Antares kantos 1.0

AUDIO SYNTHESIZER (PC/MAC)

Audio-controlled monophonic synthesizer.

Pros: Extensive synthesis tools. Deep, lush, and quasi-organic synth textures that track nuances of original signal. Most parameters can be modulated in real time via characteristics of audio input (pitch, amplitude, timbre). Can import AIFF and WAV files for use as oscillators.

Cons: Parameters can't be directly controlled via MIDI CC or a MIDI keyboard. Pitch tracking is problematic with percussive source material. Consumes considerable CPU power. Envelope attacks lack a certain "punchiness."

Antares, 831-461-7800, www.antarestech.com

$299

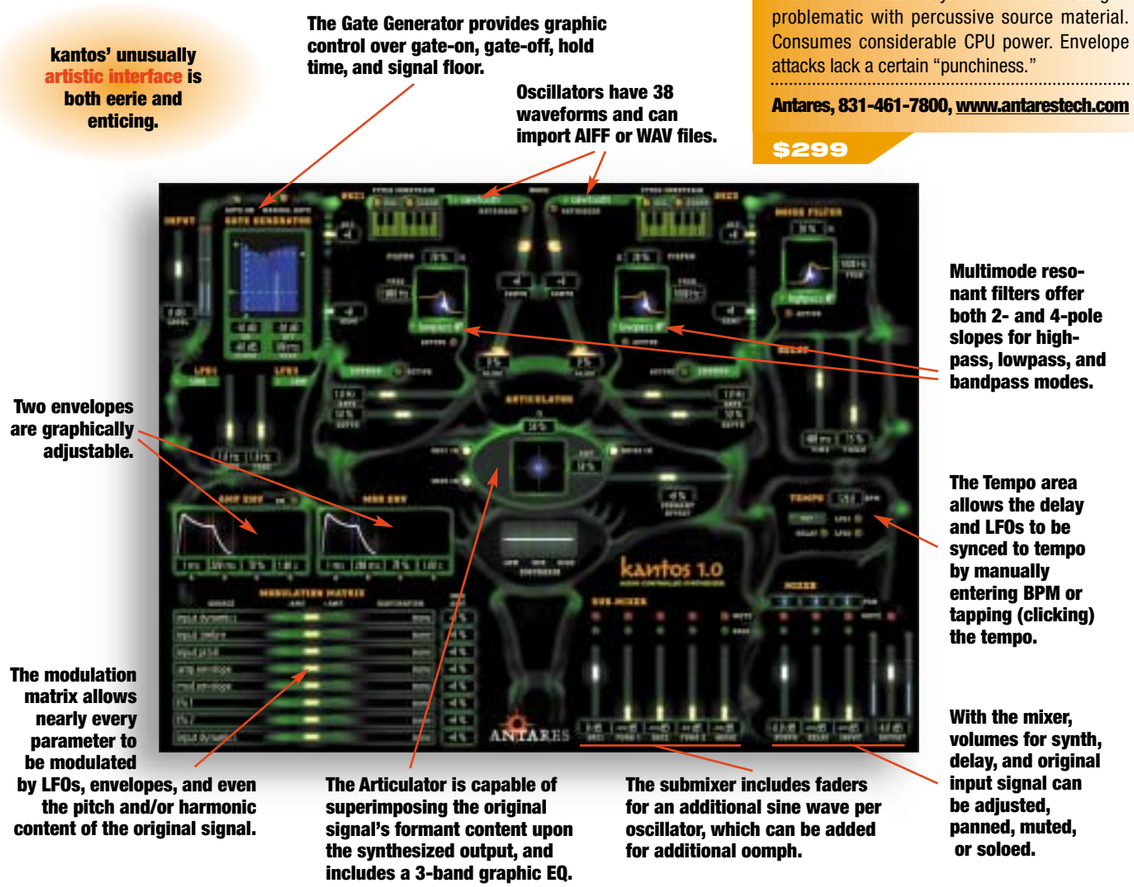

by Francis Preve

I'm a sucker for exotic audio tools — the weirder and more powerful the better. Software such as U&I's MetaSynth and Waves' Enigma plug-in have become indispensable in my studio arsenal, they allow me to create sounds that are difficult (if not impossible) to realize via standard synthesis techniques.

