Digital Synthesis

Even though they look a lot like pianos, every digital synthesizer is a computer. There are many types of synthesis advertised, but when the hype and marketing prose is translated into English, we discover they all work more or less the same way. Like CD players, digital synthesizers produce a stream of numbers at a steady sample rate, which are converted to analog form to drive speakers and produce sound. Everything we have learned about digital recording (Nyquist frequency, anti-alias filters and so forth) applies to digital synthesis.

A digital synthesizer has the same functions as a modular analog synthesizer. But instead of circuitry, each module is a subroutine within the main synthesis program. This program is essentially a loop that repeats once for each output sample. The subroutines must each be executed at least once per sample period, and many of them must execute once per note per sample period. Thus the number of notes that may be played, and the complexity of the sound available is determined by the speed of the CPU. As with personal computers, the faster CPUs are the most expensive, and this is reflected in the cost of the instrument.

The cost can be cut down by using a slow sample rate, but this is easily audible in the sound of the instrument. A state of the art synthesizer meets CD quality standards.

This is the typical signal flow, the digital equivalent of the patch:

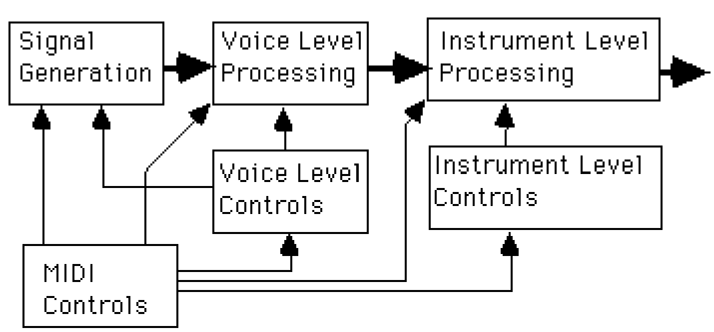

The three top boxes represent processes that deal directly with the signal and must be executed at top speed. Control processes may be updated at a more comfortable rate, so it is less expensive to add a second LFO to all voices than it is to add a second signal generator.

Signal Generation: Wavetables

The heart of any signal generation routine is Wavetable lookup. A wavetable is a section of memory that contains a representation of a segment of sound. This may be as simple as a single cycle of sine wave, or may be a complete recording of a traditional instrument playing a note. On each sample period a value is taken from the wavetable and sent to the output. If the sample rate of the synthesizer is the same as that used for the initial recording, the pitch produced will be the same. In the case of a single cycle wavetable, the frequency produced will be the sample rate divided by the number of entries in the table.

To produce a continuous tone, the table is read over and over.

To play other pitches, the values in the table are taken out of order. If every other value is output, the frequency is doubled. To go down an octave, each value could be used twice, but that would distort the wave noticeably. Modern instruments calculate a value that would fall between the entries in the table, a process called interpolation. Sophisticated interpolation routines can produce any pitch desired, including smooth glissandi and pitchbends.

Interpolation is the most time consuming process within the program loop, so the number of tones you can play this way is going to be limited in a particular instrument. This will then determine the maximum number of notes possible, depending on how many tones are used to create a note.

Variation of tone

Wavetable synthesis has one serious drawback, a very static spectrum of the sound. Most real instruments produce dynamic spectra that change over the duration of a note and never repeat in quite the same way. Simple wavetable gives the sound of a toy organ, occasionally useful, but basically boring. To get around this a variety of strategies can be found:

Very Large Wavetables

Essentially, there is enough sound in the wavetable to play an entire note on one pass, or at least to loop with a moderately long period that sounds like vibrato. Most often, these class of synthesizers use recordings of traditional instruments as the wavetable content. You may be limited to a particular set of recordings, or you may be able to make your own. These are the ubiquitous samplers and sample playback units. Sampling was first marketed by E-mu, and they have a wide range of instruments based upon the technique, but nearly every instrument on the market today has some recorded samples in it.

Frequency Modulation

We learned on the modular synthesizer how to create dynamic sounds with the simple trick of using one oscillator to modulate the frequency of another. In a digital system this can be done with elegance and control. The sounds are not inherently "natural" in the sense of sounding like traditional instruments, but a few realistic efforts exist. For sounds with a vibrancy and life of their own, the technique is unsurpassed. The Yamaha corporation had a patent on FM, and even though that patent has run out, few other manufacturers are interested in pursuing it.

Waveshaping

Waveshaping adds a step of deliberate distortion to a wavetable type generator. Usually this is accomplished by the simple expedient of using the output of the wavetable routine to find a value in a table that contains some strange transfer function. It's quick and cheap, and if the wavetable output is added to some slowly changing value before the lookup process the sound can be made to evolve over time.

