CALREC SOCCER ASSIST

In today's busy production environments, there is increasing pressure on operators to create a more immersive viewing experience. One of the many tasks the audio operator must manage is capturing all important events during live sporting occasions while maintaining a steady background level of crowd and venue noise.

Calrec's patented suite of Assist iPad® apps take advantage of Calrec's CSCP remote control protocol to help ease the pressure in this demanding working environment.

Soccer and Ice Hockey Assist allows operators to simplify the complex tasks of tracking audio during a game, providing an alternative approach for highly experienced operators while allowing those with less experience to create a quality mix very simply.

Fader Assist allows operators to remotely control fader level, PFL, cut/on and aux/main routing on all Calrec Bluefin2 consoles. The ability to remotely control these fader functions from any location within a Wi-Fi range provides a new level of flexibility and convenience for both mixing and set up.

All Calrec's Assist apps are absolutely FREE and available to download at the App Store.

This paper presents the results of subjective and empirical evaluation of Soccer Assist. Using multiple stimulus comparison, event counting, fader tracking and cross-correlation of mixes using different systems, this paper shows that lesser skilled operators can produce more reliable, more dynamic, and more consistent mixes using Assist than when mixing using the traditional fader-based approach, reducing the level of skill required to create broadcast-quality mixes.

Introduction

Mixing the audio for a live football event is a challenging task, requiring a great deal of experience and a high level of mental focus. Soccer Assist removes some of the more functional work allowing operators to focus more on the creative and aesthetic decisions to create a more complete and balanced mix.

In this evaluation of the Assist system the method of capturing the audio for a football event is presented, and the current fader-based method for mixing the audio is outlined. This is followed by an overview of the Assist system and related work. The subsequent sections present a range of tests to compare the performance of Assist and the mixes it creates against the fader-based mixes.

Capturing and Mixing the Action

The audio for many football events is captured and mixed using a tried and tested method with basic configuration and workflow largely the same across all events. Capturing the audio of on-pitch events involves around twelve shotgun microphones, placed at specific locations around the pitch as shown in Figure 1. The sound of the crowd and other ambient audio is captured with additional microphones located above or around the stands.

It is the job of the operator to ensure that the audience at home can clearly hear all key components of the broadcast and recreate, as best they can, the atmosphere of the event. To bring the viewer closer to the action the operator has to ensure that all the important on-pitch events are clearly audible.

All events on the pitch could be recreated for the listener at home by simply fading up all the pitch microphones. This is problematic due to the noise of the crowd being so high – much of it spills onto the pitch microphones. With the combined contribution of the crowd (essentially wide-band noise) from the ambient microphones and all pitch microphones, the desired audio events become masked.

During a game the operator tracks the location of the action on the pitch and fades up the relevant microphone, or microphones, to allow specific events to be heard while reducing the amount of unwanted noise. Typically one to three microphones are faded up at any one time. The operator tracks the action by monitoring a video feed from the wide angle camera mounted above the centre line.

FIGURE 1. MICROPHONE PLACEMENT FOR CAPTURING ON-PITCH ACTION

As can be seen in Figure 2, the microphones that surround the pitch are 'unfolded' onto twelve faders on the audio desk. The operators label each microphone with an arrow symbol extending from one or more lines. The lines denote the edges of the pitch and the arrows show the direction that each microphone is pointing.

The operator must project the location of the action on the pitch onto this 1D representation of the pitch boundaries. This is a challenging mental task, especially considering that the operator must anticipate where the next kick or impact will occur, and ensure that they have faded up the correct microphone, or microphones, before that event happens. Very experienced operators can mix this way without looking at the faders each time they make a change, much like an accomplished piano player does not have to look at the keyboard, but some operators do need to glance down at the faders often to make sure they are in the correct location.

It is the process of simplifying this 1D mental translation that is the key focus of this paper.

