

Spring 5-20-2021

Recording Sounds and Instruments for Sample Libraries

Camden Daly

Follow this and additional works at: https://digitalcommons.csumb.edu/caps_thes_all



Part of the [Other Music Commons](#)

Recommended Citation

Daly, Camden, "Recording Sounds and Instruments for Sample Libraries" (2021). *Capstone Projects and Master's Theses*. 1120.

https://digitalcommons.csumb.edu/caps_thes_all/1120

This Capstone Project (Open Access) is brought to you for free and open access by the Capstone Projects and Master's Theses at Digital Commons @ CSUMB. It has been accepted for inclusion in Capstone Projects and Master's Theses by an authorized administrator of Digital Commons @ CSUMB. For more information, please contact digitalcommons@csumb.edu.

Camden Daly

Dr. Lanier Sammons

MPA 475

5/18/21

Recording Sounds and Instruments for Sample Libraries

When audiences sit down to watch a movie on the big screen, the music is sometimes overlooked. While recording sessions for film can involve a large orchestra and live musicians, some tracks are filled in by digital instruments and sampled libraries. While these instruments come from digital software, the origins of these instrument libraries are anything but. These libraries sound so life-like and blend so well because they are sampled from real musicians and real instruments. To make an instrument library, one must take recordings of all the ranges, all the dynamics, and all the various techniques that the instrument can play. It is also important to record the instrument through all the different articulations: legato, staccato, tremolo, tenuto, marcato, etc. For a bowed instrument such as a cello or a violin, more articulations could be included like short harmonics, long harmonics, and pizzicato. All these techniques are worth recording for sampling purposes. The point of having an instrument library is to have full control of the instrument at its performer's fingertips, and the only way that is achieved is through the attention to detail during the recording. Even the smallest details are the difference maker in these libraries; a bow change in a violin, the breath of a flautist, or the release of a piano key have big influences in the feel of the library. These elements are essential to making a library

sound like a real performance. For this capstone project, I will describe the process to sample a grand piano, and then describe how I edited and programmed my own sample library. My own library contains sounds from around my house that I turned into a foley artist library. This library contains a unique user interface, and this instrument is a tool that I will use for my professional and personal projects.

Before going into the process of sampling specific instruments, there are some very important reasons for going through the overarching process of sampling. Many instrument libraries are made and used primarily for film composing. There are plenty of instrument libraries that are used for mainstream music, but I will focus here on instrument libraries used to compose music specifically for media like film, television, and video games. These libraries, when sampled correctly, can be very helpful for two reasons. The first is that these libraries are very good for making adjustments from recording sessions. If the orchestra has interpreted the music differently than the composer had intended, the composer can go back and use these libraries to fill in parts. Even Hans Zimmer, who is arguably one of the best film composers of all time and has transformed film music with his contemporary techniques, has used libraries to fill in parts of film scores to get his interpretation that he has envisioned. It also helps for the compositional process. The purpose of an instrument library is not always to replace the performer. Instead, a composer may use an instrument library to create samples of music for the director. A good instrument library will give the director of the project a good idea of what the music will be like without having to schedule a full recording session (“Sound Palettes”). It will also help the director get a better idea of how the music will work with the screenplay. When a composer has access to a high quality instrument library, it helps to hear all the instrument layers

of the film score before you pay massive amounts of money to use a full orchestra. Sample libraries make creating orchestral music much more accessible.

As sample libraries are often used for film composition, it is becoming more common that composers know how to make a sample library. A sample library does not need to be the most intricate, but a composer should know the skills of sampling, round robin creation, editing, programming, and scripting to make a sample library. Knowing this skill will help to get a job with professional sampling companies, allow a composer to create proprietary libraries for their own work, and help them understand how to get the most from libraries. In Hans Zimmer's online Masterclass, he makes the recommendation to sample your favorite musicians and instruments. This process helps to capture their playing style and almost allows composers to see them playing in front of their eyes while composing ("Sound Palettes"). Capturing the performer's subtle articulation, phrasing, or dynamic changes will really help a composer to know how the emotion will influence the film score.

Another important aspect of using sample libraries is that they are great tools for composers that have smaller budgets or shorter worktimes. They make orchestral music more accessible for composers that don't have access to a full orchestra. A sample library with full orchestras, pianos, percussion, or synths can make a successful film score. Especially some more expensive and high quality sample libraries can make a big difference in the way that a film composition will sound compared to stock instruments.

History of Analog & Digital Samplers

The history of capturing visual and aural media goes back over 150 years. In the 1850s,

technology came about to capture images and audio of our daily lives. The first technology to capture audio was a device called phonautograph in the late 1850s by a man named Edouard-Léon Scott de Martinville (Fabry). He captured a recording of a French folk song, “Au Clair de la Lune.” Interestingly, the records could not be played back until 2008, when audio historian David Giovannoni made the records audible (Fabry). As Giovannoni said, “the phonautograph was really the first machine to record sensory data in real time over time” (Fabry). This was the first machine that was developed for the sole purpose of recording material that humans could hear.

The phonautograph was simply a device to create a visual representation of the sound on screen that is being played. This device did not have any sort of keyboard to control the pitch. One such device that could play back acoustic sounds on a keyboard was the Chamberlin. The first instrument was made by Harry Chamberlin, who named the instrument after himself. This instrument was a keyboard that spanned just under three octaves. Each individual pitch key would be connected to one piece of tape. When the player pressed the key, the roll of tape would start playing, and when the key was released, it rewound the tape to the beginning to prepare the tape for the next time it plays (Crab). Even though this instrument was made in 1951, it still had a very impressive sound palette. Its sound palette included keyboards, harpsichords, organs, brass, winds, voices, bowed strings, plucked strings, and even two separate effects: dixieland band effects and sound effects (Crab).