That's why when Antares announced kantos last spring, I was intrigued. kantos has a serious buzz surrounding it, as it's an audio-driven synthesizer that relies on the characteristics of an input signal to control and shape its output. One thing is certain, this plug-in is definitely something new under the sun.

With the well-deserved success of their ubiquitous Auto-Tune plug-in, Antares has proven they can deliver the goods when it comes to realtime pitch detection. So the idea of a synthesizer that detects the pitch of audio and uses this to generate sounds is tantalizing to say the least.

Overview

kantos is a full-featured dual-oscillator synthesizer, but it can't be controlled directly via MIDI continuous controller or note-entry methods. Instead, kantos relies on live input or recorded audio to drive its synthesis engine. *Any* instrument — voice, guitar, saxophone, drums, even another synth — can serve as trigger and pitch information for kantos.

If this sounds powerful and unique, that's because it is. Want to double a lead vocal with a gliding synth sweep? No sweat. Want to create ethereal drones and wave-whispers that follow a guitar riff *à la* Bill Laswell? Go for it. kantos can even handle percussive material. Loops and beatboxes as sound sources yield impressive, other-worldly results, quite unlike anything you're used to in modern dance mixes. While it

Vital Stats

System requirements	*Mac*: G3 233MHz or better processor (G4 recommended) and OS as required by host application (System 8.6 to 9.X); *PC*: TBD
Platform	*Mac*: VST, MAS and RTAS; *PC*: DirectX, PC VST, and PC RTAS
Synthesis type	modeled analog
Oscillators	4 total; 2 primary oscillators generate square, sawtooth, pulse, plus 34 sampled waves, 2 fundamental oscillators generate sine waves only; can also import AIFF or WAV format files
Oscillator sync	no
Filters	3 multimode filters; 1 for each primary oscillator and 1 for the noise generator, 2-and 4-pole resonant lowpass, highpass, or bandpass; Articulator functions similarly to multi-band vocoding and includes 3-band graphic EQ
Envelopes	2 ADSR (assignable)
LFOs	2; sine, triangle, ramp up, ramp down, square, and random (sample & hold)
Polyphony	one stereo or mono voice
Multitimbral features	none
Copy protection	challenge/response verification
MIDI control	none
# of audio outputs	1 stereo pair
Built-in effects	global delay and 2 chorus modules (one per oscillator); delay time can be manually synced to tempo via BPM entry

works best with cleanly recorded monophonic passages, kantos can add spectral spice to nearly any type of source material, provided you're prepared to do a little tinkering. And there are loads of tinker-tools available within the program.

The included documentation is fantastic. Every feature is covered in detail, several times over, and there are a few tutorials that make use of different types of audio (loops, guitar, etc.). Antares is admirably forthright about what kantos will and won't do, and the manual offers quite a few suggestions for real world applications. Since kantos is so unusual, these tips are a welcome addition.

User Interface

kantos' luminous, H.R. Giger-inspired interface is breathtaking at first glance. Envelopes and filter curves are graphically adjustable, the gate functions include a visual representation of the audio stream in real time, and quite a few elements glow and pulsate when adjusted. Most parameters can also be edited via direct numerical entry for those who prefer that method. A lot of thought and creativity went into this UI, that's for sure.

Most interface elements are self-explanatory, but I still found myself scratching my head occasionally. Why? Well, for one thing, in this interface the signal flows from top to bottom instead of left to right. I've been programming sounds for decades and kantos is the first synth I've seen that uses this approach. Compounding matters slightly is the fact that the oscillators and filters flow downward in parallel — on opposite sides of the screen — into the articulator (more on this below). While this is an accurate representation of the inner workings of this plug-in, it definitely takes a bit of getting used to compared to other synth interfaces.

