INVESTIGATION OF A DRUM CONTROLLED CROSS-ADAPTIVE AUDIO EFFECT FOR LIVE PERFORMANCE

Saurjya Sarkar
Birla Institute of Technology & Science
Pilani, India
saurjya.sarkar@gmail.com

Joshua Reiss
Queen Mary University of London
London, United Kingdom
josh.reiss@qmul.ac.uk

Oeyvind Brandtsegg
Norwegian University of Science and Technology
Trondheim, Norway
oyvind.brandtsegg@ntnu.no

ABSTRACT

Electronic music often uses dynamic and synchronized digital audio effects that cannot easily be recreated in live performances. Cross-adaptive effects provide a simple solution to such problems since they can use multiple feature inputs to control dynamic variables in real time. We propose a generic scheme for cross-adaptive effects where onset detection on a drum track dynamically triggers effects on other tracks. This allows a percussionist to orchestrate effects across multiple instruments during performance. We describe the general structure that includes an onset detection and feature extraction algorithm, envelope and LFO synchronization, and an interface that enables the user to associate different effects to be triggered depending on the cue from the percussionist. Subjective evaluation is performed based on use in live performance. Implications on music composition and performance are also discussed.

Keywords: Cross-adaptive digital audio effects, live processing, real-time control, Csound.

1. INTRODUCTION

Adaptive audio effects are characterized by a time-varying control on a processing parameter. The control is computed by extracting features from the input audio and mapping their scaled versions to relevant control parameters of an effect [1]. Cross-adaptive audio effects are adaptive audio effects where the features of a signal are analysed to control the processing parameters of another signal. This opens a range of possibilities as it presents a new form of communication between musicians. The idea of another performer interfering with the way your instrument sounds is disruptive at first, but if routed carefully it can have profound impact on the way we perform music.

In recent years, we have seen a few implementations of cross-adaptive audio effects. [2] described how such effects can be used to automate the mixing process by utilizing features across multiple tracks to determine the processing applied to each track. Such intelligent mixing systems have been implanted and evaluated for automation of multitrack equalization [3] and of multitrack dynamic range compression [4]. [5-7] presented audio processing plugins with the capability to implement user-defined cross-adaptive effects. [8] developed a genetic algorithm (GA) and artificial neural network (ANN) based cross-adaptive audio effect that can control user-defined parameters to make a source audio file sound as close as possible to a given target audio file.

An advantage of audio post-production is the ability to move recorded segments and synchronize them after recording. Effects are applied during production which are time-aligned with audio events. Replicating such effects live with musicians controlling these effects while playing their instrument is often not feasible. However, it may be possible to use cross-adaptive digital audio effects to synchronize the effect applied to one source to events produced by another source. Using cues from a percussionist to synchronize effects across instruments could be highly beneficial in live performance. It can introduce live instruments to replace sampled sounds, since acting on such cues can keep effects synchronized while accommodating for human error and reaction times.

In this paper, we explore a cross-adaptive framework that enables a percussionist to orchestrate effects on other instruments in real time. In particular, we evaluate such an approach for several scenarios where musicians try to recreate, in live performance, a synchronized audio effect that otherwise would only be implemented in the studio. The effects presented here are feed-forward cross-adaptive audio effects as the features are extracted from the drum input source(s) and effects are implemented on the synth input [1]. They include short duration effects such as ducking, tremolo and filter sweep effects that can be triggered by drum cues in real time. These effects have been implemented as a live performance tool which can synchronize effects across multiple instruments using cues from the drummer. This enables the drummer to orchestrate effects across multiple instruments during performance.

Though similar effects can be implemented using MIDI triggers and sidechaining, our framework has more flexibility. With MIDI triggers or sidechaining, each drum hit would trigger the cross-adaptive effect. In our framework, the onset detector can be manipulated by the performer to trigger the effect only on louder notes or avoid re-triggering on closely spaced notes. Using an onset detector also enables the use of onsets from physical cymbals and/or drums, and does not require additional hardware.

2. IMPLEMENTATION

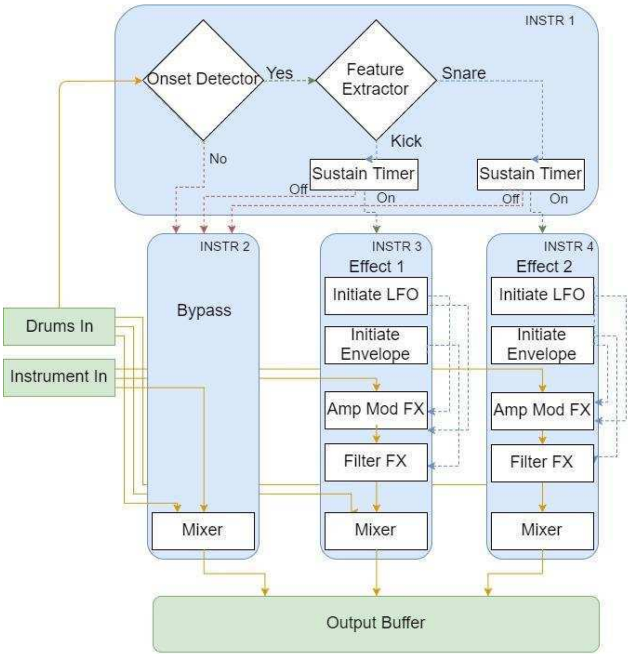

Figure 1: *A signal flowchart for the cross-adaptive audio effect framework.*

The cross-adaptive effects were implemented as a VST plugin using Csound (code available here: https://code.soundsoftware.ac.uk/attachments/download/2231/CrossAdaptiveMain.csd). The effects run within a signal flow as depicted in Fig. 1, where an onset detector runs on percussive input channels. Upon detection of an event, it initializes an envelope profile and an LFO which are sent to the desired effects acting on another instrument.