In Paul Harkins’ book entitled *Digital Sampling*, he describes in great detail the history of digital sampling and playing back recorded sounds, as well as the technology that made that ability possible. In the 1950s and 60s, there was a director working at the behavioral research

laboratory at Bell Telephone Laboratories in New York City by the name of Max Mathews. Max Mathews's work is regarded as one of the biggest historical points in terms of using technology in music. From 1957 to 1968, he was working on solving the question: "How do you take music, which is an auditory medium, and play that back through a computer that uses zeros and ones?" (Harkins 5). Computers only recognize information as a binary code, and music is heard through sound waves. The genius behind Max Mathews's innovation is that he used an IBM 7090 mainframe computer to develop the programming languages *Music I* through *Music V* from 1957 to 1968 so that computers could actually play sound. He also created a technology that created synthesized speech (Matthew). The other genius of Mathews is that he also created DACs, or digital-to-analog converters, and their complement ADCs, which are analog-to-digital converters. This technology could convert audio coming through a cable into sequences of numbers that a computer could read, and convert sequences of numbers into audio that could then be played back through speakers. Once this software and hardware was developed in the 1960s, then the use and practice of using digital technology in music skyrocketed (Harkins 17, 25).

Once the tools were developed where audio could be converted back and forth into computers, the first digital technology of sampling and playing back acoustic instruments emerged. And this first technology was called the Fairlight Computer Musical Instrument Series I and II, or CMI for short, and was developed in 1979. The Fairlight CMI was essentially a computer attached to two 6-octave keyboards. The computer itself had two microprocessors, two 8-inch floppy drives, and a QWERTY keyboard with a monitor, (Harkins 17-18). Despite its older 70s style user interface, it still uses a lot of the same technology that we use today for

sampling. The touchscreen itself came with a pen where the user could adjust the waveforms to change the shape of the note. Samples could either be recorded or loaded into the computer. The Fairlight CMI was advertised with the tagline “Orchestra for sale?” (Harkins 25). The Fairlight was used on several mainstream pop releases such as Stevie Wonder’s *Journey Through “The Secret Life of Plants”* album, Robin Scott’s “Eureka,” Peter Gabirel’s “San Jacinto,” and Britney Spears “Lucky” (Lacedelli). Even music after the 1970s and 80s used the Fairlight CMI to introduce more orchestral sounds into their music.

After the 80s, there was still more development for instrument sampling. In 1996, Steinberg released a new type of sampler instrument. This new type of software is called virtual studio technology, or VST. DAWs were invented in 1977 and entered mainstream use in the 1990s. Over the years, these VSTs that Steinberg developed started becoming a standard plugin format for DAWs. When Steinberg developed this software, they released the technical information for the plugin and that rapidly advanced the recording world (Watanabe). While the Chamberlin and the Fairlight CMI were physical pieces of sampling hardware, the Steinberg VST was the first piece of software for instrument sampling. This started the movement from physical samplers to virtual instruments. Over the last 25 years since the Steinberg VST, virtual instruments have only become more accessible and of a higher quality.

When you listen to music that has been made with real musicians and music that has been made with sampled instruments, the difference between the two can have a very wide range. And that is really depending on the detail of the sampling, the skill of the composer, and the skill of the engineer who is using the sampled instrument. As we go deeper into the details of how to

sample instruments, we will see that the level of detail that you put in exponentially increases the quality of the library.

Sampling a Grand Piano

When it comes to sampling instruments, the instrument that I most want to sample is a grand piano. I am trained as a classical pianist, and I have not found a sample library for a grand piano that has all of the detail that I am looking for. I have played piano for the last 20 years, and I have played on many types of full size and baby grand pianos. I know what response I should feel in the touch and what kinds of frequency responses you should hear from a grand piano. There are also more factors like the overall timbre, resonance, and warmth that can affect the sound of the piano. After so many years of playing piano, I am very picky on the level of detail for a sample library both in the touch and in the acoustics.. There are so many details that separate one note from another, so creating a sampled piano that meets my standards requires finding all of the different details that need to be sampled for a piano.

While initially studying film composition, I learned about the concept of instrument sampling. I also found out how intimately linked sampling is with film composition. When I was working on my first film compositions, I worked with various instrument libraries to record my compositions. And most of my libraries came from Spitfire Audio. Spitfire is based in London, England, and they are known for making very, very high-quality sample libraries for all instruments. They make instrument libraries for strings, horns, brass, percussion, keyboards, and even full orchestras. These sample libraries are not cheap by any means, but they pay attention to so much detail, and the libraries sound fantastic. They have also worked with some very big

names in the film composition industry such as Eric Whitacre, Ólafur Arnalds, the London Contemporary Orchestra, the BBC Symphony Orchestra, and even the great Hans Zimmer (“Spitfire Audio: About”). They have recorded in spaces such as AIR Studios and Abbey Road Studios, both based in London, England. They have worked with these great composers along with many others to sample their favorite sounds and musicians. The two founders of the company are two composers by the names of Paul Thompson and Christian Henson (“Spitfire Audio: About”). In one video, they go through every detail of programming a piano library (“How To Build A Piano Instrument In KONTAKT.”). Some of these steps included mapping out the samples, adjusting the ADSR envelope, velocity table, adding pedal samples, reverb, and release triggers.

Garth Neustadter is an Emmy-Award-winning film composer based in Los Angeles, California who is also involved in creating sample libraries (“Garth Neustadter: About”). Simple Sam Samples has a grand piano plug in that is a sample library of a Steinway & Sons Model D concert grand, and this grand piano was sampled by Neustadter. After purchasing this sample library, I found this was the best grand piano library that I had seen yet.

