


Some thoughts on **Octatrack**



by Merlin

Table of Contents

1. Introduction.....	2
2. Samples	3
2.1 Streaming	3
2.2 RAM memory	3
3. The sample slot lists	5
3.1 Sample editing	5
3.2 Advantages of the sample slot list.....	6
4. The flex machine	7
4.1 The flex machine and slices.....	7
4.2 The flex machine and sample locks	7
4.3 Slice locks	10
4.4 Summary and overview	11
5. Effects	12
6. Scenes.....	13
6.1 Summary and overview	14
7. Parts	15
7.1 Parts and what they store	15
7.2 Part reloading.....	16
7.3 Part reloading: Usage	18
7.4 Parts and patterns	18
7.5 Scene stacking	21
7.6 Using multiple parts	22
7.7 A small but convenient enhancement	26
7.7.1 Parts store track volumes.....	26
7.7.2 Parts store scenes.....	26
8. Volume handling.....	28
8.1 Signal flow	28
8.2 Volume and scenes	28
8.3 Correcting sample volume	29
9. Sampling.....	31
9.1 The OT's approach to sampling	31
9.2 Signal flow	32
9.2.1 Main outs	33
9.2.2 Track recorders.....	33
9.3 Sampling guidelines	34
9.4 Track recorders and the sample slot list	35
9.5 Sampling techniques	36
9.5.1 Dedicating patterns	36
9.5.2 Abusing the cue outputs.....	36
9.5.3 Abusing a track for live input.....	36

1. Introduction

This article provides an analysis of the Elektron Octatrack (OT) from a user's point of view. The aim is to look at the OT from the ground up, to explain how things work, and to develop an easy-to-understand overview of the OT's structure. Further on, several scenarios will be described in which I show/suggest how to approach the OT for common tasks.

As we will see, understanding just a couple of key concepts greatly reduces the fog and confusion which seems to haunt many owners of the OT.

Some notes:

- This article has been written in a cumulative style. If you read something on page X, I implicitly assume that everything I covered before that is known. You have been warned.
- Some information presented in this article will be experienced as "obvious" or "nothing new". This is inevitable.
- Writing this article was not exactly hard. The OT is what it currently is and it is fairly easy to write a story around things that already exist.
- I do not work for Elektron and have no insight into their technology. This article is written from a user's point of view. Anyone on the Elektron R&D team who reads this is probably in for a good laugh ;-)
- This document is an ongoing effort. Feel free to suggest and correct.

Why did I write this document? Well... having obtained the OT a few months ago, I experienced a vast overkill of its possibilities. Merging all these possibilities into a productive workflow was difficult. Therefore, I decided to write down my thoughts as this forced me to express them in an understandable way.

By doing so, this document grew as time passed, while my ability to integrate the OT's functionality into a good workflow grew with it. I therefore have to admit that I wrote this document for selfish reasons. Adapting it to a format that could be given away to the Elektron-Users forum was little extra work, however. So, if this helps people to unleash the OT, I don't see any reason for not sharing it...

2. Samples

Before diving into the OT itself, it is wise to spend some time discussing samples first. Basically, we can distinguish between two types of samples:

- Short samples (single cycle waveforms, single shot samples, and reasonably small loops), and
- Long samples that may take many minutes.

Although the above distinction is obvious, it is nonetheless important to mention as it has consequences for the way in which the OT handles them.

2.1 Streaming

Ideally, the OT should be able to handle all kinds of samples, both small and long. However, if you allow people to use extremely long samples, you would have to equip the OT with an enormous amount of memory. The solution for this is to stream long files directly from the CF-card. Although streaming is a reliable technique, it has its disadvantages.

The OT has the ability to play 8 samples simultaneously. So, if you load 8 static machines to a pattern and let them play at the same time, you are effectively sucking 8 streams of CD-quality audio from the CF-card. This puts a heavy burden on the CF-card and by doing so, you can run into timing problems if the card is cheap and not fast enough.

Therefore, although streaming can be used, my personal attitude is to avoid it as much as possible, just to be on the safe side. On the other hand, Elektron offers streaming as a valid technique, so if you use it, it should work as advertised provided that your CF-card can handle the speed.

2.2 RAM memory

An alternative to streaming is this: Equip the OT with a fair amount of internal memory. Once the OT boots up, a project is automatically loaded and a small collection of samples is read once from the CF-card and put in RAM memory.

Using the RAM memory provides several advantages. The samples are read once so no further streaming is needed. Because the samples are loaded in RAM memory, the further handling of those samples is independent of the CF-card and demands for speed and processing can be increased: Reading from and writing to RAM is simply faster than streaming from a CF-card.

The disadvantage is of course that the amount of RAM memory ultimately limits the amount of audio that can be made available to the rest of the system. Elektron equipped the OT with 80MB of memory, which corresponds roughly to 8 minutes of CD-quality audio. 80MB... hmm... doesn't seem like much!

Conclusion: Whatever drove Elektron to use only 80MB, the emphasis is clearly not on *quantity* (i.e., layering and excessive polyphony), but on *quality* and doing some serious stuff with a minimum of samples and memory.

3. The sample slot lists

In the previous section, we saw that Elektron expects you to select a relatively small number of samples and squeeze the most out of those.

Selecting a small number of samples naturally leads to the need for some kind of list in which we put the samples we want to use. Elektron provides two lists for doing this: The flex list and the static list. Since both work the same way, I will only describe the flex list, but keep in mind that the static list is used for long samples that are streamed, while the flex list is used for shorter samples that are put into RAM memory.

The flex list can be filled with up to 128 samples. When filling the flex list with samples, we therefore face two conditions:

- We cannot use more than 128 samples in our project.
- Whether we use 128 samples or less, the total sum of those samples should not exceed 80MB.

3.1 Sample editing

Selecting samples from the CF-card and putting them in a list is all fine, but we want more than that. We want very precise control over the way the sample is played. We want to “tailor” our samples so they behave the way we want. Therefore, several features should be offered:

- Defining start and end points of the sample.
- Defining loop points of the sample.
- Defining whether the sample should loop and how it should loop (repetitive start to end looping, or ping-ponging from start to end and back, etc.).
- Slicing: the ability to cut up a sample into several slices, allowing for ridiculous ways of playing back those individual slices.

All the above operations are carried out in the sample editing menu. Once you have sliced a sample or defined loop points of whatever, the decisions you made are stored *in the sample slot list*. This is an important concept; the above list of operations is stored with the sample in the sample slot list and has nothing to do with the flex machines that eventually play the sample.

This means that you can add a sample to the sample slot list, slice it, trim it, and so on without assigning it to a flex machine. You could even do it without using the sample anywhere in your project. This would obviously be pointless, but it clearly illustrates that the sample slot list lives on its own and exists throughout the whole project, whether flex machines use anything from that list or not.

To summarize, you select samples from the CF-card and put them in a list. Once a sample is in the list, you define the basic properties of that sample and thereby determine how the rest of the system should handle it.

3.2 Advantages of the sample slot list

Putting samples in a slot list and defining some basic ways the rest of the system should handle the sample has several advantages:

- **Speed:** You prepare a sample by defining and slicing it the way you want and once you have done that, the sample and its settings are available in any track holding a flex machine in any pattern in any bank throughout the project. Basically, you do the hard work once and re-use it wherever you need it.
- **Flexibility:** The same sample can be put in the list multiple times. This allows for handling the same sample in completely different ways. For example, put some kind of loop in slot 1 and create 16 slices. Now select the same sample and put it in slot 2. You can now slice it up to 64 slices, choose different start and end points, etc. Load the same sample in slot 3 and you can again make different choices. You have now defined three different ways of handling the same sample.
- **Adaptivity:** Because creating music is a highly emotional and artistic process, the sample slot list does not need to be entirely filled straight from the start. As your project progresses, more samples can be injected into the list, and unused samples can be removed as time passes. Just pay attention to the conditions described earlier: no more than 128 samples, eating no more than 80MB of memory.

4. The flex machine

So far, we have seen that samples are drawn from the CF-card and stashed into a list. Each of these samples can be prepared for further use by using the operations in the audio editing menu.

For playing samples from the flex list, the flex machine is needed. Once you place a flex machine on a track, the OT asks you to select a sample from the list. At this point you can select an already existing sample or you can select a free sample slot after which the OT allows you to select a sample from the CF-card. The selected sample is then added to the list and assigned to the flex machine.