Synthesis Overview

Because it's driven by audio rather than MIDI, kantos needs to be used as an insert or send. Audio first enters through a fairly complex gate, which can be optimized for a wide variety of signal types, both sustained and percussive. From there, the pitch-detection algorithms work their magic, controlling the pitch of the two oscillators. Each oscillator feeds its own multimode resonant filter, which combined then feed the Articulator. Modulation sources include LFOs and ADSR envelopes, as well as a modulation matrix that includes realtime response to the amplitude, pitch, and timbral characteristics of the input signal. Also included are individual chorus algorithms for each oscillator, a noise generator, and a delay line. There's a lot of horsepower under the hood of this baby.

Since kantos accepts any recorded sound as synth fodder for its engine, it's important to define the behavior of its detection algorithms, which include the gate generator, pitch-detector, and some sort of technology that extracts the harmonic content and formant characteristics for use elsewhere in the synth.

All of this begins with a gate that sports separate thresholds for both gate-on and gate-off, as well as a "floor" parameter (below which kantos will ignore all incoming audio) and a hold parameter, which determines how long the gate will stay open before closing, regardless of the gate-off setting. It's crucial that these parameters be set up according to the type of source material you're using. Otherwise, kantos will not track the signal accurately. "Incorrect" settings will often lead to spectacular accidents, but that's another story.

From there, kantos' pitch detectors kick in, analyzing a signal's frequencies, extracting the pitch, and driving the oscillators. This is where it gets a little fadery. Percussive sounds — like xylophone, timbales, and marimbas — often contain a boatload of harmonic information in their attack, making it difficult for the pitch detectors to accurately analyze the signal. Additionally, the pitch duration needs to sustain long enough for kantos to determine the fundamental frequency.

What this means in real-world terms is that sustained sounds with clean attacks and strong amplitudes work best, whereas plucky, staccato sounds (especially at fast tempos), sounds with processing (like chorus, reverb, delay, or even excessive distortion), and sounds with tons of enharmonic content can often confuse kantos, regardless of the gate settings.

At first, this may seem like a major defect in the plug-in's functionality. In practice, I found that it often adds to the character and quirky usefulness of kantos. Your mileage may vary. If you forgive kantos for this failing, you will be rewarded in other ways. It's simply a matter of your expectations.

Oscillators

Both oscillators have the customary selection of waveforms. Sawtooth, square, and pulse are accounted for, but Antares ups the ante with a fairly large array of sample-based waveforms, including voice, bass, strings, enharmonic loops evocative of classic '80s wavetable synths, industrial and synth effects, and even digideroo. The wave selection is well thought out and serves as terrific fodder for the synth engine. For additional phatness, each oscillator has its own chorus effect with independent rate and depth.

To help work around the shortcomings in kantos' pitch-detection algorithm, oscillator pitch can also be quantized to any chromatic pitch — or group of pitches — via a click-and-go keyboard. Coarse and fine-tuning, as well as independent glide for each oscillator are available, as is a retriggering function for restarting a given sample from each gate event. The combined power of these tools makes tuning kantos to your track a total breeze.

There's also a noise generator with its own filter for adding industrial grunge, nature effects, or even vocoderish whispering to your patches.

kantos allows you to import AIFF and WAV files for use as oscillator waves. There's no size restriction on the size of these files, but the manual thoughtfully indicates that smaller (under 400K) looped files will work best. I experimented with this feature and had mixed results. A classic Fairlight voice sample worked beautifully. Short rhythmic loops worked after a bit of coaxing. But larger waves from sound effects discs weren't quite as flexible. If you have a well-crafted selection of pre-looped waves, you'll be all set. If not, don't sweat it, as Antares has you covered with a few extra bundled wavetables *and* a download area on their website for even more waveforms.