Steps for Sampling

There are four steps to make a virtual instrument: recording, editing, mixing, and programming. The recording step for sampling is the first and most time consuming step for the entire sampling process. Part of the recording process is the preparation and knowing what to record. There are certain characteristics of a note that one must pay attention to. If you miss recording one characteristic of a note, the sample library will feel very weird when you play it

back. These characteristics change from instrument to instrument, so the recording engineer should understand how the instrument works before the recording session. Samples should be recorded to a grid.. Having the waveforms starting on a specific increment in the grid makes the editing and exporting process much easier later down the road. What helps lock the samples into the grid is recording with a click track. Most of the time, the tempo is around 100 to 120 BPM (Neustadter). Another tool that helps for sampling is recording loops. When you record loops, you can use the playlists to your advantage. In a sampling session, the performer records multiple takes of each individual sample. This will help save time for recording round robins, which I will go into more detail very soon.

From keyboards to strings, brass, percussion, and woodwinds, each instrument will function differently. Some of these functions and characteristics include range, articulations, and timbre. For a piano, there are essentially two parts to sample: the keys and the pedals. The first characteristic to understand is the pitch. A piano has the largest range of any orchestral and classical instrument. There are 88 semitones on a piano, and, ideally, each semitone is sampled. If time or budget is a constraint, it is possible to sample every perfect fourth or fifth. With an interval more than a fifth, then problems arise from basic acoustical physics. If the samples are spread over too far a range, then the sampler will alter the sample being played (Neustadter). While sampler programs are very good, they are not able to stretch the sample over an infinite range without warping the sample. When the sampler alters the audio file, it also alters the timbre and envelope of the audio. If it is pitched too far down, it will play the sample very slowly. The same will happen in the opposite direction. This is why sampling every semitone is very important for the quality of the library.

The second characteristic to sample are the velocity layers, which is a term related to MIDI. Velocity is known as how hard a note is played, which correlates to the dynamics of the note. The harder a key is hit, the louder the note will be. This is probably the most important characteristic to sample. In some instrument libraries from companies such as Spitfire Audio, they have up to 20 velocity layers (Neustadter). The more velocity layers that you have in your sample library, the more accurate the touch will be when you play it back. For Neustadter's piano library, he sampled 10 velocity layers, and this is the number of velocity layers that are possible to play by hand. Any more than 10 velocity layers requires a robot to get the touch exactly right (Neustadter).

When you map out the samples in a library, the velocity levels range from 0 to 127, giving 128 possible values. It is recommended to have at least two or three velocity layers. In Christian Henson's Spring Piano virtual instrument, he uses only two velocity layers, so they divide the 128 velocity values in half ("How To Build A Piano Instrument In KONTAKT"). Even with only two velocity layers, there is programming that can be done to make each transition between velocity layers seamless. Dynamic layers have a value range from 0 to 128. For these two dynamic layers, the values for the lower range are 0 to 64, and the upper range is 65 to 128. These two values from 64 to 65 is the transition point between the two layers. On one side of the transition will play quietly, and the other side will play quieter. If a performer plays consecutive notes between these two values, it will feel very awkward to play the instrument back. This is where programming can be done to make the layers feel more natural during playback. There is programming to add in transitions between each velocity layer, and this helps

greatly during playback. Setting transitions between each layer helps make the touch feel smoother and more like a real instrument.

The next characteristic to pay attention to is the articulations. This is perhaps the most important characteristic to understand before the recording process. Articulations are the way in which a specific note or group of notes should be performed beyond the basics of pitch, duration and dynamic (Dunnet). It covers techniques such as staccato, legato, slurs, and ornaments including trills, appoggiaturas and acciaccaturas (Dunnet). In a way, articulations are like musical pronunciation (Dunnet). These pronunciations change from instrumental family to family. Some instruments will be able to play some articulations, while others may not be able to play those articulations. There are also some articulations that all instruments can play, but that may have different names on different instruments. For example, strings can play a tremolo, which typically involves the performer rapidly repeating a single note, usually in eighth, sixteenth, thirty-second, or sixty-fourth notes. Keyboards usually play this effect between notes an octave apart. A percussionist calls this a roll. Brass instruments have a similar technique called flutter tongue. This is where the performer rolls their tongue, and that rapidly repeats the note, creating a tremolo-like effect. These articulations all have the same musical notation, but the terminology changes from instrument to instrument. Another example is the distinction between woodwinds and brass playing staccato and strings, which can play two variations of short other notes: spiccato and staccato. This is where a recording engineer should understand the function and family of that instrument before going into the sampling session.

The articulations for a piano are very simple. Since the function of a piano is a hammer striking a string, there are two typical articulations that a piano can make: legato and staccato.

There are other, more uncommon ways to play a piano. There are ways to play the piano muted with a finger; there are ways to get harmonics; plucking the strings is possible; and performers even use a violin bow to get a different type of note from the piano. Going to the most common notes, legato is a long note, and a staccato is a short note. When recording legato notes, the lower pitches notes on a piano will take longer to fully ring out. So, it is best practice to start sampling the lower pitches on a piano first, since the higher notes take less time to sample.

The next thing to pay attention to during recording is something that is not directly related to the performance of the note, but more for mixing in the user interface (UI) later on. This factor is the mic placements. In most instrument libraries with classical instruments, there are multiple microphones set up during the recording session. These microphone placements are usually both close to the instrument and far away from the instrument. There can also be any combination of microphone placements, where the microphones incrementally get farther away from the instrument. Microphone placement that is closer to the sound source produces a sound that is more intimate and has more of an immediate attack to the note. Farther away from the instrument, mics get a sound that is more ambient and that will have less initial attack. A control to switch between options or blend them in the user interface comes in very handy. There are some tracks that require more of an ambient sound, and other projects require a sound that is drier. In the Spitfire Audio BBC symphony orchestra library, they have 20 mic placements in the library (*BBC Symphony Orchestra Professional*).