FIGURE 2. LAYOUT OF THE MIXING DESK

Calrec Assist

While the Assist algorithm is based on the findings of "A New Technology for the Assisted Mixing of Sport Events: Application to Live Football Broadcasting" by Giulio Cengarle, Toni Mateos, Natanael Olaiz and Pauprevious Arumí, a paper presented at the 128th AES in London in 2010, in practice, Assist presents a representation of the football pitch on an iPad.

The operator inputs the locations and directions of the microphones used to capture the on-pitch action, then touches the pitch representation to indicate where the action is occurring. This basic implementation was extended to incorporate multi-touch allowing an operator to specify multiple points of interest (POIs).

Based on these POIs, the system will select the most appropriate gain to apply to each microphone in order to best capture the on-pitch audio at those locations while reducing unwanted noise from other, non-critical microphones. The approach and algorithm is similar to the process of distance based amplitude panning, although essentially applied in reverse. The system allows for the variation of the effective size of each POI. POIs automatically crossfade to avoid abrupt changes in the audio.

In order to actually affect audio, the gain values calculated by Assist are converted into fader level messages to control faders on an audio console via Calrec's CSCP, a TCP/IP control protocol.

Operational Tests

In order to subjectively evaluate the ability of Assist to create a satisfactory on-pitch mix, it was necessary to acquire representative audio source material. The following resources were captured from a football event:

- Individual audio feeds from each pitch and crowd microphone
- The full audio pitch mix as created by the audio operator
- The full audio mix of the programme, including commentary
- A video capture of the programme

FIGURE 3. ASSIST USER INTERFACE

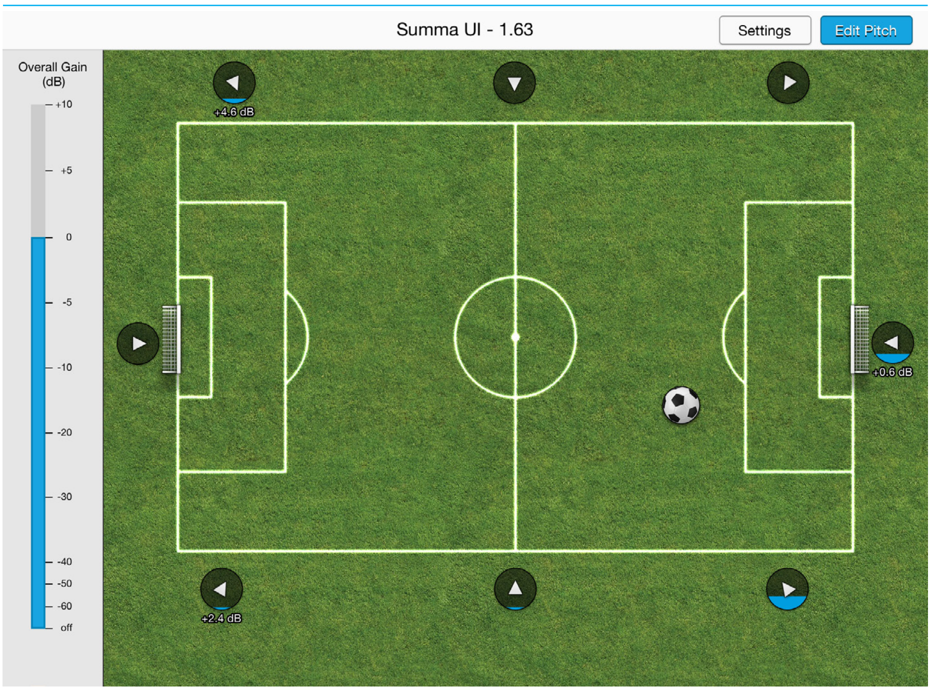

Individual microphone feeds were played back and sent to inputs of an audio console. The microphones were laid out on faders and labelled as they would be in an outside broadcast truck. The video of the game was displayed on a monitor above the console to match the location of the video monitors in a truck. Compression and EQ was applied to the crowd mix to try to match the sound of the pitch mix created by the broadcast operator.

