Transoniq Hacker

The Independent Ensoniq Mirage User's Newsletter

SPLITTING AND THE MIRAGE

by Clark Salisbury

Felicitations, samplers! This time out I want to ramble on about some stuff you can do with your factory sounds, with some (hopefully) helpful ideas that will apply to your own sampling as well.

I work in a music store (it's a dirty, rotten job, but somebody's got to... well, never mind) and one of the things I'm asked most frequently is how to deal with the rather novel approach to keyboard splitting that the Mirage uses. It's a somewhat complicated area, but as we'll soon see, complexity can be the price one pays for flexibility.

As is often the case, the Mirage organizes things somewhat differently when it comes to split-keyboard effects. As a matter of fact, "splitting" may be something of a misnomer. I tend to think more in terms of "covering and uncovering" when dealing with the Mirage. Here's what's going on. RC090317

Let's say you have a slap bass sound sampled on the lower half of the Mirage keyboard, and a guitar sound on the upper half, with the guitar sounding on all the keys above F3. The way to look at it is that the guitar sample is actually on all the keys of the keyboard, but the bass sample is covering it up on the lower half of the keyboard. This is where Ensoniq's unique "top key" system comes into play. We can say that the "top key" that the bass sample plays on is the F3 key, or key #30 (counting from the lowest up). We can now bring the split point of the Mirage down a key by selecting a new top key for the lower (bass) sample. First, make sure that you are editing the lower keyboard half. Press the 0/Program button until the display flashes L and a number from 1 to 4. Press the number (1 to 4) of the program you wish to edit. (Yes, split points can be different for each program.) If the sound is a single-sampled sound, simply call up Parameter 72 (top key) and use the cursors (up and down buttons) to choose a new top key. If you bring the top key down to 29, for example, the bass sample will play up to the 29th (or E3) key, uncovering the guitar sample on F3 and the keys above.

Simple, right? Well, not always. If the lower keyboard has been multi-sampled (as is the case with the slap-bass sound) you would have just set a new top key for the first sample on the lower half - which is not necessarily the highest sample. This is one place where Parameter 26 (wavesample select) comes in handy. If there are two or more different samples on one half of the Mirage keyboard, you have to tell the Mirage which sample to set the new top key for.

Matters become further complicated when you discover that the first wavesample isn't always the lowest sound on the keyboard. So how do you determine which wavesample is the one in the area of the keyboard you're working with? Well, one way is to count the keys up from the bottom of the keyboard to the split point. If the last key that your lower sound occurs on happens to be, say, key 30, then go to Parameter 26 (wavesample select), select #1 (with the cursor buttons), then select Parameter 72 (top key) and hit the Value button. If the display reads any number other than 30, you will know that the wavesample that you've selected is the wrong one, since its top key doesn't match the top key of the sample you want to work with. Return to Parameter 26, select wavesample #2, and repeat the process. You may have to look at all eight possible wavesamples (per keyboard half) to find the right one, but eventually the magic number (30, in this case) will appear, and you can rest assured that you've probably found the right wavesample.

"Probably?!!" you say. "What's this 'probably' stuff?" All in good time, dearie. All in good time.

First, I want to make a couple of points. One, the top key function doesn't necessarily have to be used only for changing keyboard split points. It can be used to redistribute sounds within one keyboard half. Don't care so much for the octave-plus of toms on the electric drums preset? Locate the snare drum wavesample, and set its top key for something like 20, and presto! Two octaves (almost) of snare drums. And as long as you've got the snare drum wavesample as the current sample, key in Parameter 67 (coarse tune) and use the cursor buttons to lower the snare's pitch an octave (or two, or three). Now you're getting the hang of it. Fun, huh? Now, let's do something with those pesky cymbals...

Here's an interesting idea from a friend of mine who has no fear of performing tedious operations; try locating drum sounds on just the black keys, with keyboard sounds on the white keys, so you can play all five octaves (in the key of C) and have some percussion, too!

Anyway, if this stuff doesn't keep you busy for a few months, let's return now to why I used that most hideous of words to programmers, "probably" a couple of paragraphs ago.