Part of this process of mic placements is choosing the microphones, the orientation, and the polar pattern. After discussing my goals for my sampling session with Garth Neustadter, I will have six different placements for the microphones. The first is the player perspective, which

is a mic underneath the piano facing up towards the soundboard, as well a pair of stereo mics six feet up directly over the performer's head. The next are the close mics, which are a pair of mics six inches over strings: one mic over the treble strings and the other mic over the bass strings. The next mics are the accent mics. This is a stereo pair of mics in an A-B mic orientation, four-to-five feet away from the curve of the piano and five-to-twelve feet off the ground. Ambient mics, which are a pair of mics placed in an ORTF orientation, 10-to-15 feet away, two-to-three feet apart, and eight-to-ten feet high. The last mics are the outriggers, which are two mics 30 feet apart in the sides of the room. The player perspective, inside, closer, and accent mics are all orientations that are cardioid mics directly pointing towards the piano to pick up more of the attack of the sound. The ambient and outriggers will be omnidirectional mics to pick up more of the reverberation, size, and space of the recording studio.

The next part of recording samples is something recording engineers like to call round robins. Round robins are a technique used for instrument sampling to give more authenticity and humanness to the sample library. When a musician plays successive notes on a musical instrument, every repeated note will be slightly different. Human beings are not perfect, and it is nearly impossible to play every note exactly the same. The timbre changes slightly from note to note, and this must be programmed into the sample library. Most high level instrument libraries record six-to-eight rounds of sampling (*BBC Symphony Orchestra Professional*). Every combination of pitch, velocity layer, articulation, and microphone position is rounded several times, and that number is decided by the recording engineer and programmer. This increases the total number of samples dramatically, and this gives a lot more work during the editing, mixing,

and programming stages of the sample library. Though this approach does give more work later for the sampling, this work makes the library sound less robotic during playback.

The last characteristic of the samples to record are the releases. This is the part of the note at the very end when the note is played. There is a very subtle sound as the performer releases the note. This happens for almost every musical instrument but is easier to hear with a piano. As you release the note on a piano, there is a sound that is made as the hammers come back up and make contact with the strings. These sounds are already recorded right after the note is played, so as far as the recording process goes, recording the releases does not take any extra time. It is vital to record, edit, and program the releases, because without the releases, the performer is only playing half of the note. If an experienced musician plays the sample library without the releases programmed, it will feel very awkward to the musician.

When it comes to sampling a piano, there are two halves to the piano. The first is the keys, which we just finished discussing. The second are the pedals. There are three pedals on a piano, and the most used is the sustain pedal. There is a sound that is made as the pedal is pressed and released. When Neustadter sampled his grand piano library, he found a phasing issue with the pedal. The challenge that arises is that there is a natural crossfade between the pedal-up and pedal-down samples (Neustadter). While the pedal is pressed down, the sound from the piano takes longer to decay. If there is an abrupt crossfade between the samples by the pedal up and pedal down, the volume changes wouldn't be natural.

The other challenge is that any time there is a crossfade between samples some artifacts could be produced by interactions between the samples that would be unnatural. Neustadter also said that the resonance that is heard from the pedal is most noticeable close to the piano. This

crossfade starts to make the volume unnatural as the pedal is going up and down. As the mics are placed farther away from the soundboard, the noise from the pedal is less noticeable (Neustadter). The way that Neustadter went around this challenge is to only program the pedal sound from one section of mics. The closer mics will create a more natural resonance of the pedal, whereas the resonance from the distant mics will be more negligible.

After all of the recording, the next step for sampling is the editing. Even though there can be thousands of recording samples to edit and mix, there are ways to speed up the process. The first tip for sampling editing is editing to a grid. If during the recording process, everything was recorded to a grid. Everything will be locked in with a certain time interval between the waveforms, so there are keyboard shortcuts that will split the waveform in a certain interval. This will separate each sample from one another, which will help for exporting very soon.

The next step in the editing process is to eliminate any unwanted sounds from the recording session. There are some sounds that are recorded such as bleed over, squeaky chairs, and extra noise from outside of the recording studio. In Pro Tools, the tool to do this is called Strip Silence. Strip Silence removes sounds below a cutoff level. The cutoff level for Strip Silence is determined by a specific decibel level, and that is only found by listening. When using Strip Silence for samples, there will be more use for Strip Silence in the higher pitches. Higher pitched notes will dissipate quicker, so external noises will be more noticeable in those higher registers.

The next step in editing is to set the start time on each sample. This is a crucial step in the programming stage, but it is also something that can be partially completed in the editing stage.

When you press a note, the note should play immediately. Or rather, it should play just as quickly as on a real piano. In the editing, it is crucial to edit the sample so that the sample starts right at the beginning of the waveform. In Pro Tools, enabling Tab to Transient and pressing tab will move the cursor right to the beginning of the waveform. This makes finding the start of the waveform much easier. However, it is useful to use Tab to Transient primarily when the waveform is very transient-heavy. For a piano, the attack of the note is very pronounced, and the transient is defined. Sounds from bowed strings will have less of a transient, so using Tab to Transient may not be the tool to find the start of the waveform. Once the samples all start at the beginning of the waveform, then it makes a step in the programming stage much easier later on.

Now that the editing is done, we now need to mix the samples. The most important part of mixing samples is noise reduction. Neustadter emphasizes that noise reduction is a necessity before programming. In any audio recording 101 class, every student learns that no matter how recording tools, there will be at least the slightest amount of noise floor in the track. When just one sample is played, the sound from the noise floor is unnoticeable. However, since a piano is a polyphonic instrument, a pianist can play a dozen notes at one time, which triggers dozens of discrete recordings. If multiple notes are played simultaneously, then the amount of sound from the noise floor will become very noticeable.

