APPLYING AUTOMATION TO HIGH QUALITY/LOW COST ORCHESTRA LIVE STREAMING

D. R. Chalmers

BBC Scotland, UK

ABSTRACT

Many publicly funded orchestras are looking to engage new audiences but capturing an orchestral concert well for live streaming has always been expensive and complicated so limited budgets prevent it from being a regular occurrence.

The use of robotic pan-tilt-zoom (PTZ) cameras is a way of reducing cost. These cameras do, however, require the use of a high number of different shots to avoid becoming too static and boring for the viewer. This then requires more PTZ operators and starts to negate the original savings.

This paper outlines some experiments that we have conducted with the BBC Scottish Symphony Orchestra to repurpose and adapt an existing multi-camera automation tool designed for capturing pop music performances to a different genre and style of music.

We will consider if this approach can help to keep costs down while maintaining the visual interest and complexity of a high-shot-count orchestral capture.

INTRODUCTION

BBC Scotland has an objective to expand the audience reach of our in-house symphony orchestra, the SSO, who perform regular live concerts at their permanent home in the City Halls in Glasgow. We would like to increase the number of performances that are available live to a wider audience through live streaming to the BBC website and subsequent VOD viewing via the BBC iPlayer.

Unfortunately, the process of covering a live symphony orchestra performance well for live streaming using traditional OB production techniques is prohibitively expensive to do on a regular basis on the small budgets available for live streaming.

To address this many orchestras and concert halls have already explored the use of robotic pan-tilt-zoom (PTZ) cameras as a way of reducing cost.

They can cost less to purchase, allow one operator to control multiple cameras and are small and unobtrusive so are less distracting for both performers and audiences.

They have the facility to be able to easily store and recall preset shots.

There are significant limitations, however, as the cameras are generally fixed to a single vantage point on or near the stage and can't easily move around like a jib or tracked camera without additional kit.

While there have been significant improvements in the motor control that allow some 'in vision' movement, it's still a lot harder to do dynamic movement or to react quickly to unplanned requirements.

However, by having many more framing and focusing snapshots saved as presets (up to 40 per camera is common) - it's possible to avoid the shots becoming repetitive. But - as you increase the number of presets saved per camera, you also increase the workload for the remote camera operators, who are recalling them manually in time for the next cue.

This can become very stressful for the operators, especially on long, complex classical pieces with hundreds of shot changes during the pieces.

AN OPPORTUNITY FOR AUTOMATION

We considered that a system for automated PTZ preset recall could potentially help in this context.

In some ways classical music played by an orchestra really lends itself to an automated approach because most of the performers remain relatively static (with some exceptions!); improvisation is fairly unusual, and a description of the composer's original intent already exists – the score.

There remain, however some significant challenges, specifically:

- Someone still needs to decide what visuals best represent the music at any given point.
- Matching the timing of any automation with the real-life timing of the orchestra performance can vary considerably, even from one run through to another! There is no easy solution for 'beat-matching'.
- Orchestral pieces are often quite long – up to an hour and may only be played in full once during rehearsals which doesn't give much chance to learn or adapt things.

So what is required is a PTZ automation system that enables the capture from an orchestral score of the key information; allows the straightforward preparation of camera preset cues in advance; allows easy adjustment of timing to match the actual orchestra performance and gives sufficient flexibility to allow for unforeseen changes.

To the best of our knowledge such a system did not actually exist yet, so we faced the prospect of either building one ourselves or finding something similar and adapting it.

CUEPILOT – AUTOMATED CAMERA CUE CARDS

We had a serendipitous encounter with a company at the IBC trade-show in 2018.

CuePilot (1) are a company that already produce a software-based camera cue card system that is well established in the market for high-end TV production of live music events such as the Eurovision Song Contest.

It relies on a click track or playback system generating timecode to keep the system in sync and can show the camera operators what shot they should be lining up next and how long till it's needed.

An iPad or tablet app is used in lieu of the paper camera cards normally used, and it is synchronised to a server running in the gallery.

It also can directly control the vision switcher so that the right camera shot is cut up at exactly the right moment without requiring an operator to intervene.

Our first thought on seeing a demonstration of this system was that it was extremely close to what we needed and could just require the addition of PTZ recall to make it usable in this application.