An important thing to understand here is that Elektron provides you with a user interface in which adding samples to the sample list and assigning samples to a flex machine go hand-in-hand. Keep in mind, however, that these are two separate procedures.

Now that a sample is assigned to a flex machine, you can trigger it by programming triggers on the sequencer of that track or by live recording those triggers.

4.1 The flex machine and slices

Assigning a sample to a flex machine seems easy and straightforward, but there are a couple of details to keep in mind. The samples that live in the sample list can be prepared as described earlier. This means that some samples will be sliced, while others will not. This has consequences for the way the flex machines handle them.

If a sample is sliced, you must instruct the flex machine to use those slices. Dive into the playback setup menu of the flex machine and set the slice option to `ON`. With this setting, the start parameter on the playback page can now be used to select which slice should be played.

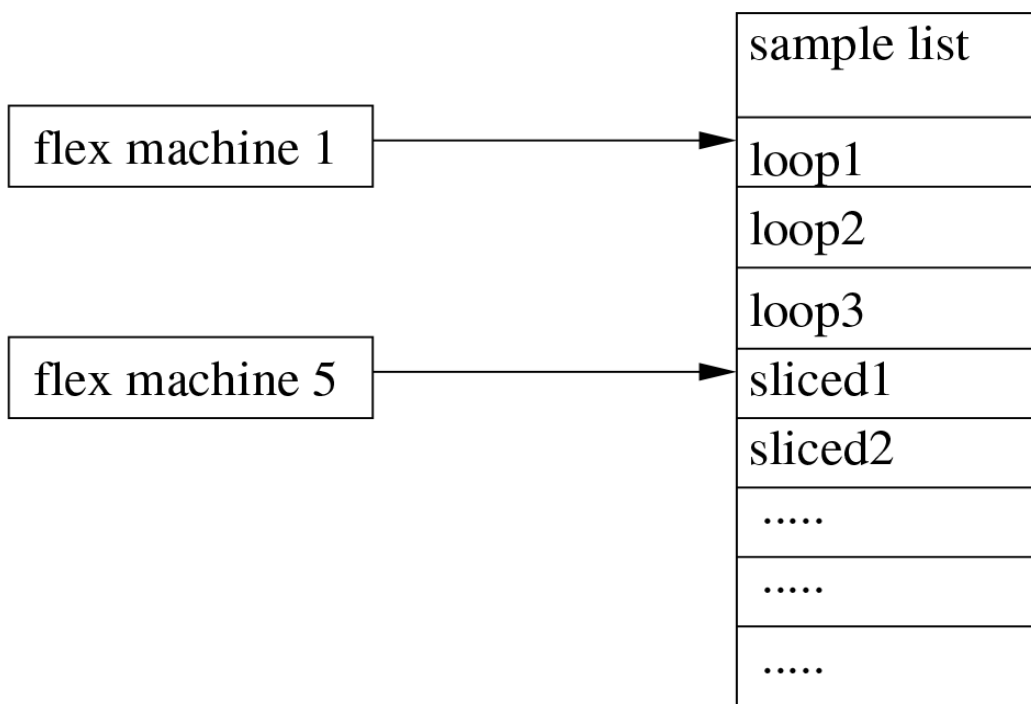
The bottom line: Because samples can be prepared in different ways, the flex machines have to be adapted/tailored to make use of what was prepared. The features you have in the sample edit menu form the basis on which flex machines operate.

4.2 The flex machine and sample locks

Before reading further, I assume that you understand the concept of parameter locks. Parameter locks form a very powerful feature when used on the flex machines. To explain their power, I will turn to a simple example. Assume that:

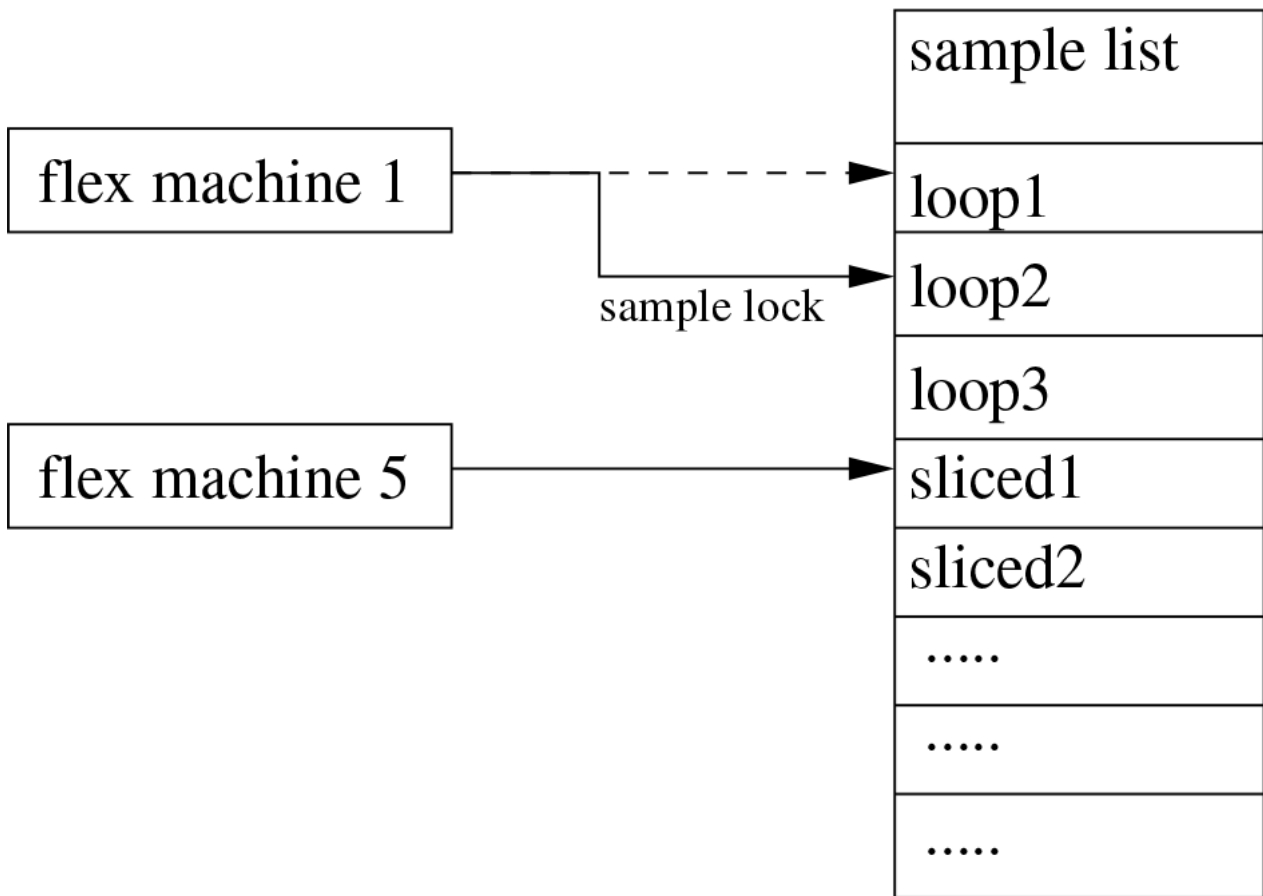
- The sample list contains three loop samples, simply called loop1, loop2, and loop3.
- The sample list also contains two loop samples which have been sliced. Let's call them sliced1 and sliced2.
- A flex machine lives on track one and has been assigned to loop1.
- Another flex machine lives on track five and was assigned to sliced1.

Schematically speaking, the pattern was set up like this:



This is a straightforward setup. Let's look at flex machine 1. Because it plays a simple loop, the sequencer only needs one trigger to play it and that trigger is placed on the first step. Nothing wrong here, but suppose I like to play a different sample. You can now hold the trigger and turn the level knob. The OT presents you with the sample slot list and you can select any of the samples in that list. Suppose I select loop2. The result is that, although loop1 was (and is) assigned to the flex machine, I have overridden this by using a sample lock on that step.