Modeling

The newest wrinkle in synthesis is Physical Modeling. In this scheme, the software is designed to recreate the actual behavior of the instrument. For instance to model a flute, you would design a subroutine that solves the equations describing the pressure variance at the blowhole. These equations would be based on (among other things) the breath pressure, the angle of inflow, and the back pressure from the rest of the instrument. That result would become a factor in the second subroutine, which may solve the equation of the resonance of the head joint. The result of that is both passed on to subroutines involving the body of the flute, and fed back to the next iteration of the blowhole subroutine. Results of the body resonance calculation likewise affect the calculations for the head joint. In this way, the entire system is interlinked just like a real flute.

To voice such a model, you specify the inflow angle and breath pressure and the open or closed state of all the keys. Modeling is very demanding of processor time, so the instruments that are just coming to market are mostly monophonic. (I find it interesting that one of the best selling of these models old fashioned analog synthesizers!)

Combination Techniques

Most synthesizers allow you to make layers of tones with a dynamically changing mix. This can create enough spectral variation to get by in complex arrangements. Roland's technique of starting with the attack of a recorded instrument and cross fading to a simple waveform for the sustaining part of the note is very popular. Emu allows the balance between two samples to be controlled by the key velocity, which reduces exact repetition of sounds. Yamaha even combines recorded samples with FM in several of its models.

Voice Level Processing

As on the analog synthesizer, generation of sound is only the beginning of the process. The sound must be shaped, in both amplitude and spectrum. To do this, the digital equivalents of Amplifiers and filters are used. The old terminology of VCA and VCF remains as DCA or TVCA (Digital Controlled Amplifier, Time Variant Amplifier) and so on. These are controlled by other familiar items, Envelope Generators and Low Frequency Oscillators. The more esoteric processors such as Ring Modulators are rare, but do turn up occasionally. There is little computation involved in amplifiers, LFOs or envelopes, so the complexity offered is usually determined by the sophistication of the user interface or marketing considerations. Filters are another story.

Analog filters hung on in synthesizers long after everything else had become digital. That is because a digital filter is computationally complex, and programming them was not well understood. That has changed. Recent models have very clean and flexible filters ranging from simple low pass to the 14 pole morphing filters found in the UltraProteous. Complex filters are usually implemented in a second processor, and require a large section of memory, so they do run up the cost of a machine, but as with analog synths, a good filter can make up for a lot of deficiencies in the sound generator.

Putting all this together, you get an instrument architecture of some sort. This one is typical (JD-990):

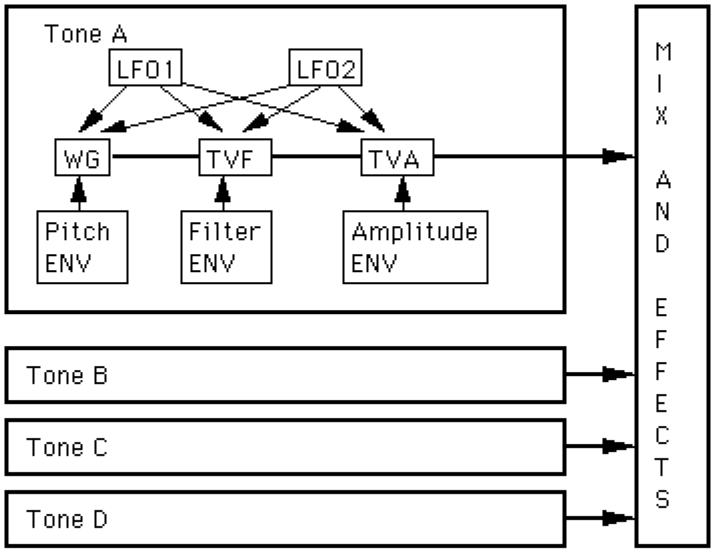

You will note that this architecture is fixed, you can't repatch things. Of course, with this simple setup little repatching is necessary. All reasonable connections exist, and you can turn some off if you don't want them.

Some instruments do allow limited patching- they give you a choice of destinations for various controls, or have a "patch cord" list which matches sources to destinations.

Instrument Level Processing

Marketing directors discovered very early on that including a reverb in a synthesizer would boost sales tremendously. This follows Elsea's first law of music merchandising "The instrument that is easiest to play and sounds best in the music store is the one that sells." A synthesizer sound always benefits from reverb, even if it's not very good. So, almost all instruments now include reverb. And while they were at it, the companies threw in delay and panning and anything else cheap and cute. These effects are applied to the mixed output of all the voices that are playing. If you can afford it, you will want to turn these off and use better quality outboard devices.

MIDI Control

Any digital synthesizer is going to conform to the expectations of MIDI. A few will outperform MIDI when played by their own keyboard, but most have descended to the lowest common denominator of the MIDI standard. Of course, not all instruments implement the same aspects of MIDI. All respond to notes with the proper pitch and most have the potential to react to velocity (although this will change from program to program within the machine), but other commands cannot be taken for granted. Study of the MIDI implementation table in the back of the manual is essential before buying an instrument.

The actual working of MIDI control is pretty straightforward. When a note on message is received, the synthesizer will begin playing a note according to its parameter settings and whatever other MIDI messages have been received to date. Some aspects of the note may be changed by incoming MIDI data as the note plays, and the note will go into its completion cycle (envelope release, etc.) when a note off comes in.