The VST plugin has three inputs, two control inputs coming from kick and snare drums and one instrument input (in this case the synthesizer) on which the effect was applied. Onset detection was applied on both the control inputs to trigger two different effects on the synth channel.

REAPER was used to run the VST plugin since it supports multitrack plugins. The effects were routed as following.

- Track 1 – Kick Drum Input. The VST is loaded on this track.
- Track 2 – Snare Drum Input. Master Send is switched off, channel 3/4 is sent to Track 1.
- Track 3 – Synth Input. Master Send is switched off, channel 5/6 is sent to Track 1.

Figure 2: *Plugin User Interface, including Envelope and LFO controls.*

2.1. Onset Detection

An onset detection algorithm [5] was applied to the amplitude envelope of the percussive signal, and transients extracted from it. We used the Csound opcode 'follow2' based on the work by Jean-Marc Jot, to extract the amplitude envelope of the percussive signal.

$$\begin{aligned} g &= 0.001^{(1/(f_s \cdot \tau_a))} && \text{if } e[n-1] < x[n] \\ &= 0.001^{(1/(f_s \cdot \tau_r))} && \text{else} \end{aligned} \tag{1}$$

$$e[n] = (1-g)\cdot x[n] + g \cdot e[n-1] \tag{2}$$

$x[n]$ is the input drum signal, $e[n]$ is the envelope extracted from signal $x[n]$. τ_a and τ_r are attack and release times, and f_s is the sampling frequency.

We then down-sample $e[n]$ to the frame rate with a frame width of 32 samples. We convert this down-sampled envelope follower to the decibel scale and compare it to a delayed copy of the same to detect transients. For an onset to be detected, all three of the following conditions should be simultaneously true for a given sample.

1. $e[n] > T_{gate}$ (3)
2. $M - e[n-d] > T_{decay}$ (4)
3. $e[n] - e[n-d] > T_{slope}$ (5)

Where M is the local maximum (the last onset trigger), d is the distance between the samples between which the slope (transient) is calculated, T_{gate} is the noise gate threshold, T_{decay} is the minimum amount by which the signal needs to decay before allowing the next onset to be detected, and T_{slope} is the minimum increase in amplitude required between $e[n]$ and $e[n\text{-}delay]$ to detect an onset. Whenever the onset detector detects an event from the drum signal, it initiates an event for the duration set by the envelope control panel.

2.2. Envelope Control

The envelope control panel, shown in Fig. 2, sets the duration and behaviour of the ASR envelope that is mapped to the parameter of the effects. The parameters mean the same as for any conventional dynamic effect. There are individual envelope controls for effects associated to different drums.

The attack time sets the time required for the control value to reach its maximum value. The sustain time sets the duration during which the control value will approach and hold its maximum value. Therefore, if the attack time is greater than sustain time, the control value will never reach its maximum value, and it will start reducing after the sustain time has passed. Thus, the actual duration of the effect is not affected by the attack time.

The release time sets the amount of time the control value takes to drop to zero from its value at the end of the sustain. The sum of sustain time and release time determine the effect duration.

2.3. LFO Control

The LFO is a low frequency oscillator used to produce tremolo and vibrato among other effects. The frequency of the LFO is set by the tempo (in BPM) and count (eg. 1/8th or 1/16th notes) inputs, shown in Fig. 2.

$$f[n] = tempo \cdot count \,/\, 60 \tag{6}$$

The frequency control button allows the frequency of the LFO to vary according to the control value '$e[n]$' set by the envelope. When turned on, the frequency of the LFO is given by:

$$f[n] = tempo \cdot count \cdot e[n] \,/\, 60 \tag{7}$$

The floor control button enables the depth of the tremolo effect to be controlled by the envelope control value. The LFO output is then given by:

$$y[n] = 1 - depth \cdot e[n] \cdot LFO[n] \tag{8}$$

The shape of the LFO can be set to sine, triangle, square bipolar, square unipolar, saw-tooth down, or saw-tooth up.

3. IMPLEMENTED EFFECTS

3.1. Amplitude Modulation Effects

Two types of amplitude modulation effects were implemented.

3.1.1. Ducking

If the amplitude modulation effect is turned on without enabling the LFO, then a simple ducking effect is produced as follows:

$$y[n] = 1 - depth \cdot e[n] \tag{9}$$

where $y[n]$ is the amplitude envelope of the output signal.

3.1.2. Tremolo

If the amplitude modulation effect is turned on with the LFO enabled, then a tremolo effect is produced. If the LFO control is switched off, then a simple tremolo is produced as per Eq. (6). If the frequency control button is turned on, then the tremolo frequency is modulated by the control envelope as given in Eq. (7). If the floor control is switched on, then a ducking tremolo effect is produced as per Eq. (8).

