

antares 

AVOX 2

Antares Vocal Toolkit

Owner's Manual

v1.1 11 March 2008

©2008 Antares Audio Technologies.
All rights reserved. Certified Isinglass-free.™
All trademarks are property of their respective owners.
Antares Audio Technologies
231 Technology Circle, Scotts Valley, California 95066 USA
www.antarestech.com
Printed in USA Rev 1.0 PN P35035-0308-M01

The Obligatory Legal Mumbo-Jumbo

The Antares AVOX 2 software and this User's Manual are protected by copyright law. Making copies, adaptations, or derivative works without the prior written authorization of Antares Audio Technologies, is prohibited by law and constitutes a punishable violation of the law.

Antares Audio Technologies retains all ownership rights to the AVOX 2 software and its documentation. Use of AVOX 2 is limited by the following license agreement.

Please carefully read all the terms and conditions of this license agreement. At the time of installation of the AVOX 2 software you will be presented with a copy of the agreement and asked whether or not you agree to it. Continuing with the installation process beyond that point constitutes such agreement.

AVOX 2 License Agreement

Antares Audio Technologies grants you a non-transferable, non-exclusive license to use AVOX 2 under the terms and conditions stated in this agreement. Use of AVOX 2 indicates your agreement to the following terms and conditions.

License

You may:

1. Use AVOX 2 on only one computer at a time. Authorization requires an USB iLok Smart Key which must be purchased separately.

You may not:

1. Make copies of AVOX 2 or of the user manual in whole or in part except as expressly provided for in this agreement. Your right to copy AVOX 2 and the user manual is limited by copyright law. Making copies, verbal or media translations, adaptations, derivative works, or telecommunication data transmission of AVOX 2 without prior written authorization of Antares, is prohibited by law and constitutes a punishable violation of the law.
2. Make alteration or modifications to AVOX 2 (or any copy) or disassemble or de-compile AVOX 2 (or any copy), or attempt to discover the source code of AVOX 2.
3. Sub-license, lease, lend, rent, or grant other rights in all or any portion of AVOX 2 (or any copy) to others.

Term of the Agreement

This agreement is effective until terminated by you or Antares. You may terminate the agreement at any time by notifying Antares and destroying all copies of the manual, and erasing AVOX 2 from all machine-readable media, whether on-line or on archival copies.

In the event of breach of any of the terms of this agreement, you shall pay the attorney's fees of Antares that are reasonably necessary to enforce the agreement plus resulting damages.

Product Installation and Required Activation

There are technological measures in the AVOX 2 that are designed to prevent unlicensed or illegal use of the Software. You agree that Antares Audio Technologies may use these measures to protect Antares Audio Technologies against software piracy.

Limited Warranty And Disclaimer

AVOX 2 AND ACCOMPANYING MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Antares Audio Technologies does not warrant that the functions contained in the program will meet your requirements. The entire risk as to the use, quality, and performance of AVOX 2 is with you.

SOME JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. THIS WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS. YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION.

Limitation of Liability

IN NO EVENT WILL ANTARES BE LIABLE FOR ANY DAMAGES, INCLUDING LOSS OF DATA, LOST PROFITS OR OTHER SPECIAL, INCIDENTAL, CONSEQUENTIAL OR INDIRECT DAMAGES ARISING FROM THE USE OF AVOX 2 OR ACCOMPANYING MATERIALS. THIS LIMITATION WILL APPLY EVEN IF ANTARES OR ITS AUTHORIZED AGENT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. YOU ACKNOWLEDGE THAT THE LICENSE FEE REFLECTS THIS ALLOCATION OF RISK. SOME JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

Whew! Now that that's over, let's get on to the good stuff.

Contents

Chapter 1	Getting Started	3
Chapter 2	Introducing AVOX 2	5
Chapter 3	Harmony Engine Vocal Modeling Harmony Generator	9
Chapter 4	THROAT Physical Modeling Vocal Designer	11
Chapter 5	DUO Vocal Modeling Auto-Doubler	19
Chapter 6	CHOIR Vocal Multiplier	21
Chapter 7	PUNCH Vocal Impact Enhancer	23
Chapter 8	SYBIL Variable Frequency De-Esser	25
Chapter 9	ARTICULATOR Vocal Formant and Amplitude Modeler	27
Chapter 10	MUTATOR Extreme Voice Designer	37
Chapter 11	WARM Tube Saturation Generator	43
Chapter 12	ASPIRE Aspiration Noise Processor	47
	Index	51

Welcome!



On behalf of everyone at Antares Audio Technologies, we'd like to offer both our thanks and congratulations on your decision to purchase the AVOX 2 Antares Vocal Toolkit, a suite of uniquely powerful DSP tools for creating spectacular vocal tracks. As an AVOX 2 owner, you are entitled to receive notification of any software updates, technical support, and advance announcements of upcoming products. But we can't send you stuff unless we know who and where you are. So please, register.

At Antares, we are committed to excellence in quality, customer service, and technological innovation. With your purchase of AVOX 2, you have created a relationship with Antares which we hope will be long and gratifying. Let us know what you think. You can count on us to listen.

Again, thanks.

The Whole Antares Crew



Chapter 1: Getting Started



How To Use This Manual

The plug-ins that make up the AVOX 2 suite have been designed with clear, consistent, friendly user-interfaces. If you come to AVOX 2 with previous vocal processing experience, you will find that most of their functions are self-evident and are extraordinarily easy to use. However, since some of the AVOX 2 tools do things that have never been done before (or do them in new ways), a few aspects of their user-interfaces may not be immediately obvious.

Consequently, we encourage you to at least scan each of the individual plug-in's chapters for descriptions of their various parameters. In particular, Harmony Engine offers new (and extremely intuitive) ways to specify vocal harmony, while THROAT, ARTICULATOR and ASPIRE offer entirely new approaches to vocal processing, providing capabilities that have simply never before existed in a DSP plug-in. We strongly suggest that you read their chapters in their entirety to gain an understanding of what they can do.

If you just can't wait to get some sound going, feel free to load up some vocal tracks and play around with the various AVOX 2 plug-ins. (Be sure to check out Harmony Engine's and THROAT's collection of factory presets. They'll give you a good idea of what they're capable of as well as providing inspiration for your own creations.) Then come back and dig into the details.

The Contents of This Manual

Chapter 1: Getting Started

The chapter you are reading. Provides information on installing and authorizing AVOX 2.

Chapter 2: Introducing AVOX 2

This chapter provides an overview of the tools that make up the AVOX 2 plug-in suite.

Chapter 3: Harmony Engine Vocal Modeling Harmony Generator

Barely a chapter, here's where we explain why there's a separate manual for Harmony Engine.

Chapters 4–12

These chapters are reference information for all of the controls used in each of the other AVOX 2 plug-in interfaces.

Installing AVOX 2

AVOX 2 plug-ins are designed to function in a wide variety of digital audio applications. Please refer to your specific host application's user manual for more information on installing and using plug-ins.

Authorizing AVOX 2

Authorization is the process by which AVOX 2 is allowed to run on your computer. Detailed instructions covering the available authorization options will be found in the the file "Authorization Read Me" which is included on the installation CD ROM or with your software download.



NOTE: When initially installed, this software will run for ten days without authorization.

So even if you can't authorize it right away you can still use your software in the meantime. (During this period, click the "Continue"

button whenever you are presented with the Trial Period screen at launch.) But don't procrastinate too long. After those ten days are up you will no longer be able to launch AVOX 2 until it's authorized.

Technical Support

In the unlikely event that you experience a problem using AVOX 2, try the following:

1. Make sure you have the latest version of the software. You can download and install the latest version of the software from the following web page:

<http://www.antarestech.com/download/update.shtml>

2. If you are having problems authorizing your software, be sure that you have the latest version of the PACE Interlok drivers. You can download and install the latest version for your operating system from the following web page:

<http://portal.knowledgebase.net/article.asp?article=174703&p=5764>

If your problem is not resolved after taking the above actions, try the following:

1. Make another quick scan through this manual. Who knows? You may have stumbled onto some feature that you didn't notice the first time through.
2. Check our web page for tips, techniques, or any late-breaking information:

<http://www.antarestech.com>

3. Consult our searchable knowledgebase at:

<http://www.antarestech.com/support/index.html>

4. Call your local Antares dealer.

Chapter 2: Introducing AVOX 2



The heart of any great song is a great vocal sound. With the AVOX 2 Antares Vocal Toolkit, we've combined ten state-of-the-art vocal processing modules that open up an entirely new world of vocal processing capabilities. AVOX 2 gives you the power you need to create stunning vocal tracks in any musical style as well as design unique vocal effects for audio post-production applications.

In the following chapters we'll deal with each AVOX tool in turn, but first let's take a brief look at what's in AVOX 2 and how you might use the tools in combination.

AVOX 2 Overview

The AVOX 2 Antares Vocal Toolkit includes:

Harmony Engine Vocal Modeling Harmony Generator

Harmony Engine is a real-time harmony generating plug-in that puts professional-quality vocal harmony arrangements within reach of any songwriter, producer, musician or engineer.

With four independent harmony voices, a variety of powerful harmony generating modes, humanization features for natural sounding performances, and a flexible real-time preset system for harmony and vocal type, Harmony Engine provides incredibly easy-to-use tools to quickly and easily produce virtually any vocal arrangement you can imagine.

Whether you're an experienced vocal arranger, a songwriter looking for that perfect backup vocal, or a composer experimenting with unique vocal effects, Harmony Engine gives you entirely new ways to create the harmony parts you hear in your head. In fact, experimenting with different harmonies is so easy (and, dare we suggest, fun), you may find yourself using Harmony Engine to explore harmonic alternatives you may have never otherwise considered.

THROAT Physical Modeling Vocal Designer

THROAT is a revolutionary new vocal tool that lets you process a vocal through a meticulously crafted physical model of the human vocal tract. THROAT begins by neutralizing the effect of the original singer's vocal tract and then gives you the ability to specify the characteristics of the modeled vocal tract.

THROAT's controls allow you to modify the voice's glottal waveform as well as globally stretch, shorten, widen or constrict the modeled vocal tract. For even more detailed control, THROAT's graphical Throat Shaping display allows you to individually adjust the position and width of five points in the vocal tract model, from the vocal chords, through the throat, mouth and out to the lips. Finally, THROAT's Breathiness controls let you add variable frequency noise to the model, resulting in a range of vocal effects from subtle breathiness, to raspiness, to a full whisper.

While THROAT has been designed to allow subtle modifications to a voice's vocal quality, the range of the controls also allows the creation of vocal tract models well beyond the limits of physical human anatomy, offering the possibility of vocal characteristics that are simply unattainable by any other means.

DUO Vocal Modeling Auto-Doubler

DUO automatically generates a doubled vocal part from an existing vocal. In addition to allowing programmable variation in pitch, timing and vibrato depth, for even more realism, DUO also includes a simplified version of THROAT's vocal modeling to provide timbral variation for the doubled part. DUO's output section gives you independent control of the level and stereo position of the original and doubled voices.

CHOIR Vocal Multiplier

Neither a harmonizer nor a conventional chorus effect, CHOIR is a unique processor that actually turns a single voice into 4, 8, 16, or 32 distinct individual unison voices, each with its own pitch, timing and vibrato variations. The voices can, if desired, be spread across the stereo spectrum. Assign an instance of CHOIR to each of four voices singing four-part harmony (or voices from Harmony Engine) and voilà, instant choir.

PUNCH Vocal Impact Enhancer

PUNCH is a unique processor that, as its name implies, lets you give your vocal more dynamic impact, allowing it to cut through a dense mix with clarity and power.

SYBIL Variable Frequency De-Esser

SYBIL tames vocal sibilance with threshold, ratio, attack and decay controls as well as a variable sidechain high pass frequency to match any vocal performance.

ARTICULATOR

Vocal Formant and Amplitude Modeler

The alien offspring of a vocoder and a modern-day version of the venerable talk box, ARTICULATOR lets you extract the formant and amplitude information from a vocal (or other dynamic source) and apply it to any other audio track or to a built-in broadband noise generator.