Occasionally, when working with changing split points, you may locate what you think is the appropriate top key number, but if you lower the top key number either nothing happens, or else you find you now have a totally different sound between your lower and upper keyboard sounds. What has happened is this: you have "uncovered" another higher-numbered wavesample on the lower keyboard half. To proceed, then, locate the number of the offending wavesample, and lower its top key also. That should take care of that. But wait, there's more...

You can use this method to actually get more sounds on one disk than you're supposed to.

HOW TO GET 24 SOUNDS ON A DISK

Probably the easiest way to check this feature out is to boot up the disk labeled "This disk contains 12 lower and 12 upper sounds..." etc. You should have received it with your Mirage. Load Lower Sound 1. Sound L1 is a nice tubular bells type sound. L2 is an (ahem) electric piano, L3 is an organ, and L4 is a Martian clavinet. The way Ensoniq gets four distinct sounds for each of the four programs (L1 - L4) is through multi-sampling and by using Parameter 27, initial wavesample.

First, the tubular bells sound is sampled and stored in wavesample #1 and #2. (If you've read your manual thoroughly, you may remember something about Mix Mode, Parameter 28. When Mix Mode is activated, as it is with this disk, multi-sampled sounds in even and odd numbered wavesamples are paired, i.e. #1 with #2, #3 with #4, etc.) The electric piano is stored in wavesamples #3 and #4, the organ in #5 and #6, and the "clavihetron" in #7 and #8. When the presets are being organized, then, Parameter 27 (initial wavesample) is invoked. For preset L1, initial wavesample is set to #1. Since we're in Mix Mode here, wavesample #2 is paired with #1, and they act together as a single sample. Top key for this sound is set to 61, effectively covering the entire keyboard with the tubular bells sample. Now on to preset L2.

Since we want to distribute our electric piano sound across the entire keyboard, we'll set initial wavesample (Parameter 27) to #3 (remember, wavesamples 3 and 4 are paired). Then we set the top key (Parameter 72) for 61 and viola! Five octaves of electric piano. We can repeat this process for programs L3 and L4, thus yielding four distinct sounds from one lower keyboard disk load. And we can repeat the process with lower keyboard L2 and L3, of course. And we can repeat it again to get 12 sounds on the upper keyboard. But wait - if we have all of our lower keyboard sounds set with a top key of 61, how can we hear what's in the upper keyboard programs?

Well, we can't. Not unless we "uncover" the upper sounds with one of our lower presets.

Boot up Lower Sound #3 from the synth sounds disk, and key in program L4. On the lowest key of the keyboard, you should find a rock band endlessly thrashing away at an E minor vamp. The rest of the keyboard will now have whatever sound was on the upper keyboard half before you loaded the lower sounds from the synth disk. The idea here is that this sound (Lower 3, L4) has its initial wavesample set to #7 (paired with #8) and a top key of 1, thus "uncovering" the upper keyboard sounds. Now you can go ahead and load the upper keyboard sounds. The only catch is that you'll have that damn rock band on your lowest key for all the upper sounds - but it seems a fairly small price to pay for the added flexibility of having so many sounds on a single disk. As a matter of fact, it seems to me that if any of you die-hard samplers out there want to try this procedure on your own, but with Mix Mode (Parameter 28) set to off, you should be able to get 12 lower and 12 upper split keyboard sounds by sampling a new sound to each of the eight wavesamples per keyboard half. RC090317

In other words, sample a bass for wavesample #1 and a piano for wavesample #2, then arrange them into a "split" sound in Program L1. Then sample a zylophone for wavesample #3, and a flute for wavesample #4, and arrange them into Program L2. See what I mean? Time consuming? Yes. Rewarding? Maybe. Will it teach you something about the Mirage? You bet.

LICKING ENVELOPES

By Clark Salisbury

The first sample I ever did on the Mirage was my own voice - a rich baritone, worldly, yet with subtle overtones of youthful poignancy. I plugged in my trusty Shure SM57, hit "sample lower" and presto! "SF" (sampling finished) was flashing in the display and I prepared to explore new areas of sonic enlightenment. I pressed a key. My voice sort of popped out of the speakers like the last of the toothpaste from the tube, darkening and decaying to silence rather quickly. This was not quite the ethereal choir I had expected. "So what gives?", I thought.