Schematically speaking:

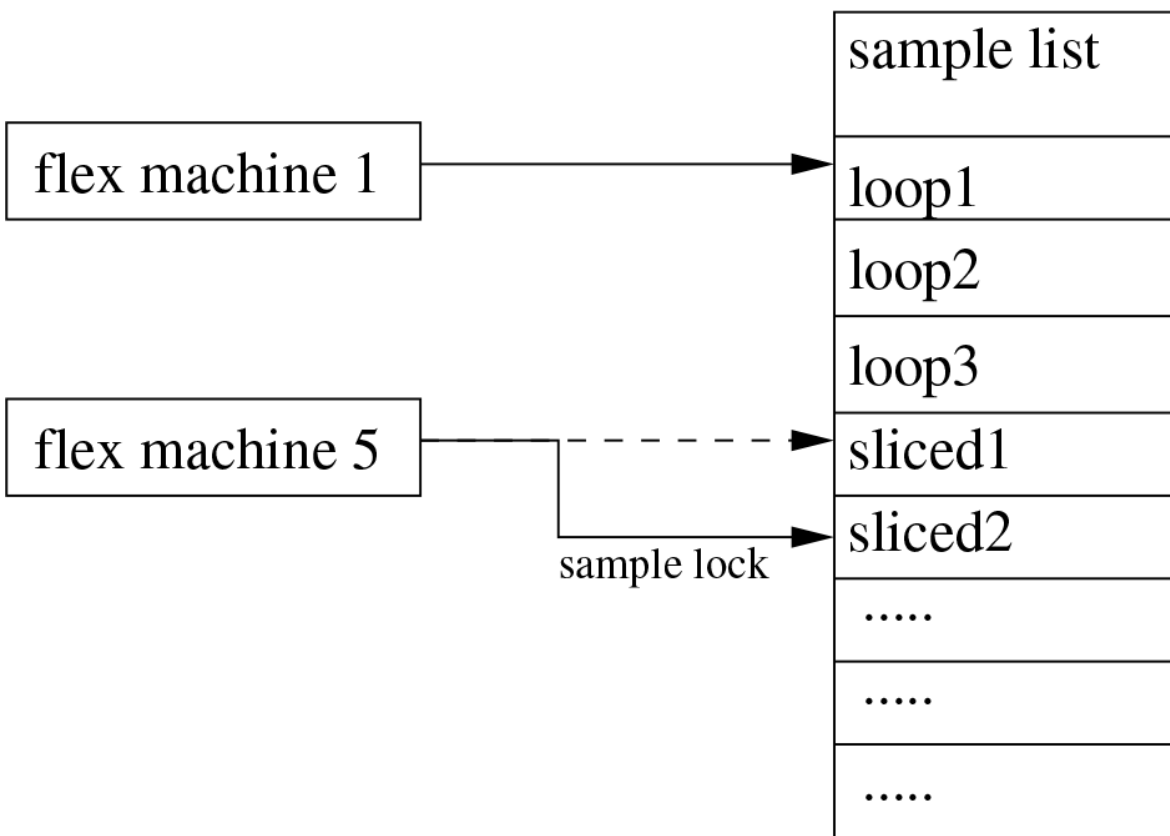


And there it is. By using sample locks, I can play *any* sample from the sample list at *any* step on *any* track that contains a flex machine in *any* pattern in *any* bank. That is a pretty powerful way of dealing with samples. This approach shows how different the OT works when compared to other samplers. Instead of coming up with excessive layering and polyphony, Elektron put the power in the *accessibility* of samples.

4.3 Slice locks

Again, look at the example as used in the previous section. Flex machine 5 has been assigned to a sliced sample. For using the slices, the flex machine was set up as described earlier. In the previous section, we saw that by using sample locks, you can select whatever sample is in your sample list. So far so good, but just as the sample can be locked, so can the start parameter on the playback page. The start parameter determines which slice is played. Therefore, you can not only determine which sample you want to play, but also which slice you want to play *in that specific sample* by locking the start parameter. As an example, assume that we lock flex machine 5 to sliced2.

Schematically, this looks like:



Because the sliced2 sample also contains slices, I can now lock the start parameter to select whichever slice I need.

The bottom line: Any flex machine configured for dealing with slices can play any slice of any sample that contains slices. You select the sample by using sample locks and you select the slice by locking the start parameter on the playback page.

Let's take this to its logical extreme and do a little math. The sample slot list can contain up to 128 samples. Each of those samples can contain 64 slices. We therefore have $128 \times 64 = 8192$ slices available.

Add to this that we can do the same thing with the static sample list and static machines, and this number doubles to a massive 16384 available slices.

4.4 Summary and overview


Up to this point, the overall structure of the OT is, in my view, still relatively easy to grasp:

- Samples are selected from the CF-card and put into a sample list, which is available project wide.
- Samples that are in the list can be prepared for further mangling. This is done in the sample edit menus.
- Flex machines play samples from the sample list. Sample- and slice-locks can be used to access any slice of any sample in the list wherever and whenever you need.

5. Effects

In the previous section, we saw how samples can be prepared and played back in many different ways. Although this already creates a lot of welcome possibilities, it would be nice if we could wrench the samples through some effects as well. In contrast to previous models like the Machinedrum and the Monomachine, Elektron provides us with two effect slots in which we can *insert* an effect. In other words, the sound that is produced by a machine is fed into two FX processors of our choice.

Schematically speaking:



An advantage of this approach is that we can *select* which effects we want to apply to a specific sample. That is a lot more flexible than being forced to use a fixed array of effects.

6. Scenes

In this section I will describe scenes, but *not* by wasting a lot of pages on how to set them up, as I would end up rewriting Elektron's manual. Instead, I'll look at scenes from a slightly more conceptual point of view and provide you with a simple model of how to look at them.

Basically, a scene allows you to use the crossfader to control several parameters on several pages on several tracks at once. The idea of attaching several parameters to a crossfader is not exactly new. For example, if we use Ableton Live with Akai's APC40 controller, we obtain the same functionality. The crossfader on the APC40 can be attached to several parameters scattered all over the place and a simple fade changes them all.

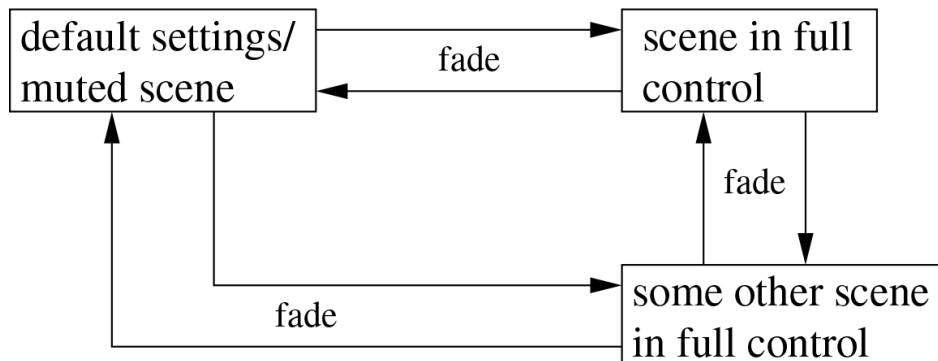
On the OT, Elektron extended this functionality to a higher level: Complete collections of parameters can be stored in a scene and no less than 16 scenes can be loaded to both scene slots living on the left and right side of the crossfader. At this point it is important to look at a few scenarios:

- **Fade from default to a scene:** Assume that the crossfader is on the left and that scene slot A is muted. This effectively disables the scene and the pattern plays according to its default parameter settings. A scene has been loaded to scene slot B. Now, when I move the crossfader, I am fading from the default settings to the scene settings. Fading back brings me back to the default settings of the pattern.
- **Fade from scene to scene:** While the crossfader is at scene slot B, I can load another scene to scene slot A. Fading back to the left now moves the parameters from scene B to scene A.
- **Muting a scene:** Finally, when I mute a scene, the OT immediately reverts back to its default settings or gradually moves back if I fade to it.

To summarize, there are three key things you can do with scenes:

1. Fade from default to scene.
2. Fade from scene to scene.
3. Fade from scene to default.

Schematically speaking:



Metaphorically speaking:

A scene allows you to undertake an expedition from the base camp (default settings) into new territory (a scene) and move further and further away as you keep loading new scenes and fade to them. Returning to base camp is always possible by muting a scene and fading to it. Of course, you can also skip the fading and mute/unmute a scene while the crossfader is there. In that case, you cause a sudden jump from default to a scene or vice versa.