3.2. Filter Effects

Filter sweep effects were implemented using low pass, high pass and notch filters whose center or cutoff frequencies were dynamically controlled by the ASR envelope value. The filter center/cutoff frequency was updated at the frame rate.

3.3. Effect Performance

Fig. 3 shows synthesizer and drum samples used to show the spectral and temporal changes applied by the effect. The drum sample was composed of a kick followed by a snare drum hit, which were routed to different effects.

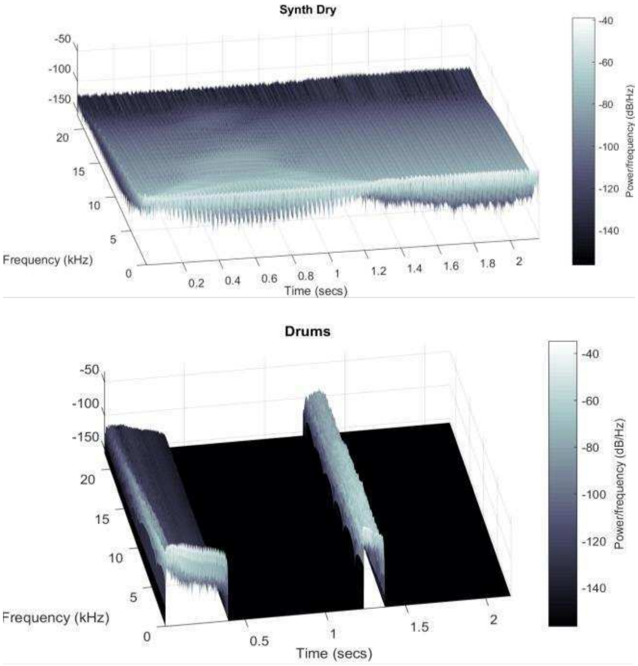

Figure 3: *Spectrograms of dry synthesizer sample (top) and drum sample (bottom).*

For the filter sweep effect, shown in Fig. 4, the kick trigger was associated to a low pass filter with cutoff frequency varying from 20 Hz to 22 kHz, as per the control envelope, and the snare trigger was associated to a high pass filter with cutoff frequency from 20 Hz to 20 kHz.

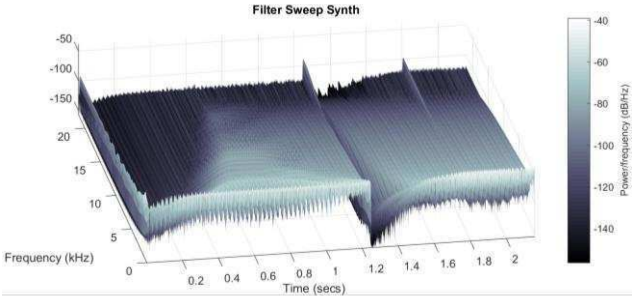

Figure 4: *Filter Sweep Effect Spectrogram.*

For the tremolo effect, the kick trigger was associated with a sine tremolo at quarter note triplets and the snare trigger was associated with a sawtooth tremolo at quarter notes, shown in Fig. 5.

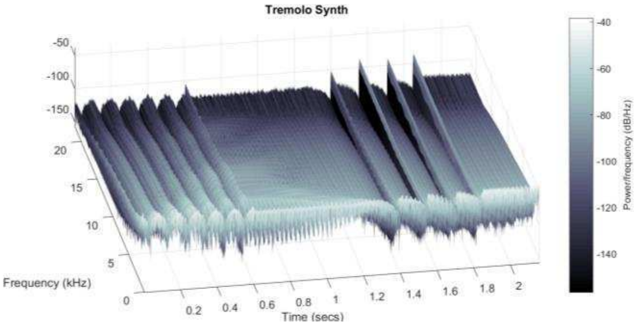

Figure 5: *Tremolo Effect Spectrogram*

.

A frame glitch was observed in the spectrogram at turn-off and turn-on time of each effect (seen prominently in Fig. 5 at 1.8 sec) but this was not audible during performance.

4. PERFORMANCE EVALUATION

To evaluate the performance of this effect and its implications on live performance, a performance study was conducted with 5 amateur musicians (2 drummers, 3 keyboardists). Participants were asked to replicate drum-synchronized effects on synthesizers from popular songs, with and without the effect. Their experiences were contrasted and analysed to assess the effectiveness of cross-adaptive audio effects as a live tool, and their applicability in modern music.

4.1. Test Setup

A two microphone setup was used to take input from the drums, one for the snare and one for the kick. Gate thresholds of -10 dB and -18 dB were used for the kick and snare microphones respectively. The high thresholds allowed the drummer to selectively trigger the effect on specific notes using heavily accented notes.

A MOTU Hybrid MKIII sound card was used as an interface to provide input to Reaper. The native MOTU ASIO drivers were used with a buffer of 256 samples and sampling rate of 96kHz. The cross-adaptive effect was applied on a VST synthesizer representing the instrument.