Fortunately, the developers understood our requirements and offered to implement the PTZ recall function.

They added support for the Sony VISCA over IP protocol (a common protocol used by a number of PTZ camera manufacturers).

The software now allowed the user to specify whether a camera was manual or PTZ and to set up a list of preset shots for each PTZ camera which would send the correct VISCA commands to recall that preset when triggered.

Since the software could 'look ahead' and see what the next preset for each camera was going to be, and since it was controlling the vision switcher (so knew which camera was currently 'on air') it could recall the next preset as soon as the previous one was no longer needed.

TESTING

Once the developers had implemented the PTZ recall feature we arranged to run an initial test with a full orchestra. Following the success of the initial test we used the system in production for three further live streams of SSO concerts.

These were live streamed to the BBC website as well as other platforms and made available afterwards on the BBC iPlayer VOD service.

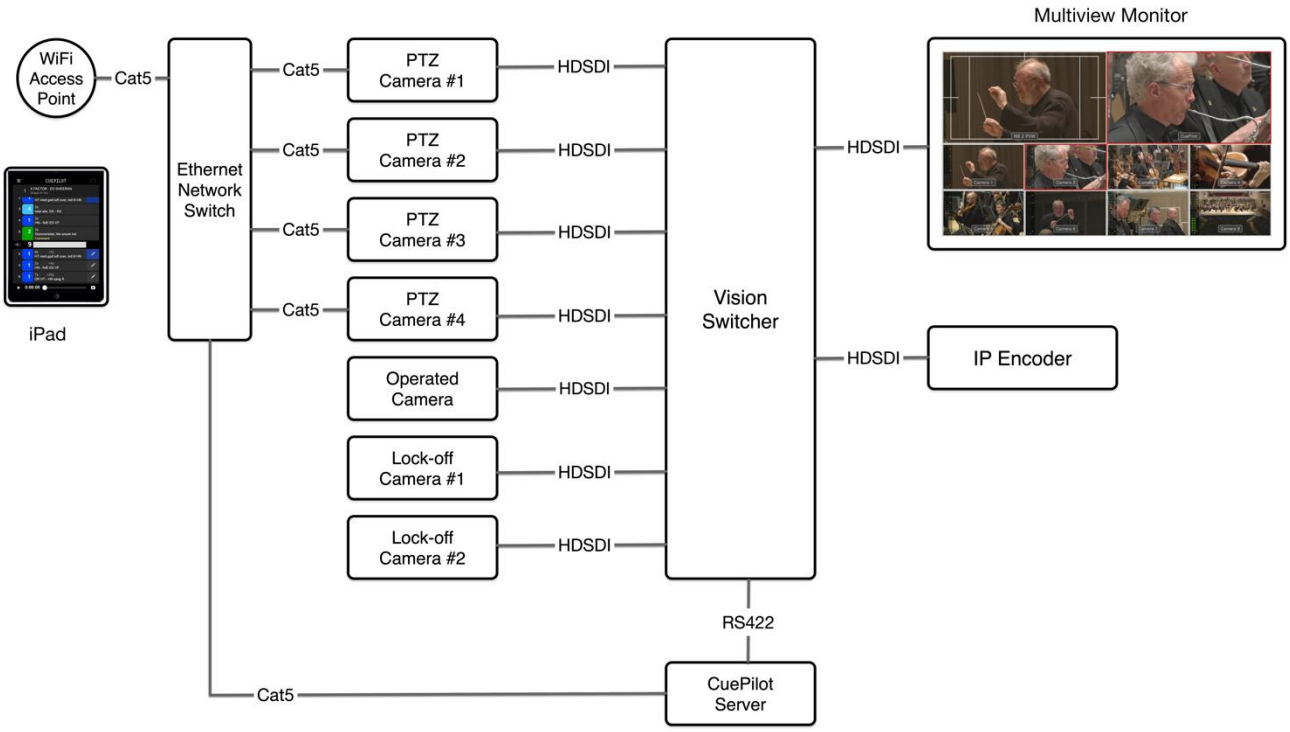

Figure 1 – CuePilot test system setup

The setup for the initial test is shown in figure 1. We had four broadcast quality PTZ cameras remotely controlled by a single control panel and two locked off cameras – a wide shot of the stage and a reverse shot of the conductor. Then we had a single EFP style shoulder-mount camcorder on a tripod at the rear of the hall with a long zoom lens (48x) controlled by a dedicated camera operator.