6.1 Summary and overview

So far, the overall structure of the OT is still fairly straightforward:

- We draw samples from the CF-card, place them in a list, and prepare them for further use.
- We load machines to tracks and assign (some of) our prepared samples to them.
- We set up effects to mangle sounds coming from the machines.
- We set up scenes to allow for changing several parameters at once, thereby allowing us to drift away from the default settings (and back again) in a convenient way.

A crucial thing to understand here is how things depend upon each other. First you select samples, then you set up the effects as you need them and once this foundation stands, you build scenes for performing on that platform.

7. Parts

The part system is probably the most complicated aspect of the OT. Although the part system is very powerful and allows for very intriguing setups and possibilities, it is also a huge source of confusion for both newcomers and for more experienced users. Describing what parts are is difficult because they can be used in so many different ways. Therefore, simply describing what a part is will not help, because this results in a rather abstract and formal description which will keep many people puzzled about what parts are and how to use them.

I therefore decided to take a different route. Instead of coming up with a formal definition, I will explain what parts are and what they do by describing a couple of simple experiments. By simply following this step-by-step approach, our understanding of parts increases and (hopefully) evolves to a level where all the fog has disappeared... First though, we have to dive into a little dry theory concerning a simple question: Which information is stored where?

7.1 Parts and what they store

Question: What information does the OT need to play back a pattern as we intend?

Let's see... First of all, it needs to know which samples we want to use. The samples are stored in the sample list as described earlier. That list, however, is available project wide. It is therefore not *specific* to our pattern. Therefore, I ignore the sample list.

Next: The OT wants to know which machines have been loaded to which tracks and how we have configured them. So here we have the first valid piece of information the OT needs: machines and their settings.

Next: The OT also wants to know what effects we use on those machines and how we have configured them. So here is number two: effects and their settings.

Next: We have machines and their effects, and we have also set up several scenes that we use for performing. Scenes are therefore number three.

Next: All tracks play their audio on their own volume. Volume settings of the tracks are therefore number four.

Finally: The OT wants to know when to trigger the samples and whether parameter locks were used to add some bells and whistles, so here is the fifth piece of information the OT needs.

To summarise: The OT needs five chunks of information to play a pattern as we want:

1. Triggers and parameter locks
2. Scenes
3. Effects and their settings
4. Machines and their settings
5. Track volume settings

The next question: Where should all this stuff be stored? The most straightforward way would be to store all of this stuff in a pattern but, and this is where it gets interesting, Elektron decided to do something different. Instead of storing all this information inside a pattern, Elektron has spread it over patterns and what Elektron calls parts. So, what is stored where? Well:

- **Patterns** store triggers and parameter locks.
- **Parts** store machines (+settings), effects (+settings), scenes, and track volumes.

So why this distinction? At first glance this seems to complicate things unnecessarily. Well, time for some action...

7.2 Part reloading

First things first, let's build a simple pattern to work with. The pattern I built is set up as follows:

- **Track 1** holds a bass drum. To give it bite I put a compressor on the track followed by a delay, which will be used when constructing a buildup.
- **Track 2** holds a clap. To sculpt it, I put a filter and a reverb behind it.
- **Track 3** holds a loop with some hi-hat stuff. Effects are filter and delay.
- **Track 4** holds a bassline loop with filter and lo-fi applied to it.
- **Track 5** holds a pad with filter and reverb. Further on, I set up a couple of scenes for creating some movement.

Now then, try the following:

1. Press **[FUNCTION]** + **[PART]** to enter the parts menu. You will see four parts. The top one is highlighted and shows a small star. That star is important, more on that later.
2. Press **[FUNCTION]** + **[EDIT]**. A menu appears which allows you to save the part. Save and leave the menu.
3. Play the pattern. If all went well, nothing has changed and the pattern plays as usual.

4. While playing, tweak your pattern by randomly changing parameters on whatever track. Also try replacing effects with different ones, and remove some machines or assign different samples to them. Be brutal here; change the pattern so heavily that the original pattern is completely destroyed.

Note: Do not record parameter locks and do not add or remove triggers.

Basically, only change whatever is stored in a part but leave intact what is stored in a pattern (trigs and parameter locks).

Congratulations: You've just turned a nice sounding pattern into a complete mess.

The next steps:

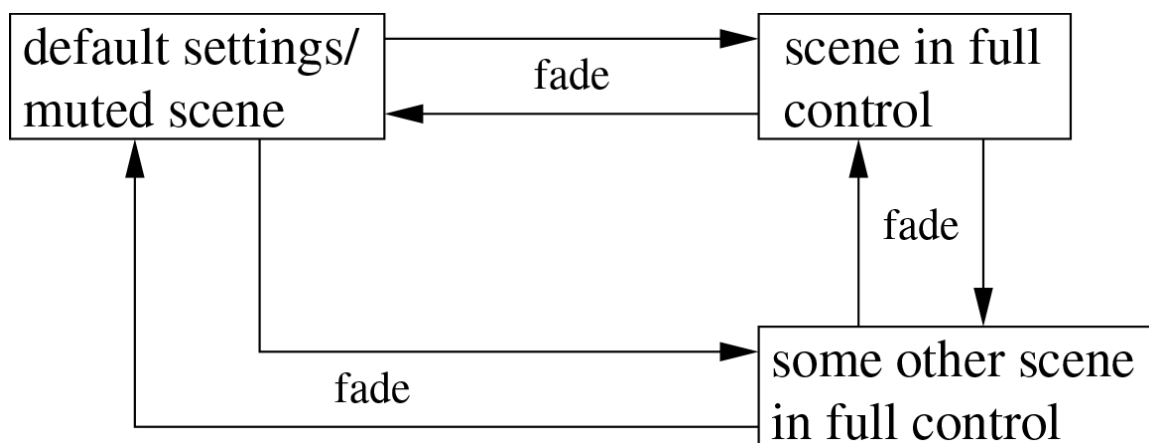
5. While playing your destroyed pattern, hit **[FUNCTION] + [RELOAD]**. Once you do that, all the destruction you caused is undone and the pattern is instantly reverted to the state it was in when you saved the part.

6. Repeat the above: destroy your pattern, change parameters, replace machines, whatever, and hit **[FUNCTION] + [RELOAD]** again to return home.

Metaphorically speaking:

By saving a part, you have set up a base camp. Once set up, you can go out for a walk and see what is out there by experimenting with different machines, effects, and scenes. Delete, change, and replace whatever you want. Once you get lost, hit **[FUNCTION] + [RELOAD]** and you are instantly zapped back home.

Schematically:



As you can see, reloading a part has some similarity with using scenes. Both can be used to drift away from the standard settings and both allow you to return home safely. Scenes, however, are only used for changing parameters and allow *fading* between parameter values.

Parts, on the other hand, are far more radical. They don't allow fading, but they do allow replacement of machines, effects, and scenes while keeping the possibility of reverting everything back to the last saved state.

So, what was that little star next to the part number about? Easy. Once you save a part, the current state of the pattern is exactly the same as what was saved in the part. Once you make a change (by turning a parameter knob or replacing an entire machine or whatever) the current state has changed from what was saved. The star indicates that modifications have been made to the originally saved state. In other words, when you see that star, what you are currently hearing is in some way different from what you saved to the part.

7.3 Part reloading: Usage

There are several situations in which reloading a part can be handy. Some suggestions:

Painless experimentation: This was described in the last section. After you've saved the part you can replace, throw away, and sculpt whatever you want, but the base camp is always within reach.


Incremental development: Assume you have a pattern running with a saved part. As you listen critically to it, you find that the hi-hat is a little too loud. You correct it by reducing its volume. You also find that the bass drum you used is completely rubbish and should be replaced. So, you find a suitable bass drum and dial it in. As time passes, you keep on fine-tuning the pattern until it sounds profoundly better than what you started with. Once you're satisfied, save the part! That is, move your base camp to this better sounding location.

7.4 Parts and patterns

In the last section, we saw some basic usage of parts and gained some understanding about how they work. This is not the big picture though. In this section, I'll focus on the relation between a part and a *collection* of patterns.