Two plug-ins that can reduce noise are the Waves X-Noise and the Izotope RX. Both plug-ins are useful in their own ways. By first glance, the Izotope RX may seem like the more efficient option. The Izotope RX has more controls and a graphic display in the user interface. The graphic display helps to see how much noise reduction is being applied to the audio track. Izotope RX also has a very helpful feature called “learn.” When playing the audio while having

this selected, it listens to the audio and automatically adjusts the settings to separate the noise from the good audio. Having this learning feature in the interface is a very helpful starting point for the engineer. The Waves X-Noise does have controls like the threshold, reduction, ADSR, and resolution. It does also have a graphic display to see what levels are being reduced. As for cost, the Izotope RX is much more expensive than the Waves X-Noise, but Izotope delivers all the tools for the price. Out of the two plug-ins, both are great options for mixing. This is of course just a starting point, because an engineer should always trust their ears to make the best mixing decisions possible. There is a balance with using noise reduction. There needs to be enough use of the noise reduction to eliminate extra from the noise floor, but too much reduction will compromise the timbre of the note (Neustadter).

There are also difficulties in noise reduction among different registers. Low registers on a piano are easier to apply noise reduction to without affecting the note itself. There are two types of samples where a mixing engineer needs more time for noise reduction. In higher registers and in softer dynamics, the level between the sample that is desired and the noise floor is much closer. This makes it easier to remove parts of the desired sample while taking out the noise floor. A mixing engineer will need to take more time in these frequencies in order to keep the timbre of the notes.

The next step of mixing is to balance the levels. During the recording process, higher notes will have a tendency to play quieter than low notes. For example, if you play the lowest A on a piano pianissimo, and then play the highest C pianissimo, the high C will still play slightly softer than the low A. Even with a talented performer playing what they think as pianissimo, the high notes will still play slightly quieter than a low note with that same dynamic. This is where

during mixing, it is important to balance the volume levels across all the pitches and all dynamic layers. For each dynamic layer, each pitch across the entire instrument should have about the same volume. There is a difference between what we perceive as loudness versus a measured amplitude. Our ears are a very subjective way of measuring dynamics. We can think that one sound is louder than another, but because our ears are so subjective, the sound could be at the same amplitude. The amplitudes between each dynamic layer should not be completely flat, because then the instrument will not feel natural.

The last step of mixing is exporting and labeling all of the samples. During the exporting process, all of the samples are already separated from each other, and exporting the samples is a simple step in most DAWs. Exporting hundreds, even thousands of samples can take a good amount of time, but using key commands can save lots of time. After the key command, then you can label the sample. The labeling process is also very important, not just for finding the samples, but also for programming the samples into the library later on. For labeling samples, it is most helpful to name the samples by library name, instrument, dynamic layer, and pitch with octave number (“How To Build A Piano Instrument In KONTAKT”). In Christian Henson’s video on programming a piano library, his approach produces names of samples such as “WPP M149 pp A1.” The WPP M149 is an organizational system that Spitfire Audio uses to organize their libraries. The lowercase “pp” is the dynamic layer, which stands for pianissimo. P would be for piano, F for forte, FF for fortissimo, and so on. The A1 is the pitch. For a piano, the pitches will range from A0 all the way up to C8 with the number indicating the octave. It is very important to label the pitch keys correctly because in our sampler, it will look for the pitch key in the name to center the key (“How To Build A Piano Instrument In KONTAKT”).

Now we get to the programming stage. This is where the tips and tricks in the editing and mixing process will pay off. The first thing to do with programming is to pick a software. All of the plug-ins mentioned below can be used for any DAW, such as Pro Tools, Ableton Live, Logic, Cubase, Reaper, and others. One option for using the samples is the Structure plug-in. Structure is a plug-in that comes free with Pro Tools. Structure has a very simple user interface (UI). The user can drop in samples into the piano keys to associate the sample with a specific key. Structure has the ability to transpose the samples by octaves, a semitone, or to fine tune the sample by cents. It also can bend the pitch of the notes, add polyphony, and adjust the available range of the instrument (*Structure: User Guide*). Overall it is a good plug-in and it does its job, but does not offer detailed UI options.

Sforzando is another option. It has the piano keyboard along the bottom of the user interface, as well as an area above to use during programming. It is a good program where you can put in samples and has a few more functions than Structure, but like Structure, it does not have anywhere to put in code or scripting for more advanced features in the user interface.

Logic EX24 is another sampling program. The UI for the Logic EX24 is more complicated than the other two plug-ins. The appearance for the EX24 looks more like a stage mixing console. It has many more knobs and grids for loading in samples, but it is more difficult to understand what each part of the UI is at first glance. It is not the most welcoming for a first glance.

The last plug-in on this list is Kontakt by Native Instruments. The Kontakt player is the most used among professional sampling companies. Companies such as Spitfire Audio release sample libraries in Kontakt. Kontakt is perhaps the most dense program to learn but the most

versatile at the same time. When opening Kontakt, it is purely a blank slate. The UI for the library is made by the engineer. There are two downsides of Kontakt. The first is that it is not cheap. There is a free version, but most sample libraries in Kontakt only play for 10 minutes before stopping in the free version. To play the libraries without any stoppage, you must purchase the full version of Kontakt, which costs \$400 ("Samplers: Kontakt 6 | Komplete"). The second is that there is a big learning process for using Kontakt, and that is with the scripting. Making the interface in Kontakt requires using the script editor and coding to make the parts of the interface. While Kontakt might be one of the most complicated samplers in DAWs, it is definitely worth the time of learning the software. Learning the mapping editor, wave editor, scripting editor, and the coding language helps tremendously in making a library that sounds and feels realistic, as well as looking very professional and welcoming to the user.