MUTATOR Extreme Voice Designer

MUTATOR provides a combination of tools for creating unusual, weird, or downright wacky voices. With high-quality pitch shifting, throat modeling, pitch-tracking ring modulator-based "mutation," and a mind-bending Alienization function, MUTATOR is the perfect tool for unique special vocal effects and post-production sound design.

WARM Tube Saturation Generator

Based on Antares' legendary (at least among those of us who work here and, based on the volume of emails we've gotten since we discontinued it, a fair number of others as well) Tube plug-in, WARM, warms up your vocals with Antares' world-renowned tube modeling technology, while being so processing-efficient, you can put an instance on pretty much every track in your project and still barely see your CPU meter twitch. (And yes, that's probably the longest single sentence in this manual.)

ASPIRE Aspiration Noise Processor

ASPIRE is the world's first tool for modifying a voice's breathiness independently of its harmonic content. Whether reducing a bit of vocal rasp or adding a bit of smokiness, ASPIRE lets you modify the amount and quality of a voice's aspiration noise (the component of every voice that's the result of air passing over the vocal chords) without otherwise affecting the vocal's harmonic characteristics.

AVOX 2 Processing Guidelines

While it might seem to be convenient to consolidate all of the AVOX 2 functions in one huge integrated plug-in, some of the AVOX 2 functions use quite a bit of processing power. By organizing them as separate plug-ins, you can choose to instantiate only those you need for each particular vocal track, thereby using only as much of your computer's DSP power as you absolutely need.

While it's not likely that you'll be using all ten plug-ins on every track, it is likely that you'll often be using more than one. When that's the case, keep the following in mind:

General Guidelines

- In general, you should progress from the cleanest possible source vocal through successive stages of processing. Harmony Engine, THROAT, ARTICULATOR, MUTATOR, ASPIRE, DUO and CHOIR in particular are designed to process pitched monophonic voices and will function much more effectively with clean input signals.
- If a track requires de-essing, use SYBIL as the first of the AVOX 2 plug-ins.
- Next in line would be either of the "Voice Creation/Modification" plug-ins, i.e., THROAT, MUTATOR and/or ASPIRE.
- Once you've tweaked the voice exactly as you want it, you can double, multiply, or harmonize it with DUO, CHOIR or Harmony Engine.
- Finally, WARM and PUNCH work well with pretty much any input, so they would typically be the last of the AVOX 2 plug-ins in the chain.
- Effects like reverb and conventional chorus or flanging should typically be applied after processing by the desired AVOX 2 plug-ins.

Special Cases:

- While you will usually choose between DUO or CHOIR, you can use them together for a really huge vocal section. Assign DUO to your track and pan the original and doubled voices to opposite tracks. Then assign separate instances of CHOIR to each of the two tracks. Humongous!
- If you will be using Harmony Engine to generate harmonies from a single vocal line, start with SYBIL (if necessary) and then THROAT and then feed its output to Harmony Engine. If your host supports it, assign or export each of Harmony Engine's outputs to individual tracks and use DUO or CHOIR on each track to create a vocal group of the size of your choice. If desired, PUNCH can be used on THROAT's output or on the individual Harmony Engine outputs prior to DUO or CHOIR.
- If you have a vocal with problematic vibrato (either too much or too little), you can use DUO for vibrato modification. Assign DUO to the track and use only the doubled output. Set all parameters to their minimum effect and use the Vibrato function to adjust the performance's vibrato depth. See Chapter 5 for more details.
- For maximum control of the timbre of doubled voices (at the cost of increased CPU requirements), assign DUO to your vocal and pan the original and doubled voices to opposite tracks. Then assign an independent instance of THROAT to each part and create a unique timbre for each voice.

But all that being said, always feel free to ignore any of the above suggestions. The AVOX 2 tools offer entirely new realms of creative possibilities. There is no wrong way.

Now on to the details.

Chapter 3: Harmony Engine

Vocal Modeling Harmony Generator



If you've purchased a boxed copy of AVOX 2, you will have noticed that it includes a separate manual for Harmony Engine. And if you've downloaded AVOX 2, you'll find the Harmony Engine PDF manual included with the plug-in.

There are a couple of reasons for this. First, Harmony Engine is by far the deepest plug-in in the AVOX 2 bundle, so there is a lot to say about getting the most out of it. And, from a more practical perspective, since Harmony Engine is also sold as a stand-alone product, we've already done the work of formatting the manual and then printing a whole bunch of them, so it's just a lot more efficient to include one of them here.

Bottom line, read the Harmony Engine manual, try out the factory presets, and play with the included tutorial sessions.





THROAT's controls allow you to modify the voice's glottal waveform as well as globally stretch, shorten, widen or constrict the modeled vocal tract.

Chapter 4: THROAT

Physical Modeling Vocal Designer



THROAT is a revolutionary new vocal tool that lets you process a vocal through a meticulously crafted physical model of the human vocal tract. THROAT begins by neutralizing the effect of the original singer's vocal tract and then gives you the ability to specify the characteristics of the modeled vocal tract.

THROAT's controls allow you to modify the voice's glottal waveform as well as globally stretch, shorten, widen or constrict the modeled vocal tract. For even more detailed control, THROAT's graphical Throat Shaping display allows you to individually adjust the position and width of five points in the vocal tract model, from the vocal chords, through the throat, mouth and out to the lips. Finally, THROAT's Breathiness controls let you add variable frequency noise to the model, resulting in a range of vocal effects from subtle breathiness, to raspiness, to a full whisper.

While THROAT has been designed to allow subtle modifications to a voice's vocal quality, the range of the controls also allows the creation of vocal tract models well beyond the limits of physical human anatomy, offering the possibility of vocal characteristics that are simply unattainable by any other means.

In order to understand what THROAT is doing and how you can use it to process your vocal tracks, it helps to know how our throats actually work to create what we perceive as unique vocal qualities.

Vocal production starts with the vocal chords. Air from our lungs is forced through the vocal chords, causing them to vibrate. The contour of this vibration is the glottal waveform. The actual shape of the waveform is affected by each individual's anatomy as well as the pressure applied to the vocal chords. From there, the voice is propagated through the throat, the mouth and out through the lips. It is the shape of these structures, both their length and width, that create the resonant characteristics that combine with the glottal waveform to define a unique vocal identity.

With THROAT, for the first time, you have individual control over each of the elements that go into creating a distinct vocal character. Whether you are a producer or engineer looking to subtly enhance a

singer's performance, or a sound designer in pursuit of a totally new vocal effect, THROAT will give you creative capabilities that have simply never before existed.

Controls SET UP

In order for THROAT to do the best possible job of modeling, it needs to know some basic things about the source audio. The following three controls are used to characterize the vocal that you will be processing:

Vocal Range

Use this control to select the range of the track you will be processing. Choices include Soprano Voice, Alto/Tenor Voice, Bass/Baritone Voice and Instrument (a general setting for anything that isn't actually a vocal). Matching the appropriate algorithm to the input results in faster and more accurate pitch detection and more accurate modeling.

To select vocal range, click on the Vocal Range pop-up and then select the desired range from the pop-up list.



NOTE: *Choosing the wrong Vocal Range (or just forgetting to set it at all) can result in compromised performance. Pay attention.*

Source Glottal Waveform

The glottal waveform is the waveform produced by the vibration of the vocal chords. The range of an individual's possible waveforms is defined by their particular anatomy. Within that range, the waveform can change pretty dramatically depending on the performer's singing style. Imagine, for example, the progression from a breathy whisper, to a straight-ahead pop vocal, to hard rock or full-on operatic aria.

A lot of subtle factors influence the glottal waveform, but for ease of setting, we have characterized the choices as loudness. Typically, as loudness increases, so does the pressure applied to the vocal chords and with that change in pressure a corresponding change in glottal waveform.

When setting this control, select the loudness level (soft, medium, loud, intense) that most closely matches the performance you are processing. If you're not sure where your performance lies on this scale, don't worry. There is no need to obsess over subtle distinctions. If in doubt, just pick one. Once you have set up a model, you can always come back and try another setting and see if it makes a (positive) difference.

To select vocal type, click on the Vocal Type pop-up and then select the closest stylistic character of the vocal performance from the pop-up list.

Source Throat Precision

If you've already been poking around the THROAT interface, you might be asking yourself, "What the heck is throat precision?" And a good question it is.

As it happens, this control works a bit differently than the previous two. Here, you're not being asked to tell THROAT something about the input, but instead to tell THROAT how precise to be in its attempt to calculate the characteristics of the input based on the type of modeling you intend doing (the choices being subtle, medium, or extreme).

So why not always leave it at "subtle"? Another good question. The answer is that THROAT's most precise analysis results in extremely accurate characterization of the source throat, which works well for models that are within the general range of human anatomy (i.e. subtle modifications). However, for more extreme models, the "subtle" setting can in some cases result in artifacts — most often a sort of "whistling." In those cases, such artifacts can often be reduced or eliminated by choosing a different setting for this control.

As a result, the strategy for this control should be to always start with "subtle" (which is the default) and, if a particular model results in artifacts (and you don't actually like those artifacts), change the precision one step at a time (i.e. to medium and then extreme) until you get the desired effect.

To select source throat precision, click on the Precision pop-up and then select the appropriate model type from the pop-up list.

The next two controls allow you to add a variety of breathiness effects to your modeled voice:

BREATHINESS

Breathiness Mix

This control lets you select the amount of breathiness component mixed into your modeled voice. With a setting of 0, there will be no breathiness (apart from what might be in the original vocal). At a setting of 100, the model will be all breathiness, with none of the original vocal characteristic present at all. Intermediate settings will produce mixes of the original vocal and the breathiness component.

The sonic character of the breathiness will depend both on the articulation of the original vocal and, more dramatically, the setting of the Frequency control described below.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Breathiness Frequency

This control lets you set the high pass frequency of the breathiness component (i.e., the frequency above which the breathiness will be

present). This frequency determines the audio range and character of the breathiness effect.

At high settings, the effect will be rather whispery and ethereal (depending, of course, on the original vocal and the Mix amount). At lower settings, the effect is more of a raspiness (again, depending on the mix). Experimentation is the best way to become familiar with the possibilities.

Command (Mac)/Control (PC) click the control to reset it to its default value of 4000 Hz.

The remaining controls are used to define the model vocal tract:


THROAT MODELING

Model Throat Length

The Throat Length control allows you to globally lengthen or shorten the modeled throat. Values above 1.00 represent a lengthening of the throat while values below 1.00 represent a shortening of the throat.

The actual values represent the percentage change in the throat length. For example, a value of 1.20 represents a 20% increase in throat length, while a value of 0.80 represents a 20% decrease in throat length.

Changes made to this control are reflected on the Graphic Throat Display described below. If you have used that display to create a custom throat contour, this control will preserve the overall contour while scaling it by the selected amount.

 **NOTE:** While this control gives you the ability to radically change the throat length, keep in mind that the variation in the length of human vocal tracts is rarely more than about 25% in either direction. If you are looking for a “realistic” vocal characteristic, start with modest settings of this control. (As a visual reference, this range is indicated by color on the control scale.) More extreme settings can produce dramatic results, but probably not what anyone would call “realistic.”

Command (Mac)/Control (PC) click the control to reset it to its default value of 1.00.

Model Throat Width

The Throat Width control allows you to globally widen or constrict the modeled throat. Values above 1.00 represent a widening of the throat while values below 1.00 represent a narrowing of the throat.

The actual values represent the percentage change in the throat width. For example, a value of 1.20 represents a 20% increase in throat width, while a value of 0.80 represents a 20% decrease in throat width.

Changes made to this control are reflected on the Graphic Throat Display described below. If you have used that display to create a custom throat contour, this control will preserve the overall contour while scaling it by the selected amount.