I'll start with the same example I used earlier. Assume that the pattern lives on A1, that it is using part 1, and that the part has been saved. This gives me a clean point to start from. Before going further, it is important to understand what the memory of the OT looks like at this point.


The following diagram shows the current situation:



The arrow illustrates that the pattern is using the part. Remember, patterns store triggers and locks, while the part stores the underlying infrastructure: machines, effects, scenes, and track volumes.

What happens when we copy pattern A1 to A2? Well, it's obvious that all the triggers and parameter locks that live on pattern A1 are now copied to pattern A2. But what happens to the part? Is a hard copy of the part made as well? Nope, the part is left unchanged. Instead, pattern A2 now *also* uses part 1.

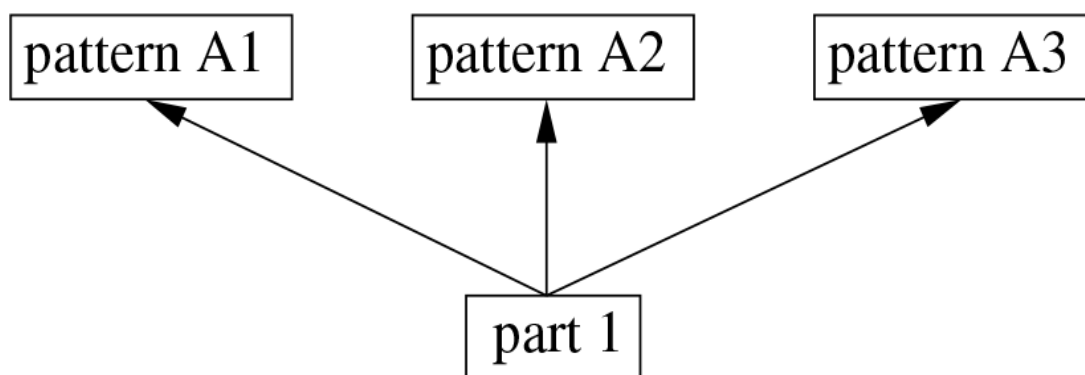
Schematically:



Think for a moment about the implications of this. If I change something to the part while pattern A2 is running, that change will be reflected in pattern A1 as well, because it uses the same part. This is a huge source of confusion to many users as they (wrongly) think that a pattern stores all of its information locally. This is not the case, as I described earlier. We can therefore conclude something very important: *Any changes you make to a part will be reflected in all patterns that use that part.*

Let's expand a little further on this. I move back to pattern A1 and copy it to pattern A3.

Schematically:



At this point you should be able to figure out where this is going. I can create several variants of a pattern by copying and pasting at will and all these patterns will use the same underlying machines, effects, and scenes. Any changes made to the machines, effects, or scenes in any of those patterns will be reflected in all other patterns that use this part.

So, what is the purpose of this? Easy: convenience and coherency. For example, suppose that the resonance of my filter in track five is a little too aggressive, so I tame it a bit. If every pattern stored its information locally, I would have to change this on pattern A1, then move to pattern A2 and change it again, then move to pattern A3 and change it for the third time... which profoundly annoys me.

With multiple patterns using the same part, I only need to make changes once and all patterns using that part follow with it. This saves me a lot of dull work and guarantees that all patterns form a coherent collection of variations, all using machines that sound the same, effects that behave the same, and scenes that work the same over all these patterns.

This expands to *performing* with multiple patterns. Suppose A1 is running and that I make changes to various parameters, track volumes, etc. When I switch to A2 (or any pattern using this part), the parameters I have changed will stay changed (as long as the part is not reloaded), allowing for a natural switch from A1 to A2. If every pattern stored these settings locally, switching a pattern would cause all kinds of parameter jumps.

The only real difference between patterns A1 and A2 lies in the triggers and the parameter locks. Although it may look a little silly at this point, remember that parameter locks can be used to play *any* slice of *any* sample in the sample list, so triggers and locks alone already create enough mayhem to turn pattern A2 into something that is very different from A1 and A3. To top it off, consider a simple real-world example:

- Pattern A1 contains my basic pattern. I can use scenes to perform some tricks on this pattern and also use the part reload function I described earlier.
- Pattern A2 contains a copy of pattern A1, but here I've added some extra triggers on my bass drum at the end of the pattern. I've also used some parameter locks to play the bass drum in reverse. Additionally, I've used some locks to retrigger a couple of hi-hats and, finally, I've used locks to add dirt to my bassline on track four.
- As pattern A1 ends, I instruct the OT to move on to A2. A2 contains the buildup and I patiently wait for it to develop. As A2 ends I direct the OT to pattern A3.
- Pattern A3 serves as a break; all triggers on the bass drum were deleted. Instead, I've put one trigger on step 1 and locked it to a sample in the list which contains a massive noise sound that slowly increases in volume and washes away everything. As this pattern approaches its end, I direct the OT to pattern A1 and everything starts over again.

7.5 Scene stacking

In the last section, we saw how multiple patterns can use the same part. Scenes are stored in the part. This means that once a scene is active, it will *remain* active even if I move on to the next pattern. In other words, 16 scenes can be used freely over all patterns that use the same part. This can be used to move from scene to scene and from pattern to pattern in a natural way. A slight problem emerges, however. I can prepare 16 scenes for use with my patterns, but how do I remember which scenes do what? Ok, I could write it down, but there is a simple way of letting the movement through patterns go hand in hand with the movement through scenes. I call it *scene stacking*.

The idea is easy:

1. When the first pattern runs, it runs without being affected by any scene.
2. I move the crossfader to the right and load scene 1 to the left. Fading to the left now puts scene 1 in full control.
3. I copy scene 1 to scene 9, and load scene 9 to the right. Crossfading to the right now has no effect, since scene 9 holds the same parameter values.
4. I make all the changes I want to scene 9. Because the parameters of scene 1 are still stored in scene 9, I can move naturally from 1 to 9. The parameters that changed in scene 9 are stacked *on top of* what was already changed in scene 1.
5. I copy scene 9 to scene 2, which is loaded to the left. Again, I make changes *on top of* what was changed in scene 9.
6. Copy 2 to 10, modify, and fade.
7. Copy 10 to 3, modify, and fade.
8. Copy 3 to 11, modify, and fade.
9. ...

By setting up scenes like this, I can move through them in a natural way. Scenes 1 to 8 are always on the left while scenes 9 to 16 are always on the right. Fading from scene to scene is easy:

- If the crossfader is on the left, I press the right scene button, look at the LEDs (9-16), and I load the scene that comes after the one currently loaded.
- If the crossfader is on the right, I press the left scene button, look at the LEDs (1-9), and I load the scene that comes after the one currently loaded. Because moving through scenes has now become easy, moving from one pattern to the next one has become easy as well. As I move through the patterns, I fade from scene to scene. Those scenes have been set up in such a way that they create a fluent progression through the patterns.

7.6 Using multiple parts

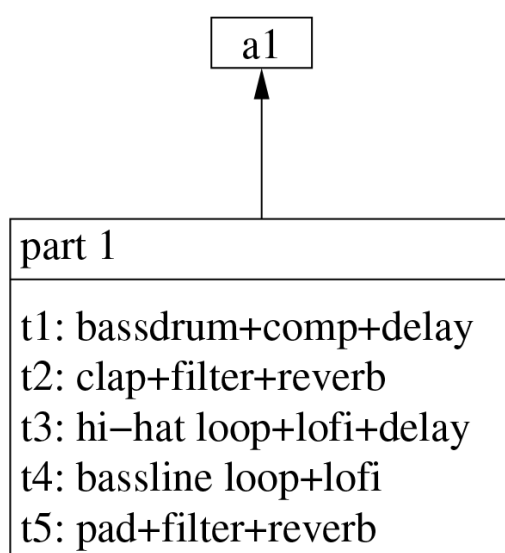
As we saw in the previous section, parts store machines, effects, scenes, and volume settings. Several patterns can use a particular part, and their machines can be sequenced in various ways by playing around with triggers and parameter locks.

So far, however, we have only used one part and as you have probably figured out already, the OT can hold up to four parts per bank. Four parts means we have four individual “platforms” (machines, effects, scenes, and volume settings) that we can feed to our patterns.