Filters & Articulator

kantos is no slouch in the filter department either, featuring resonant lowpass, highpass and bandpass mode for each of the three filters (one for each oscillator and one for the noise generator). Better still, each filter operates in 2- or 4-pole mode and is graphically adjustable. While I would have liked to see a notch or phaser mode included, it's not a heartbreaking omission.

The oscillator/filter combos feed the "Articulator," which is where the serious sonic exotica occurs. kantos' Articulator superimposes the formant and timbral characteristics of the input signal on the synthesized audio. Two parameters — amount and Q — are adjusted via a biaxial interface element (with yet another luminous orb indicating its position in the X/Y axis). A third parameter, Formant Offset, shifts the input's detected formants up or down for even more sonic mayhem.

Also included in this section is a 3-band (low, mid, high) graphic equalizer for emphasizing various frequency ranges. The actual ranges aren't indicated or separately adjustable, but the inclusion of this EQ helps shape the overall sound of the Articulator's output.

What does it sound like? Kinda like a vocoder. Sorta like a speech synthesizer. Different than both, but definitely useful and otherworldly. Setting the amount and Q parameters to their maximum values yielded the most dramatic results, but it can also be used subtly for adding a little something extra to patches, depending on your source input.

LFOs & Envelopes

Patch animation is covered nicely via two LFOs, two general-purpose envelopes, one of which can be used to modulate the amplitude of the synthesized tone.

The envelopes are of the standard ADSR-type, with time parameters adjustable in millisecond increments. Well not *exactly* millisecond, as 173ms was rounded to 171ms and 982ms rounded to 987ms, but it's more than adequate for most applications.

The LFO shapes include sine, triangle, ramp up, ramp down, square, and random waveforms with rates ranging from 0.1Hz to 20Hz. Both LFOs can be synced to tempo, but this value must be entered manually or tapped, as it's not detected automagically under VST. This isn't a major deal, but I would have liked to see the addition of various note values instead of strictly whole-note resolution (one cycle = one measure).

Delay

A simple delay algorithm provides everything from echoes to metallic static flange effects. Delay time and feedback parameters are included, and the time can be synchronized to tempo, again by directly entering BPM or tapping the tempo with your mouse. The tap-tempo function worked adequately, but I preferred simply to enter the BPM based on my sequence tempos.

As with the LFOs, tempo synchronization applies strictly to quarter-note values. So if you want different rhythms, you'll need to get a calculator and do the math yourself. (You can download the free Music Production Calculator, www.hitsquad.com, for this kind of work.)

Realtime Modulation

With all of the amplitude, pitch, and formant detection going on under kantos' hood, Antares decided to expose these features for use as modulation sources. Each of the three characteristics (and/or the envelopes and LFOs) can be used to modulate nearly every significant parameter available, including oscillator pitch, filter cutoff and Q (for each filter), articulation parameters, chorus rate/depth, delay time, delay feedback, LFO frequency . . . the list goes on. Suffice to say, the most common modulation destinations are available, and then some.

It's quite something to have the overall amplitude or pitch of the input signal modulate the articulator or delay feedback. When applied to acoustic or vocal articulations, the results have a strangely organic quality, despite their obviously electronic origins. This is truly inspiring stuff.

Mixers

Once processed (or mangled, as the case may be) the oscillator signals are sent to two mixer sections. First, the submixer accumulates the outputs from each oscillator and the noise generator, as well as a discrete sine wave generator for each oscillator that derives its pitch directly from the pitch-detection algorithms. The sine wave generator is a nifty touch as it allows you to thicken the synth's output by reinforcing the fundamental frequencies of the input signal *after* pitch-quantization has been applied via the individual oscillator settings. ➥

The combined output of the submixer is then routed to the final mixer, which has one fader each for synth output, delay output, and direct input signal respectively.

