good reasons for this:

For all normal program material intended to be heard, the multiband section should be in limiting. This way the known and desired sound of the processor is definitely having effect on the signal. This NOT like normal "Production"-style or "Bus" compression where one is merely interested in reducing the dynamic range and preventing overload - that's what the "Production" compressor is for. The multiband limiter is used to ensure that whatever is intended to be heard is heard, and heard in the way one wants it to be heard. In conjunction with their integral AGCs they allow the aggressive processing necessary to increase apparent loudness at the same time as regulate and compensate for a wide variety in source levels. Under normal circumstances, this can result in jaw-dropping total gain-reductions. Just remember they are in effect the combined gain-reductions of a "leveller" or pre-processing AGC plus that of the limiting stages, which with other solutions would be separate, but total to no less.

Secondly, be aware that the sound whilst in limiting is likely to be radically different to the unprocessed sound, once it drops out of limiting. So, unless this is a desired effect (and it can be a valid and interesting one), ensure the signal stays in limiting if you intend it to be heard!

Setting Up the Output.

Once a satisfactory spectral balance has been achieved using the individual bands' trims and possibly post-EQ, it is necessary to match the multiband limiter's output with the rest of the world. It is presumed that a desired output level from the box has been decided, and that the Output Limiter has been set up to capture at that point. Let's say it is +8dB. The Output Limiter should be set to that threshold with zero-overshoot settings (default, but 0.5mS attack, 50mS release). With normal program material running through the multiband, and gain-reductions plumbing untold depths, the "Output Level" control on the Multiband Limiter should be carefully adjusted such that the Output Limiter is active by just a few dB, say 3dB, or 5dB at most. The output of the Multiband, although each band is of a known level, is a mix of the three, and so is instantaneously uncertain within limits. Running it a few dB into the Output Limiter reduces this uncertainty without radically altering the carefully crafted sound.

No-free-lunch department: The downside to all these games is that highly Loud processing although initially mightily impressive can become tiresome, and eventually aggravating with time. There is an inverse relationship between aggressiveness and long-term listenability.

Be aware too that heavy compression is not kind to music which has been low-rate perceptually encoded (e.g. <128k MP3); the encoding artifacts start to stand out, adding to an already probably grating experience.

System

Sample-rate selection, stereo/mono selection, input settings protection from presets, security (passwords etc.), skin selection (”look” of the GUI).

To a large extent, the selections here are of the “set-and-forget” variety, and mirror the System setup features and functions accessible through the AP3’s front panel. There are five groupings to the System control screen:

SAMPLE RATE — Allowing selection of 44.1kHz or 48kHz internally generated sample rates, and the option of “Auto-Follow”, under which the AP3 follows the sample rate of a valid applied AES/EBU signal.

STEREO — Through three buttons, this determines if the AP3 is going to operate as two discrete channels (“Dual Mono”), or if being made stereo, from which of A or B channels the common stereo parameters will be taken (“Stereo from A”, “Stereo from B”).

INPUTS — “Inputs follow presets”, if unchecked, allows the AP3 to ignore input source (Mic/Line/AES) and input gain settings within a preset. The signal processing setup within a preset may be of interest, but the preset may have been made using a different source/gain than that which is hard-wired into the AP3.

SECURITY — Allows access unto password-world (“Passwords” button). It is presumed that access to the GUI grants full access to all features and functions of the connected AP3, internal password-limited access notwithstanding; indeed the password and access levels may be altered freely from the GUI. Clicking “Password” brings up a dialog box displaying the existing pass-code for the connected AP3; it may be altered by directly entering a new number between 0 and 255 in each field. The level of access outside password protection (user access) may be selected by clicking on the appropriate “radio” button.

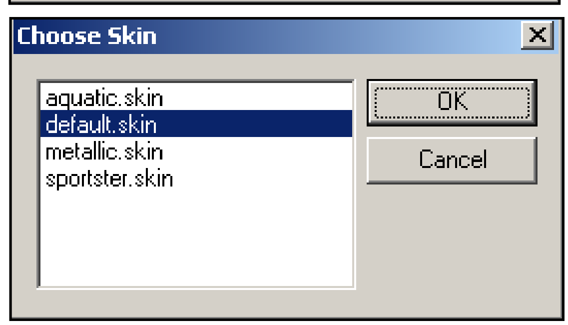

INTERFACE OPTIONS — Clicking “Skin” shows a selection of differing colors and styles available for the GUI. Selecting one and “OK” will cause the screen’s look to change in a few seconds.

Side Bar Region

To the right of the control area is a vertically disposed row of buttons:

Take

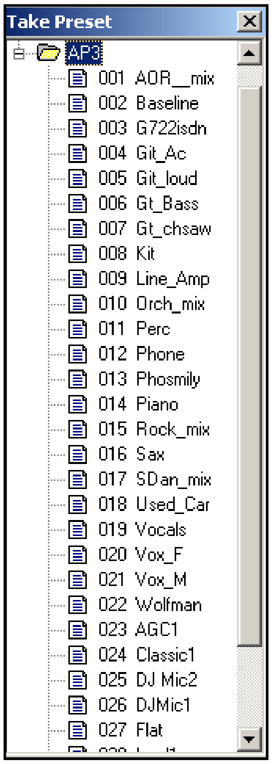

When left-clicked, a "Windows" box appears, showing folders that contain presets, typically "Favorites" and "AP3", this last being presets already in place in the presently connected device. Opening a folder (if not already so) lists the presets in order. Double-clicking on a preset brings it immediately into use in the AP3. The box stays open until deliberately closed, allowing differing presets to be double-clicked upon readily in succession. This is a very direct means of comparing presets. ***It is important to remember to save the present settings as a preset before invoking others, or they will be lost.***

Save

A small "Windows" dialog box appears, which prompts for a name under which to save the present settings of the AP3. The preset will be saved into the next available empty slot. Alternatively, by nudging the preset number ("PRE #") up or down, it is possible to save the present settings in place of an existing preset's settings. Needless to say, this erases, as in loses, as in destroys, the overwritten preset's contents.

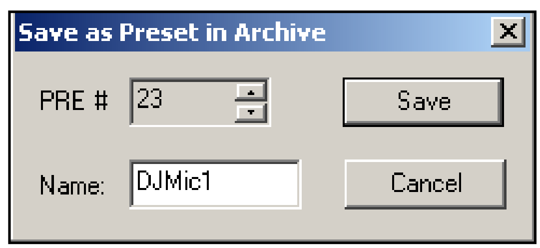

Presets

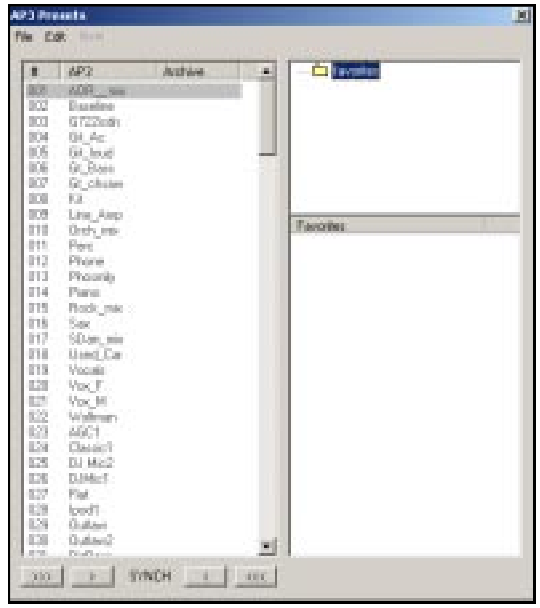

A three-panel "Windows" box appears. The large panel on the left has three columns, being the preset number, AP3 preset contents, and the GUI's mirror archive contents for that AP3; often the last two are the same or similar. If they are dissimilar, it has to be decided which is in error. The two lists may be synchronized using the "Synch" buttons at the bottom; with these it is possible to move one selected preset from one column to the other (single arrow) or all of them (triple arrows). Treat With Respect! "Windows" file drag-and-drop techniques work to move presets around, too.

The upper and lower right hand panels give access to the GUI's presets library; the contents of folders in the upper panel are shown in the lower panel; drag-and-drop allows movement between the library folders and the AP3 and its mirror archive in the GUI.

Devices

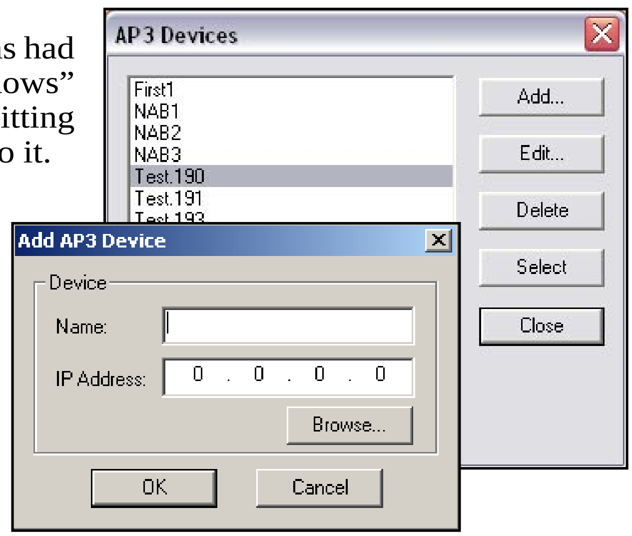

A list of AP3 devices with which the GUI has had cause to become familiar is shown in a "Windows" box. Highlighting the desired AP3's name and hitting "Select" causes the GUI to attempt to connect to it.

In the event no AP3's are listed, or one is adding a further device, or a new one for the first time, "Add" brings up a small dialog box which asks for the new prospect's name and IP address. "Browse" searches out likely victims connected to the same LAN as the GUI; this is the easy way if the AP3 has already been given an IP address.