In this section, I will extensively cover an example which shows how multiple patterns, scenes, and parts can be used to construct a transition between two patterns. To get started, assume the following (very basic) setup:

- A pattern lives on track one.
- The pattern uses part 1.
- The part is saved: basecamp is secured.

Rather than describing what kind of sounds were loaded to which track, I've laid them out in the diagram below:



I can play around a little with this pattern by changing some parameters, muting and unmuting, manually triggering samples, etc. Remember that I can always return to the basecamp by using the part reload function as described earlier.

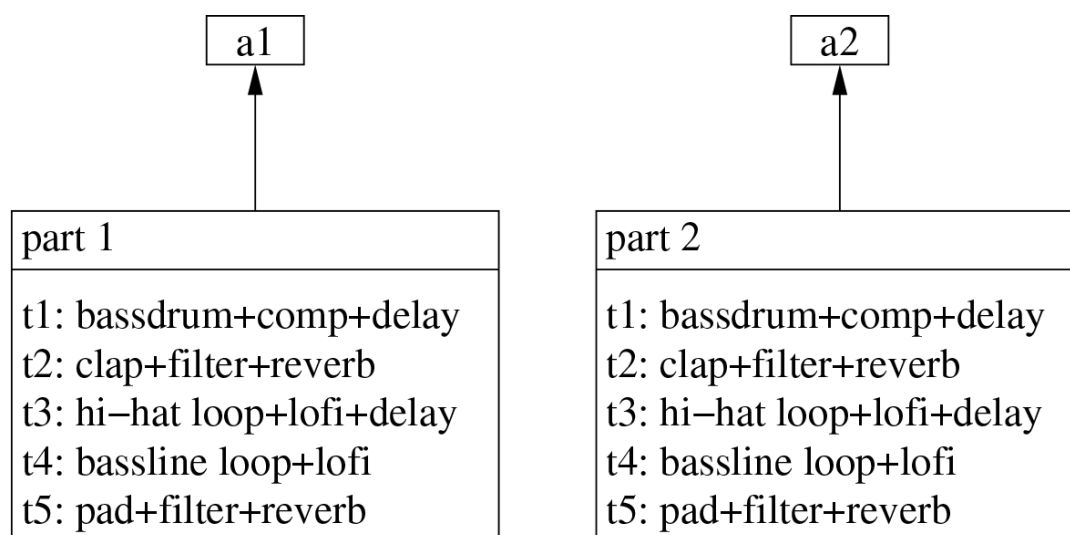
Time for a change: I want to construct a buildup. I do this by setting up a scene on scene slot B. The scene is constructed easily:

- The bass drum feeds a delay, which results in a bass drum bouncing all over the place.
- The filter on the clap is opened, resonance is increased a little, and reverb is added.
- The rate reduction and distortion parameters on the hi-hat are used for creating artifacts and destruction.
- The distortion parameter is increased on the bassline to let it go nuts.

With this scene loaded, a simple and slow fade of the crossfader from left to right pumps more and more madness into the pattern. As I keep moving the fader, the buildup grows more and more... to what?! Well, eh... Injecting a buildup is fine, but sooner or later a buildup has to be followed by a move to a new direction. I can move to a new pattern or fade to a new scene as was described in the last section, but the problem is that whatever I do, the machines and effects will remain the same. For really making a move into a new direction, some more radical things have to be done:

1. I copy pattern A1 to A2. As described earlier, we now have two patterns both using the same part.
2. While in pattern A2, I turn to the part menu and copy part 1 to part 2.
3. Copying from part 1 to part 2 still doesn't mean that A2 is using part 2. So, I select part 2 and hit [YES].
4. I save all parts.
5. I look at the OT's screen and verify that A2 is indeed using part 2.

After the above operations, the memory of the OT looks as follows:

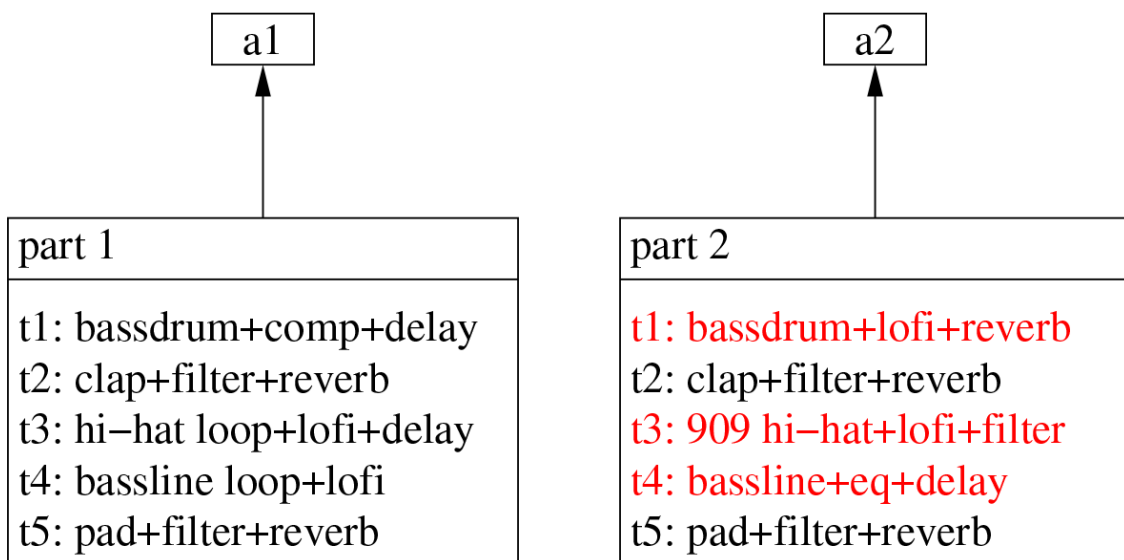


As you can see, I now have two patterns, *both using their own part*. Since A2 is a copy of A1 and part 2 is a copy of part 1, both patterns will sound exactly the same. Yet, modifying part 2 has now become fully independent of part 1.

Time for some changes: I want A2 to sound far more aggressive than A1. In order to achieve this, I modify several elements in part 2:

- The bass drum is replaced by a more aggressive one. I also replace the effects: lo-fi and reverb are now in place.
- The hi-hat loop is replaced with an aggressive 909 hi-hat + lo-fi + filter.
- The bassline is replaced by a really deep and pumping one + EQ + delay.
- Clap and pad are left unchanged.

I also add some triggers and parameter locks on various tracks and tweak and fine-tune until I have the loop I want. To make these changes permanent, I save all parts. After the above operations, the OT memory looks like this:



The differences between part 1 and part 2 have been colored red. Right, all has been set. Let's perform:

1. I switch to A1 and hit **[PLAY]**.
2. I load my buildup scene to scene slot B.
3. Time for the buildup. I slowly move the crossfader and the pattern starts building and building...
4. At a moment of my choice, I hit **[PATTERN] + [A2]**. At the moment I do that, I already know that when A1 ends, it will be followed by the more aggressive sounding A2.
5. I wait until A1 ends and on the moment A2 kicks in, I mute scene slot B. By muting scene B, A2 plays without being affected by the scene.
6. A2 now plays cleanly: The new bass drum and bassline really punch away and the 909 hi-hat cuts through the air. The original pad and clap from part 1 are still present as they were. This new pattern still has some remnants of its predecessor, but the newly added elements pushed this pattern into a new direction.

I am now left with A2, which uses its own part. I can therefore set up scenes that are tailored for that pattern, goof around with parameters and the part reload function, mute, unmute, etc. I can also copy A2 to other patterns and build variations, knowing that all these variations use the same part. Effectively the game starts all over again.

7.7 A small but convenient enhancement

In the previous section I constructed a pattern, used a scene to create a buildup, and moved to the next pattern, which uses its own part. Once A2 kicked in, I muted the scene and A2 ran completely clean.

Running clean in this case means that the aggressive bassline I installed on part 2 is *immediately* heard on full volume at the moment A2 kicks in. Nothing wrong with this, but suppose that I don't want that. Suppose I don't want to hear the bassline at the moment A2 kicks in and that I want to fade it in myself as A2 plays on. How do we do that?

At first glance, I could just mute the bassline track the moment A2 kicks in. Although this is possible, it requires me to be very precise in my timing. To make things worse, I already need to be careful with timing because I have to mute scene B at the exact moment that A2 kicks in. This is obviously not convenient. Let's explore some solutions.