As many of you may have guessed, what gave was that I didn't start from a "blank vanilla" setting (as the Advanced Sampling Guide calls it.) Rather, I had booted the Mirage with the piano disk, and then sampled without changing any of the piano program settings. So what happened was that my voice sample played back all right but was processed by filter, envelope, and amplifier settings that had been programmed for a piano sound. This is not what I had in mind.

So, if you had hoped to avoid having to learn anything about analog synthesizer techniques by buying a Mirage, surprise! I'm here to tell you about envelopes generators, filters, and amplifiers, those pesky processors from the analog world.

An envelope is fairly easy to understand. If a sound is broken down to its components we find three basic elements, each of which is a manifestation of some sort of change over a period of time. They are pitch (or frequency), volume (or amplitude), and tone (or timbre). Let's start with amplitude.

Different types of sounds change differently across time in terms of their amplitude. A piano, for example, has a percussive attack (beginning) and will decay (die down) at a more or less steady rate if you continue to hold the key down until it finally reaches silence. A flute, on the other hand, reaches its full amplitude somewhat more slowly (less percussively) than a piano. However, it will sustain at a more or less even volume for as long as it is blown into (also unlike a piano) and the note decays quite abruptly when one stops blowing.

Attack, decay and sustain are three of the components that make up a basic envelope. There's one more part that we use to create a working envelope. It's called release, and simply put, it is the length of time it takes for a sound to decay after you have released the key(s) on the keyboard. Here's the analogy.

When you play a note on an organ, the note disappears (decays) instantly once you release the keys. When you play a note on a piano, however, the note will decay while you are holding the key and die rather quickly once you release it. But a little bit of the sound hangs over, especially on the lower notes. The sound lingers usually for only a fraction of a second, but that minute final decay heavily colors the way in which we percieve the overall sound. I've read of studies that suggest that a sound's envelope (attack, decay, sustain, and release times) are just as important, if not more important, than a sound's timbre when it comes to how we percieve different sounds. If you think this unlikely, then set up a piano sound on the Mirage and set its attack time (Parameter 50) to, say, 30 or so and its release time (Parameter 54) to 0 (long, bowed attack, no release). Pretty weird, huh?

Anyway, the four parts of our basic envelope are attack time, decay time, sustain level, and release time. We can say that every sound has some sort of amplitude envelope. A piano, for example, has an envelope that we could graph in this way:

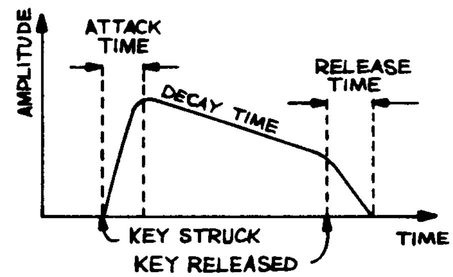

We could graph an organ in this way:

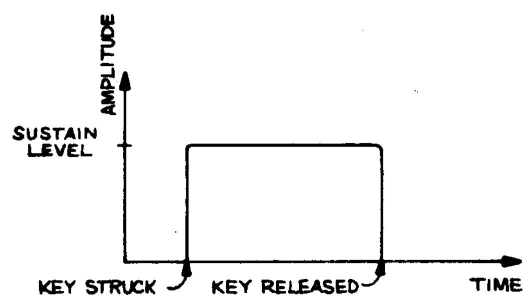

And a trumpet, whose attack is somewhat louder than the sustain portion of the sound, would look like this:

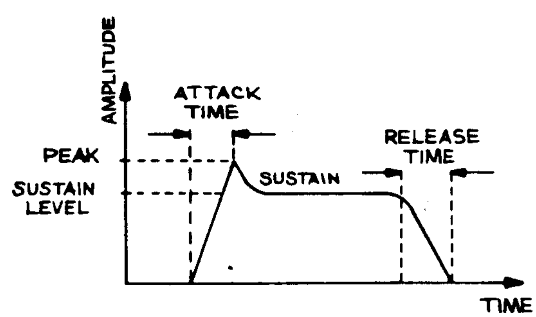

So the four parts of our basic amplitude envelope are attack time, decay time, sustain level (the volume at which a sound sustains, if it sustains) and release time. The common abbreviation is ADSR. RC090317