Q Save (Quick Save)

This saves the immediate present AP3 settings in a temporary buffer, allowing continued adjustments to be made without losing a known working point. "Q Save" works in conjunction with . . .

Compare

. . . which toggles between the settings in the "Q Save" buffer and the immediate settings; an arrow indicates which of the two is active on the AP3 at any instant. This arrangement greatly facilitates incremental adjustments while building a "sound" or a new preset.

Bypass

Allows a "sanity check" by bypassing the signal processing in the box altogether; the input and output gain controls remain active.

Title Bar Region

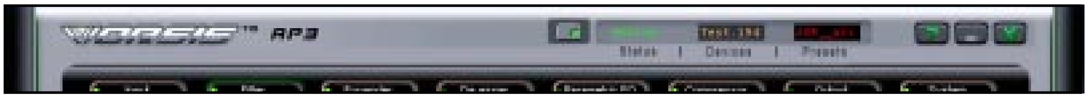

Along the top edge of the AP3 GUI screen (in line with the "Vorsis AP3" product label to the left, and the Windows "About"/"Minimize"/"Exit" icons to the right) are indicators and controls for the management of devices and presets. The AP3 GUI is capable of controlling multiple AP3s ("Devices") and managing the "Presets" within them. Whether a connection is made ("Status"), and which device and which preset are presently under command are indicated. Double-clicking these indications will gain access to respective management screens.

Status

This indicates: "Online", meaning the GUI is in communication with and directly reflects an AP3; "Offline", meaning the GUI and AP3 are disconnected; "Trying", meaning the GUI is attempting to find the AP3 on the LAN and connect to it. An adjacent button toggles between "Online" and "Offline".

Devices

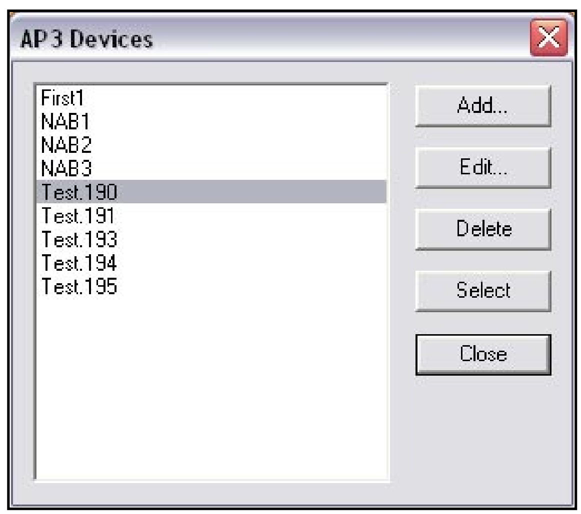

The name of the AP3 to which the GUI is connected (or with which it is attempting to connect) shows here in amber.

Double-clicking on the name area brings up the "Devices" "Windows" box. A list of AP3 devices with which the GUI has had cause to become familiar is shown in a "Windows" box. Highlighting the desired AP3's name and hitting "Select" causes the GUI to attempt to connect to it.

In the event no AP3's are listed, or one is adding a further device, or one for the first time, "Add" brings up a small dialog box which asks for the new prospect's name and IP address. "Browse" searches out likely victims connected to the same LAN as the GUI.

Presets

The name of the last invoked preset is indicated here. Green indicates that the AP3's settings for this preset agree with those in the GUI's mirror archive; red indicates a discrepancy. Double-clicking on the name brings up a "Differences" - parameter-by-parameter, setting-by-setting - comparison list between the AP3 and the GUI mirror archive.

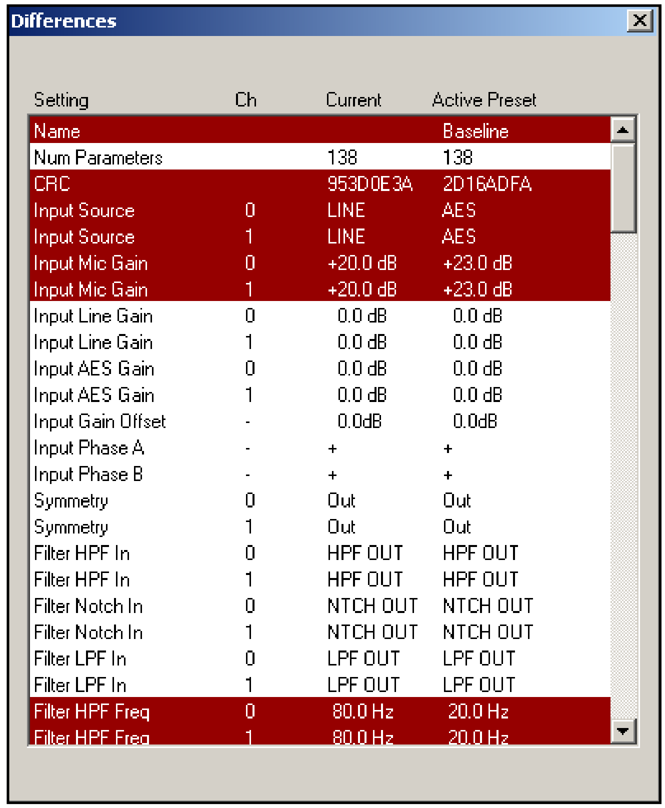
Differences

Setting	Ch	Current	Active Preset
Name			Baseline
Num Parameters		138	138
CRC		953D0E3A	2D16ADFA
Input Source	0	LINE	AES
Input Source	1	LINE	AES
Input Mic Gain	0	+20.0 dB	+23.0 dB
Input Mic Gain	1	+20.0 dB	+23.0 dB
Input Line Gain	0	0.0 dB	0.0 dB
Input Line Gain	1	0.0 dB	0.0 dB
Input AES Gain	0	0.0 dB	0.0 dB
Input AES Gain	1	0.0 dB	0.0 dB
Input Gain Offset	-	0.0dB	0.0dB
Input Phase A	-	+	+
Input Phase B	-	+	+
Symmetry	0	Out	Out
Symmetry	1	Out	Out
Filter HPF In	0	HPF OUT	HPF OUT
Filter HPF In	1	HPF OUT	HPF OUT
Filter Notch In	0	NTCH OUT	NTCH OUT
Filter Notch In	1	NTCH OUT	NTCH OUT
Filter LPF In	0	LPF OUT	LPF OUT
Filter LPF In	1	LPF OUT	LPF OUT
Filter HPF Freq	0	80.0 Hz	20.0 Hz
Filter HPF Freq	1	80.0 Hz	20.0 Hz

"SYNCH" allows the presets within the AP3 and the GUI to be made the same. The red color indicates a discrepancy between the actual current AP3 settings and the GUI's mirror archive for the preset; this usually reflects that settings have been twiddled with since the preset was "Taken". At this point a choice should be made: if the original preset is still valid a re-Take should be performed; but if the current modified version is the valid one, the current settings should be saved back into the originating preset.

To paraphrase: There are many ways during normal usage in which one or some of - (a) the actual current settings of the AP3, (b) the preset within the AP3 from which the settings originated, and (c) the GUI's archived version - can become different. A red preset label indicates this. Determine which of the three is actually correct, then make the red flag go away and synchronize the system by:

a. Current settings correct:

- Save the current settings back into the originating preset.
- Copy the newly refreshed AP3 preset into the GUI archive.

b. Originating AP3 preset correct:

- Re-"Take" the preset into the AP3.
- Ensure the GUI archive accurately reflects that preset, copy over if necessary.

c. GUI mirror archive version correct:

- Copy the archived preset over into the AP3 preset.
- "Take" the preset into the AP3.

Notes on “Online” and “Offline” Working

Most often operationally it will be required that the GUI act directly and instantaneously upon the AP3 selected under “Devices”: changes on the main control screen take immediate effect on the signal processing and any selected presets are immediately rendered active. This is “Online”.

“Offline”, however, is more of a supervisory and system management mode, and for those who feel the need to set up complex audio processors without listening to them. Changes can be made to preset settings and configurations, and manipulation of the presets selection list can be made without affecting signal processing within the AP3 itself, which can continue to be processing away obliviously. (Preset list manipulation can be undertaken “Online”, too, but only by those with nerves of steel. One click from potential disaster . . .)

Importantly, any adjustments made on the GUI “Offline” will not take effect until the AP3 is rendered “Online” again.

It is easy for the AP3 itself and the GUI to lose agreement; either the AP3 can have been adjusted using the hardware front-panel, or adjustments made on the GUI, either “Offline” or when the AP3 and GUI have had a tiff and weren’t talking to each other for some other reason. Extensive tools to reconcile such wayward units are available under “Presets” in the Side Bar screen region.

Invoking a preset - by double-clicking a preset from the box brought up by clicking the “Take” button - does one of two things, depending whether the GUI is in “Online” or “Offline” mode:

- “Online” - the double-clicked preset is immediately sent to the AP3 and becomes active.
- “Offline” - the double-clicked preset brings that preset onto the main control and display screen for perusal / editing. Any data entry / control which involves activity within the main graph will still operate upon the preset, but not on the AP3’s audio.