Routing:

- Kick Input -› Channel 1 of plugin
- Snare Input -› Channel 3 of plugin
- Synth Input -› Channel 5 & 6 of plugin

Table 1: *List of Songs for Performance Test.*

Name	Artist	Start/Stop	Effect
Closer	The Chainsmokers feat. Halsey	1:10/1:31	Level controlled by ASR synchronized with kick and Snare
In the name of love	Martin Garrix feat. Bebe Rexha	0:49/1:03	Kick triggered tremolo and snare triggered mute
All we know	The Chainsmokers feat. Phoebe Ryan	1:14/1:36	Kick triggered fixed duration tremolo

Table 1 lists the three song sections and the respective effects that were selected as test cases. Each song section was performed with 2 groups of musicians with alternating order of performing manually or with the effect. Table 2 shows the order and total time spent practicing and recording for each of the 13 experiments that were conducted. The test conditions and setup were identical for all sessions and symmetrically alternated between different songs and different approaches.

Table 2: *Log of Conducted Performance.*

Expt. No.	Song No.	Drummer	Key-boardist	Method	Time Taken
1	1	A	C	Effect	4:41
2	2	A	A	Manual	12:15
3	2	A	A	Effect	4:23
4	3	A	B	Effect	2:17
5	3	A	B	Manual	7:40
6	2	B	B	Manual	6:43
7	2	B	B	Effect	3:35
8	1	B	B	Manual	7:15
9	1	B	B	Effect	4:42
10	3	B	A	Effect	1:54
11	3	B	A	Manual	7:10
12	1	B	A	Manual	5:15
13	1	B	A	Effect	2:40

4.2. Performer Background

Keyboardist A is a classical pianist with 12 years of training and enjoys listening to ethnic music.

Keyboardist B has played piano for 20 years, sang as an Alto in multiple choirs over time and listens to classical music and orchestral movie soundtracks.

Keyboardist C has played piano for 5 years and occasionally plays other instruments including keyboards, guitars, harps, and ukuleles.

Drummer A studied music at GCSE and A-level on clarinet and started playing drums, bass and singing in bands at 16. He likes music that is simple and elegant, like a tight funk groove over a flashy show of chops.

Drummer B played keyboard since childhood and has played drums for the last 8 years. He plays mostly rock-oriented genres like Blues-Rock, Grunge, Indie, Nu Metal, Punk-Rock.

4.3. Test Case Analyses

The plugin setting for each experiment is available at - https://code.soundsoftware.ac.uk/attachments/download/2230/Table%203.docx.

4.3.1. Song 1 (All we know)

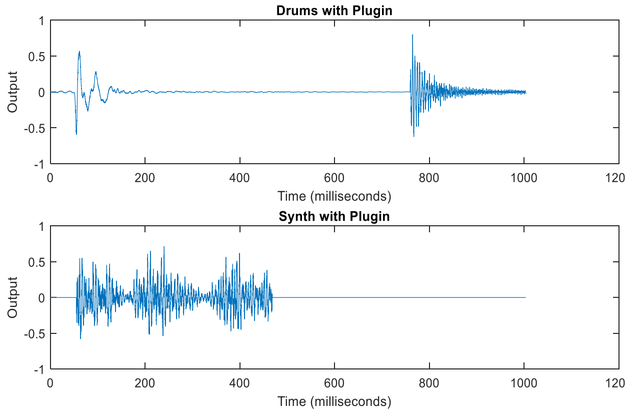

Figure 6: *Song 1 Recording with effect (Expt. No. 4).*

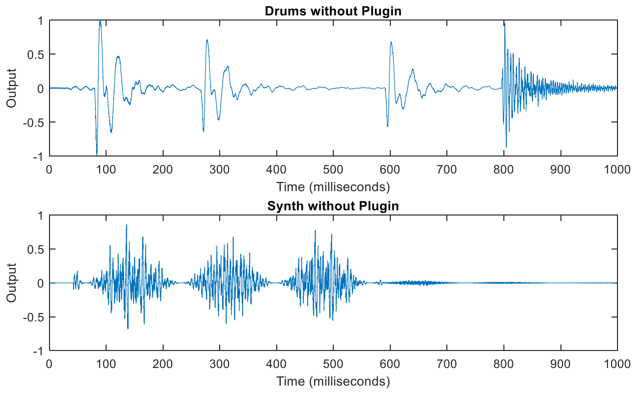

Figure 7: *Song 1 Recording without effect (Expt. No. 5).*

The effect to be replicated was to trigger 3 cycles of tremolo synchronized with the kick from the drummer. The duration of the synthesizer note was supposed to be exactly 3 cycles.

With Effect:

As seen in the drum track recordings shown in Fig. 6 & 7, the drummer had to modify the groove by removing a few double kicks to make the plugin trigger the effects on the synthesizer in the desired manner. The keyboardist did not need to control the onset time and duration of the chord, but only keep holding the correct chord throughout.

In Fig. 6, we see that each synth note was synchronized to each drum onset with a delay of 2.3ms from the edge of the transient of the kick drum hit (see Fig. 12 & 13). The duration of the note was fixed by the sustain duration. This was set to 3 cycles of the tremolo at quarter notes at 90 BPM.