To put this programming into practice, I created a sample library of Foley sounds. Foley artistry is the practice in film and TV of taking everyday sound effects and dramatizing them to create more interest for the audience. I recorded my own Foley sounds to program into Kontakt so that I can create my own Foley artist projects. For this library, I compiled 11 objects from my house and sampled them. The two biggest groups will be pitched and unpitched sounds. Each one of the pitched sounds will need a subgroup. For the pitched sounds, there are three different sounds that I turned into pitched instruments: a PVC pipe, a wine glass, and a wine bottle opening. There are three sounds for the wine glass: rubbing on the rim, tapping with a finger, and tapping with a pen. All these objects are percussion-like sounds. For the library, I recorded four velocity layers for each sound. To load in the samples, I loaded in each dynamic layer one at a time. There are 128 values for dynamic layering, so four dynamic layers break down each

velocity layer to 32 values. The first layer will be from 0 to 31, then 32 to 63, 64 to 95, then finally 96 to 127. In Kontakt, organization of types of samples is established by “grouping.” Grouping can be used for many different purposes, depending on how the programmer decides to use the grouping. Groupings can be used for controlling mic positions, as well as being able to select what sound or articulation to play at one moment. In Neustadter’s piano library, he creates groups and routes those groups to volume knobs. In Spitfire Audio’s library, they use groupings to separate articulations for each instrument. In the Spitfire Audio BBC Symphony Orchestra, they also use keyswitches to flip between articulations: legato, staccato, marcato, tremolo, etc. They load in articulation in separate groups, and the key switches allow the engineer to press a note on the MIDI controller to flip between articulations. This helps save processing power by reducing the number of tracks during the recording session.

For my library, I am making groups and subgroups for the articulations. For the unpitched sounds, I programmed them from lowest pitch to highest pitch. For any sound that we call “unpitched,” we say that because there isn’t a specific fundamental frequency from that sound, meaning that it doesn’t have a note that you can match on a piano keyboard for example. For example, a snare drum can be repitched, but it won’t be heard as a middle C or any other pitch. Any recorded sound can be stretched from one single hit, but it won’t sound very musical or pleasant to play.

There are more objects sampled in this group, and a few sounds have two-to-four articulations. These unpitched sounds were laid out in one group on one map. These sounds make up a unique type of drum kit; so sounds like a low thump on a desk will be on the lower end of the keyboard, and taps on metal and glass objects will be higher on the keyboard. These

taps sound like a hi hat, and the sound of candy pouring in a jar sounds like the ride of a drum set. These sounds will be helpful in the future for making either Foley sounds or a drum pattern for music.

The first step for loading in the samples is to pull the samples into the mapping editor. The mapping editor is a graphic of a keyboard with a grid above (see fig. 1). Along the horizontal axis are the pitches. For pitched instruments, as you go farther left, frequency decreases. As you go farther right, it will play higher frequencies. On the vertical axis, this controls the dynamics and velocity layers. As samples are loaded into the mapping editor, samples must have the correct values for the velocity layers. As above, the values for the four dynamic layers are 0 to 31, then 32 to 63, 64 to 95, then finally 96 to 127. 0 to 31 is the quietest layer and 96 to 127 is



fig. 1

the loudest layer. Changing the range of the velocity layer will adjust the height of the samples in the editor, which then allows the samples to be played with a specific pressure on the note. The samples with the loudest dynamics will have the highest range of velocity values, which means that you can only play the loudest samples when you hit the keyboard the hardest. Conversely, when you press the keyboard very softly, it will play very quietly. And this is exactly what we want for our virtual instrument.

After setting the velocity layers, we then need to set the root key. Setting the root key centers the pitch of the note with the corresponding key on the mapping editor. Without centering the root key, the note will play at a different speed causing the pitch to be transposed. For some sounds that I transformed into pitched instruments, I started with one single pitch. I imported the sample onto the note that makes the true pitch of the sound. Once the pitch is stretched across the mapping editor, it will help if the true pitch of the sound is matched as best as possible to the note on the piano. This process is then repeated for any sound that is desired to be turned into a pitched instrument. There is a limit on how far a single pitch can be stretched across the mapping editor. A rule of thumb is a perfect fourth or fifth, but various instruments and different timbres will make this range change slightly.

After the work in the mapping editor, all the samples are loaded into the library. The next step is to go into the wave editor. Since I went through the specificities of recording and editing the samples to a grid, all the samples will have the same gaps between the start of the .wav file and the waveform. While in wave editor, this is where the start times can be set (see fig. 2). This takes a bit of trial and error, but any musician can tell whether the sample is



fig. 2

playing too quick or too slow. Once the number of desired samples for the correct touch is found, then this number can be applied to all the samples. To double check the touch of the library, it is helpful to play the library at various dynamics to check all the dynamic layers.

The last step of programming is scripting. This is the most complex part of the programming, and it is the unique part of programming. This is where the user interface is created, along with inserting faders, knobs, reverb controls, adding in backgrounds, and many other features in the UI. This step of programming has two essential purposes. The first is that programming the UI gives the engineer or composer more control over the timbre of the instrument. The second, which may just be more important, is for aesthetics. The UI should look welcoming as soon as the user opens it. If a composer sees a library that is very professional looking, it gives him or her more inspiration to use the library. This then brings more creativity for the composer or performer.

The first part of making the user interface is making the background. To make the background, the code has a few crucial parts. The first part of the code sets the height and width of the background in pixels, and this is determined by the pixel size of the image that you insert. The last line of this code shows the library where to look for the image file on the computer that makes up the background.

The next part of coding the library is the knob graphics for the reverb. The graphics for the knobs were made in a Japanese program called KnobMan (“Kontakt Tutorial: How to Make Knobs with Custom Graphics”). What Kontakt needs is a picture for each position that a knob can be in. In KnobMan, it is very easy to make a .jpeg file with the knob positions. This .jpeg will have a vertical line of knob graphics, and each knob will have more of the dial filled in than

the last. In the graphics that I made for my library, I created a .jpeg with 20 dial positions, which is sufficient to make the animations in the library smooth. While my knobs and graphics are rather simple, KnobMan can make some very complex and professional looking graphic designs.