Every fader in each of the mixers also includes mute and solo controls with visual indication of their status. This is handy as you can then quickly isolate each characteristic of your sound for further tweaking.

In Use

During the time I worked with kantos, I found myself encountering combinations of acoustic and electronic sounds that almost defied description. Had I heard them in other artists' tracks, I would definitely have wondered how they were created. Applying kantos to melodic vocals definitely demonstrated the power and accuracy of Antares' pitch-detection expertise in their best light. Once the gates were set up correctly, kantos did a terrific job of tracking even subtle nuances in the performances. I tried both male and female vocals with equal success. Using kantos in this manner was perfect for supporting a vocal with a simple, slightly "vocoded" pad lead. Spoken word material yielded a strange yet cool speech synthesis emulation. More ambitious sonic experiments led to interesting methods of replacing the lead vocal for breakdowns or psychedelic verses on dance tracks.

Other acoustic instruments produced mixed results. Sustained guitar solos, trumpet passages with clear tones, and saxophone samples were successfully tracked (for the most part), with some lovely melodic ambiences being generated as a result. Marimba and kalimba loops didn't fare as well, though limiting the oscillators' pitch-quantization to relevant keys went a long way toward helping kantos deliver useful results.

Drum loops, on the other hand, proved to be fantastic fodder, despite my initial expectations about kantos' ability to track the pitch. We're talking about a drum loop after all! I explicitly chose pitches I wanted kantos to generate (via the oscillators' keyboard interface) and let 'er rip. Everything from synchronized, percolating sweeps to ghostly, syncopated undercurrents was possible. Creating rhythmic drones with this plug-in gave me hours of tranced-out pleasure. Downtempo. Uptempo. It all worked, once the gates were set up accordingly.

Applying kantos to synth parts was a bit different. Single oscillator patches with medium decays worked fairly well, as long as the decays weren't too abrupt and the release times didn't cause notes to overlap. While it sounded quite cool when layered with the original part, the inherent processing delays involved in analyzing the signal characteristics made it slightly problematic to get really punchy envelopes, regardless of kantos' ADSR settings. The kantos manual suggests several ways to replace bass synth parts on the fly by tracking the original bass line with a sine-wave patch, then processing the result with kantos. I tried it. It worked okay, but not really mind-blowingly accurate or in your face. Then again, if you think about it, this isn't a huge deal since driving kantos with another synth is a mobius loop of redundancies. If you want to track a plucky arpeggio with a burbling ethereal drone, it's more than possible — just be sure to adjust the pitch quantization to match or compliment the original part.

While kantos definitely consumes more CPU power than the average plug-in, it worked quite well overall. I *did* encounter a few mysterious CPU spikes when running multiple instances of kantos. The CPU meter was hovering around 50% on the dual-G4, but Cubase's overload light came on unexpectedly, interrupting the sound. It was impossible to determine whether this was a bug or simply the side-effect of a power-hungry plug-in, but it didn't crash either machine, so it's not really "dangerous," per se.

Conclusions

Make no mistake — kantos is capable of taking sounds to places you've never heard before. How you react to kantos will depend on your overall approach to music and your needs from a synth plug-in. If you crave weirdness and possibilities from your instruments, you're gonna love kantos. It's a happy-accident generator if there ever was one. Doubling instruments and vocals with textures that follow the nuances of the performance is a truly unique sound that isn't easily duplicated by other products, if at all.

On the other hand, if you expect kantos to track everything you throw at it perfectly and behave like a well-tempered sonic servant, you'll be disappointed. The pitch-detection and gate algorithms require certain kinds of content to function at their best. This is not a be-all do-all synth plug-in. It's a new way of approaching recorded material.

kantos is a unique addition to the world of plug-ins. It's not for everyone, but for some musicians and producers, it will be a deep well of inspiration.

Francis Preve is a remixer/producer who has worked with Orbital and Utah Saints, among others. Get the facts at www.fap7.com.