Five subjects, who were technically competent engineers with audio mixing experience but who lacked experience in mixing live football events, were asked to create a mix of a short, three minute segment of the match using the faders on the console. Each subject was then asked to make a mix of the same segment using Assist to mix the pitch audio. Subjects were allowed one pass at the mix using each method to avoid the operator being able to remember the flow of action which would allow them to better predict how to mix the audio. The subjects were allowed to familiarise themselves with the operation of both methods prior to the recording of the mixes.

Capturing the Mix Data

For each mix, playback of the audio and video was started at a known location. The recording would end after a pre-defined period. This ensured that there was no difference in the overall length of each mix, and that the start and end positions of each mix were the same.

Mix data consisted of individual post-fader output from each source on the console, as well as a full stereo mix of the pitch and crowd. Fader level data from the audio console was captured at a rate of 10 Hz.

Listening Tests

In order to get a subjective opinion on the performance of Assist compared to fader based mixes, a multiple-stimulus listening test process was use. Participants were presented with a statement and asked to compare mixes based on this rather than comparing the various sonic or artistic qualities. The statement was: “The mix captures all important on pitch events.”

The samples consisted of a 30 second segment from the five fader mixes and the five Assist mixes created during the operational tests. An anchor was provided as a source in the test and was a mix consisting of the crowd noise only, with no contribution from any pitch microphones. All mixes were accompanied with the appropriate video content.

Eight subjects participated in the listening tests. Some had a background in audio equipment design, some from a practical audio background, and others had no professional audio experience at all.

The audio mixes and accompanying video were looped. Each mix was assigned to a random audio channel before each listening test started. Participants were free to jump between the different mixes as required using the solo buttons on the fader strips.

Listening Test Results

As can be seen in Figure 4, the Assist mixes consistently score higher than the fader based mixes and generally have a narrower confidence interval. Opinion