NOTE: Similar to the Length control above, this control gives you the ability to radically change the throat width. Again, if you are looking for a “realistic” vocal characteristic, start with modest settings of this control. (And again, this range is indicated by color on the control scale.) More extreme settings can produce dramatic results, but probably not what anyone would call “realistic.”

Command (Mac)/Control (PC) click the control to reset it to its default value of 1.00.

Model Glottal Waveform

As was explained above in the Source Glottal Waveform section, the glottal waveform is the waveform produced by the vibration of the vocal chords. We used the Source Glottal Waveform control to help THROAT neutralize the effect of the original vocal’s glottal waveform. The Pulse Width and Voice Type controls let you define the glottal waveform you want to model.

Glottal Pulse Width

This control allows you to select the pulse width of the modeled glottal waveform. If you are at all familiar with analog synthesizers, you can think of this as being vaguely similar to the variable pulse width control on a square wave oscillator (and if you’re not familiar with analog synthesizers, don’t worry, just move the slider and listen to what happens).

You will find that the most dramatic timbral changes are usually found at the extreme ends of this control's range. If you're looking for realistic, stay in the middle 80% of the range.



IMPORTANT NOTE: *Keep in mind that this control interacts with the Glottal Voice Type control below. Each Voice Type sets a default Pulse Width that is associated with it. So if you set this control and then select a new Voice Type below, the value of this control will change to the new voice type's associated pulse width default. For that reason, it is usually best to select the Glottal Voice Type first and then make any desired adjustments to the pulse width.*

Command (Mac)/Control (PC) click the control to reset it to its default value of 64.

Glottal Voice Type

This control is the converse of the Source Glottal Waveform control described above. We used that control to tell THROAT the characteristic of the original performance. We use this control to tell THROAT what kind of characteristic we would like to model.

The choices here are the same as the choices for the Source Glottal Waveform control (i.e., soft, medium, loud, and intense). If you want to preserve the stylistic character of the original vocal, set this control to the same voice type as you set the Source Glottal Waveform (e.g., if you set Source Glottal Waveform to "soft," set this control to "soft").

As mentioned above, this control interacts with the Glottal Pulse Width control. Each of the Voice Types sets a default Pulse Width that is associated with it. Once a Voice Type is set, the pulse width may then be adjusted separately.



NOTE: *Despite the value names (soft, loud, etc.), this control does not necessarily change the actual level of the signal (although the "intense" setting does sometimes result in some level gain). Its purpose is to model the glottal waveform that would result from the various styles of singing. If you want to adjust the level, use the Output Gain control described below.*

To select vocal type, click on the Vocal Type pop-up and then select the desired stylistic character from the pop-up list.

GRAPHIC THROAT DISPLAY

Okay, here's where the fun really starts.

The Graphic Throat Display allows detailed tweaking of the model throat. For the purpose of this control, the model throat is divided into four sections defined by five boundary points. For reference, these sections and boundary points are illustrated on the head graphic to the left of the Throat Shaping display. Point (1) starts at the vocal chords and the sections progress through the throat and the mouth out to the lips, which are at point (5).

The Throat Shaping display itself consists of two elements: the original throat plot and the model throat plot. The original throat plot is colored blue and is not user-changeable. It represents the original length and width of the four throat sections and serves as a reference relative to which changes to the model throat are made.

The model throat plot is colored red and includes five control points that can be grabbed and moved, effectively adjusting the length and width of individual sections of the throat.

Additionally, the background of the display is divided into two sections by color. The lighter blue central area represents the range of adjustments that are consistent with typical human anatomy. As long as all control points and plot lines fall in this area, the results of your model will be more or less "realistic." The darker blue area represents the range of adjustments that exceed the dimensions of typical human vocal tracts. When any control point or plot line falls in this area, the result may or may not sound like something that you might recognize as a human voice. The more points and/or plot lines fall in this area, the more extreme the effect.

When you open a new instance of THROAT, the default state of the Throat Shaping display is with the original and model throat plots exactly superimposed, indicating no difference.

If you adjust the Throat Length and/or Throat Width sliders, you will see the entire red model plot move to reflect the changes. Length changes are self-evident. For width changes, any part of the model plot that lies above the original plot indicates a widening of the throat while any part of the model plot below the original plot indicates a narrowing of the throat.

Here are some things to consider when working with the Throat Shaping display:

- Point (1), which represents the vocal chords, serves as the anchor of the vocal tract. It can be adjusted for width but not length.
- When you move points (2), (3), or (4) horizontally, you are adjusting the relative lengths of the adjacent sections. For example, if you move point (3) to the right, you are lengthening the section between (2) and (3) while shortening the section between (3) and (4). The overall length of the throat model remains the same.
- Only by moving point (5) can you change the overall length of the model (apart from using the Throat Length slider, of course).
- As we've mentioned, if you're looking for "realistic" vocal characteristics, you would do well to start with relatively small adjustments that result in all control points and plot lines remaining in the central light blue area.
- In addition to the plot point positions, watch the contour of the plot lines connecting them. It's possible to place the points in positions relative to each other that cause the plot lines to bow out towards the edges of the display (or even pin against an edge of the display). This will almost always result in artifacts of one sort or another. (Of course, if it's artifacts you're looking for, they may be just what you want.)
- When you have created a custom model contour, the Model Throat Length and Width controls will adjust the overall throat length and width while retaining (and scaling) your custom contour.
- Extremely striking effects can be created by moving plot points in realtime. You can do this manually (for one point at a time) or,

much more powerfully, you can use your host's automation capabilities to program movements of all five points simultaneously. For the purposes of automation, each point is represented by two parameters, one for horizontal position (length) and one for vertical position (width).

In addition to the original and model plots, when THROAT is processing audio the display will also contain real-time representations of the original and modeled throat contours. As with the plots, the blue contour is the original throat and the red contour is the model throat.

As you begin to get familiar with the Throat Shaping display, trial-and-error will no doubt be the first order of the day. However, with a little experience, you will soon be able to predict what effect a particular plot adjustment will have. Checking out the factory presets, with a particular eye towards model plot shapes, should help you on your way.

Reset

Clicking the Reset button cancels any custom contour you have programmed, but retains any global Stretch and Width settings set by the Model Throat Length and Width controls. To reset those controls, Command (Mac)/Control (PC) click them to set them to their default values.

Output Gain

This control lets you adjust the output level of the modeled vocal over a range of +/- 24 dB. As you will discover, some models result in substantial level changes. This control is used to bring them back up or down to the desired level.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0 dB.

Level Matching

As mentioned above, some model settings result in substantial level changes. The Level Matching function attempts to compensate for level differences between the original and processed versions by automatically applying gain adjustments to the modeled version.



VERY IMPORTANT NOTE: This function is provided for ease of A/B comparisons with the original vocal. It inserts some processing into the modeled vocal path that can have a very small but nonetheless real effect on the overall audio quality. Once you have a model you like, turn the Level Matching function Off and use the Output Gain to adjust the proper level.

Click the Level Matching button to toggle its state. The button will display On or Off as appropriate.

Bypass

This control is used to (you guessed it) bypass the plug-in. It has been designed to provide artifact-free bypass switching so that you can use it to seamlessly enable THROAT only where desired on a track.

Click the Bypass button to toggle its state. The button will display On or Off as appropriate. Keep in mind that “On” means that the Bypass function is on, i.e., the plug-in is bypassed.

Usage Tips

- THROAT is a monophonic processor.
- For best modeling performance, THROAT needs to be able to detect the pitch of the original performance. To do that, THROAT requires a clean, pitched, monophonic signal. If THROAT can not reliably detect the pitch of the input, either because of a noisy or effected signal or because the input is not a single monophonic voice, it will not fail, but will fall back to an alternative (but not quite as accurate) modeling method.
- As has been mentioned several times already, for natural, “realistic” results, always start with relatively small adjustments to the various throat modeling controls, whether the Length and Width sliders, or the Graphic Display control points.

The actual range of variation in human anatomy is quite small relative to the overall range offered by THROAT. When you’re first getting started, try limiting control changes to around +/- 15%. As you get a feeling for the effect of various settings, slowly increase the ranges until things start sounding weird (not that that’s necessarily a bad thing).

- If you just want to start playing, try the following brief tutorial.


Start by assigning THROAT to a vocal part and set the above contour on the Throat Shaping display:

Be sure you have set the proper Vocal Range and Source Glottal settings and leave all the other controls at their default values.

Now, while listening to your vocal, slowly move the Model Throat Length slider from 1.00 to 1.25. Note the effect. Depending, of course, on the source vocal, you might typically hear the vocal quality become subtly darker, more male, older, etc.

Then try moving it back to 0.75. Again note the effect. In this case the result might be subtly more bright, female, young, etc.

Next, return the Model Throat Length slider to 1.00 and repeat the above with the Model Throat Width.



Finally, experiment with various combinations of the two controls. As you become familiar with their effects, try more extreme settings of the controls.

- Speaking of weird, it may be useful to understand just why extreme models are perceived as “unrealistic” or strange.

As mentioned above, the range of human vocal anatomy is actually relatively small. In order that we can differentiate between the voices of the many people we encounter in our lives, the auditory processing function in our brains is extremely sensitive to very small differences in vocal timbre within the typical range of variation. But as a result of this, we have no frame of reference for vocal timbres that fall outside of the range of common human anatomy. We can’t picture the person who would sound like that. So we think more in terms of “filtered,” “tubular,” or whatever. Trust us, if there were people with vocal tracts of the more extreme dimensions THROAT is capable of, this is what they’d sound like.

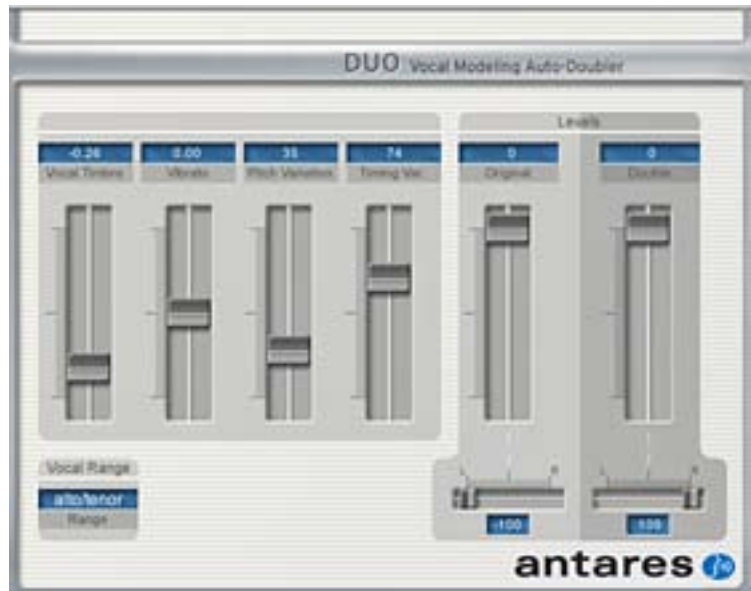
- While THROAT has been designed for voice, experiments here have shown that it can produce quite striking effects when used on other instrumental tracks. Try it with drums, guitars, bass, in fact pretty much anything. (In particular, applying it to a drum loop and then following that up with PUNCH has resulted in some truly twisted beats. Give it a try.)
- Play, play, play.

A few words about the THROAT factory presets

Unlike synth presets, which will always sound the same for all users, the THROAT presets are heavily dependent on the vocal tracks that you use them on. Consequently, they serve primarily as starting points for various general effects, with the expectation that you will tweak them to work best with your particular audio.

Specifically, they all contain the default settings for Vocal Range and Source Glottal Waveform controls. After calling up a preset, you should be sure set each of those controls to reflect your track.

Also, although some presets are named to give a general idea of their intent, feel free to experiment with any preset on any source track. There are no rules.



DUO makes use of Antares' unique vocal modeling and vibrato processing technologies to create a doubled part that actually sounds like another singer.