Programs

The sound you get out of a synthesizer is determined by the values used for all of the subroutines in the processing loop. These values are known as parameters, and even a simple instrument will have over a hundred of them. A complete collection of the parameters is known variously as a program, patch or voice, depending on the manufacturer. In any case, programming a synthesizer means adjusting the parameters to get a sound you like. The program can then be stored in a bank of memory and recalled by a MIDI message.

Optionally, programs can be stored on data cards, memory cartridges, cassette tape, floppy disks, or send out via MIDI to another synthesizer or a computer.

According to Elsea's law, all of the memory banks will have programs in them when you buy the machine[1]. Since merchandising studies show that few musicians actually program their synthesizers, many of an instrument's stock programs will be in (cheap) read only memory.

The User Interface.

Most instruments are modal, with three or four distinct operations. When you change modes, most of the buttons take on new meanings (these are often indicated by color coded labels). Typical modes are:

- Play one channel at a time with one voice
- Play a lot of channels with different voices(Multi play)
- Edit the voices
- Edit Multi play setups
- Change the behavior of the instrument.

The last category can include such things as tuning, MIDI channels received and transmitted, maybe the viewing angle of the display. This kind of information is remembered when the power goes off, so anything you change will be a surprise to the next user!

SAVING

Most instruments work in a manner similar to computer applications, where you can call up a document and make changes in it, but the changes will not be permanent until they are saved. Programs are not saved to a disk however, they are saved in a battery operated memory bank. (On some instruments, or with the help of a computer based librarian, entire banks can be saved onto disks.) You don't have to save them in the same place they came from, the save procedure will always give you a chance to pick the destination.

Generally, changing edited programs without saving will result in the loss of all changes.

EDITING

Originally, digital synthesizers had one knob per parameter, just like the analog machines. By the second generation, manufacturers had decided that since you only turn one knob at a time, only one was necessary[2], along with a digital display that told you what you were changing. (So much for unconventional ways of playing.) This means you can't just grab and twiddle, you must navigate a complex maze of menus, pages and submenus, and study the manual to see if what you are about to do will have any effect. Maybe that's why customers quit programming the things.

The displays are typically organized in groups, and each display may contain several items. Here is part of the Yamaha SY-35 setup:

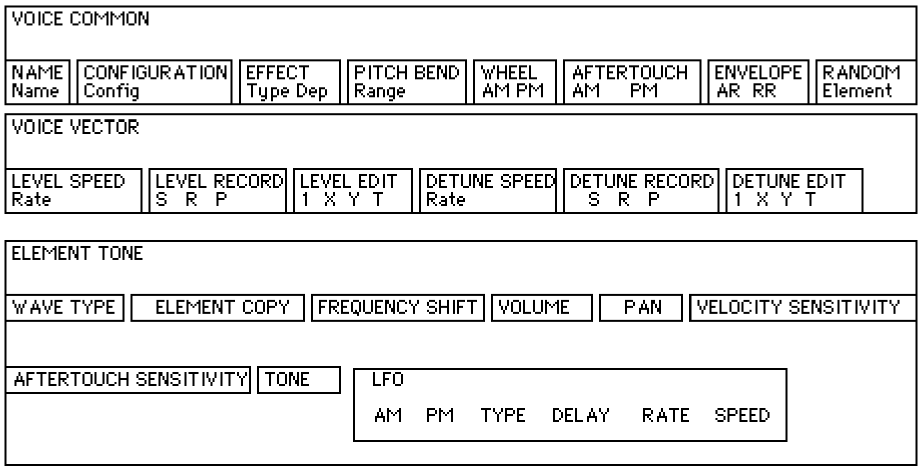

In this case, you get to a group such as Voice Common by pushing a button[3], which leaves one of the displays showing. A display may have one to 10 items in it. Somewhere in the display is a cursor that shows what will change when you press the data increment or decrement buttons. There are other buttons to move the cursor. (You change displays by getting the cursor under the display name and changing that.)

Sometimes, a function may apply to one of several similar parts of the voice. For instance Element Tone: Wave Type display really looks like this:

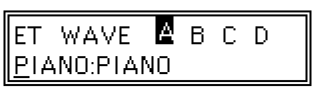

There are four tones to this voice: A,B,C and D. One of those letters in the display will be highlighted, indicating which tone you are adjusting. There are buttons to push to change this.

Sometimes menus will have submenus. On the Yamaha TX-81 you'll see a question like

EDIT EG?

and if you hit the YES button, you will descend one more level. Once you have stepped through all of the options of the envelope generator you will be returned to the upper level.

Relief is in Sight

More recent (and more expensive) instruments have larger displays that behave much like Windows dialogs. There will be a few buttons under the display, and the function of the buttons is indicated in the display itself. Some new machines now feature calculator style keypads for entering data, and there is at least one new instrument with a touch sensitive display, and another that allows you to plug in a PC type keyboard.

Source: http://www.co-bw.com/Audio_Digisynths.htm