I/O Schematic Drawings & Load Sheets

Chapter Contents

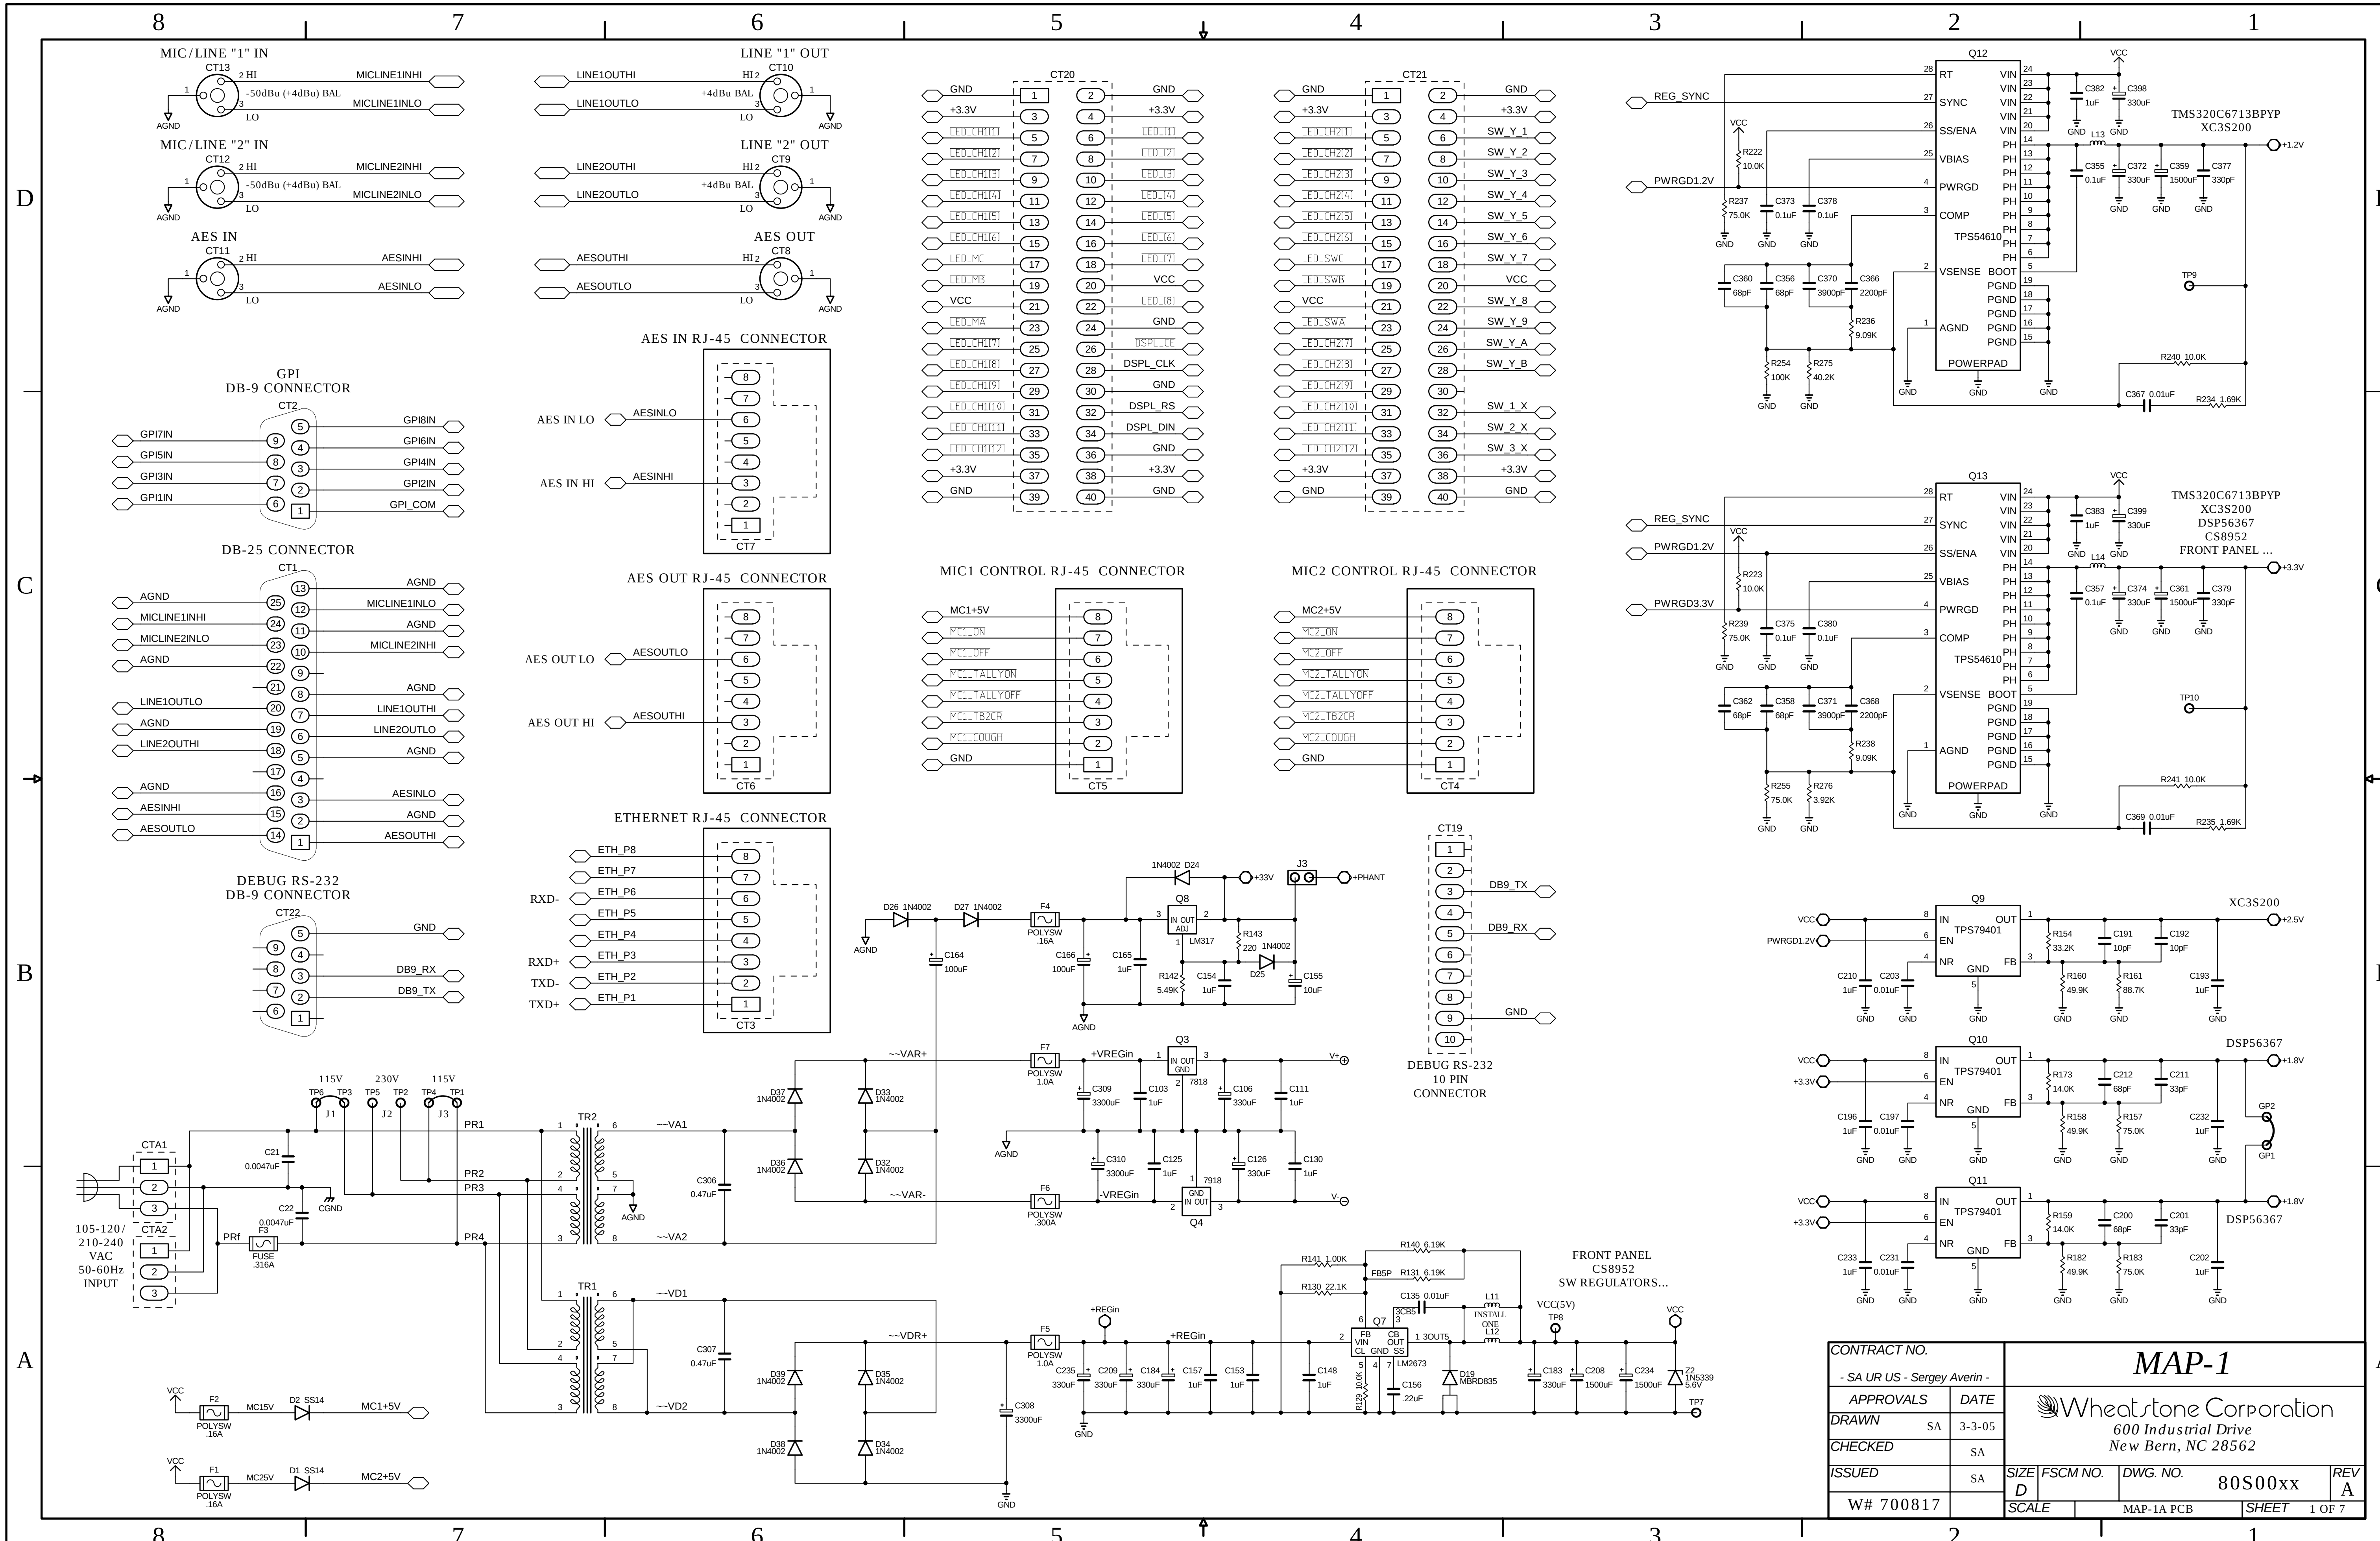

AP3 Mic Audio Processor Schematic

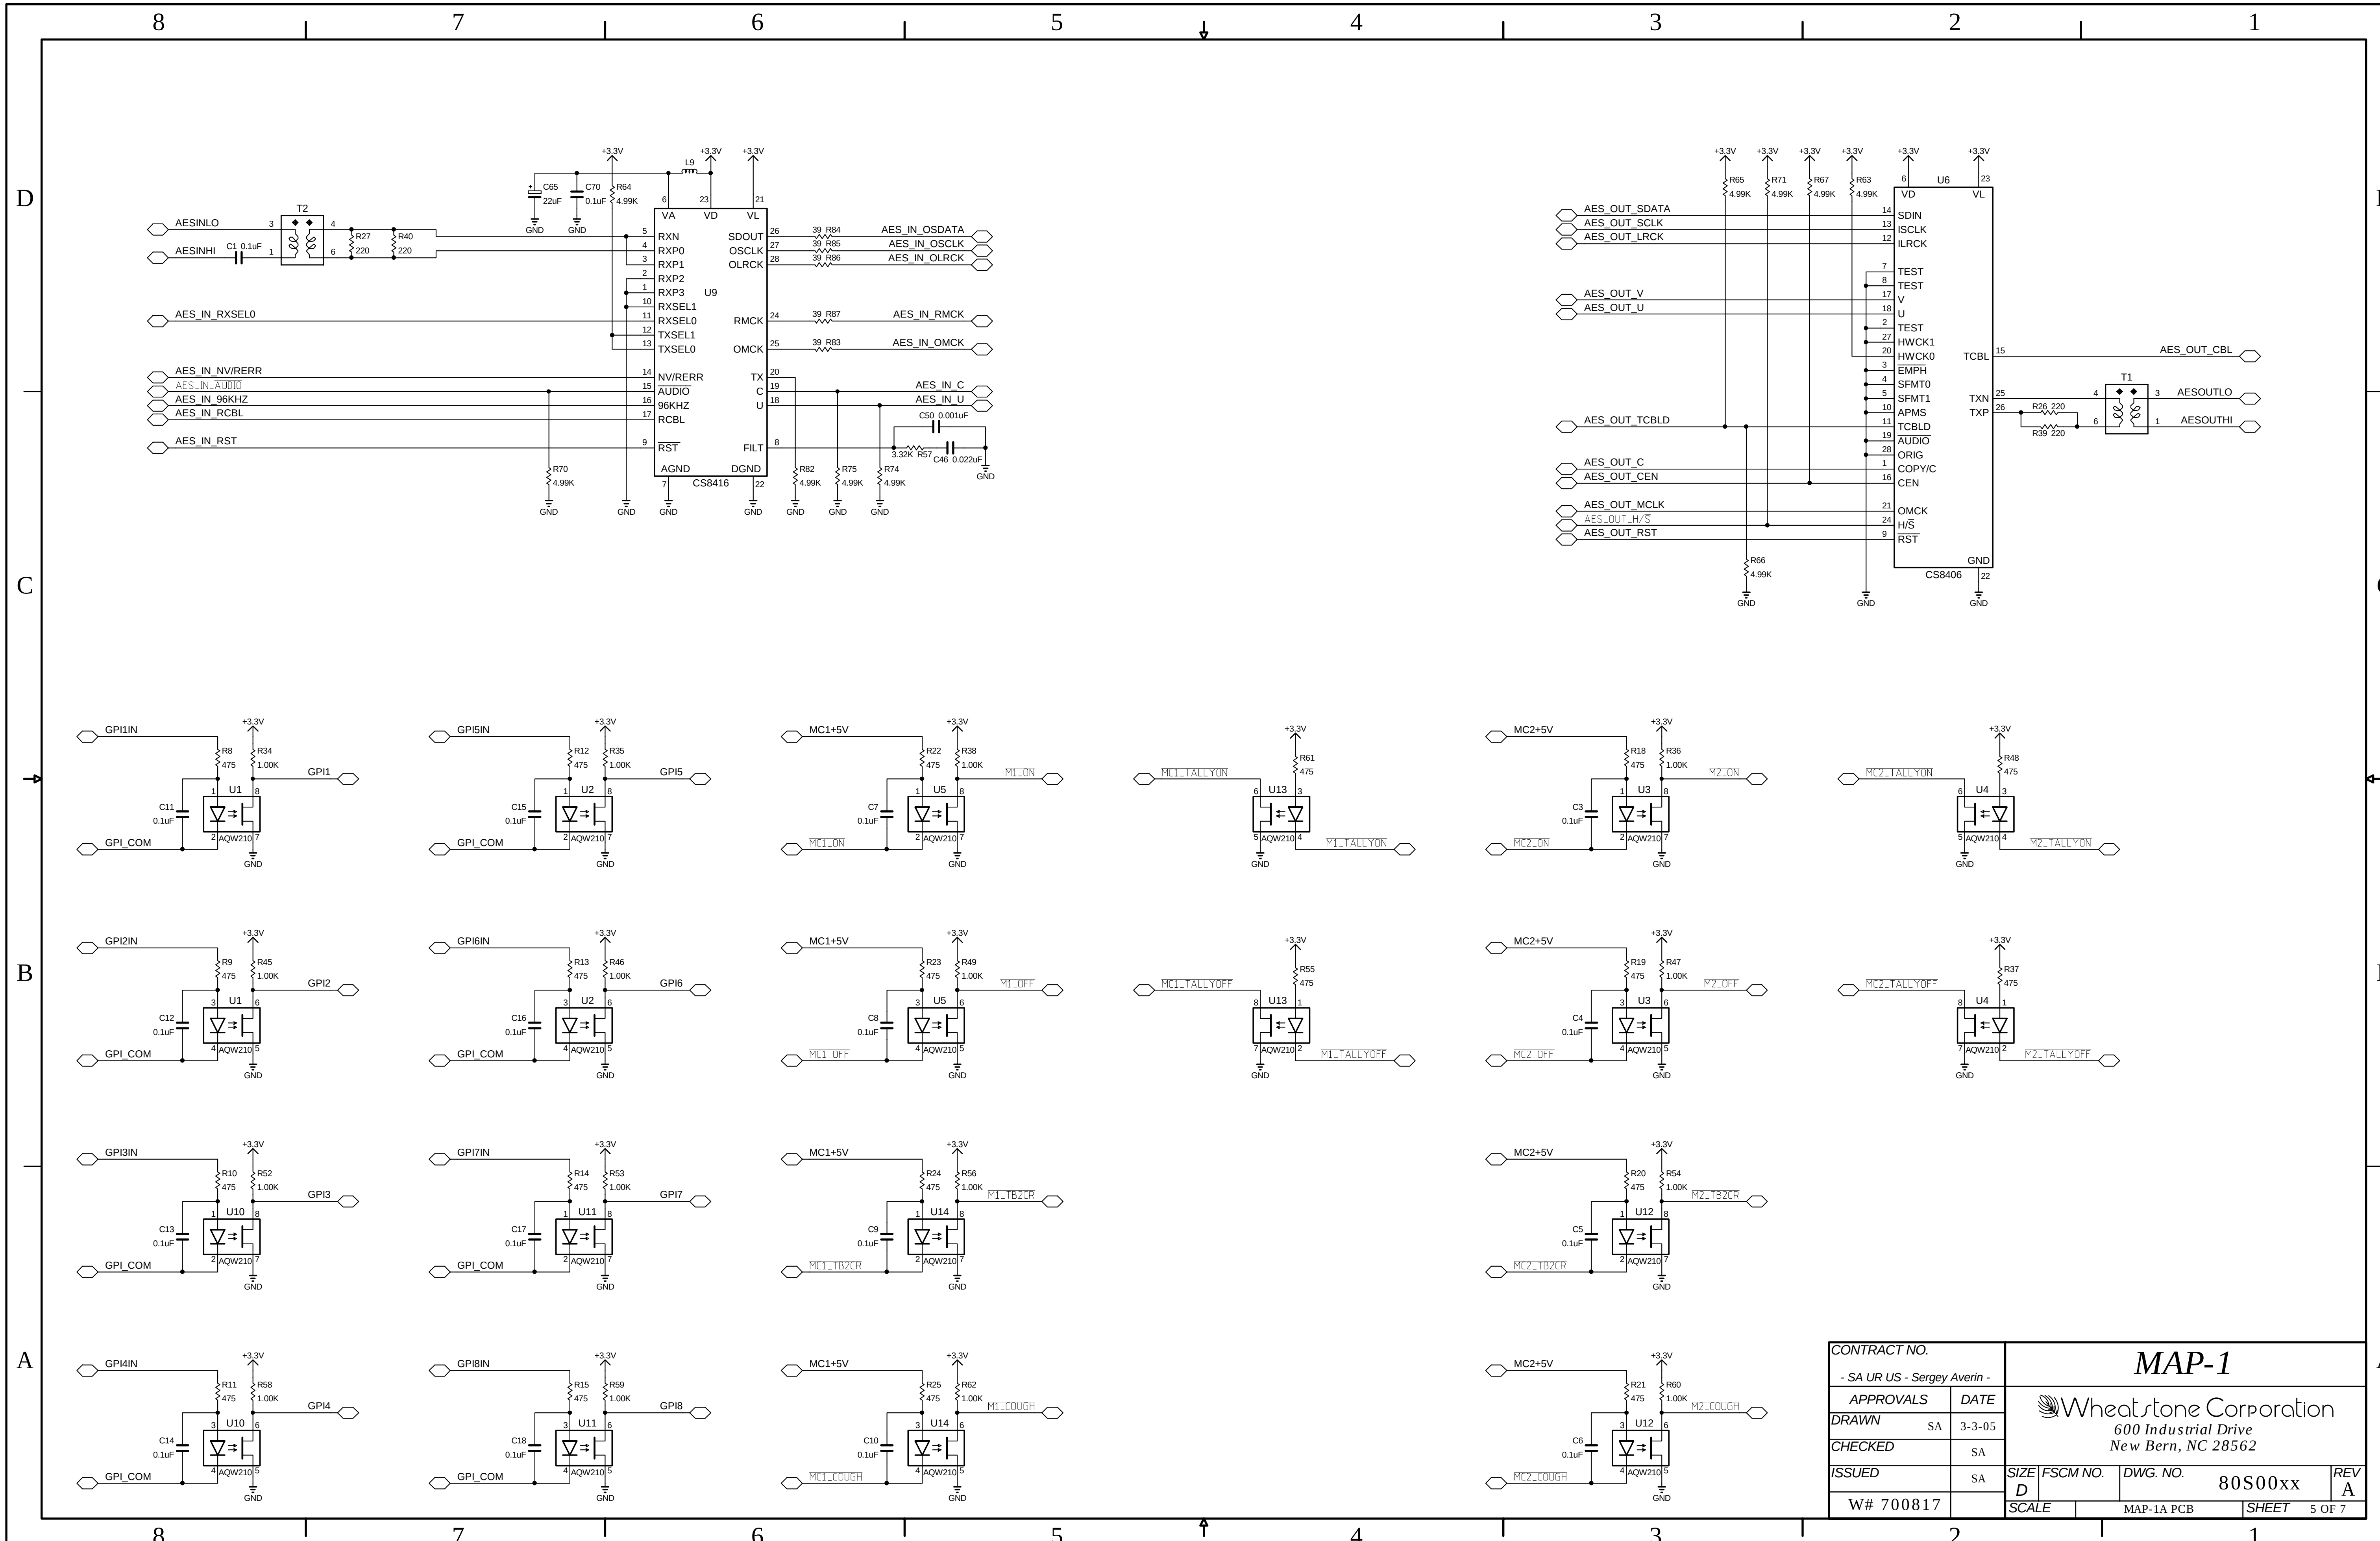

AP3 Mic Audio Processor Schematic

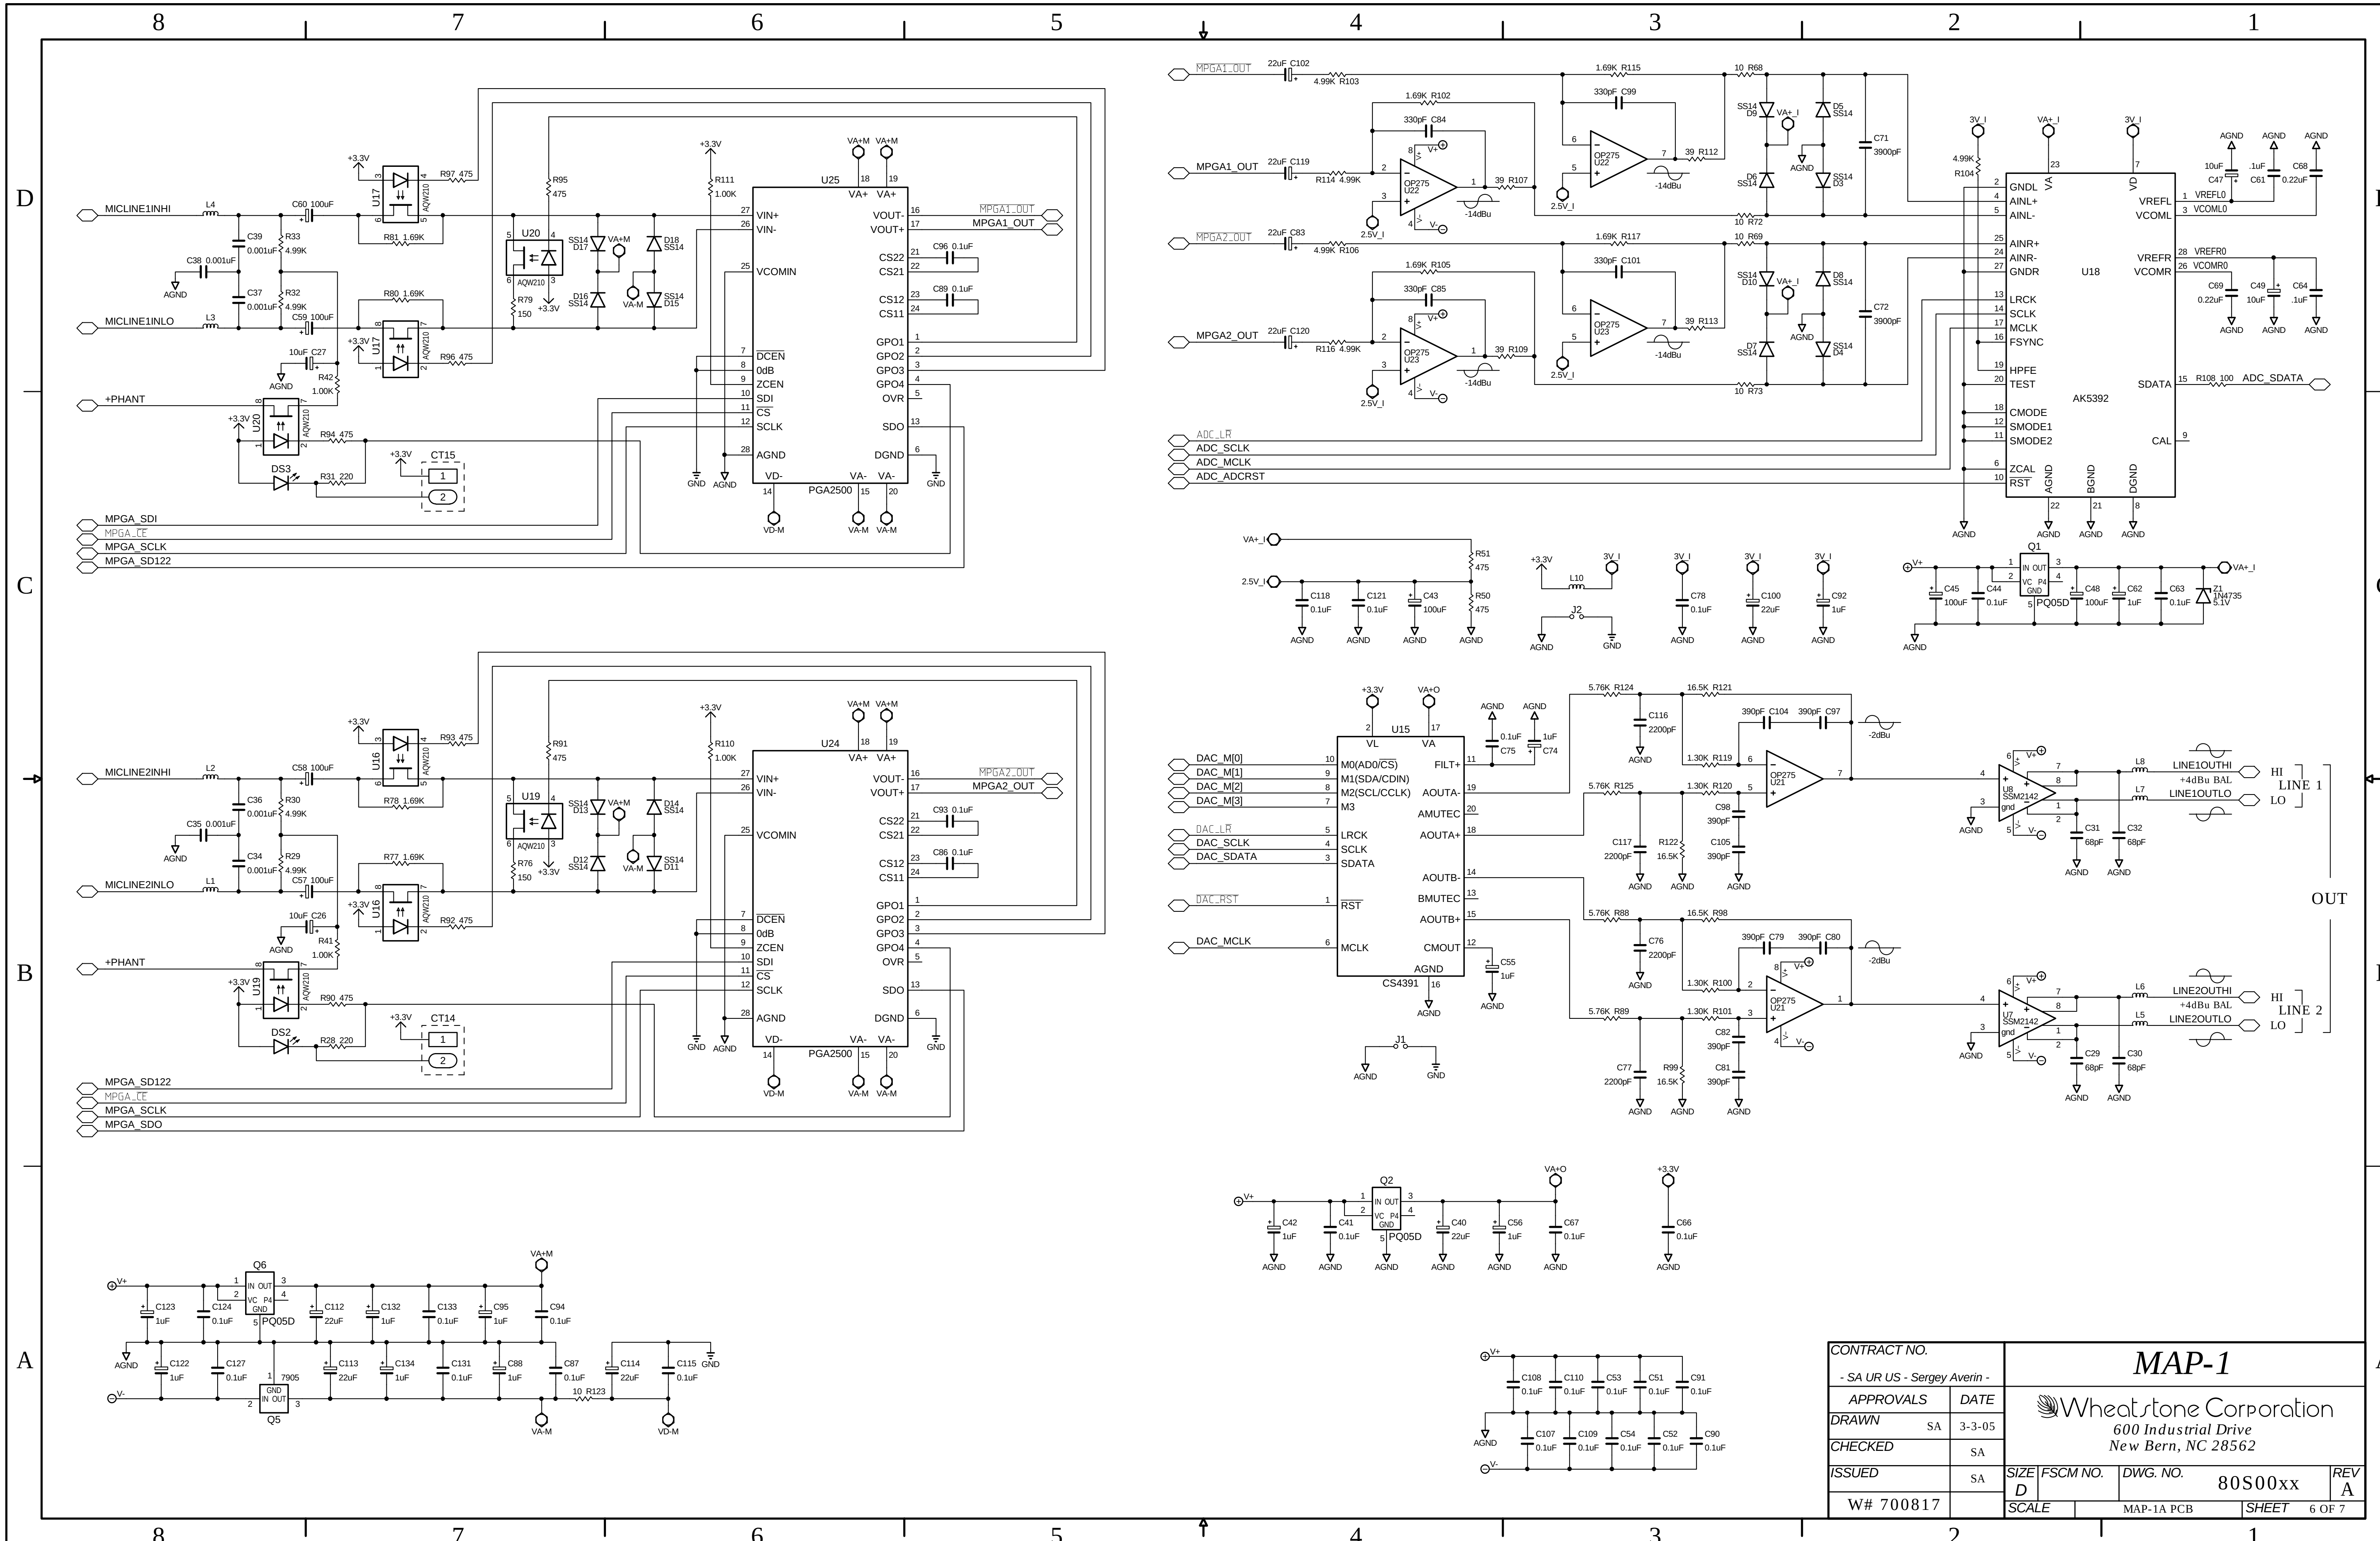

AP3 Mic Audio Processor Schematic