Without Effect:

In Fig. 7 we see the kind of issues that performers faced. First, onset times of the synth and the drum always had a difference due to human error, ranging from 23ms to 98ms. Second, note durations varied between 478 and 580 ms, consistently greater than the desired value of 48 ms. This can be attributed to reaction times since performers depended on the audio cue to complete 3 cycles of tremolo and then release the note.

Moreover, sometimes the phase of the tremolo effect was not perfectly synchronized because of the above-mentioned errors, which can be disorienting for the drummer.

4.3.2. Song 2 (In the name of love)

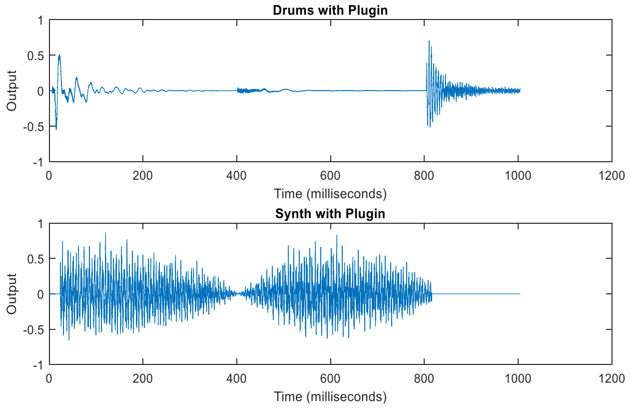

Figure 8: *Song 2 Recording with effect (Expt. No. 7).*

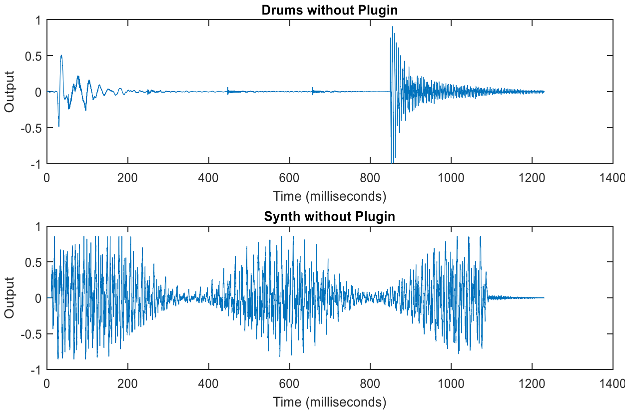

Figure 9: *Song 2 Recording without effect (Expt. No. 2).*

This song required the synthesizer to start a note with full depth tremolo on each kick drum hit, and stop the note at every snare drum hit. This was implemented by using the kick-based trigger to enable output and start an LFO to control the amplitude modulator for a long duration (longer than the length of one bar), and setting the snare-based trigger to mute the output. The manual performance used a regular note-synchronized tremolo effect and required the keyboardist to control the note onset and release.

With Effect:

This song required the drummer to play normally (as seen in Fig. 8 & 9, the drum tracks are identical), but the keyboardist had to hold the notes slightly before the drummer and hold it longer than the song required. This ensured that the note onset and release were controlled by the drum triggers alone. Since the song was slow, this was not difficult and the effect worked very well, see Fig. 8. A constant delay of 2.3 ms was observed between the peak of each drum transient and the note onset/release. This delay was limited by the latency of the sound card and driver.

Without Effect:

Fig. 9 shows that, when manually trying to replicate this song, there were similar onset differences as in the previous experiment. The slower tremolo of this song caused phase errors to be less than in the previous experiment. But due to the slow tempo and abrupt stop at every snare hit, it was very noticeable when the keyboardist released the note later than the snare hit.

4.3.3. Song 3 (Closer)

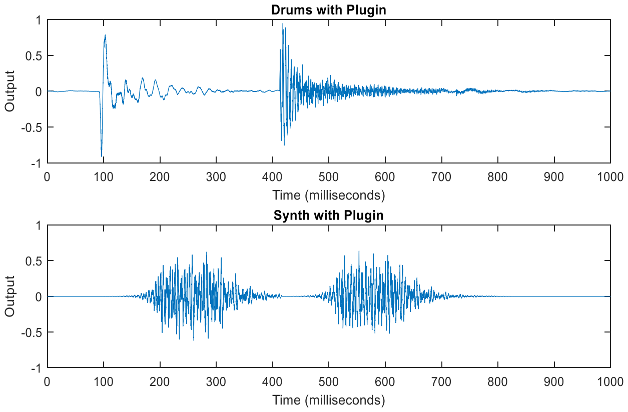

Figure 10: *Song 3 Recording with effect (Expt. No. 9).*

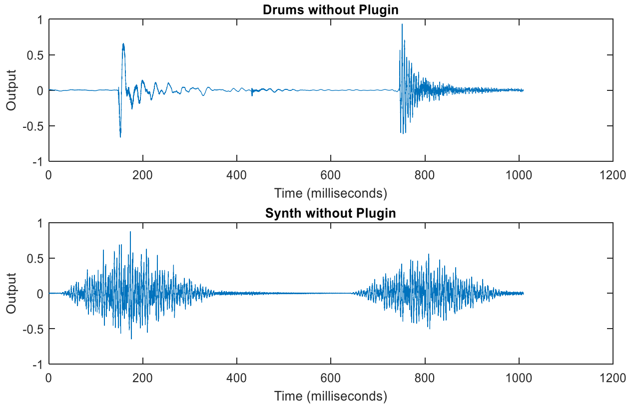

Figure 11: *Song 3 Recording without effect (Expt. No. 13).*

In this experiment, each synthesizer note had ~0.1 second attack and release time, and were played staccato. Each of the synth notes were played only with each kick and snare notes.