FIGURE 4. MEAN MIX RATING WITH 95% CONFIDENCE INTERVAL

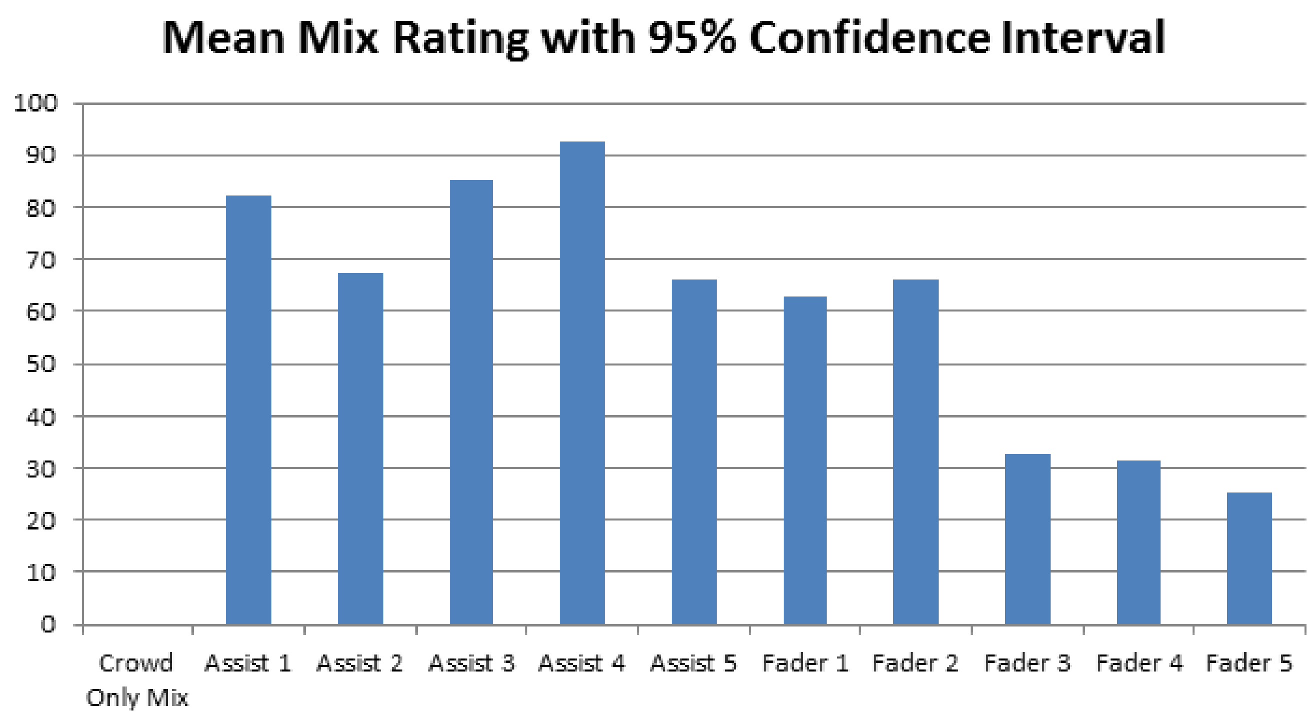

FIGURE 5. EVENT TYPE

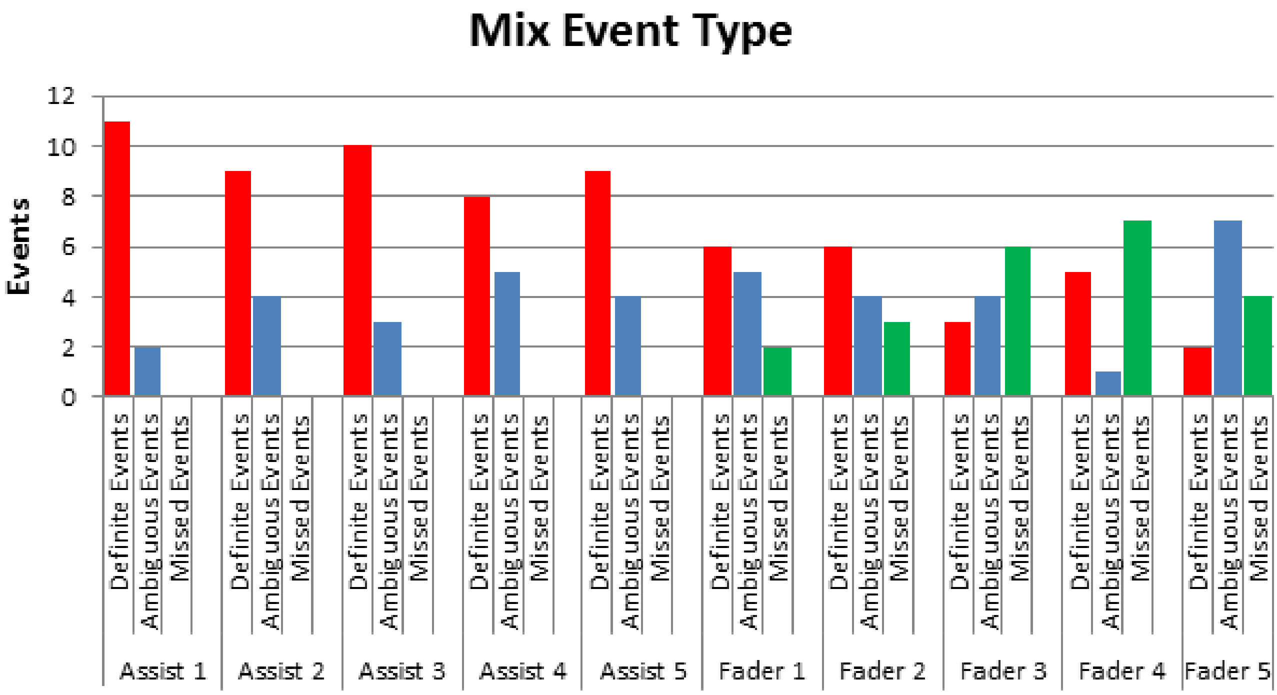

of the fader based mixes had greater variation between subjects but was generally lower than the Assist mixes.