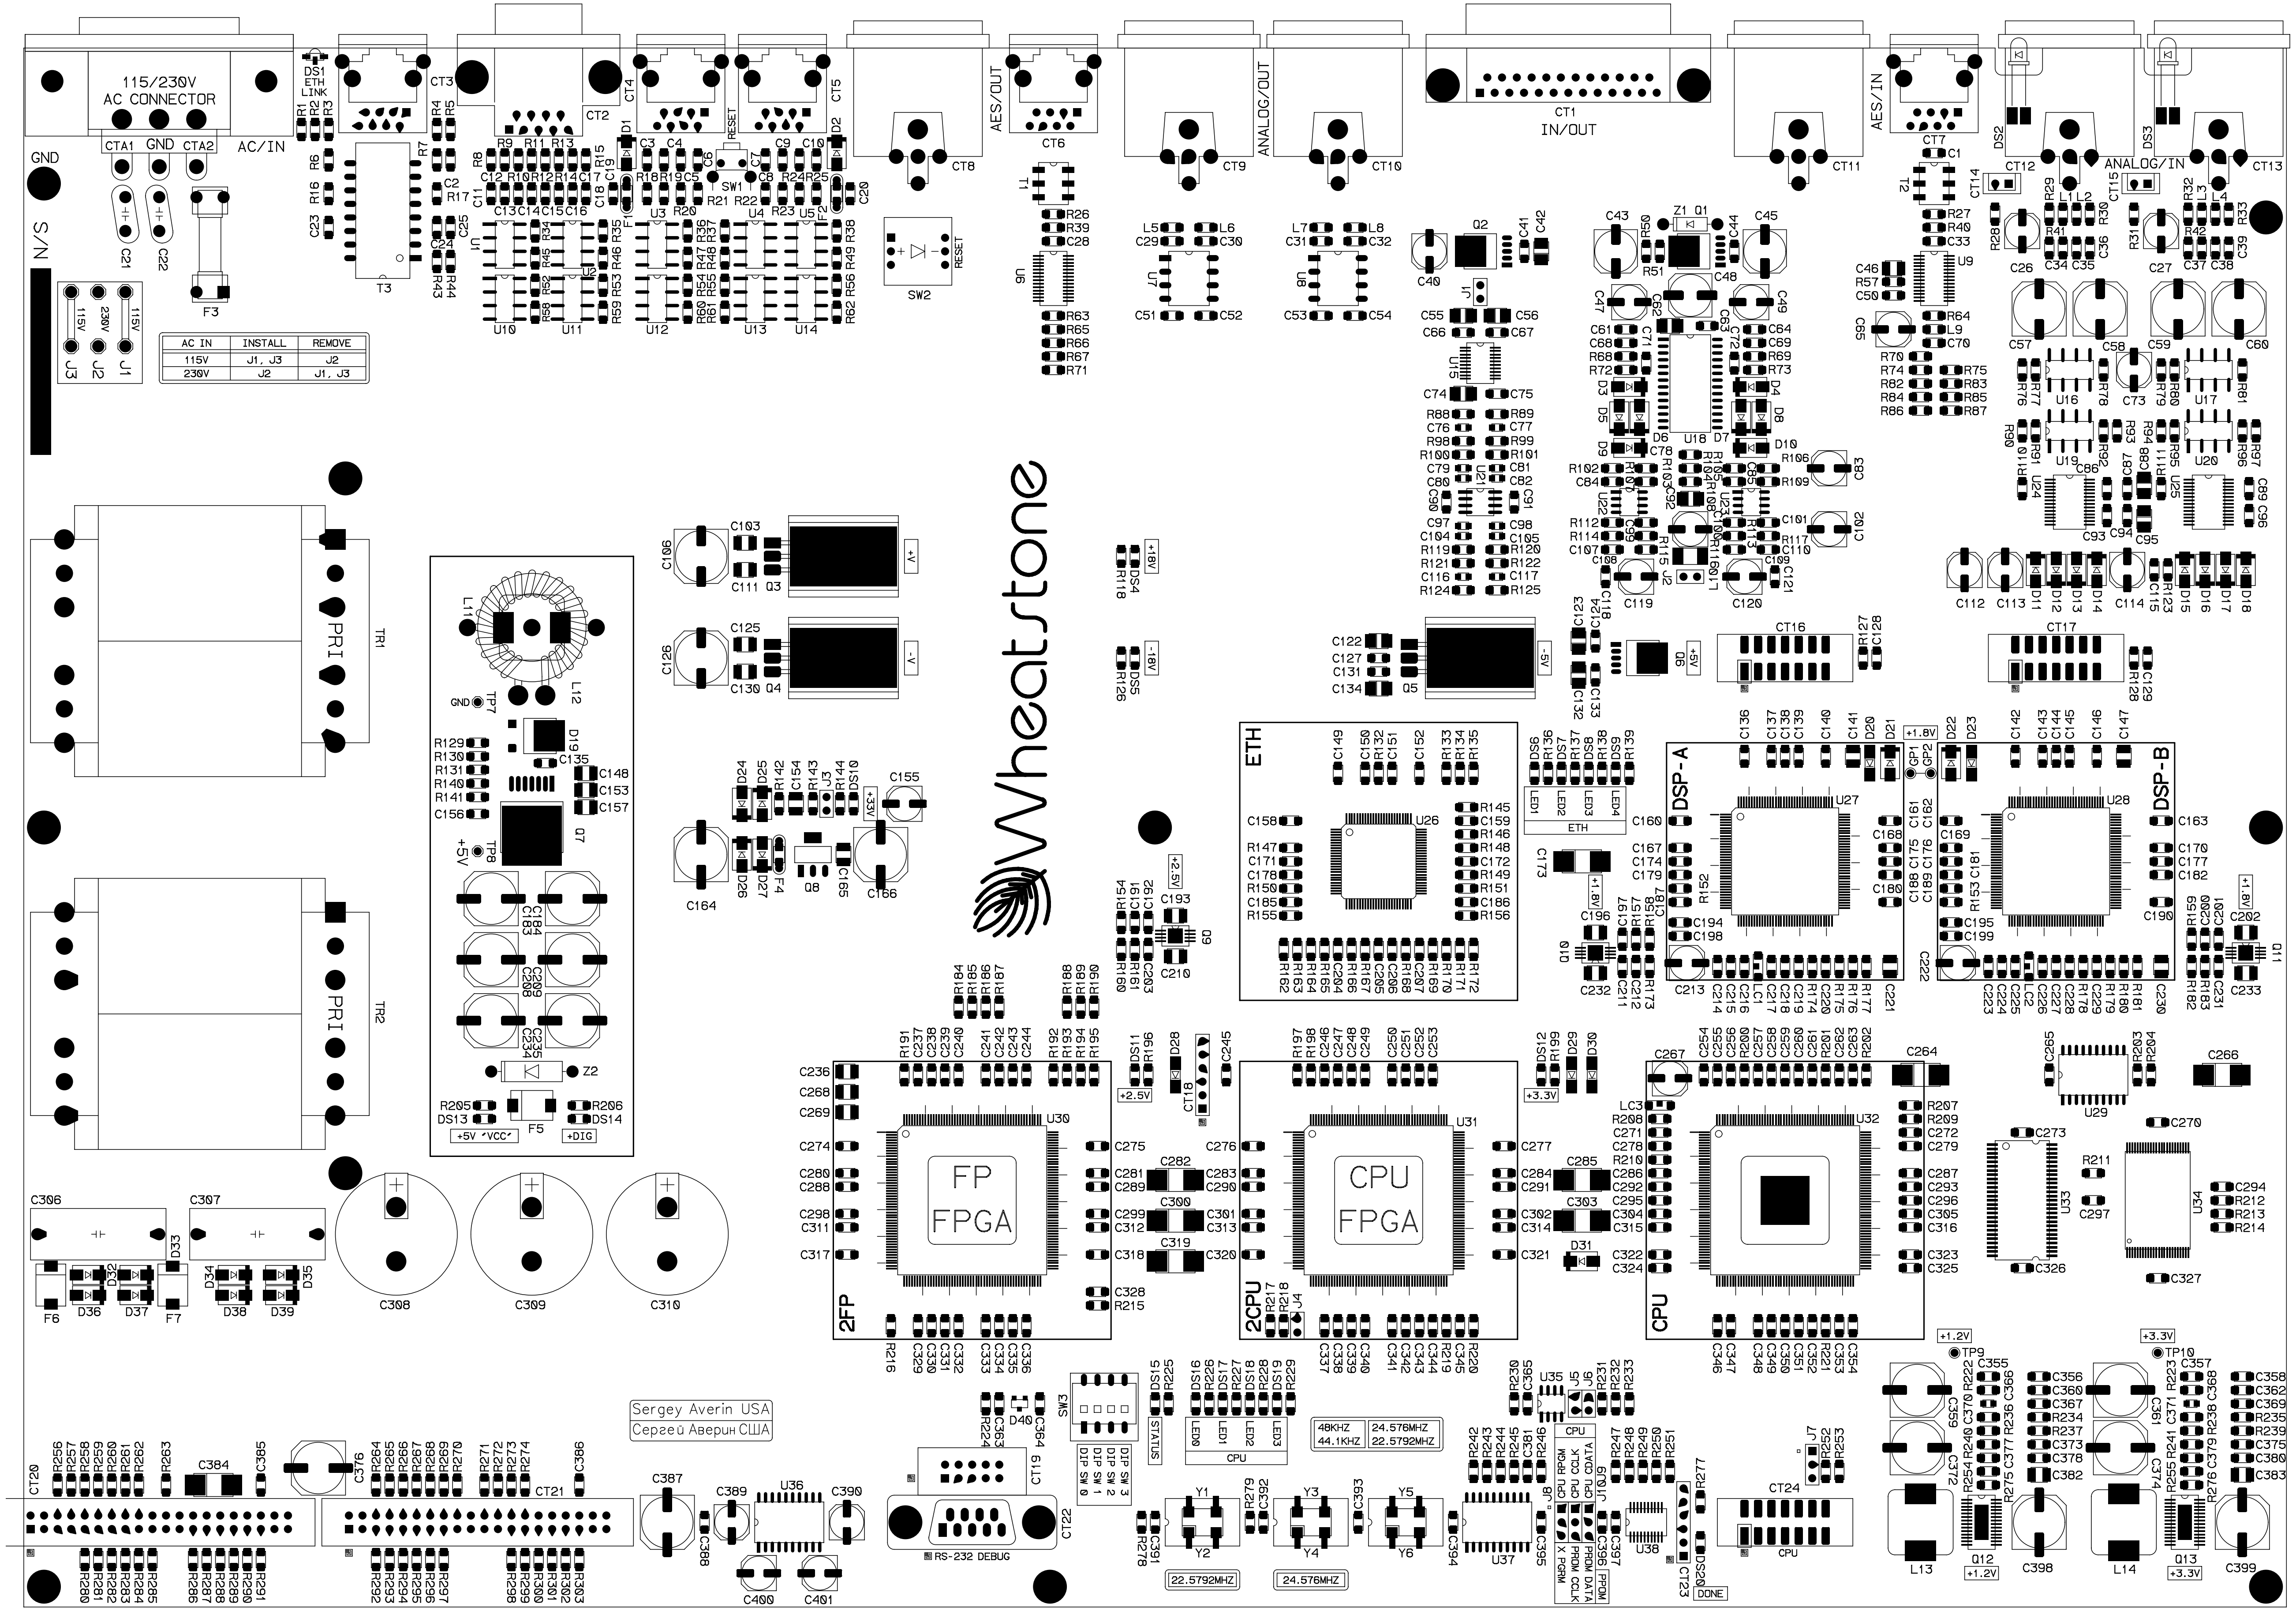

AP3 Mic Audio Processor - Load Sheet

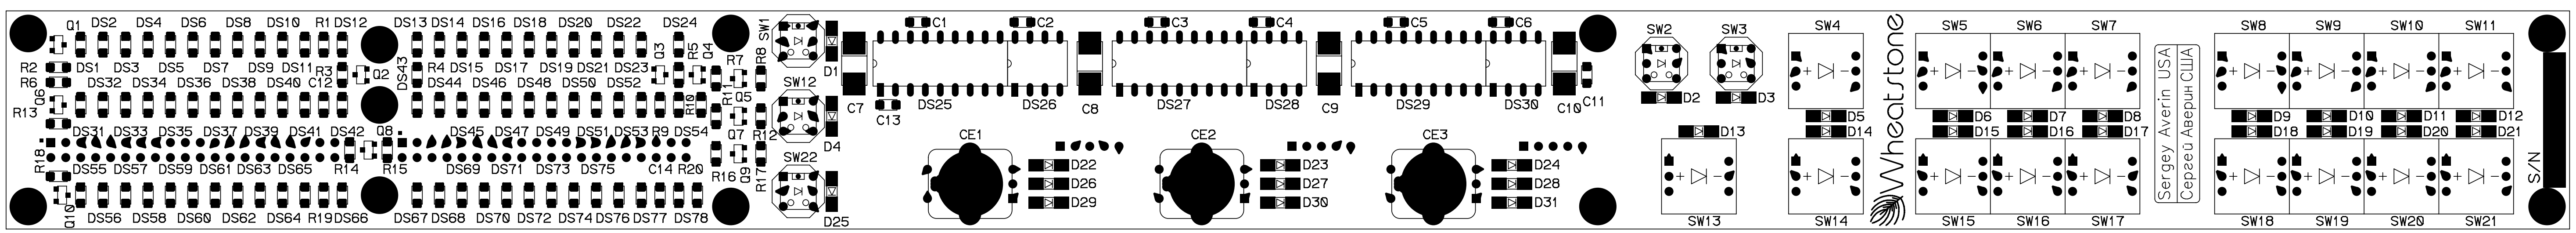

PSW-1 Processor Switch Card - Load Sheet

Appendices

Appendix 1

Appendix 2

Appendix 3

Appendix 1

Contents

Using the AP3 or "Walk softly, and carry a Vorsis AP3."

Introduction

The Vorsis AP3 is a highly complex piece of audio processing gear. Typical industry hubris is such that most users will not get even close to reading this until panic ensues, or perhaps ever.

Nonetheless, it is strongly recommended that some low-stakes time is spent playing with the AP3 and its accompanying GUI software - which really makes operation of the unit easier - so becoming familiar with it. Such preparation will ensure that not only will there be good results under fire, but will definitely pay off in the long run by making for far better crafted presets and better sounds.

Lofty words dispensed with; here's the:

Five-Minute Processor

Power up the AP3. Connect input source and destination.

Stereo, or two channels of Mono? Press the "Channel 1", "Channel 2", or both together held for two seconds for Stereo.