In more traditional forms of synthesis, we would normally start with our raw, full volume sound and use a handy device called an envelope generator to shape its dynamic movement. The envelope generator is a simple electronic device used to output a control signal which is (hopefully) more or less analagous to the acoustic envelope that we are trying to synthesize. In other words, if we start with a raw violin-type waveform (sound) and want to create a slow, bowed type of attack, we would set the attack control on our envelope generator fairly high. The effect is that when we depress a key on our synthesizer, the sound doesn't abruptly jump to full volume. The rate at which it builds is determined by the attack-time setting. Likewise, we can effect changes in the decay rate, sustain level, and release rate by manipulating the D, S, and R controls on our envelope generator.

At this point I should note that the envelope generator does not act directly on the waveform in question; it merely sends a generic signal which is, in this case, used to control a device called a VCA, or voltage-controlled amplifier. A VCA is a simple amplifier-type circuit with one important distinction. Its output volume (actually, the amplifier's gain) can be controlled by a voltage. This is how it works.

The VCA has three connections - audio input, audio output, and control signal input. This last will accept voltages within a certain range, say 0 to 10 volts. If we connect our violin waveform to the audio input and 0 volts to the control input, we will get no sound at the output. If we apply 10 volts to the control input, we get our violin wave at the output, full volume. Neat, huh? The tricky part is that if we apply a voltage that changes slowly from 0 to 10 volts, the output will give us our violin wave, but it will build slowly from 0 to full volume. Conversely, if we apply a control voltage that begins at 10 volts and decreases gradually to 0 volts, our violin wave will start out at full volume and gradually decay to silence. So where do these rising and falling voltages come from? You guessed it! From our friend the envelope generator. The basic processing circuit now consists of three parts. Our waveform generator (in the Mirage it is called a digital oscillator or DO), our VCA, and our envelope generator (EG). The output of the DO is connected to the audio input of the VCA. The output of the envelope generator is connected to the control input of the VCA, and the output of the VCA goes (more or less) to the input of our monitoring system so we can hear our sound. One word though - if you look in your parameter list in the Mirage in hopes of finding something called VCA, forget it. For one thing, the Mirage uses DCAs (digitally controlled amplifiers) instead of VCAs. The only difference between DCAs and VCAs is that DCAs are controlled by numbers rather than voltages. And, since I find VCAs not only more familiar to most people, but also easier to explain, I hope you'll bear with my somewhat obtuse approach to outlining their function. Also, since this configuration is permanently "hard-wired" within the Mirage, and a DCA has no controls of its own (only inputs and outputs) there are no parameters to change that are directly a part of the DCA. The envelope parameters present us with all the controls we have for the DCA (never mind about keyboard scaling and envelope generator velocity sensitivity - that's for later).

There is one other EG control that I've left for last, primarily because in most "traditional" synthesizers it isn't included in the EG section at all. It's called "peak" and in other synthesizers it would correspond to EG depth or EG intensity. Its function is to control the intensity with which the EG controls the DCA. In other words, if you want to synthesize a slow, bowed attack for the violin waveform but don't want the waveform to reach maximum volume by the time it has completed the attack portion of its cycle, you can use the peak control of the EG to determine how much effect the EG has on the DCA. If the peak control is set for its maximum value, the violin waveform will eventually reach its full volume. If the peak control is set for half of its maximum value, the violin waveform will eventually reach half of its original volume. Now the applications of this may not be immediately apparent, but they are there. Trust me.

At this point, some of you may be saying, "Well, that's just fine, Clark. But, since the Mirage can sample any sound, doesn't it sample the sound's amplitude envelope and everything? I mean, why do we even need to mess with synthesizing all these envelopes and stuff, anyway? And when are you coming over for dinner?"