Chapter 5: DUO

Vocal Modeling Auto-Doubler



The DUO Vocal Modeling Auto-Doubler automatically generates a doubled vocal part from any existing monophonic vocal. Unlike conventional doublers that simply apply pitch and delay variation to the original part, DUO makes use of Antares' unique vocal modeling and vibrato processing technologies to create a doubled part that actually sounds like another singer. There is no easier or quicker way to create a realistic doubled vocal part.

Controls

Vocal Range

Use this control to select the range of the track you will be processing. Choices include Soprano Voice, Alto/Tenor Voice, Bass/Baritone Voice and Instrument (a general setting for anything that isn't actually a vocal). Matching the appropriate algorithm to the input results in faster and more accurate pitch detection.

To select the desired Vocal Range, click on the Vocal Range pop-up and then select the desired range from the pop-up list.



NOTE: *Choosing the wrong Vocal Range (or just forgetting to set it at all) can result in compromised performance. Pay attention.*

Vocal Timbre

Behind this unassuming slider lies a compact version of the THROAT vocal modeler. When this control is set to 0, the doubled voice will have exactly the same character as the original voice. As you move the slider up or down, DUO sends the doubled voice through a progressively more extreme vocal model. I.e., near 0, the timbre is very close to the original voice. As the control approaches .40 or -.40, the vocal timbre undergoes a fairly radical change.

As you might imagine, this slider simultaneously controls a number of vocal modeling parameters. We've linked them together to give you quick and easy access to a range of useful timbres.

Although the exact effect of any setting will depend a good deal on the vocal being processed, in general:

- Selecting a positive value will result in a lengthening of the vocal tract and hence an older/darker/more male vocal quality.
- Selecting a negative value will result in a shortening of the vocal tract and hence a younger/brighter/more female vocal quality.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Vibrato

The Vibrato control allows you to select the amount of variation in vibrato depth applied to the doubled voice.

Selecting a positive value increases the doubled voice's vibrato depth relative to the original voice.

Selecting a negative value decreases the doubled voice's vibrato depth relative to the original voice.



IMPORTANT NOTE: *Unlike vibrato controls on most vocal processors, this function does not simply use an LFO to apply a synthesized vibrato. DUO actually analyses the original vocal part to identify its vibrato and modify its depth. As a result, the modified vibrato will still reflect the overall contour of the original vibrato. As an additional consequence however, the Vibrato control will only have affect if the original vocal contains vibrato. It will not add vibrato to a performance that does not originally have it.*

The Vibrato control may be used by itself to modify the vibrato of a performance. See the Usage Tips below for details.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Pitch Variation

The Pitch Variation control lets you select the amount of random variation in pitch applied to the doubled voice. The higher the value, the larger the maximum amount of allowable variation.

Command (Mac)/Control (PC) click the control to reset it to its default value of 10.

Timing Variation

The Timing Variation control lets you select the amount of random variation in timing applied to the doubled voice. The higher the value, the larger the maximum amount of allowable variation.

DUO's timing variation is completely independent of pitch variation. To get a good idea of the effect of timing variation, use the level controls to listen only to the doubled voice. set Vibrato to 0 and Pitch Variation to 10. Set Timing Variation to maximum and process a rhythmic vocal performance. Pretty weird.

Command (Mac)/Control (PC) click the control to reset it to its default value of 10.

Original Level

Sets the level of the original voice.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Original Pan Position

Sets the original voice's location in the stereo spectrum. This control only functions in Stereo or Mono-> Stereo modes.

Command (Mac)/Control (PC) click the control to reset it to its default value of -100 (far left).

Double Level

Sets the level of the doubled voice.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Double Pan Position

Sets the doubled voice's location in the stereo spectrum. This control only functions in Stereo or Mono-> Stereo modes.

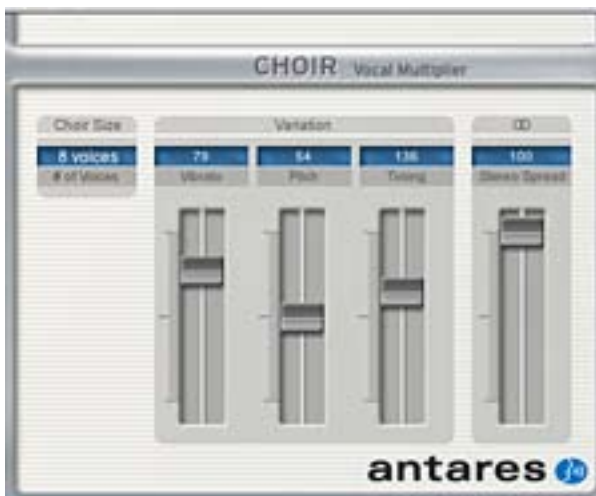
Command (Mac)/Control (PC) click the control to reset it to its default value of 100 (far right).

Usage Tips

- DUO is available in mono and stereo versions and, depending on the capabilities of your host program, a mono -> stereo version (which in most cases is the preferred routing). In the case of the stereo version, DUO processes only the left channel. DUO is dramatically more effective with the original and doubled parts panned apart, so if at all possible, try always to use it with stereo output.
- For best performance, DUO requires a clean, pitched monophonic signal. If DUO can not reliably detect the pitch of the input, either because of a noisy or effected signal or because the input is not a single monophonic voice, it will apply Vocal Timbre and Timing variations only.
- Experiment with various Pitch and Timing settings. Different combinations can result in distinctly different stylistic effects.
- DUO's Vibrato function can be used by itself to modify the vibrato depth on a recorded track. If you have a track with too deep a vibrato (or too shallow, although that's a much rarer problem), assign DUO to the track and set the mixer so that only the doubled track is heard. Set Vocal Timbre to 0 and Pitch and Timing to their minimum values. While listening to the track, adjust the Vibrato slider to reduce or expand the vibrato depth as desired. Notice that while you are modifying the vibrato depth, the expressive contour of the singer's original vibrato is retained.

Chapter 6: CHOIR

Vocal Multiplier



CHOIR is a unique processor that turns a single monophonic voice into 4, 8, 16, or 32 distinct individual unison voices, each with its own pitch, timing and vibrato variations.

The voices can, if desired, be spread across the stereo spectrum. When multiple instances of CHOIR are assigned to individual harmony parts, the result is an amazingly realistic large vocal ensemble.

Controls

Choir Size

This control allows you to select the number of individual voices that will be generated from the original voice. Choices are 4, 8, 16, and 32 voices.

Vibrato Variation

The Vibrato control allows you to select the range of variation in vibrato depth applied to the generated voices. Each voice is individually assigned a vibrato variation relative to the original voice. The higher the value of this control, the greater the possible vibrato variation.

Command (Mac)/Control (PC) click the control to reset it to its default value of 30.

Pitch Variation

The Pitch Variation control lets you select the range of the random variation in pitch applied to each generated voice. Each voice

is individually assigned a pitch variation relative to the original voice. The higher the value of this control, the larger the maximum amount of allowable variation.

Command (Mac)/Control (PC) click the control to reset it to its default value of 30.

Timing Variation

The Timing Variation control lets you select the range of the random variation in timing applied to each generated voice. Each voice is individually assigned a timing variation relative to the original voice. The higher the value of this control, the larger the maximum amount of allowable variation.

Command (Mac)/Control (PC) click the control to reset it to its default value of 100.

Stereo Spread

This control selects the extent to which the generated voices are spread across the stereo spectrum. At a setting of 0, all of the voices

appear in the center of the stereo soundstage. As the value is increased, the voices spread out from the center until, at the maximum value, they appear across the entire stereo soundstage.

This control only functions in Stereo or Mono->Stereo modes.

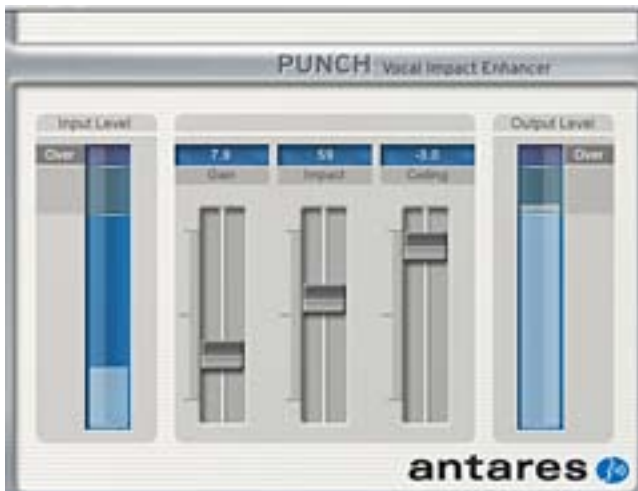
Command (Mac)/Control (PC) click the control to reset it to its default value of 100.

Usage Tips

- CHOIR is available in mono and stereo versions and, depending on the capabilities of your host program, a mono -> stereo version (which in most cases is the preferred routing). In the case of the stereo version, CHOIR processes only the left channel. CHOIR is dramatically more effective with the voices panned across the stereo spectrum, so if at all possible, try always to use it with stereo output.
- The real power of CHOIR comes not as a processor for one voice, but as a processor to assign to each of a number of harmony parts. Even if you have only one singer, have them overdub the basic harmony parts and then process those parts through THROAT to give each the character of different voices. Then assign an instance of CHOIR to each of those parts and create your vocal ensemble.
- Alternatively, create harmony parts by processing a single vocal part through a harmonizer and then assigning each individual harmonizer output to an instance of CHOIR.
- Keep in mind that the greater the number of voices, the greater the CPU usage. If you will be using multiple instances of CHOIR, it might be wise to limit each instance to 4 or 8 voices.
- Along those same lines, even if you have tons of CPU power, more is not always better. Match the choir size to the style of your music. Not every song needs the Mormon Tabernacle Choir.
- Experiment with different combinations of the Pitch and Timing controls. They can create dramatically different vocal ensembles. With Timing at its minimum and substantial Pitch Variation, you have a choir that's rhythmically tight but a bit loose with intonation. Conversely, reversing those settings gives you a group that's solidly in tune, but rhythmically loose. And there are of course many variations in between. Match the performance style to the style of your music.
- For best performance, CHOIR requires a clean, pitched monophonic signal. If CHOIR can not reliably detect the pitch of the input, either because of a noisy or effected signal or because the input is not a single monophonic voice, it will apply Timing variations only.

Chapter 7: PUNCH

Vocal Impact Enhancer




PUNCH is a processor that is designed to optimize the level of a vocal track to allow it to cut through a dense mix with power and clarity.

It provides a combination of compression, gain, limiting, and overload protection with a user interface designed for speed and simplicity.

Controls

Input Level

This meter displays the input level of the signal to be processed.

 **NOTE:** This display is for reference purposes only. None of the following controls affect the input level. The result of all processing will be reflected on the Output Level display described below.

Gain

This control allows you to increase the gain of your vocal. It will typically be set in combination with the Impact control below.

The effect of this control will be reflected on the Output Level display.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.0.

Impact

This control lets you select the amount of “punch” that is added to the vocal. As you increase the value of this control, level variations in the vocal performance are equalized and their level raised. The effect of the control will be reflected on the Output Level display.

Command (Mac)/Control (PC) click the control to reset it to its default value of 0.

Ceiling

This control allows you to attenuate the signal after all other processing.

Although PUNCH allows you to create a full-level signal without risk of ugly distortion, it will usually be wise to apply at least a little attenuation here to allow yourself some room for any further processing that might apply gain. The default setting of -3.0 dB is a good starting point, and you can always come back and add in more if necessary as you continue processing and the mix takes shape.

Command (Mac)/Control (PC) click the control to reset it to its default value of -3.0 dB.