With Effect:

The effect had an exponential ASR envelope to control amplitude of the synth output. Due to the exponential nature of the attack curve, the output amplitude grows very slowly initially, staying inaudible for a while. Hence, the perceived onset of the note from the effect was delayed compared to the trigger time of the effect, especially for short attack times (<0.2 seconds). As seen in Fig. 10, performers perceived a lag of about 70 ms while using the effect. This was disruptive for the performers and made it difficult for them to use the effect. We observed that during performance, notes that occurred with the snare hits did not seem to have the distinct lag, although the processing complexity was the same. This was probably because the snare drum itself has longer sustain, thus masking the duration when the effect output is inaudible and slowly blends in as the effect becomes louder and the snare fades away. The song had multiple, fast chord changes. This made the effect more difficult to use since the keyboardist needed to anticipate the drum onset and hit the note prior to that.

Without Effect:

This experiment was trivial as the only variable was the onset times. The difference in onset times in this experiment was greater than the previous experiments since the song was faster and had frequent chord changes. We see in Fig. 11 an instance where the keyboardist's onset occurs significantly before the onset of the drummer.

4.4. Performer Analysis

4.4.1. Drummer A

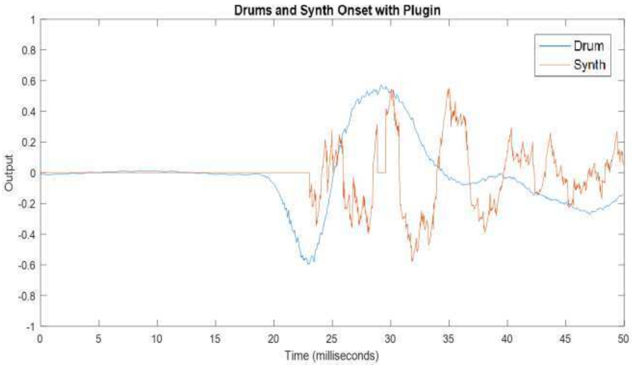

Figure 12: *Drummer A onset detection example.*

Drummer A used the heel down technique while playing the kick drum. Thus, his kick drum hits were soft and had long rise and short decay times. As seen in Fig. 12, the onset detector was triggered at the end of the first half oscillation of the kick drum diaphragm. Drummer A also hit the snare softly with a rim click, which caused the snare to sustain for a shorter time.

4.4.2. Drummer B

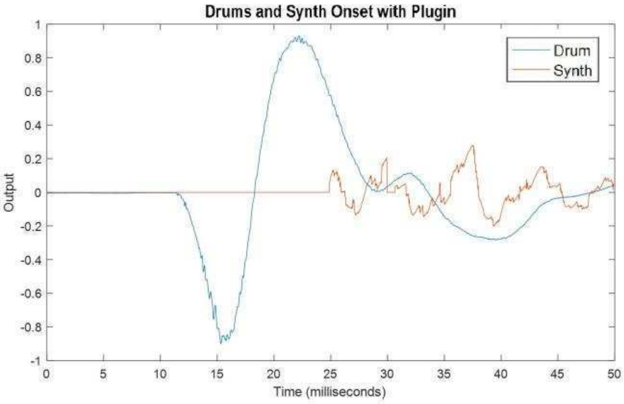

Figure 13: *Drummer B onset detection example.*

Drummer B used the heel up technique while playing the kick drum, thus producing loud hits with short rise and long decay times. As seen in Fig. 13, the amplitude envelope increases until the second half oscillation of the diaphragm, when the onset detector was triggered. Drummer B also hit the snare drum heavily, causing the snare to rattle for a longer duration. This reduced the perceived lag of the effect in experiments 9 and 12 (as seen in Sec. 4.3.3).

4.5. Performer Feedback

A survey was conducted for the performance test and participant responses and comments were recorded. Questions were designed to be answered on a 1 to 5 scale and comments were taken for each question.

4.5.1. Understanding the effect

All participants responded that they understood how the effect worked and felt that the effect was very intuitive. One comment stated that effects with shorter attack times were easier to anticipate and accommodate for while performing. Some performers reported latency with the song Closer which was disorienting and made it difficult to use the effect (see Sec. 4.3.3).

4.5.2. Ease of use

All drummers believed that the effect did not make their performances easier. This can be attributed to the fact that the drummers are constantly concerned with listening to the output from the keyboards to verify if they are playing correctly to trigger the effect as the keyboardist desires. Another factor may be that they were asked to replicate certain effects as they were in a song. So, they were constantly evaluating whether they are performing as expected for the test.

In all cases except the song Closer, all the keyboardists believed that the effect made their performance a lot easier.

4.5.3. Impact on performance

All performers responded that they had to change their performance technique mildly or moderately to make the plugin work. Drummers also noted that with the kick drum, using dynamics to control the effect was very difficult. Thus, when using the effect, they could not change the kick pattern, and potentially lost some of the groove in a song.