Press "IN" button, or ensure it is unlit (and the AP3 is in bypass)

"IN" Switch
Lit = Processing
Dark = Bypass

Press "INPUT" button.

Rotate the left-most encoder to select the input type (mic, line, or AES).

Press "GAIN" button.

Rotate the right-most encoder until "0.0 dB" shows on the right-most display.

Adjust the left-most encoder until the input source material is peaking about +10dB on the input meter - between 0 and +20 is fine, but avoid clipping.

If necessary, readjust "Out Gain" to suit the destination.

The AP3 should now be passing audio and be roughly in accord with the environment's level structure.

Press the "PRESETS" button.

Cycle through the presets using the left-most encoder until one named appropriately to the program material shows; press the "TAKE" button until the light goes out.

Press the "IN" switch, if not already lit (and the box processing).

If necessary, go back to "In Gain" and "Out Gain" to adjust for levels.

Listen.

If unsure, peruse the presets some more and pick another to try.

The AP3 is now on-the-air! And probably well enough until it is possible to spend more than five minutes on optimizing the processing.

About Listening

There is lots of listening to be done setting up a box as complex and flexible as the AP3. It is critical, particularly when processing is about to be introduced into a chain which previously may have had none, or comparatively little. Going overboard in blazing glory, or underwhelming through timidity, are common errors with new processing, and best avoided by examining the new sounds as analytically and completely as possible in the contexts of the application and expected source program material.

It is ***firmly recommended*** that a pair of ***Good*** headphones be acquired, if not already owned. ("Good" means did-not-come-free-with-anything; a pair of "Good" headphones costs hundreds of dollars in the same way that a pair of "Good" loudspeakers costs thousands. "Good" means they don't get lent out to anybody else. Ever. Especially not air staff.) A better investment cannot be made, and at minor cost in relation to the processing. It can be quite chastening how iffy most studio monitoring systems can sound in comparison to a pair of "Good" neutral headphones, and how much more detail one is capable of discerning. This is not to suggest all listening and setting up of the processor should be done with these headphones; in contrast, one should strive to listen to the results of presets under development on as many different monitoring systems and in as many different environments as possible, to ensure sonic transportability. Something may sound just grand in one place, but have - for example - the bass seem totally overblown in another. A reference pair of "Good" headphones will help avoid the easy excesses of a new processing system. Eventually, cross-experience will lead one to be able to reasonably anticipate a sound's transportability simply from the headphones, in a manner impossible from more highly colored monitoring.

Processing settings are not once-and-for-all entities; as long as one may have spent fine-tuning a setup for a given application, deliberate or accidental revisiting is inevitable. Say, one day listening to the station in the car, a sneaking need to tweak back the mid-band compressor a hair will surface. Or the realization that the expander on the main vocal may be set a tad high. Processor settings are in constant evolution to suit the subtly altering purposes and evolving sensibilities in their creator. Listen often and long, and be prepared to adjust settings to accommodate the results. To quote Cole Porter: Experiment!

Appendix 2

Contents

Parameters, Units and Ranges.

Approximately following the AP3's signal path, these are the values and ranges appropriate to each type of processing.

System Level

Headroom:	20dB
Nominal Operating Level:	-20dBFS

Input Gain

a. Microphone

Type:	Electronic Differential
Input Impedance:	> 2kohms at 1kHz
Optimum Source Impedance:	200 ohm
Phantom Power:	32V via 6.8kohm per leg
Maximum Gain:	70dB
Minimum Gain:	-10dB
Gain Increments:	1dB

b. Analogue Line Input

Type:	Electronic Differential
Input Impedance:	> 10kohm (bridging)
Optimum Source Impedance:	< 1kohm
Maximum Gain:	20dB (i.e. -20dBu input)
Minimum Gain:	-10dB (i.e. +10dBu input)

c. Digital Line Input

Type:	AES3/EBU
Maximum Gain:	+12dB
Minimum Gain:	-12dB
Default Gain:	0dB
Gain Increments:	0.5dB

Filters

a. High-Pass Filter

Filter class:	24dB/oct. Butterworth
Frequency:	20Hz – 1kHz (80Hz)
Defaults:	Off

b. Low-Pass Filter

Filter class:	24dB/oct. Butterworth
Frequency:	1kHz – 20kHz (4kHz)
Defaults:	Off

c. Notch

Frequency:	20Hz – 20kHz (60Hz)
Depth:	> 20dB
Defaults:	Off

Equalization

Four identical, independent bands of parametric equalization:

Frequency:	20Hz – 20kHz
Bandwidth:	0.2 – 3 octaves
Lift / Cut:	+ / - 14dB
Defaults:	All off
	1. 50 Hz, 1 oct., +6dB
	2. 800Hz, 0.5 oct., -4dB
	3. 2.8kHz, 0.66 oct. , +3dB
	4. 16kHz, 0.4 oct., +8dB

De-Esser

Recursive-style true (not broadband) De-Esser.

Threshold:	-10 to -60dBfs (-20dBfs)
Attack:	0.1mS – 100mS (5mS)
Release:	50mS – 500mS (100mS)
De-Ess Frequency:	20Hz – 20kHz (4kHz)
De-Ess Bandwidth:	0.15 – 4 octaves
Defaults:	Off

Expander ("Downward expander")

Threshold:	-60 – 0dBfs (-40dBfs)
Open:	0.1mS – 100mS (0.1mS)
Hysteresis:	3dB
Hang:	0 – 1S (100mS)
Close:	50mS – 3Sec. (300mS)
Ratio:	1:1 – 1:20 (1:2)
Depth:	0 – 40dB (14dB)
Defaults:	Off

AGC

A common set of parameters control the behaviour of the three independent band-split AGC systems.

Attack:	50mS – 500mS (200mS)
Release:	100mS – 3Sec. (1.5Sec.)
Defaults:	Off

Compressor - "Production"

A three band frequency split compressor/limiter, each band having individual parameters:

Threshold:	-30 to -70dBfs (-40dBfs)
Attack:	0.1mS – 300mS (1mS)
Release:	33mS – 1Sec. (150mS)
Ratio (soft-knee):	1:1 – 20:1
LF Crossover Frequency:	20Hz – 1kHz (250Hz)
HF Crossover Frequency:	1kHz – 20kHz (4kHz)
Drive (Overall "Threshold"):	0 – 100% (50%)
Trim (HF, MF, LF):	+/- 18dB (0dB)
Makeup Gain:	-20 to +48dB (0dB)

Multiband Limiter - "Airchain"

An alternative to the above three-band compressor, this is a three-band peak limiter, each band having individually set parameters:

Drive	+8 to -20dB (0dB)
Attack	0.2mS – 1S (0.5mS)
Release	33mS – 1S (50mS)
Backoff	0dB to -10dB (0dB)
Clip	Hard / Soft (Soft)
Trim	+/-18dB (0dB)
Freeze	0dB to -60dB (-40dB)
LF crossover frequency:	20Hz – 1kHz (250Hz)
HF crossover frequency:	1kHz – 20kHz (4kHz)

Output Limiter

A zero-overshoot-capable peak limiter.

Threshold:	-50 to -10dBfs (-10dBfs)
Asymmetric Threshold:	100 – 150% (100%)
Attack:	0 – 100mS (5mS)
Release:	10mS - 330mS (50mS)

Appendix 3

Contents

For the most part there are no user-replaceable parts in the Vorsis AP3. A complete list of available components is shown on the next page. Contact Vorsis technical support for further information.

Vorsis (600 Industrial Drive, New Bern, North Carolina, USA 28562) may be reached by phone at 252-638-7000, fax 252-637-1285, electronic mail "techsupport@vorsis.com".

REPLACEMENT PARTS —AP3 PROCESSOR		
COMPONENT	**DESCRIPTION**	***WS P/N***
MAP-1 LOADED CARD	PROCESSOR LOADED CARD ASSEMBLY	***"008413"***
PSW-1 LOADED CARD	SWITCH LOADED CARD ASSEMBLY	***"008411"***
CABLE	50 COND FLAT RIBBON CABLE	***"150007"***
POWER CORD	7 1/2' BLACK POWER CORD	***"150017"***
RIBBON PLUG	40 PIN RIBBON PLUG	***"250053"***
I/O CONNECTOR	DB-9 RIGHT ANGLE PC MOUNT CONNECTOR	***"220016"***
I/O CONNECTOR	RIGHT ANGLE 25 PIN PC MOUNT CONNECTOR .318 FEMALE	***"220120"***
I/O CONNECTOR	FEMALE XLR, PCB	***"260003"***
I/O CONNECTOR	MALE XLR,PCB	***"260004"***
I/O CONNECTOR	RIGHT ANGLE SHIELDED RJ-45 CONNECTOR	***"260049"***
POWER CORD CONNECTOR	POWER CORD CONNECTOR WITH PRINTED CIRCUIT TERMINALS	***"230071"***
NKK SWITCH	JB15 SWITCH W/BRIGHTER GREEN LED AND SILICON GASKET	***"510289"***
NKK SWITCH	JB15 SWITCH W/BRIGHTER YELLOW LED AND SILICON GASKET	***"510291"***
SWITCH	SINGLE POLE MOMENTARY SWITCH W/HOLES FOR LED, NO LED INSTALLED	***"510293"***
SWITCH CAP	WHITE SWITCH CAP	***"530004"***
LUMA BUTTON	GREEN LUMA-3 LIGHT PIPE BUTTON	***"530277"***
ENCODER	11MM ROTARY ENCODER WITH THREADED BUSHING	***"560002"***
ENCODER KNOB	11MM BLACK PUSH-ON KNOB	***"520105"***
ENCODER CAP	PLAIN BLUE CAP FOR 11MM COLLET KNOB	***"055200"***
DISPLAY	4 SEGMENT GREEN ALPHA NUMERIC DISPLAY	***"600016"***
MANUAL	OWNER'S MANUAL	***"008409"***