Output Level

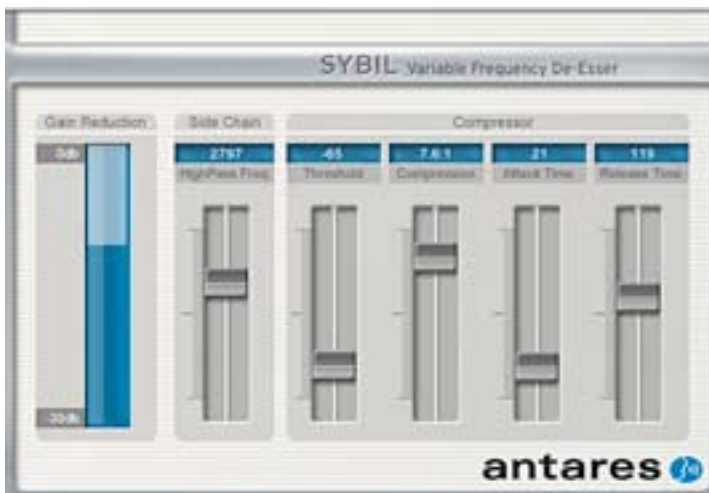
This meter displays the level of the signal after all processing by PUNCH. The Gain, Impact, and Ceiling controls should be set in combination such that no clipping occurs.

Usage Tips

- PUNCH is available in mono and stereo versions.
- Since the point of PUNCH is to optimize the impact of your vocal track in the mix, it makes sense to set the controls while listening to the track in the context of the entire mix.
- Every vocal performance is unique. Finding the optimum settings for PUNCH is very much a matter of experimentation. Luckily there are only two controls that really matter, so the prospect is not exactly daunting.
- The Impact control has been purposely designed with a wide enough range to produce some pretty odd effects at its extreme. If you're looking for that sort of thing, check it out.
- While PUNCH has been designed for the voice, it can perform its magic effectively on pretty much any recorded part. It will even do interesting things to entire recorded mixes. Check it out.

Chapter 8: SYBIL

Variable Frequency De-Esser




SYBIL is a digital equivalent of a traditional vocal de-esser.

When recording spoken or sung material, the sibilant sounds (Ss, Ts, CHs, and SHs) in the track sometimes (depending on the individual performer) appear louder than the rest of the signal. The effect can sound unnatural and often irritating. The solution to this problem is to compress only the sibilants, thereby lowering

their level relative to the rest of the track. Processing a signal this way is called de-essing.

The diagram below shows how a compressor and a high pass filter are traditionally configured to accomplish de-essing.



SYBIL uses a digital algorithm to implement the de-esser function. While the details of the algorithm are quite complex, the resulting effect is functionally equivalent to the diagram above.

Controls

Gain Reduction Meter

The Gain Reduction Meter displays the amount of compression taking place. If the frequency and threshold controls are set properly, the meter will display little gain reduction during vowel sounds and soft consonants, and substantial gain reduction during sibilants. For this reason, the meter is a useful tool when setting the other controls.

High Pass Frequency

This control sets the frequency of SYBIL's side-chain high pass filter. The goal is to set the frequency such that the filter passes any sibilance (thereby keying the compressor), but not any of the desired signal.

Command (Mac)/Control (PC) click the control to reset it to its default value of 8000 Hz.

Threshold

This control sets the threshold level of SYBIL's compressor. The range is from 0 dB to -80 dB. In most cases, the default value of -48 dB is a good place to start.

Command (Mac)/Control (PC) click the control to reset it to its default value of -48 dB.

Compression

This control sets the compression ratio of SYBIL's compressor. The range is from 1.0:1 to 99:1. In most cases, the default value of 2.4:1 is a good place to start.

Command (Mac)/Control (PC) click the control to reset it to its default value of 2.4:1.

Attack Time

This control adjusts the speed with which SYBIL's compressor responds to peaks in the signal coming from the high pass filter. The range of the control is from 3 milliseconds to 100 milliseconds.

Command (Mac)/Control (PC) click the control to reset it to its default value of 10 milliseconds.

Release Time

This control adjusts the time it takes the compressor's gain to increase 6 dB after the signal coming from the high pass filter drops below the threshold. The range of the control is from 3 milliseconds to 200 milliseconds.

Command (Mac)/Control (PC) click the control to reset it to its default value of 20 milliseconds.

Usage Tips

- SYBIL functions solely as a mono processor.
- If the high pass frequency is set too low, non-sibilant components of the signal will be compressed and the vocal will have its highs attenuated. If it is set too high, some sibilance will still remain. The trick is to find that ideal point where only the sibilance is affected. When set correctly, you can't tell that there is processing going on. If you can hear something happening, SYBIL needs to have its settings tweaked.

Chapter 9: ARTICULATOR

Vocal Formant and Amplitude Modeler



The alien offspring of a vocoder and a modern-day version of the venerable talk box, **ARTICULATOR** lets you extract the formant and amplitude information from a vocal (or other dynamic source) and apply it to any other audio track or to a built-in broadband noise generator.

Perfect for talking guitars, singing synths, whispering voices, and a wide range of special effects.

So what's going on here?

ARTICULATOR is unique among the AVOX 2 plug-ins in a number of basic ways:

First, it requires two different audio sources, what we call the Audio Signal and the Control Signal (more about those in a bit).

Second, while every other AVOX 2 plug-in is designed to process and affect a vocal track,

ARTICULATOR is designed to use a vocal track to control the processing of a second track, ideally a track with broadband harmonic content like a rich synth patch, a processed guitar, or the built-in noise generator (but typically not another vocal track).

Finally, because different hosts offer different (or no) capabilities for routing multiple tracks to a single plug-in, the version of ARTICULATOR for each plug-in format features a unique Input Routing subsection that is designed for the specific capabilities of that format.

Signal Selection

As mentioned above, ARTICULATOR requires two audio sources (although the built-in Noise Generator can serve as one of them). They are:

The Control Signal: This signal serves as the source of the formant and amplitude information that will be applied to the Audio Signal. A good Control Signal will have a lot of variation in both loudness and resonant harmonic content. A voice, whether singing, speaking, or just making weird rhythmic vocal noises makes an ideal Control Signal.

The Control Signal will always be mono.

The Audio Signal: This is the signal that the Control Signal's formant and amplitude information is applied to. A good Audio Signal will be rich in harmonic content and sustained in nature. Examples include synth pads and patches rich in overtones, distortion-rich guitars sounds, or even entire instrumental mixes. The built-in Noise Generator also makes a good Audio Signal, but its character will pretty much always be in the nature of whispering. (It can also be mixed with an external Audio Signal.)

The Audio Signal can be mono or stereo, depending on the capabilities of your particular host.

Signal Routing

Since few current plug-ins need to make use of two independent audio signals, some ingenuity is required to use ARTICULATOR in some hosts.

The easiest way to route the required signals is through the use of a sidechain input. In this scenario, you simply instantiate ARTICULATOR on the Audio Signal's track and route the Control Signal through the sidechain input.

Unfortunately, not all hosts currently support sidechain routing. If yours doesn't, you will have to use one of the alternative methods outlined in the Input Routing section below.


Controls

Input Routing Section

As mentioned above, the RTAS, VST and Audio Units versions of ARTICULATOR each include an Input Routing section unique to each format's capabilities. In the case of RTAS and VST, these sections are simply informational in nature, providing an indication of what's going where. In the case of Audio Units, the Input Routing section includes controls that let you select the routing that is supported by your particular host. Details below.



RTAS Since Pro Tools™ supports sidechain routing, setting up the RTAS version is quite straightforward. Simply select the Audio Signal by instantiating ARTICULATOR on the desired mono or stereo audio track and then assign the Control Signal via the sidechain input. The RTAS Input Routing section includes a red "LED" that lights to confirm that the sidechain input has been enabled in Pro Tools.



VST The version of VST that is most commonly in use at the time this is being written (and is used by ARTICULATOR) does not support sidechain input (but see the note below). Consequently, the methods for getting both signals into the VST version of ARTICULATOR vary depending on what kind of track it's instantiated on. The Input Routing section will change to reflect each of the following three possible options:

Instantiation on a mono track: In this configuration, instantiate ARTICULATOR on the track of the Control Signal. ARTICULATOR's internal Noise Generator



functions as the Audio Signal. (This is useful primarily for whispering and modulated noise effects.)

Instantiation on a stereo track: In this configuration, ARTICULATOR treats the right channel of the stereo track as the Control Signal and the left channel of the track as the Audio Signal.

Keep in mind that this will usually require that you create the stereo track by combining (via whatever method your DAW provides) the two independent tracks you wish to use as Control Signal and Audio Signal (being careful to ensure that they are hard panned to the appropriate channels).



In this configuration, the Audio Signal is limited to being mono (although it can be combined with the Noise Generator in Stereo Noise mode).

Instantiation on an LSR Group Channel:

Instantiation on an LRS Group Channel: In this configuration, ARTICULATOR treats the surround channel of the group as the Control Signal and the left and right stereo channels as the Audio Signal.

Again, this will require that you create the LRS Channel signals by combining the independent tracks you wish to use as Control Signal and Audio

Signal (again being careful to ensure that they are hard panned to the appropriate channels).

This is the most flexible of the three options, in that it allows you to process a stereo Audio Signal.



NOTE: A new version of VST that does include sidechain support was released just as AVOX 2 was being released. It's not yet clear how quickly this version will be supported in all hosts. There is the possibility (but no guarantee) that a future version of ARTICULATOR may add sidechain support for VST. If so, it would be noted in an accompanying Read Me file.



AUDIO UNITS

Some Audio Units hosts support sidechain input and some don't. Consequently, the Input Routing section for the AU version provides controls that allow you to set up ARTICULATOR to take advantage of whichever capabilities your host provides.

The Input Routing section will change to reflect each of the two possible options:

Instantiation on a mono track: In this configuration, you can select either the sidechain input or the audio track itself as the Control Signal. If your host supports sidechain input, you will typically always

select it as the Control Signal. If your host does not offer sidechain input, select Audio Track.

Depending on your choice, the resulting Audio Signal source will appear in the Audio Signal display as follows:

- If you select **Sidechain**, the Audio Source will be the audio track that ARTICULATOR is instantiated on.

- If you select **Audio Track**, ARTICULATOR's internal Noise Generator functions as the Audio Signal. (This is useful primarily for whispering and modulated noise effects.)



Instantiation on a stereo track: In this configuration, you can select either the sidechain input or the left channel of the audio track or the right channel of the audio track as the Control Signal. If your host supports sidechain input, you will typically always select it as the Control Signal. If your host does not offer sidechain input, select either the left or right channel of the audio track.

Depending on your choice, the resulting Audio Signal source will appear in the Audio Signal display as follows:

- If you select **Sidechain**, the Audio Source will be the stereo audio track that ARTICULATOR is instantiated on. This is the most flexible configuration in that it allows for a stereo Audio Signal.
- If you select **Audio Track L** as the Control Signal, Audio Track R will be selected as the Audio Signal.
- If you select **Audio Track R** as the Control Signal, Audio Track L will be selected as the Audio Signal.

In these latter two cases, you will typically be required to create the stereo track by combining (via whatever method your DAW provides) the two independent tracks you wish to use as Control Signal and Audio Signal (being careful to ensure that they are hard panned to the appropriate channels). In these configurations, the Audio Signal is limited to being mono (although it can be combined with the Noise Generator in Stereo Noise mode).

Input Mixer Section

The Input Mixer controls let you set the relative levels of the Audio Signal and the internal Noise Generator as well as optionally applying compression to the Audio Signal and selecting mono or stereo noise.



AUDIO LEVEL

This control sets the level of the audio signal (duh!). Its range provides for both attenuation and gain. The range is from -inf (i.e., off) to +12 dBs of gain.

To adjust Audio Signal Level, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the amount of attenuation or gain in dBs.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0 dB.



COMPRESSOR

This control lets you engage a compressor with fixed parameters designed to maximize the Audio Signal's level. It would typically be used in the case where the Audio Signal has quite a bit of dynamic variation that you would like to even out so that the Control Signal's dynamics can be applied to it.

Clicking the Comp button will toggle the compressor on and off. The button will turn blue to indicate that the compressor is on.



NOTE: Engaging the compressor usually results in an overall increase in gain. You may need to adjust the Audio Level control to compensate.