The drummers responded that they had to change their groove for the songs Closer and All We Know, while the keyboardists responded that they had to change their performance only for Closer.

4.5.4. Creative possibilities

We received a mixed response when we asked musicians about the possibility to improvise while using the effect. Some performers who could use the effect comfortably believed there was room for improvisation, especially for the drummer to perform a solo while the keyboardist is passive. Others felt that using the effect successfully required musicians to know in advance what the other musician is playing, thus making improvisation difficult.

Most participants said that the effect is well suited for electronic music genres. One comment stated that it could be applied to a broad variety of genres if used creatively.

Some participants believed that this effect might restrict the way musicians compose their music. One participant noted that this need not be true since the effect is a performance tool rather than a composition tool, and might open up interesting opportunities to bands that may not have included keyboards otherwise.

4.5.5. Musical Expression

One drummer responded that the effect opens new dimensions of expression for a drummer but it also gives the drummer more responsibilities. Another drummer noted that constraints on the ability to change the kick/snare pattern in the groove is limiting. And being in control of effects across multiple instruments is interesting but also frustrating for the other instrumentalists.

One keyboardist responded that the effect is more useable when the song is rhythmic, enhancing tightness while simplifying the keyboardist's job. Another responded that it is not possible for a keyboardist to express oneself using the effect, but the drummer has more freedom and ability to express and make the song richer.

5. LATENCY ANALYSIS

The plugin has very little latency (<2.6 ms). For the live implementation, the latency of the sound card, driver and the buffer size used determine the overall latency of the effect. Fig. 14 shows the time difference between the onset of the gated drum note and the onset of the effect.

Offline rendering gives a latency of 0.72 ms (32 samples) from the peak of the transient (can be seen in Fig. 14 by comparing the Drum and Filter Synth plot at Time = 0), which is the minimum latency due to a 32-sample frame length and detection commencing at the peak of the onset transient.

During live testing, a time difference of <12.3 ms (~550 samples) was observed from the beginning of the transient (excitation due to drum hit) and ~2.5 ms (~110 samples) from the peak (onset). This is because physical drums were observed to have an attack transient of 4-10ms and the sound card used with ASIO driver at 256 samples buffer size has a latency of 2.6 ms. Except the experiments related to song 3, none of the performers reported issues due to latency as the effect latency was <10 ms which shows no significant difference in responsiveness of an instrument [9].