These knob graphics are for the reverb control. To make the reverb control, the first thing to do is to put a reverb insert as one of the send effects. There are many different effects in Kontakt, including delays, EQs, and modulation. To put in a reverb, I selected a convolution reverb. After inserting an effect, a factory preset can be selected. There are many presets to select such as piano reverbs, drum reverbs, vocal reverbs, various room sizes, and plate reverbs. For this Foley library, I selected the preset 01 Real Rooms > Music Studio A. For a piano library, I would select a piano reverb as one of the presets, but that creates a reverb too big for a Foley library. For the scripting, there are two things to do. The first bit of script is to tell Kontakt to establish a knob for reverb. The second step is to establish where on the UI the knob is going to be placed. This bit of code also tells Kontakt what effect this knob is controlling, which is the reverb insert that I put in. This also will establish what the text on the knob will be, which most people put as “reverb.”

In many instrument libraries, including orchestral and piano libraries, there are controls for the microphone placements. This will help control the timbre of the instrument. This is where the groups come into play. For this library, I used the groups to separate each sound and object that was sampled. For a classical instrument such as a piano, each group will have the samples from one microphone position. To make a slider, we use scripting that says “declare ui slider.” This means make a fader on the user interface. On another line of code on the bottom of the script editor, we route the buses into the fader. This way when the faders move up and down, this

will turn up or down the volume of the groups. For a piano library, this is the type of coding that I would use. For a Foley library, I routed the groups to buttons. These buttons effectively work as switches and give the ability to select one group at a time. This Foley library is not a library where I want to play multiple sounds at once, unlike a piano library where I want to hear multiple microphone positions simultaneously.

For the scripting of these buttons, there are two parts of the code. The first part of the scripting is to establish the button. The coding “declare ui button” creates the button on the user interface, then this section of code creates the name of the control. The second part of this code “on ui control” must be placed at the end of the scripting and must be placed in between lines of code that read “end on.” This section of code is similar to the function of the reverb control. This line of code establishes the order of the buttons in the code. If there are multiple faders or knobs, this where Kontakt needs to see a numbering system for the controls. This ordering is “mic1button”, and the button would be replaced with knob or slider, depending on what is preferred on the UI. This section of code also routes the buses of the groups into the button. There are 16 buses for the groups, and Kontakt reads the values for the 16 buses from 0 to 15. So the first bus would be route 0; bus number two will be route 1; and so on. This is another spot where it is essential to pay attention to what bus is assigned to what group. If the group numbers are not routed correctly, then the buttons, faders, or knobs will not control the volume of the group.

There is one more part of the script that is important, and that is changing the units for the volume. When a control is put on the UI, the units of the volume don't read as decibels; they just read as a value from 0 to infinity. While it will get louder and quieter as you move the control, it still doesn't look right to users if the values aren't in decibels. To make the control read in

decibels, the code that needs to be typed is “Knob Unit DB.” This will then make the numbers read from negative infinity decibels to positive 12 decibels. This is normally how mixing engineers read values when mixing.

In summary, the steps for sampling any instrument start with understanding the articulations of the instrument. The next step is to record those articulations at multiple dynamic layers with microphone positions to control the timbre and space of the instrument. Then, the final steps are to edit and mix the samples and then program the samples into a software for playback. Overall, I’m very pleased with how the library turned out. Regardless of the instrument that is sampled, the recording, editing, mixing, and programming processes are exactly the same. While my library may not be the most professional looking, I am still pleased with the way that it sounds. There are an infinite number of sounds that can be added to a library, so it can constantly be updated. For the aesthetics and looks of the user interface, that is a matter of small tinkering. It is a matter of how the engineer would like the interface to look. This included what background to have, where the controls are placed, and what affects the virtual instrument will have. All of this just takes time to perfect the look and function of the library. More expensive libraries like those from Spitfire Audio take months to sample the instrument, and even more months to edit, mix, and program their instrument libraries to make their libraries achieve perfection.

Works Cited

BBC Symphony Orchestra Professional. Version 1.2.0. Spitfire Audio. Software.

Crab, Simon. "The 'Chamberlin', Harry Chamberlin, USA, 1951." *120 Years of Electronic Music*, 120years.net/the-chamberlin-harry-chamberlin-usa-1951/

Crofts, Mike. "Miking Up a Concert Piano with DPA." *SoundOnSound*. Nov. 2018.
www.soundonsound.com/techniques/miking-piano-concert-dpa

"David Fray records Bach concertos for 2,3,4 pianos" *YouTube*, uploaded by Warner Classics. 31 Oct. 2018. www.youtube.com/watch?v=c0f-e1jZ100

Dunnet, Ben. "Articulations." *Music Theory Academy*.
www.musictheoryacademy.com/how-to-read-sheet-music/articulation/.

Fabry, Merrill. "What Was the First Sound Ever Recorded by a Machine?" *Time Magazine*. 1 May 2018. www.time.com/5084599/first-recorded-sound/

"Garth Neustadter: About." *Garth Neustadter*. www.garthneustadter.com/#!/about.html

Harkins, Paul. *Digital Sampling: Design and Use of Music Technology*. New York, NY:

Routledge, 2020.