So 7 camera sources in total fed into a vision switcher which was controlled by the CuePilot server.

We also used iPads running the CuePilot CueApp to show the upcoming shots to the camera operator and to the remote PTZ operator.

For the subsequent concerts we increased the complexity of the system by dropping one static camera and adding two additional PTZ cameras, taking the total camera count to eight, with six under CuePilot control. This allowed us to position some of the cameras in among the performers onstage for more close-up shots which add to the visual experience.

We also used a larger ATEM vision switcher with 2 mix/effect busses. This allows the CuePilot system to drive one M/E bus and present the output of that bus to the other bus which serves as the main programme output. This allows an ‘emergency manual cut’ option if, for some reason the CuePilot software failed or got out of sync during a performance.

The following sections outline some of our findings:

Preparation

Prior to a performance it is necessary to prepare the run-order in the CuePilot tool.

The director had already marked up the score with their choice of shot for the entire piece, specifying a shot number, a camera number and a preset number in each case. An example of this is shown in Figure 2.

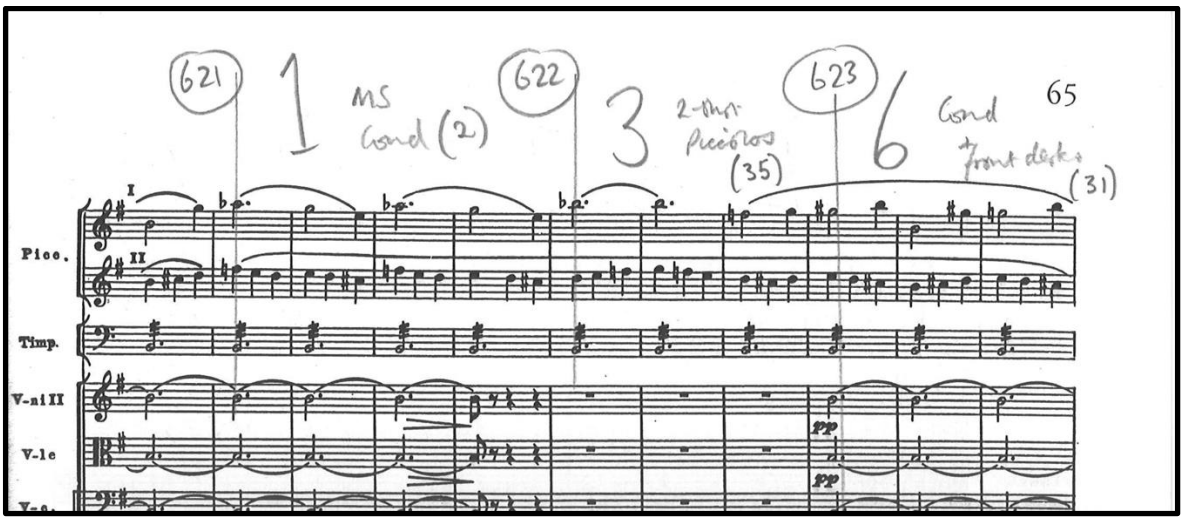

Figure 2 – example of marked up score

This produced a list of preset shots per camera and these were then entered into the CuePilot system to produce a timeline of cues.