NOISE LEVEL

This control sets the level of the internal Noise Generator. The range is from -inf (i.e., off) to +12 dBs of gain.

To adjust Noise Generator Level, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the amount of attenuation or gain in dBs.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of -inf (i.e., off).



STEREO NOISE

ARTICULATOR's internal Noise Generator can provide either a single noise signal, or two uncorrelated noise signals panned to the

opposite extremes of a stereo output.


Clicking the Stereo Noise button will toggle it between mono and stereo noise. The button will turn blue to indicate the selection of stereo noise.



NOTE: *Stereo noise is only available from stereo or mono->stereo versions of ARTICULATOR. If you have instantiated a mono version, this button will not appear.*

Formant Extraction Section

The controls in this section allow you to optimize the formant extraction process for your specific Control Signal.



VOICE TYPE

Use this control to select the range of the track you will

be processing. Choices include soprano voice, alto/tenor voice, baritone/bass voice and instrument (a general setting for anything that isn't actually a vocal).

Matching the appropriate algorithm to the input results in more accurate formant extraction.

To select voice type, click on the Voice Type pop-up and then select the desired range from the pop-up list.



NOTE: *Choosing the wrong Voice Type (or just forgetting to set it at all) can result in compromised performance. Pay attention.*



TRACKING

In order to accurately analyze the Control Signal, ARTICULATOR requires a periodically repeating waveform,

characteristic of a voice or solo instrument. The Tracking control determines how much variation is allowed in the incoming waveform for ARTICULATOR to still consider it periodic.

If you are working with a well-isolated solo Control Signal (e.g., tracking in a studio or from a multitrack master) you can typically set the Tracking control to 25 and forget it.

If, on the other hand, your Control Signal is noisy or not well-isolated, or you are dealing with a particularly breathy or guttural voice, it may be necessary to allow more signal variation (higher Tracking numbers). However, under these circumstances tracking is not guaranteed and a too "relaxed" setting may negatively affect the accuracy of the formant extraction.

To adjust Tracking, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. Remember, higher numbers allow more variation, while lower numbers require less.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 25.

Formant Modulation Section

Another way of looking at the Control Signal's formant information is to think of it as a real time model of the vocal tract of the person singing or speaking. The controls in this section allow you to modify the extracted formant information by modifying the geometry of that modeled throat.



THROAT LENGTH

This control lets you increase or decrease the length of the modeled throat that is applied to the Audio Signal.

Values above 1.00 represent a lengthening of the throat while values below 1.00 represent a shortening of the throat. The actual values represent the percentage change in the throat length. For example, a value of 1.20 represents a 20% increase in throat length, while a value of 0.80 represents a 20% decrease in throat length.

The effect of this control is so dependent on the combination of each specific Control Signal and Audio Signal that the only guideline for working with it is to just check out its entire range and discover what works best for your particular needs.

To adjust Throat Length, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the percentage of change.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1.00.



THROAT WIDTH

As you might imagine, this control is similar to the Throat Length control described above, but instead modifies the width of the modeled throat.

Values above 1.00 represent a widening of the throat while values below 1.00 represent a narrowing of the throat. The actual values represent the percentage change in the throat width. For example, a value of 1.20 represents a 20% increase in throat width, while a value of 0.80 represents a 20% decrease in throat width.

Again, your best bet in working with Throat Width is just to try a variety of settings to hear what works best for your project.

To adjust Throat Width, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the percentage of change.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1.00.

Amplitude Extraction Section

In the schematic signal flow, this section represents the process of extracting the Control Signal's amplitude envelope (which is a fancy way of saying "the way its loudness changes over time").

ARTICULATOR does this all by itself with no need for any input from you.

Amplitude Modulation Section

Here's where you decide how much effect the Control Signal's amplitude envelope will have on the Audio Signal.




ENVELOPE AMOUNT

This control sets the amount of amplitude modulation (i.e., change in loudness) applied to the Audio Signal by the Control Signal's amplitude envelope.

When this control is set at 0, the Control Signal's amplitude envelope has no effect on the Audio Signal. In this case, the loudness dynamics of the Audio Signal are preserved and only the formant information from the Control Signal is imposed on the Audio Signal.

When this control is set at 100 (maximum), the full effect of the Control Signal's amplitude envelope is applied to the Audio Signal.

And as you might imagine, intermediate settings of this control apply intermediate amounts of amplitude modulation.

 **NOTE:** In order for the Control Signal's amplitude envelope to have maximum effect on the Audio Signal, the Audio Signal should be a sustained sound (like the Noise Generator or a synth pad) or at least one in which audio is constantly present.

If there are periods of silence in the Audio Signal, its output during those periods will be silence, regardless of the Control Signal's envelope. If the Audio Signal is constant, but has a varying dynamic range, you might, depending on your goal, consider engaging the Compressor in the Input Mixer to even it out.



TIP: The easiest way to understand the effect of this control is to do the following:

1. Do whatever your DAW requires to set up ARTICULATOR with a speaking or singing voice as the Control Signal.
2. In the Input Mixer, set Audio Level to $-\infty$ and Noise Level to a comfortable level.
3. Set the Formant Extraction controls as appropriate for your Control Signal.
4. In the Output Mixer, set Control Signal Level to $-\infty$ and Audio Signal level to a comfortable level.
5. While playing the Control Track, vary the Envelope Amount control from 0 to 100 and listen to the result.

To adjust the Envelope Amount, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 100 (maximum amount).

EQ Section

The EQ Section provides one band of traditional parametric EQ for fine-tuning the Audio Signal (as well as the Noise Generator output, if used).



FREQUENCY

This control selects the frequency of the EQ (i.e., the frequency that will either be increased or decreased via the Gain control). The range is from 100Hz to 4000Hz.

To adjust Frequency, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1500Hz.

**Q**

The Q control allows you to set the slope (and, consequently, the bandwidth) of the EQ's peak or dip. Its range is 0.10 to 3.10. High Qs result steep slopes (and narrow bandwidth), while lower Qs result in more gentle slopes (and wider bandwidth). (The

easiest way to understand this is to watch the green EQ plot while adjusting the Q and Gain controls. All will become clear.)

To adjust Q, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0.71.

**GAIN**

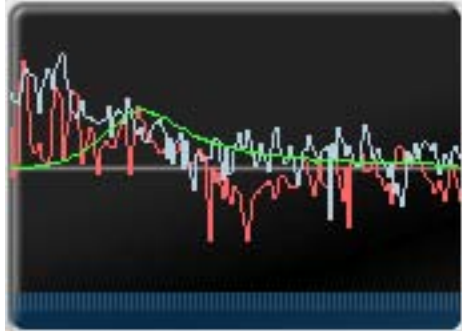
The Gain control lets you set the amount of boost or cut that will be applied to the Audio Signal at the selected Frequency. The range is from -12.0dB to +12.0dB.

To adjust Gain, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0.0.

**BYPASS**

The Bypass button allows you to bypass the effect of the EQ section. When the button is red, EQ is bypassed. Click the button to toggle its state.

**DISPLAY**

The EQ Section Display provides a real-time graphic representation of the Control Signal (red plot), the Audio Signal (blue plot) and the EQ curve (green plot).

Output Mixer Section

Here's where you can set the output levels and pans of both the processed Audio Signal and the Control Signal.

**AUDIO SIGNAL LEVEL**

This control applies gain or attenuation to the processed Audio Signal plus Noise (if any). The range is from -inf (i.e., off) to 12 dBs of gain.

To adjust Audio Signal Level, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the amount of attenuation or gain in dBs.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0 dB.

**AUDIO SIGNAL PAN**

This slider allows you to set the Audio Signal's location in the stereo spectrum.



NOTE: *Audio Signal Pan only functions in Stereo or Mono->Stereo versions of ARTICULATOR. If you have instantiated a Mono version, the slider will be disabled.*

Command (Mac)/Control (PC) click the control to reset it to its default value of 0 (absolute center).



CONTROL SIGNAL LEVEL

This control applies gain or attenuation to the Control Signal. The range is from -inf (i.e., off) to 12 dBs of gain.

To adjust Control Signal Level, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the amount of attenuation or gain in dBs.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of -inf (off).



CONTROL SIGNAL PAN

This slider allows you to set the Control Signal's location in the stereo spectrum.



NOTE: *Control Signal Pan only functions in Stereo or Mono->Stereo versions of ARTICULATOR. If you have instantiated a Mono version, the slider will be disabled.*

Command (Mac)/Control (PC) click the control to reset it to its default value of 0 (absolute center).

Usage Tips

Although we've already mentioned this a few times, it bears repeating that to obtain the most interesting results from ARTICULATOR, it's most important to select appropriate Control and Audio Signals.

Control Signals

- First and foremost, the Control Signal should be dynamic, both in formant information and volume (as would be case with speaking or singing).
- Since the actual pitch of the output is provided by the Audio Signal, if you are singing the lyrics to a song, it's not necessary to sing the actual notes of the melody for the Control Signal. It's only necessary to enunciate the words clearly and distinctly, with whatever expressive gestures are appropriate for the track. Even a monotone will work.
- Don't limit yourself just to words. All manner of vocalizations and vocal effects can be used to provide a dynamic Control Signal. Try vocal beat box sounds or the kinds of vocal wah wahs that simulate a resonant synth filter.
- Interesting effects can result from using the amplitude envelope of an intricately rhythmic Control Signal like congas, tablas, or even an entire drum track.

Audio Signals

- The most important characteristic of a good Audio Signal is that it has rich harmonic content. The formant information from the Control Signal essentially accentuates certain frequencies while attenuating others. But you can only hear this effect if the Audio Signal actually contains harmonic information at those frequencies.

Using a sine wave, for instance, which contains only a fundamental frequency with no harmonics, would provide almost no raw material for the Control Signal to act on. Most solo instruments (e.g., trumpet, flute, violin, etc.), as well as solo voice, are not much better.

Again, better candidates are rich synth pads, distortion-rich guitar sounds, thick chords (from any combination of instruments), and even whole mixes.

- An exception to the point above is that if you are only interested in using the Control Signal to provide amplitude modulation, the harmonic content of the Audio Signal will not be an issue.
- If you do want the Control Signal to provide the rhythmic articulation, the Envelope Amount control should be set at 100 and the Audio Signal should be sustained in nature.
- If the Audio Signal itself is rhythmic and you want to preserve that rhythm, set the Envelope Amount control to 0, in which case the Control Signal will affect only the harmonic articulation.


Chapter 10: MUTATOR Extreme Voice Designer



MUTATOR provides a combination of tools for creating unusual, weird, or downright wacky voices.

With high-quality pitch shifting, throat modeling, a pitch-tracking ring modulator for “mutation,” and a mind-bending Alienization function, MUTATOR is the perfect tool for unique vocal effects and post-production sound design.

Controls



Vocal Range

Use this control to select the range of the

track you will be processing. Choices include soprano voice, alto/tenor voice, baritone/bass voice and instrument (a general setting for anything that isn't actually a vocal).

Matching the appropriate algorithm to the input results in faster and more accurate pitch detection and more accurate pitch shifting and modeling.

To select vocal range, click on the Vocal Range pop-up and then select the desired range from the pop-up list.



NOTE: Choosing the wrong Vocal Range (or just forgetting to set it at all) can result in compromised performance. Pay attention.

Voice Design Section

The following three controls allow you to create the basic vocal characteristics of your voice.



SHIFT PITCH

As its name suggests, this control is used to shift the pitch of the input voice either up or down.

MUTATOR's pitch shifting automatically includes formant correction (i.e., it preserves the resonant characteristics of the original voice), however you can further modify the vocal characteristics using the Throat Length and Throat Width controls described below.

The range of the control is four octaves — two octaves up and two octaves down. The control is calibrated in 48 discrete steps of a semitone each.