7.7.1 Parts store track volumes

A part stores machines, effects, scenes, and the volume at which the tracks feed their audio to the main outs. Therefore, setting the track volume on the bassline's track to 0 and saving the part will "mute" the track once A2 kicks in. I can then select the track and increase the volume as the pattern plays on.

There is a downside to this, however. Every time the part loads, the volume of that track is set to 0. Every time I use the part reload function, the track drops to 0. So as long as the part doesn't reload, I am safe, but this is not ideal.

7.7.2 Parts store scenes

Remember that a part stores machines, effects, scenes, and volume. The fact that a part stores scenes is very important here. While A1 is playing, the scenes in part 1 can be used to do whatever you want. For creating the buildup towards A2, I loaded scene 9 in scene slot B and slowly faded to it. So, once A1 ends, the fader is pushed full to the right and scene 9 is in full control.

Then A2 kicks in and loads its own part. The fact that A2 loads its own part implies that it *also loads its own scenes!* Elektron went one step further though: A part also stores which scenes are loaded to scene slots A and B and also whether a scene is active or muted. Let's use that to our advantage:

1. I move to A2 and *delete* scene 9.
2. I load a clean scene 9 to scene slot B. By starting with a clean scene, I make sure that no unexpected parameters are affected by that scene.
3. I configure scene 9 in such a way that the volume of the aggressive bassline is set to 0.
4. I save the part to make this permanent.

And there it is. Once A2 kicks in, I don't need to mute scene B at that exact moment anymore. In fact, I don't need to do anything: At the moment the OT switches from A1 to A2, part 2 is loaded and thus scene 9 of part 2 is loaded as well. This new scene kills the volume of the bassline. Moving back the fader to the left (where no scene lives at all) fades in the bassline to its normal level. The pattern now runs cleanly.

Metaphorically speaking:

I have left my previous basecamp for good and arrived at my new home. Since I really liked the clap and pad, I took them with me, but someday when I move to the next basecamp, I might leave them behind...

8. Volume handling


Once a machine plays a sample, several parameters affect its volume. Understanding which parameters live where and how they interact with scenes will avoid several volume handling pitfalls.

8.1 Signal flow

The signal that a machine generates is first wrenched through the amplifier envelope. On the amp page, you will find the `VOL` parameter, which is a bipolar parameter, meaning that its range goes from -64 to +63. The parameter defaults to 0. Generating tremolo effects can be done by letting an LFO modulate this parameter. Mimicking side chain compression is done by setting up an LFO in such a way that it decreases volume on those moments a bass drum kicks in. Parameter locks can be used for that as well.

After the amp envelope, the signal moves through two FX processors. Finally, the track volume parameter determines the overall volume of the whole track (sample + FX).

Schematically:



As you can see in the diagram, the OT offers both pre- and post-FX volume handling on each track.

8.2 Volume and scenes

On top of the regular volume parameters, the OT offers two additional parameters which are exclusively available to the scenes:

Track volume: `XLEV` can be set to min or max. The parameter is shown once you hold one of the scene buttons. Min equates simply to 0 and max equates to the volume that is currently set (that is, if the scene is muted). It therefore allows you to “mute” a track when the scene is active or to fade in a track to the “usual” track volume.

Perhaps using parameter locks on the amp volume parameter does the job. It does indeed, but now I have to remember that the bass drum should be locked on every step I program in patterns that use that bass drum. Although this is a solution, it is still not very convenient.

Effectively the problem is this: everything in your patterns was ok, until that misbehaving bass drum was dialed in. The problem is caused by the sample itself, and should therefore be corrected on the sample level and *not* by disrupting existing settings which also affect other elements. Basically, the sample should adapt to the system and not the other way around.

Fortunately, the OT offers a simple way to correct “misbehaving” samples. On the attributes page in the audio editor, a `GAIN` parameter is found. This parameter allows you to boost or decrease the sample volume with a massive 24dB in both directions. Simply start the pattern and set the volume to taste in the attributes page: Many problems with inconsistent track volumes or parameter locks are now history, as that parameter is considered an attribute of the sample and stored in the sample slot list.

9. Sampling

So far, I have completely ignored the OT's sampling abilities. I had a simple reason for doing so, however. Explaining the overall structure of the OT is hard enough and talking about sampling would not have changed the overview. In order to avoid making things more complicated than necessary, I have therefore avoided the sampling issue as a whole.

9.1 The OT's approach to sampling

When I look at samplers of the past, I can see that most of them have the ability to record new material from the outside world and play it back. Indeed, this is obvious, but the point is that on most samplers these two features have been completely separated. Most samplers offer the user a "sampling mode" and a "playback mode" and for good reasons. When recording new material, you enter sampling mode, hit start, and make the recording. Once the recording has taken place, some facilities are offered for trimming and preparing the sample before it can finally be used for playback. Over the years, several enhancements have been added. For example, midi can be used to synchronize the sampling and playback to external gear, etc.

Then, in 2004, Ableton Live 4 hit the markets. Ableton took a different approach to sampling: It offers a system in which sampling and playing back samples can be done without stopping the sequencer or switching between sampling and playback modes. In other words, Ableton integrated both of these functions in a system under the direct control of a sequencer and a user manipulating it.

The biggest difference between Live 4 and Live 3 lay in the fact that Live 4 was the first version that allowed the use of VST synths and effects, thereby connecting the world of sampling to the world of producing, and blurring the distinction between both.

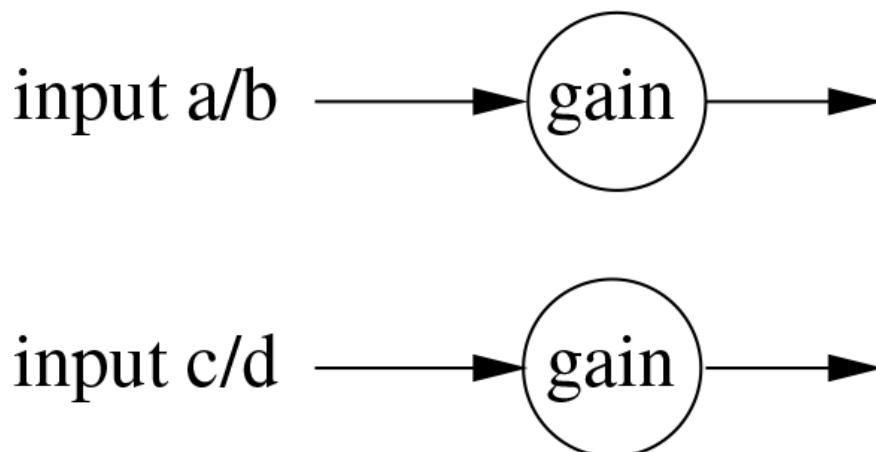
At that time, Elektron must have already been thinking about integrating samplers in their products. In 2005, the Machinedrum UW was introduced, and this unit takes more or less the same approach to sampling as Ableton: both sampling and playback can be done without stopping the sequencer. By doing so, sampling new material has become just as trivial as playing back samples. This approach allows for a seamless integration of sampling in a performance, and that was a big step forward from what already existed at that time.

9.2 Signal flow

Before diving into the sampling process itself, I should mention a few words about signal flow. Perhaps a little dull, but it is important because understanding the signal flow will avoid problems with volume. First things first, the OT offers two stereo line inputs, simply named input A/B and input C/D. When external equipment is connected to the OT, a small problem appears: some synthesizers deliver a very hot signal while others deliver a signal that is too soft. The OT, therefore, should offer a facility to boost or suppress an incoming signal *before* that signal hits the rest of the system.

The OT offers simple `GAIN` parameters for both A/B and C/D in the mixer page. The gain allows for boosting an incoming signal by 12dB or suppressing it to 0. The role of these parameters should be clear: you use the gain for correcting the volume of incoming signals so that the rest of the OT receives a healthy signal. Ideally, the gain should be configured once and be left untouched throughout the rest of the project. Set and forget. The gain parameters therefore cannot be automated, sequenced, or whatever. They exist for providing the OT with healthy signals from the outside world and that's it. Any tricks for adjusting volume, fading in or out during a performance, etc., should be made elsewhere.