“How To Build A Piano Instrument In KONTAKT.” *YouTube*, uploaded by Christian Henson, 2 Jun. 2019, www.youtube.com/watch?v=b9mqaby0Axs&list=PLPb9C4eRs8MgEVi8GXYpwadIwwRJmocQN

“How To Make A Sample Library.” *YouTube*, uploaded by Spitfire Audio. 25 Sep 2016. www.youtube.com/watch?v=Yvco_DWMV_g&list=PLPb9C4eRs8MgEVi8GXYpwadIwwRJmocQN&index=3

“How to Mic and Record a Grand Piano.” *DPAMicrophones*. 1 Nov. 2015. www.dpamicrophones.com/mic-university/how-to-mic-a-grand-piano.

“How To Sample A Piano (in 2 minutes 10 seconds) - PB001.” *YouTube*, uploaded by Christian Henson. 21, Dec. 2018. <https://www.youtube.com/watch?v=GUAq5R4eAVE>

Hugill, Andrew. “The Digital Musician.” Taylor & Francis Group. 2012. *ProQuest EBook Central*. ebookcentral.proquest.com/lib/csumb/reader.action?docID=981765

“Is this the perfect Kontakt starter template? (Also: How to swap effects and add round robins).”

YouTube, uploaded by David Hilowitz Music. 25 Apr. 2020.

www.youtube.com/watch?v=swaoGvPY2mM&list=PLPb9C4eRs8MgEVi8GXYpwadIw wRJmocQN&index=5&t=18s

“Kontakt Tutorial: How to Make Knobs with Custom Graphics + FREE PVC Tube Kontakt Library” *Youtube*, uploaded by David Hilowitz Music. 2 Jan. 2019. www.youtube.com/watch?v=hVp07BaFUyI

“Kontakt Tutorial: How to Make Mic Level Knobs (using Instrument Busses) + FREE KONTAKT SAMPLE.” *YouTube*, uploaded by David Hilowitz Music. 11 Jan. 2020. www.youtube.com/watch?v=S0cYRHhML60

Lacedelli, Stefania Zardini. “The Fairlight CMI: The Secret Composer of the Music You Love.” *Science and Media Museum*. 22 Jul 2020. blog.scienceandmediamuseum.org.uk/fairlight-cmi-playlist/

Leavell, Chuck. “Miking a Grand Piano.” *Sterling Audio*. 2015, sterlingaudio.net/miking-grand-piano/

Mantione, Philip. “The Fundamentals of Sampling Instruments and Libraries” *Pro Audio Files*. 2 Feb. 2018. theproaudiofiles.com/fundamentals-sampling-and-instruments-libraries

Matthew, Terry. "The Man Who Made His Computer Sing: Max Mathews & the Invention of Digital Music." *5MagNet*, 4 Jun. 2019, 5mag.net/features/max-mathews-inventor-digital-music/

McGuire, Sam, and Roy Pritts. *Audio Sampling: A Practical Guide*. Taylor & Francis Group, 2008. *O'Reilly*, learning.oreilly.com/library/view/audio-sampling/9780240520735/01_Title.html

Martin, Dave. "How to Record a Piano." *Sweetwater Audio*. 19 Nov. 2019. <https://www.sweetwater.com/insync/how-to-record-a-piano/>

Neustadter, Garth. Personal Interview. 9 Sep 2020.

"Part 2: Building A Sample Instrument." *YouTube*, uploaded by Spitfire Audio. 20 Jun. 2020. www.youtube.com/watch?v=bY8BwXT8Afs&list=PLPb9C4eRs8MgEVi8GXYpwadIwwRJmocQN&index=2

Pitman, Toby. "Introduction To Scripting in Kontakt, Part 1." *AskAudio*. 1 Oct. 2012. ask.audio/articles/introduction-to-scripting-in-kontakt-part-1

Rogers, Neil. "Ambient Mic Placement." *SoundOnSound*. Feb. 2015. www.soundonsound.com/techniques/ambient-miking-room-improvement

Rosenberg, Nathan. "Miking the Grand Piano." *Recording Magazine*. Music Maker Publications.

recordingmag.com/resources/recording-info/mics-miking/miking-the-grand-piano/

Russ, Martin. *Sound Synthesis and Sampling*. Elsevier/Focal Press, 2004.

learning.oreilly.com/library/view/sound-synthesis-and/9780240516929/xhtml/17_chapter4.xhtml

"Samplers: Kontakt 6 | Komplete." *Native Instruments*. [www.native-](http://www.native-instruments.com/en/products/komplete/samplers/kontakt-6/)

[instruments.com/en/products/komplete/samplers/kontakt-6/](http://www.native-instruments.com/en/products/komplete/samplers/kontakt-6/).

"Sound Palettes." *Hans Zimmer: Film Scoring*. Season 1, episode 7, *Masterclass*, 8 Sep 2017.

www.masterclass.com/classes/hans-zimmer-teaches-film-scoring/chapters/sound-palettes?action=preview&controller=chapters&course_id=hans-zimmer-teaches-film-scoring&id=sound-palettes&logged_in=true

"Spitfire Audio: About." *Spitfire Audio*. <https://www.spitfireaudio.com/about/>

Structure: User Guide. Air Music Technology, [https://aed625a7d6b8aa15a975-2b](https://aed625a7d6b8aa15a975-2bf75360321f1a0597a021ce2af6519b.ssl.cf2.rackcdn.com/743/documents/AIR_Music_Tech_-_Structure_v2_-_User%20_Guide.pdf)

[f75360321f1a0597a021ce2af6519b.ssl.cf2.rackcdn.com/743/documents/AIR_Music_Tech_-_Structure_v2_-_User%20_Guide.pdf](https://aed625a7d6b8aa15a975-2bf75360321f1a0597a021ce2af6519b.ssl.cf2.rackcdn.com/743/documents/AIR_Music_Tech_-_Structure_v2_-_User%20_Guide.pdf).

Watanabe, Yuya. "Insights into VST and its virtual instruments." *Steinberg*. 17 Nov. 2015.

www.steinberg.net/en/products/vst/vst_story.html