Since an octave is equal to twelve semitones, the values of this control are equivalent to the following musical intervals:

- 1 = minor 2nd
- 2 = major 2nd
- 3 = minor 3rd
- 4 = major 3rd
- 5 = perfect 4th
- 6 = aug 4th/dim 5th
- 7 = perfect 5th
- 8 = aug 5th/min 6th
- 9 = major 6th
- 10 = minor 7th
- 11 = major 7th
- 12 = octave
- 13 = minor 9th
- 14 = major 9th
- 15 = minor 10th
- 16 = major 10th
- 17 = perfect 11th
- 18 = aug 11th
- 19 = perfect 12th
- 20 = minor 13th
- 21 = major 13th
- 22 = minor 14th (rarely used)
- 23 = major 14th (rarely used)
- 24 = 2 octaves

To adjust pitch shift, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the number of semitones shifted.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0.



THROAT LENGTH

This control uses the vocal tract modeling technology from THROAT to modify the vocal characteristics of the input voice by actually varying the length of the voice's model vocal tract. It can be used in combination with the Shift Pitch control to modify gender and/or vocal quality.

Values above 1.00 represent a lengthening of the throat while values below 1.00 represent a shortening of the throat. The actual values represent the percentage change in the throat length. For example, a value of 1.20 represents a 20% increase in throat length, while a value of 0.80 represents a 20% decrease in throat length.

In addition to simply changing vocal timbre, increasing throat length is useful when the original input is female and you are shifting it down and want it to sound male. Conversely, decreasing throat length is useful when the original input is male and you want to shift it up to sound female or childlike.



NOTE: While this control gives you the ability to radically change the throat length, keep in mind that the variation in the length of human vocal tracts is rarely more than about 20% in either direction. If you are looking for a "realistic" vocal characteristic, start with modest settings of this control. More extreme settings can produce dramatic results, but probably not what anyone would call "realistic."

To adjust Throat Length, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the percentage of change.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1.00.



THROAT WIDTH

As you might imagine, this control is similar to the Throat Length control described above, but instead modifies the width of the modeled throat.

Values above 1.00 represent a widening of the throat while values below 1.00 represent a narrowing of the throat. The actual values represent the percentage change in the throat width. For example, a value of 1.20 represents a 20% increase in throat width, while a value of 0.80 represents a 20% decrease in throat width.



NOTE: As with the Throat Length control above, if you are going for "realism," limit yourself to modifications of about +/- 20%. For special effects, anything goes.

To adjust Throat Width, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the percentage of change.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1.00.

Mutate Section

Having created your basic vocal characteristic with the Voice Design controls, the Mutate section lets you turn that voice into a variety of, how shall we put it, alternative entities (i.e., weird creatures).



MUTATION

The Mutation control allows you to select among 24 different varieties of mutation, as well as being

able to select "0," which, not surprisingly, will result in no mutation effect.

Although the sound of each mutation setting depends to a great extent on the characteristics of the specific voice being processed, as a general rule, odd-numbered settings will typically result in more enharmonic overtones being added to the voice (i.e., weird growly or grunginess) than even-numbered settings. When first experimenting, try starting with "8" and then work out from there.

Beyond that, there are really no rules for predicting how a particular setting will affect a particular voice. Trial and error is the method of the day.



NOTE: For the technically curious, what's going on here is a unique variation of the technique called *ring modulation*, a staple of early analog synthesizers. What makes MUTATOR's mutation process special is that MUTATOR is constantly detecting the pitch of the input voice and shifting the relative pitch of the modulating signal to match, thereby preserving the basic mutation quality regardless of pitch changes in the input. A setting of "8" sets the modulation signal to the same frequency as the input voice. "4" sets it an octave lower and "16" an octave higher. Intermediate settings result in more complex ratios and more complex harmonic structures.

To select a Mutation variety, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the mutation number.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0 (no mutation).



MUTANT MIX

This control allows you to set a mix between the original unprocessed voice and the voice after being processed by the Mutation control. This mix is then sent on to the Alienize section.

When this control is set to the extreme “Original” end of its range, only the original unprocessed voice will be sent. At the extreme “Mutated” end, only the processed voice will be sent. In the middle of the range, a selectable mix of the two will be sent.

To set the Mutant Mix, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the percentage of the Mutated signal in the mix.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 100 (all mutated signal).

Alienize Section

The Alienize section lets you give your speech the quality of an alien language by chopping it up into small segments and playing each segment in reverse. (The effect is kind of hard to describe in words, so just try it. You’ll see.)



ALIENIZE

The Alienize switch lets you activate and deactivate the Alienization process. Click anywhere on the switch to toggle its state between On and Off.

It turns blue when Alienization is On and gray when it is Off.



DIALECT

This control sets the length (in time) of the speech segments that are reversed. Low numbers result in short

segments, while progressively higher numbers result in progressively longer segments. At very low numbers, the segments are so small that you can still easily understand the original spoken or sung words. At higher numbers intelligibility becomes progressively worse, until at very high numbers, each segment can include one or more entire words.



NOTE: It should be noted that Alienize is a real-time function, i.e., it doesn’t simply analyze and reverse sections of a pre-recorded track. You could, in fact, connect a microphone and have MUTATOR Alienize your voice as you speak or sing. Try it. It’s fun.



ANOTHER NOTE: Because the Alienize function needs to store up a small bit of your speech before it can start Alienizing it, there will be a short delay between the beginning of the audio and when the alienized voice begins to sound (the exact delay depends on the selected Dialect setting). Conversely, depending on exactly when your audio ends, the alienized voice may continue for a very brief time after you stop while the last alienized segment is completed. (Think of it as Alien Time-shift.)

To set the Dialect, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. The numeric display will indicate the relative segment length (in arbitrary units).

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 10.

Input/Output Section

Here's where you can adjust the output level, bypass the entire plug-in, and get a visual representation of the before and after versions of your audio.



OUTPUT VOLUME

This control applies gain or attenuation to the output signal. Since MUTATOR's various functions can substantially increase or decrease the level of the input, you can use this control, in combination with the output level meter, to insure that your output is at an appropriate level without clipping.

To adjust output volume, click on and move the fader to the desired position. The numeric display will indicate in dBs the exact amount of gain or attenuation applied.

Command (Mac)/Alt (PC) clicking anywhere in the fader track will reset the control to its default value of 0dB.


OUTPUT LEVEL METER

Use this meter in combination with the output level control to set the optimum output level. Ideally, the meter should peak in the upper area of its range without repeatedly pegging at 0 dB.



BYPASS

The Bypass control lets you completely bypass the plug-in. When the Bypass button is blue, the plug-in is processing normally. When it is red, it is bypassed. Click the button to toggle its state.



AUDIO DISPLAY

The Audio Display gives you a visual indication of the original and processed versions of your audio (and is just kind of cool to watch). The blue plot is the waveform of the original audio, The red plot is the processed audio. For hours of fun, try changing the various controls while watching their effect on the red plot.

Usage Tips

The effects of the various MUTATOR functions are so dependent on the nature of each individual voice, it's hard to give useful general guidelines. If there was ever a plug-in that rewarded just messing around, this is it.

With that caveat:

- If you're looking to create the voice of a demon, try shifting pretty much any voice down 8-12 semitones (depending on how low the voice is to begin with), set Throat Length to 1.10, and set Mutation to "2". Then tweak as desired.
- To turn a woman into a child, shift the voice up 5 semitones, set Throat Length to 1.30, and set Throat Width to 0.60. Then tweak as desired.
- The effect of Alienization Dialect settings are entirely dependent on the specific speech it's applied to. However, for an interesting musical effect, use a long cycle period and try matching the period to the tempo of your song.



WARM is an easy-to-use plug-in that accurately models the sonic effects of recording audio through a high quality analog tube preamplifier.

Chapter 11: WARM

Tube Saturation Generator



Based on Antares' legendary (at least among those of us who work here and, based on the volume of emails we've gotten since we discontinued it, a fair number of others as well) Tube plug-in, WARM, warms up your vocals with Antares' world-renowned tube modeling technology, while being so processing-efficient, you can put an instance on pretty much every track in your project and still barely see your CPU meter twitch.

So what exactly is WARM?

WARM is an easy-to-use plug-in that accurately models the sonic effects of recording audio through a high quality analog tube preamplifier. It offers the choice of two different tube models as well as a unique OmniTube function that allows the application of the tube effect to the entire signal (as opposed to only the transients, as is typically the case with a tube preamp).

WARM's extended parameter ranges have been designed to allow the creation of a wide variety of sonic effects beyond the classic subtle tube "warmth." Additionally, WARM has been designed to be extremely DSP efficient, so that even a relatively modest native system will be capable of running WARM on an almost unlimited number of tracks simultaneously.

What's the big deal with tubes anyway?

For most of us in audio recording, distortion is a bad word. We typically want the cleanest possible recordings, free of any audible distortion. But for certain types of audio there is an exception: the distortion that comes from transient clipping in a high quality tube preamplifier.

When high quality tube pre-amps are operated in their linear range, there is virtually no signal distortion and their audio qualities are essentially identical to that of high quality solid state pre-amps. However, it is common for transients to exceed the preamp's linear voltage range, resulting in distortion. And here's where things get interesting. The distortion characteristics of a vacuum tube pre-amp are vastly different than those of solid state amplifiers. Specifically, the overtones produced by tube distortion are generally harmonically related to the fundamental input pitch, whereas solid state overdrive distortion generally produces non-harmonic overtones. This tube distortion is often described as adding a certain "warmth" to a sound. This is in contrast to what is often described as the "brittleness" of the solid state sound.

As a result, many people pay large sums of money for tube preamps designed to provide this desirable but elusive quality. It was the best sonic qualities of a number of these preamps that served as the sources for WARM's digital models.




Controls


Input Level

This control applies gain to the input signal before it's passed to the tube model. It is typically used to ensure that the level of the input is high enough to exceed the clipping level in the tube model. As a starting point, adjust the gain such that the input level meter needle spends most of its time in the top two (gray and purple) meter segments. Further adjustments can be made in combination with the Drive setting (see below).

To adjust gain, click on and move the fader to the desired position. The numeric display will indicate in dBs the exact amount of gain applied.

Command (Mac)/Alt (PC) clicking anywhere in the fader track will reset the control to its default value of 0dB.

 **NOTE:** Since the Drive control is capable of a maximum of 12dB of gain, the input signal transients must be at least above -12dB to produce any tube effect. If the input meter never enters the top segments of the meter, no tube effect will be audible with even maximum Drive.

 **ANOTHER NOTE:** Unless you are going for some special effect, you do not want to have clipping at the input stage. This is not the "good" clipping that happens in the tube model. Input clipping is that nasty digital kind.



Tube Selector

WARM includes two different tube models. The Velvet (blue) tube models the subtle effects of a classic high-quality tube preamp. While this model is perfect for adding warmth to vocals, it's just as useful for acoustic guitar, horns, strings, synth pads, in fact pretty much anything.

The Crunch (red) tube models the deep, warm distortion of an overdriven tube amplifier. Use it in moderation to impart a subtle (or not-so-subtle) effect to almost any track. Or crank it up for that "industrial" vocal sound, electric guitar, bass, or synths.

To select a tube, simply click the appropriate button.




Drive Amount

The amount of tube saturation effect applied to your audio is controlled by the Drive control in combination with the Input Gain control described above.

The Drive control determines the amplification factor of the modeled tube pre-amp with the numeric display indicating the amplification in dB. At 0 dB, no distortion occurs, even for full amplitude (+1 or -1) signal levels. These levels represent the “rails” of the amplifier. As the Drive is increased, the amplification is increased. Any regions of the signal that increase beyond the rails generate saturation distortion.

Because the maximum Drive amount is +12 dB, generating the tube effect requires the original signal to be at a level greater than -12 dB. If this is not the case, you should adjust the Input Gain control to increase the level of the sound. (Be certain that Input Gain is not increased so much as to cause clipping.) It may be necessary to go back and forth between Drive and Input Gain a few times to get exactly the effect you want.