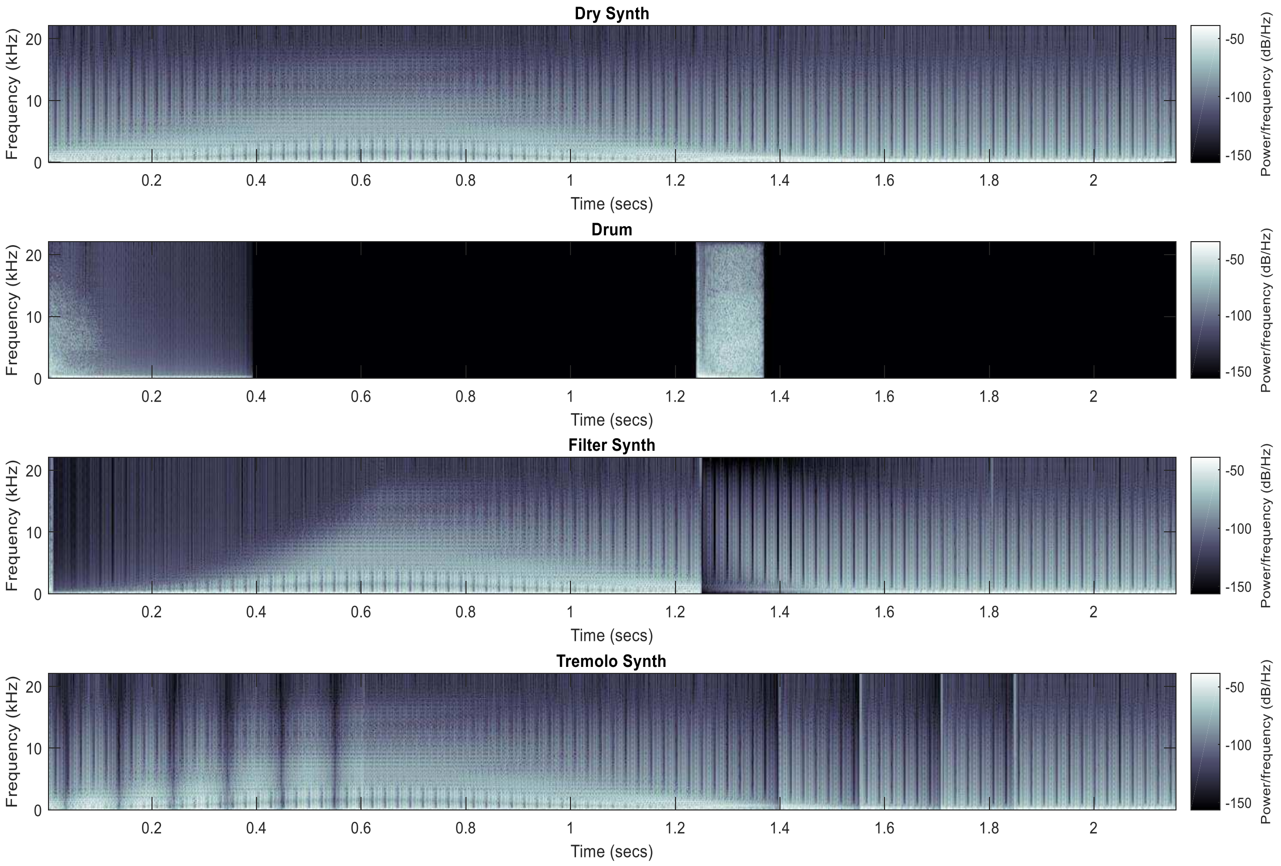

Figure 14: *Temporal view of offline test spectrogram.*

6. RESULTS

Performance tests showed that the effect works better for LFO synchronization and fast attacks. Effects with longer attack times do not provide immediate feedback to the performers, thus giving a sense of lag. This longer attack time setting was a larger contributor to lag than any signal processing latency. Fig. 15 compares the time taken to achieve final recording for each of the experiments. D-A K-B implies the experiment conducted with Drummer A and Keyboardist B. The time required for performers to achieve desired performance while using the effect was consistently shorter than when performing with manual effects. This might be biased for the selected songs and tasks since the effects in the song were difficult to replicate manually and the effect was particularly useful for the given situations.

Sound samples from the performance tests are available at: https://code.soundsoftware.ac.uk/attachments/download/2232/Performance%20Examples.rar

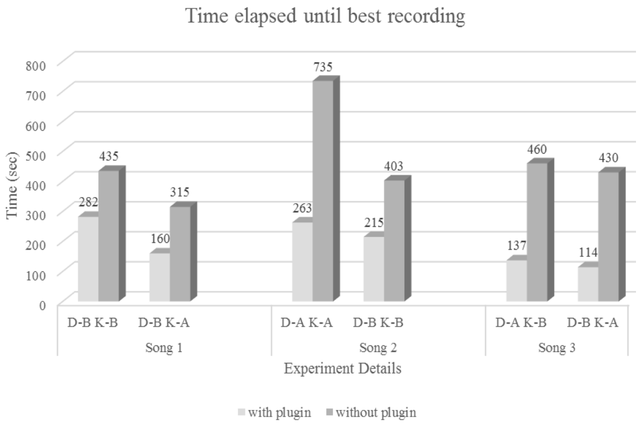

Figure 15: *Experiment duration comparison*.

7. CONCLUSION

We investigated whether the kind of effects created in the studio for popular music using modern digital audio tools can potentially be recreated live using cross-adaptive architectures. We implemented a cross-adaptive audio effect for live performance that enabled the drummer to orchestrate effects across instruments using drum cues. The idea to have effects synchronised to the drum cues is very intuitive and our experiments have reflected the same. Results showed that the cross-adaptive architecture was successful for achieving tasks in several scenarios based on application of effects in post-production, although such limited evaluation may not uncover the most significant challenges in their use.

The use of cross-adaptive effects not only has a drastic impact on the performance of musicians but also potentially affects the way music is composed when using such effects. Better feature recognition techniques to distinguish higher level drum cues like grooves and rolls would enhance the functionality of these effects, allowing them to be less intrusive and more powerful.

8. REFERENCES

[1] V. Verfaille, U. Zolzer, and D. Arfib, ‘Adaptive digital audio effects (a-DAFx): A new class of sound transformations’*, IEEE Transactions on Audio, Speech and Language Processing*, Vol. 14 No. 5, pp. 1817–1831, 2006

[2] E. Perez Gonzalez and J. D. Reiss, Automatic Mixing, *DAFX: Digital Audio Effects*, Second Edition (ed U. Zölzer), John Wiley & Sons, Ltd, Ch. 13, p. 523–550.

[3] S. Hafezi and J. D. Reiss, 'Autonomous multitrack equalisation based on masking reduction,' *Journal of the Audio Engineering Society*, Vol. 63 No. 5, May 2015.

[4] Z. Ma et al., 'Intelligent multitrack dynamic range compression,' *Journal of the Audio Engineering Society*, Vol. 63 No.6, June 2015.

[5] O. Brandtsegg, ‘A Toolkit for Experimentation with Signal Interaction’, *18th Int. Conference on Digital Audio Effects (DAFx-15)*, Trondheim, Norway 2015.

[6] M. Wright, A. Freed, and A. Momeni, ‘Opensound control: State of the art 2003,’ *New interfaces for musical expression (NIME)*, pp. 153–160, Singapore 2003.

[7] O. Campbell et. al., ‘ADEPT: A framework for adaptive digital audio effects’, *2nd AES Workshop on Intelligent Music Production (WIMP)*, London 2016.

[8] I. Jordal, O. Brandtsegg, G. Tufte, ‘Evolving neural networks for cross-adaptive audio effects’, *2nd AES Workshop on Intelligent Music Production (WIMP)*, London 2016.

[9] R. H. Jack, A. McPherson, T. Stockman, ‘Effect of latency on performer interaction and subjective quality assessment of a digital musical instrument’ *Proc. Audio Mostly*, Norrköping, Sweden 2016.