All subjects correctly identified the anchor (the crowd-only mix). Assist mix 4 was consistently rated highly, with all subjects scoring it above 80%, with half of the subjects rating it as the best mix. It was the most consistently rated of all the mixes.

Event Capture
It was noticed during the listening tests and while observing the mixing processes, that Assist seemed more capable of capturing all of the on-pitch events compared to a fader based approach, where a number of key events were conspicuous in their absence. To determine this, the number of definite, ambiguous and missed audible events were counted in each of the same mixes that were presented for the subjective listening tests.

The audible events were defined as visible events that created an expectation of accompanying audio. These were noticeable hard kicks and passes, body contact and ball bounces; there were twelve clear events in total in the examples used. Each mix was listened to three times while viewing the corresponding video. A simple scheme was used that related the numbers one, two and three to definite, ambiguous and missed events.

1. Definite event: a clearly audible accompaniment to the visual event that met the listener's expectations.
2. Ambiguous event: an audible accompaniment to a visual event that was either barely audible, or that didn't fully meet the expectation of the listener.
3. Missed event: a visual event with no accompanying audio event. Listener expectation was not met at all.

As can be seen in Figure 5, there is a clear distinction between the Assist mixes and the fader mixes. The main statement that stands out is that across all the Assist mixes, not a single event was missed; including the more 'difficult' events such as ball bounces. The level of definite events was also consistently higher than in the fader based mixes.

There was much more variation across the fader mixes, with some performing much better than others. We can also see that in three cases, the number of missed events exceeded that of the number of definite events.

Mix Correlation
Observing the creation of the mixes, and listening to the stereo mixes afterwards, it was noticed that the Assist mixes seemed to be quite similar to one another. This is presumed to be the case due to the way that Assist encourages operation. While it can be used in a number of ways, the

most obvious and simple method is to use the location of the ball as it moves around the pitch as the POI. Every subject in each test who followed the ball, moved the POI along a similar path, resulting in a similar output from the processing.

With the fader-based mixes, subjects were observed to take a variety of approaches. Some subjects chose to open only the one or two microphones nearest to the action and open and close different microphones as the action moved around the pitch. One subject chose to open more microphones at once, sometimes five or six, while keeping them all at a lower level to try to reduce the chances of missing specific events with the compromise that an event captured may be at a lower level than if they had use fewer, more relevant microphones.

As the subjects also all found it difficult to transform the two dimensional location of action on the pitch to the one dimensional fader layout, they were often slower to update the faders to reflect the 'correct' location. Some were faster at this than others. These differences in approaches and difficulties in location could contribute to the lower mix correlation between fader based mixes.

Further Enhancements
Following feedback from the research sessions, Calrec made further improvements to the system and trialled these changes with real-world users. These refinements more accurately emulate the decisions and fader moves of experienced engineers.

Ball tracing: it was discovered that accurately tracing the location of the ball did not product the best result. For example, when the ball is kicked from one end of the pitch to the other, tracing the path of the ball causes all mics from one end of the pitch to the other to quickly fade in and out in sequence.

A preferred method was to simply tap one end of the pitch to set the first POI at the source of the kick, then tap the other end of the pitch to set a POI where the ball will end up. The crossfade system creates a smooth and pleasing transition between the two locations, more accurately emulating the fader moves of an experienced engineer.

Individual mic selection: sometimes it is necessary to leave a mic open at a certain level, or to leave it faded out if it is faulty or picking up undesirable noise. This is easy with manual fader control – just don't touch the fader. When all mics are under control of a single algorithm, this requires more thought. Should this be required then a double tap on a mic in the UI will stop the app affecting that fader. Another double tap will bring it back under control again.

Gain selection: on occasions where a particular mic needs a bit of extra gain, or should be favoured against other microphones, a quick swipe up on a mic in the UI will give additional gain until the user swipes down, returning it to normal.

Manual override: it is important that the user can put Soccer Assist down and take over manual control on the faders at any time. Touching a real fader always overrides the App. As the App is an external controller it can be switched off at any time, leaving the desk exactly where it was.