Well, that's a very kind offer and I'm free next Tuesday night. And, yes, the Mirage will sample any sound, envelope and all. But let's take a look at a piano sample, just for fun.

First, go out and sample the C below middle C on a nice Bösendorfer piano. Remember, now, we want to capture as much of the harmonic content as possible, so let's use a nice high sample rate. All set? Good. Now, you should have about 2 seconds of sample time, so go ahead and sample your note. What? You say the note must take about 4 or 5 seconds to decay? No problem. Sample the first 2 seconds and then loop the sample. What? Now you say the sample sustains forever, like an organ, and doesn't stop until you release the key? Easy. Find the decay control on your amplitude EG (Parameter 52) and the sustain control (Parameter 53). Set sustain (53) for 0 (since you don't want the piano to sustain unnaturally) and vary the decay control until the piano sample has a nice, natural sounding decay. And, as long as you are at it, add a little bit of release time (Parameter 54) because when you release the keys on a real piano, it takes a split-second for the sound to die away. Oh, and if you can't find a Bösendorfer, you can try these ideas with the piano sample you received with your Mirage. And, if you want to do that, but are having some trouble keeping straight which sample is which on the factory disk read last month's article. And, as long as you're doing all this, you might as well start playing with the attack control which will give you some fairly bizarre bowed-piano effects. What the heck, play with all the controls indescriminately.

Okay, now that we've got that out of our system, let's take a look at the filter and the other envelope generator. Don't worry. The hard part is behind us.

A filter is just what it says. It filters out parts of the sound that we don't want, just like a coffee filter filters out coffee grounds. The filter can be very useful for producing rich, full-bodied sounds. Become a filter-achiever.

Anyway, the filter in the Mirage is what's called a low-pass filter (24 db-per-octave, resonant filter for you techno-weenies out there) and its job is to pass the low frequencies, and filter out higher (brighter) frequencies. The effect is that the filter can be used to darken the sound.

The filter in the Mirage is a voltage-controlled filter. It has three connections; an audio input, a control-signal input, and an audio output. (Beginning to sound familiar?) It acts just like a VCA, except that when voltage from the EG is applied to the control input, the sound at the audio output is closer to its maximum brightness, and when less voltage is applied, the sound becomes progressively darker. As a matter of fact, if you filter a sound too heavily, you will get nothing at the output, because you have filtered out all the frequencies. This is sort of like using tin foil for a coffee filter. Caveat Emptor.

Like a VCA, the filter can change across time, gradually brightening and/or darkening a sound. We control these changes with the filter EG, setting the attack, peak (intensity), decay, sustain, and release parameters for the desired effects.

Be aware, though, that the filter and amplifier work together. If you set the filter for a fast decay, the amplifier decay control may seem to have no effect, since, if the sound has already been filtered completely out, there is nothing left for the amplifier to work on.

There are a couple of extra controls on the filter that are not found on the amplifier. The one labelled "Filter Cutoff Frequency" (Parameter 36) sets the initial brightness (or darkness) of the filter. If you wish to use the filter to darken a sample over time, do not set this control too high. It sets the lowest point that the filter can go to - regardless of other settings.

"Filter Q" (or resonance) sets the amplitude of the resonant peak at the filter cutoff point. I like to think of it as "quackiness." Try it and see. It's Parameter 37.

"Filter Tracking" (Parameter 38) allows the filter cutoff to be determined by key number. The effect is that the higher on the keyboard you play, the brighter the sound. This is the way many acoustic instruments act.

Anyway, I realize that this is a lot of information to digest at one sitting. I suggest sitting down with your Mirage, your parameter chart, and this article, and try varying some of these parameters on your factory samples as you go. Before too long you should get the hang of it, and when you do you'll have a basic understanding of the principles of voltage control. These principles apply not only to the Mirage, but to 95% of all synthesizers manufactured to date, so I feel it's pretty worthwhile stuff. Meanwhile, if you get hung up, you can write to me c/o the Hacker. And for further reading, check out SYNTHESIZER BASICS from GPI Publications (publisher of KEYBOARD MAGAZINE). Until next time, "..the knee bone's connected to the thigh bone..."