 **NOTE:** *If your audio was recorded at an exceptionally low level, it may be that even maximum Input Gain and maximum Drive will still not result in a level high enough to generate distortion. In that case, either re-record your audio at a higher level (if possible) or use your waveform editing program to digitally increase the level (keeping in mind that this may negatively affect the signal quality).*

To adjust the Drive amount, click on the fader handle and move it left or right to the desired setting. The numerical display will indicate the amplification factor in dB and the tube graphic’s will “glow” will change to provide an additional visual indication (i.e., “eye candy”).

Command (Mac)/Alt (PC) clicking anywhere in the scale will reset the Drive value to 0dB



OmniTube

Without OmniTube engaged, WARM functions exactly like a tube preamp, i.e., the

regions of the input signal that exceed the clipping level (typically transients) are affected and all other regions are passed with no change.

With OmniTube on, all regions of the input signal are subject to tube distortion. To accomplish this, when you engage OmniTube, a compressor is inserted into the signal path after the Input Gain control and before the tube model. The compressor is set to compress the signal and then apply sufficient makeup gain to ensure that the Drive control can drive the entire signal above the clipping level. After the tube effect is applied to the entire signal, an inverse gain function restores the signal’s original dynamics.

The effect of OmniTube is highly dependent on the character of the input signal, which tube is selected, and the setting of the Drive control. Experimentation is the order of the day.

To toggle OmniTube on or off, simply click on the OmniTube button. The color of the switch will change to indicate its current state.



NOTE: *Despite the fact that OmniTube employs a compressor, it is important to remember that the final output signal is not compressed and that the original dynamics are maintained. If you want to actually compress your signal, you’ll need a separate compressor somewhere else in the signal chain.*



Output Attenuation

This control applies attenuation to the signal after it's passed through the tube model. It is typically used to ensure that no clipping occurs as a result of the gain applied by the Drive control. As a starting point, set the attenuation to 0dB (i.e., no attenuation). If the Clip indicator lights, apply just enough attenuation to avoid the clipping and cause the output level meter needle to spend most of its time in the top two meter segments. (The top of the level meter's range is 0dB.

To adjust attenuation, click on and move the fader to the desired position. The numeric

display will indicate in dBs the exact amount of attenuation applied.

Command (Mac)/Alt (PC) clicking anywhere in the fader track will reset the control to its default value of 0dB.

Usage Tips

At its most basic level, WARM is designed to add an extremely accurate model of a tube preamp's sound to warm your digital tracks. But we have purposely designed WARM to have a range of controls that allow you to push well beyond that model. Here are a few other things to try:

- Pass your entire final mix through WARM. Instantiate a stereo version of WARM (or more if you are mixing for surround) on your mix bus. Start with a little Velvet tube with OmniTube off. Experiment with the Drive amount and/or adding OmniTube for a unique timbral effect. Or switch to the Crunch tube for a more dramatic effect. The result will depend a lot on the nature of your mix, but will rarely be less than "interesting."
- Run your vocal through the Crunch tube with OmniTube on and a high Drive level. Instant industrial vocal.
- Although we usually warn against input clipping, for extreme effects try purposely clipping the input (for digital distortion) and then passing the distorted signal through the tube model. This will usually sound pretty ugly, but in some cases it might be just the ugly you're looking for.
- Experiment and let your ears be your guide.

Chapter 12: ASPIRE

Aspiration Noise Processor



ASPIRE is the world's first tool for modifying a voice's breathiness independent of its harmonic content.

Whether reducing a bit of vocal rasp or adding a bit of smokiness, ASPIRE lets you modify the amount and quality of a voice's aspiration noise without otherwise affecting the vocal's harmonic characteristics.

What's aspiration noise?

As we mentioned back in the chapter on THROAT, when you speak or sing, air from your lungs is forced through your vocal chords, causing them to vibrate. The contour of this vibration is the glottal waveform. The actual shape of the waveform is affected by each individual's anatomy as well as the pressure applied to the vocal chords. From there, the voice is propagated through the throat, the mouth and out through the lips. It is the shape of these structures, both their length and width, that create the resonant characteristics that combine with the glottal waveform to define a unique vocal identity.

But going into a bit more detail, we find that the vocal waveform is actually a combination of two different components. One, the one we usually think of as the primary element of vocal


timbre, is the harmonic content generated by the process described above. However, there is a second component that is generated by the turbulence of the air rushing past the vocal chords. This component, aspiration noise, does not have harmonic content and can roughly be thought of as "breathiness."

Aspiration noise presents itself in a variety of ways, from a subtle bit of breathiness to a full-blown rasp. Of course, whether it's a problem or not depends on the amount of noise and the style of a performance. A whole lot of roughness that would work well for screaming punk or roots blues might well be a problem for mainstream pop or a smooth ballad (of course, it also might not). Up to now, there wasn't much you could do about it. Now, with ASPIRE, there is.

What's ASPIRE do?

Quite simply, ASPIRE analyzes a vocal in real time and separates the aspiration noise component from the harmonic content. It then allows you to either increase or decrease the amount of aspiration noise, as well as letting you affect the character of the aspiration noise by applying a parametric EQ to the noise component.

Controls



Voice Type

Use this control to select the vocal type

of the track you will be processing. Choices include soprano voice, alto/tenor voice, baritone/bass voice and instrument (a general setting for anything that isn't actually a vocal). Matching the appropriate algorithm to the input results in faster and more accurate analysis.

To select vocal type, click on the Voice Type pop-up and then select the desired type from the pop-up list.



NOTE: Choosing the wrong Voice Type (or just forgetting to set it at all) can result in compromised performance. Pay attention.



Tracking

In order to accurately analyze the input, ASPIRE requires a periodically repeating waveform,

characteristic of a voice or solo instrument. The Tracking control determines how much variation is allowed in the incoming waveform for ASPIRE to still consider it periodic.

If you are working with a well-isolated solo signal (e.g., tracking in a studio or from a multitrack master) you can typically set the Tracking control to 25 and forget it.

If, on the other hand, your signal is noisy or not well-isolated (as might be more common in a live performance situation) or you are dealing with a particularly breathy or guttural voice, it may be necessary to allow more signal variation (higher Tracking numbers). However, under these circumstances tracking is not guaranteed and a too "relaxed" setting may introduce distortion and popping.

To adjust Tracking, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise. Remember, higher numbers allow more variation, while lower numbers require less.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 25.



Display

The display provides a graphic representation of the harmonic content of the input signal (the white plot), the aspiration noise component and its relative level (the red plot), and the current setting of the parametric EQ (the green plot).



Increase/Reduce

Click the appropriate button to either Increase or Decrease the amount of aspiration noise. The appropriate controls will become active (and the remaining controls will be disabled).



IMPORTANT NOTE: ASPIRE acts on the aspiration noise naturally present in a person's voice. It does not synthesize aspiration noise from scratch. There are people with extremely pure voices that contain virtually no audible aspiration noise. In those relatively rare cases, ASPIRE will have no audible effect.



REDUCTION

This control allows you to select the amount of reduction of the aspiration noise. A setting of 0.0

provides no reduction. Maximum reduction is -12.0 dB. The red aspiration noise plot on the graphic display will move to reflect the setting of the control.

To adjust Reduction, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of -6 dB.



INCREASE

This control allows you to increase the amount of the aspiration noise (although see the


section on EQ below for an explanation of how this function can also be used for selective reduction). A setting of 0.0 provides no increase. Maximum increase is 12.0 dB. The red aspiration noise plot on the graphic display will move to reflect the setting of the control.


To adjust Increase, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 6 dB.

EQ Section

The EQ Section provides one band of traditional parametric EQ.

 **NOTE:** It's important to remember that this EQ affects just the aspiration noise component, not the main harmonic component. If you want to apply EQ to the main voice, you'll need to use another EQ somewhere else in the processing chain.

 **ANOTHER NOTE:** Although the EQ is only active in the "Increase" mode, various combinations of Increase amount, Frequency, Q, and Gain can result in selective shaping of the aspiration noise. Specifically, using negative Gain amounts, you can use this mode to decrease the aspiration noise in a specific frequency range while increasing it (or leaving it unchanged) in another.



FREQUENCY

This control selects the frequency of the EQ (i.e., the frequency that will either be increased or decreased

via the Gain control). The range is from 0Hz to 4000Hz.

To adjust Frequency, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 1500Hz.



Q

The Q control allows you to set the slope (and, consequently, the bandwidth) of the EQ's peak or dip.

Its range is 0.10 to 3.10. High Qs result steep slopes (and narrow bandwidth), while lower Qs result in more gentle slopes (and wider bandwidth). (The easiest way to understand

this is to watch the green EQ plot while adjusting the Q and Gain controls. All will become clear.)

To adjust Q, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0.70.



GAIN

The Gain control lets you set the amount of boost or cut that will be applied to the aspiration noise component

at the selected Frequency. The range is from -12.0dB to +12.0dB.

To adjust Gain, click on the knob and move the cursor up to turn the knob clockwise, or down to turn it counterclockwise.

Command (Mac)/Alt (PC) clicking on the knob will reset it to its default value of 0.0.

Usage Tips

- If you'll be using any of the more radical AVOX 2 plug-ins on a voice, ASPIRE should come first in the signal chain (or second, if you will also be using SYBIL).
- As in many of the AVOX 2 plug-ins, we've provided control ranges that allow you to create effects well beyond what would normally be considered "realistic." In particular, large Increase amounts in combination with high Gain and Q at various frequencies can result in some rather unusual effects.
- As mentioned above, ASPIRE will only have an audible effect if a voice actually contains natural aspiration noise. If you are working with one of those rare super-pure voices and need to add an element of breathiness, use the Add Breathiness function in THROAT to add synthesized simulation of aspiration noise.

Index

A

- Antares Audio Technologies ii, 1
- ARTICULATOR 6, 27
 - Controls 28
 - Amplitude Modulation Section 33
 - EQ Section 33
 - Formant Extraction Section 31
 - Formant Modulation Section 32
 - Input Mixer Section 30
 - Input Routing Section 28
 - Output Mixer Section 34
 - Signal Routing 28
 - Signal Selection 27
 - Usage Tips 35
- ASPIRE 6, 47
 - Controls 48
 - Display 48
 - EQ Section 49
 - Increase/Reduce 48
 - Voice Type 48
 - Usage Tips 50
 - What's aspiration noise? 47
 - What's ASPIRE do? 48
- Authorizing AVOX 2 3

C

- CHOIR 6, 21
 - Controls 21
 - Choir Size 21
 - Pitch Variation 21
 - Stereo Spread 21
 - Timing Variation 21
 - Vibrato Variation 21
 - Usage Tips 22

D

- DUO 6, 19
 - Controls 19
 - Double Level 20
 - Double Pan Position 20
 - Original Level 20
 - Original Pan Position 20
 - Pitch Variation 20
 - Timing Variation 20
 - Vibrato 19
 - Vocal Range 19
 - Usage Tips 20

H

- Harmony Engine 5, 9

I

- Installing AVOX 2 3

L

- License Agreement iii

M

- MUTATOR 6, 37
 - Controls 37
 - Alienize Section 40
 - Input/Output Section 41
 - Voice Design Section 37
 - Usage Tips 41

O

- Overview 5

P

- Processing Guidelines 7
- PUNCH 6, 23
 - Controls 23
 - Ceiling 23
 - Gain 23
 - Impact 23
 - Output Level 24
 - Usage Tips 24

S

- SYBIL 6, 25
 - Controls 25
 - Attack Time 26
 - Compression 26
 - Gain Reduction Meter 25
 - High Pass Frequency 25
 - Release Time 26
 - Threshold 26
 - Usage Tips 26

T

- Technical Support 4
- THROAT 5, 11
 - Controls 11
 - Breathiness 12
 - Graphic Throat Display 14
 - Set up 11
 - Throat Modeling 13
 - factory presets 17
 - Usage Tips 16

W

- WARM 6, 43
 - Controls 44
 - Drive Amount 45
 - Input Level 44
 - OmniTube 45
 - Output Attenuation 46
 - Tube Selector 44
 - So what exactly is WARM? 43
 - What's the big deal with tubes anyway? 43