Schematically:



Now that we have a healthy signal, we can do two things with it:

1. Send it straight to the main outs.
2. Feed it to various track recorders.


9.2.1 Main outs

Sending the signals to the main outs is probably the easiest part: The `DIR` parameters in the mixer page are used for that. By increasing the `DIR` parameter, you determine how loud the incoming signal is fed *directly* to the main outs. Hence the name “`DIR`”. All other things that happen inside the OT are completely separated from this, with one important exception:

Feeding an incoming signal directly to the main outs is all fine, but during a performance I’d like to fade the live inputs in and out, when I need them. During a performance, however, I am probably already goofing around with various track parameters, scenes, etc. Every time I want to fade the live inputs in or out, I have to turn to the mixer screen to control them. Not very convenient...

Fortunately, the volume of the live inputs can be controlled with scenes. While in the mixer page, hold the scene button and you will see that the `DIR` parameters turn to “`XDIR`”. I can now set the volume of both inputs to maximum or minimum. This allows me to fade the live inputs in or out with the crossfader without opening the mixer page. Live inputs are now fully integrated in a performance. I can use the scenes to control parameters on all tracks and mix the live inputs in or out as I need them. Effectively I can use the scenes for controlling ten “channels” of audio: eight tracks from the OT itself and another two live inputs.

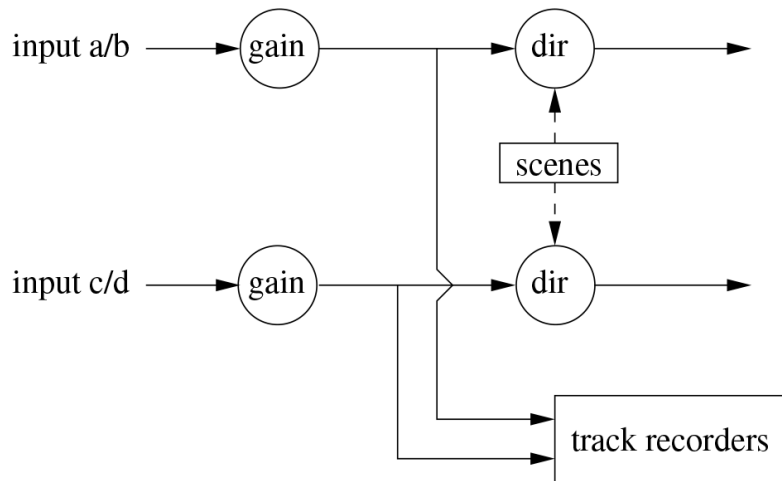
Schematically:



9.2.2 Track recorders

Live inputs can be sent directly to the main out, but we can also feed them to the track recorders for sampling.

Schematically:



Designing the signal flow in this way is pretty clever, as it allows for fading the live inputs in or out, irrespective of the sampling process. I can therefore sample without hearing the inputs, or vice versa, or any other combination.

9.3 Sampling guidelines

The OT features eight fully independent recorders on each of its tracks. The recorders are built “on top” of a track and are always available, regardless of what happens on a track. Each of these recorders can record from any combination of three sources: The live inputs A/B and/or C/D and/or a third source, which can be set to any track, the main output, or the cue output.

Recording from any of these sources can be triggered either by the sequencer or manually, and I will not dive into the details of that here, since I would end up rewriting Elektron’s manual. I will, however, present you with a couple of brief guidelines:

Sequenced sampling: By setting up recorder triggers, the recording process is under full control of the sequencer. Triggering recorders from the sequencer has an advantage over manual sampling: You guarantee that sampling is always synchronized to the sequencer. Thus, trimming or further preparation of the sample is usually not needed. It can be used directly in the rest of the system without stopping the sequencer.

One-shot triggers: If you want to sample *once* without overwriting the sample the next time the pattern loops, you can use one shot triggers. They trigger the recorder once but no more. The one-shot triggers can be re-armed at will, so this is a nice technique if you want sequenced sampling, but only when you need it. This more or less equates to triggering a recorder in Ableton Live.

Manual sampling: By triggering the samplers manually, you throw away the timing advantages of sequenced sampling, but the inadequacies of your own timing can lead to happy accidents.

Quantized sampling: This is a form of manually triggering the recorders, but by setting the `QREC` and `QPL` parameters in the record setup page, the recorders are not triggered instantly, but quantized to the OT's sequencer. This more or less equates to triggering a recorder in Ableton Live.

9.4 Track recorders and the sample slot list

So where do these recorders write their audio? Well, every track recorder has its own dedicated record buffer. This means that a recorder on track one always writes its audio to record buffer one. This is hardwired into the system.

So where do the record buffers live? Answer: In the sample slot list. Think for a moment about the implications of this.

As described earlier, the sample slot list is available project wide. In other words, once a recorder writes something to its record buffer, that sample is instantly available on any track holding a flex machine in any pattern in any bank as long as the track recorder doesn't overwrite the sample. And even if it does, that new sample is, again, immediately available everywhere.

Finally, a couple of handy details about the record buffers:

- **Editing:** There is no difference between a record buffer and a regular sample slot. So, anything that is put into a record buffer can be fine-tuned with the audio editor. This includes slicing.
- **Saving to CF-card:** If you want to secure a sample that is currently in the record buffer, you can dive into the audio editor and select the file page. There you find the option `SAVE AND ASSIGN SAMPLE`. By selecting that, the contents of the record buffer are written to the CF-card. You decide how it should be named and afterwards, the OT automatically assigns that sample to a free sample slot. This is handy: If you want to save a recorded sample from destruction and immediately make it available in your project, this is the way to go.

9.5 Sampling techniques

I will now describe a couple of sampling strategies.

9.5.1 Dedicating patterns

The fact that a recorded sample is instantly available anywhere creates many possibilities. One could, for example, set a couple of recorder triggers in a pattern and immediately move on to the next pattern which doesn't have these triggers. In this way, a pattern could be regarded as a recorder fly-by: move to the pattern for recording what you need and immediately leave to the next pattern where (some of) these samples are used, amongst others.

Although sacrificing a full pattern to hold just a couple of recorder triggers, may sound like a waste of patterns, we have enough of them, so this is not a big deal. Of course, one could also use one-shot triggers for setting up recorder fly-bys.

9.5.2 Abusing the cue outputs

Every track in the OT can be sent to the master output, but also to the cue outputs. The cue outputs are independent of the master output. As described earlier, a track recorder can record the live inputs and a third source. This third source can be any track, the master outs and the cue output.

So, set up a track recorder to record the cue output. Whatever combination of live inputs and tracks you need can be sent to it by cueing them, irrespective of what happens on the master outs. If, for example, you need a recording of the bass drum, hi-hats, claps, and input C/D, just cue these four, trigger the recorder the way you want, and you obtain a sample of what you selected.

If you need some form of pre-listening, plug in some headphones and assign the headphones to the cue output in the mixer page.

9.5.3 Abusing a track for live input

This sample technique is nasty and not described in the manual as a separate technique, but it works and is, in my honest opinion, very usable, so I present it here.

As described earlier, track recorders and tracks are independent. A flex machine running on a track will not bite the track recorder on the same track and vice versa. There is, however, a way of setting up a track recorder and a flex machine in such a way that they do bite each other in a useful way. I set it up as follows:

- Set up a track recorder and let it record the live input. Record length should match the pattern length.
- Place a full trigger on step 1. This allows the recorder to keep recording, boldly overwriting the last recording every time the pattern re-loops. Although the recorder does its job, I cannot hear the live input. Set up a flex machine on the same track.
- Assign the flex machine to its own record buffer. In other words, if the recorder of track four was used, we now put a flex machine on track four and assign it to record buffer four.
- Trigger the flex machine on step 1.

Indeed, this is the same step on which the recorder was triggered. We have now set up a strange situation: the track recorder is writing to the buffer while the flex machine is reading from it at the same time. So, what happens? Well, and this is documented in the manual, in such a situation, you will hear the live inputs. This yields several advantages:

- Removing the recorder trigger breaks the read/write situation. As a result, the flex player simply plays the contents of the track recorder. Live input has gone.
- Whether you hear the live input or the sample, they are both affected by the track effects.
- The track volume can be used to fade in or out the track as usual, whether it plays the buffer or passes the live input. Oh, and did I mention scenes?