Clark Salisbury is Product Specialist with Portland Music Co. in Oregon, and is also a partner in "The Midi Connection," a Portland-based consulting firm. He has been actively involved in the composition, performing, and recording of electronic music for over five years, and is currently involved in producing and marketing his own pop-oriented compositions.

USING MIX MODE FOR QUICKER CHANGES AND 48 SOUNDS/DISK

By Tom Darling

Midwest District Sales Manager
Ensoniq Corp

The two questions I hear the most are, "How many sounds can be stored on a disk?" and "Can I change sounds faster than seven seconds?" Well, have no fear, hackers, I have good news for you.

The answer to both these questions can be found in learning how to use the Mix Mode (Parameter 28) to your advantage. This Mix Mode is a much-overlooked feature of the Mirage, which, if used properly, can allow you to store different wavesamples in each of the Mirage's digital oscillators. Check out Sound Disk #2 and you'll see how by calling up Programs L1, L2, L3, and L4, you get four different synth sounds instantly - and they play on all 61 keys. This is because each of the program variations don't necessarily have to be just different envelope and filter settings on the same sample but can actually hold completely different samples. When you load a lower sound from Sound Disk #2, you are really loading four different synthesizer samples - each of which is stored in one of four lower program variations.

Confusing? A little. Let's get into the "how to" of this. As I hinted above, the secret to this is the Mix Mode. By storing a different sample into each program variation and by using Mix Mode, you can have eight lower and eight upper sounds for each of the three lower and upper sound positions on the disk - 48 sounds in all!

First, load the MASOS Operating System and Sound 3 from the MASOS disk. This configures the Mirage lower and upper memory for eight samples each. This is an even division of memory for each sample. You may wish to change this as you go. Now, call up Lower Wavesample 1. (Make sure Parameter 77 is turned ON.) Now, set Parameter 27 to 1. This is Initial Wavesample. It tells the Mirage which one of the eight wavesamples to play first. Now, we use the Mix Mode. Turn the Mix Mode ON (Parameter 28). Make the sample. We have 32 pages of memory available for each lower wavesample. This is plenty for most synthesizer samples since a synthesizer produces its waveform in the beginning of its sound and then basically just replicates it out over the rest of memory. In fact, many times I've gotten great samples from synths in just one page of memory.

Using the Mix Mode like this, the Mirage will put your first wavesample into Oscillator 1. Oscillator 2 is still empty. Now, go to Lower Wavesample 2 and set the Initial Wavesample (Parameter 27 again) to 2. Make your second sample. This second sample will go to Oscillator 2. Set the top keys of these samples to be whatever you like. You now have the following: If you sampled, say, a brass sound into Wavesample 1 and a clavinet sound into Wavesample 2, you'll have both sounds activated when you call up Sound 1 Lower, Variation L1. Now set Parameter 35 to 00. With the Mod Wheel back you will hear the brass (Oscillator 1), and with the Mod Wheel forward you'll hear the clavinet (Oscillator 2).

With Parameter 35 set between 01 and 31 the key velocity will control the mix of brass and clavinet. Layering can also be achieved by setting Parameter 35 to 00 and putting the Mod Wheel halfway forward. Be sure to set Parameter 34 to around 31 so you get an even volume balance.

By continuing through all eight lower and upper wavesamples in this manner, you can really have eight completely different sounds, lower and upper, at command instantly by just calling up the four lower or upper program variations and then using the Mod Wheel. Amazing eh? Well, there's more - don't forget that when you use the Mod Wheel for a mix of the two samples, you actually get a whole new sound. Also, you should try detuneing Oscillator 2.

Well, that's the long answer to the two short questions. With each seven-second lower or upper load you can actually bring in eight sounds for instant recall. And, each disk can actually store up to 48 different samples.

On a slightly different subject, I've developed a certain technique for sampling synthesizers that you may find quite helpful. I usually do one-page loops. I set the sample time to 30, (33k sample rate) and sample a C note, one octave below middle C. I always allocate as little memory as I think I can get away with. This makes the one-page loops less apt to have clicks.

I hope you found this article useful. I look forward to writing more in the future. Take care and best of sampling!