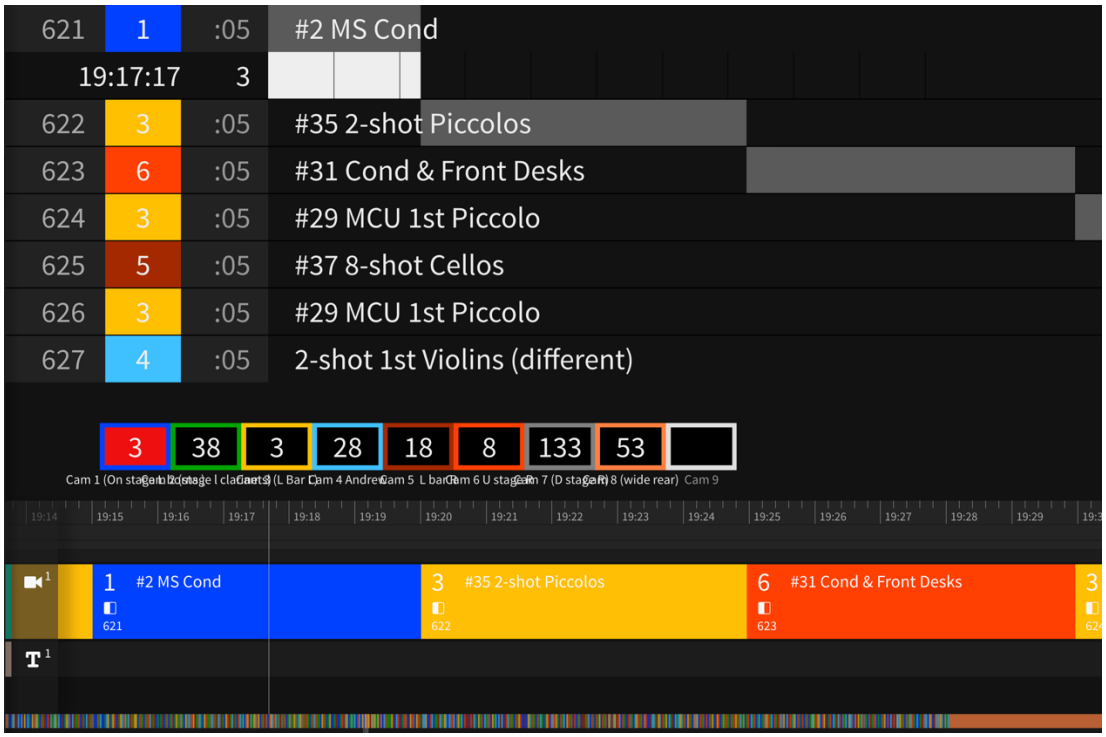

Figure 3 – corresponding CuePilot timeline

The main challenge came when we started to try to map the shot timings against a commercial recording of the same piece imported into the tool.

It became clear that the method normally used in CuePilot of setting a master tempo for the entire piece and then using beats/bars markers to navigate through the piece was not really usable. The tempo variations within the piece required multiple time signature and tempo changes which, while possible within the tool, was a very time consuming process.

We decided at this point to abandon any attempt to match the exact timings of the recording and instead to just try to match the bar numbering on the score.

For the later productions we went even further and didn't even try to set the correct timing – all cues were added with a default 4 second gap. This worked well for the director, who was getting their timing info directly from the score, but less well for the camera operators who had lost any sense of how long it was till the next cue.

Synchronisation

When we started trying to replay the cues based on the prepared timing we very quickly realised that it wasn't going to work. The operator was spending all their time worrying about keeping in sync and no time at all thinking about the actual shots.

We called a halt to the testing and after some discussion the developers retired to their hotel and quickly implemented a new mode which they called 'step mode'.

In this mode the timeline still played but, if it reached a cue point and the operator had not pressed the 'next cue' button, it would pause.

Equally if the operator pressed the next cue button early the timeline would jump ahead to the next cue.

It's a very intuitive approach as you are still taking the shots when you want them but all the other parts of the directing and vision switching process are handled for you.

In testing this proved to be very effective. The operator was able to keep in sync - as long as they knew where they were in the score.

Manual cut-aways

Another feature added during testing was the ability to manually over-ride the currently cued camera shot manually by pressing a number key matching the camera number.

So the operator could decide not to take the next cue but to tell CuePilot to cut up one of the currently available camera shots instead.

This was useful if a particular shot had become unusable for some reason, or if a camera failed or lost connection.

Interestingly, while it was easy to cut away from the timeline of prepared cues, it proved more difficult to switch back again in a seamless way. This is an area that needs further development.

Crewing

We settled on a crewing arrangement where one person just followed the score and called out shot numbers, while the other person operated the CuePilot software, stepping through the shots when the right time had been reached. A third person would be tweaking the recalled preset shots, doing fine adjustments on focus, exposure and framing.

In the later concerts we noticed that while the director was concentrating on the cue timings it was hard for them to be across the different camera options and make decisions about skipping unusable shots etc. They ended up still looking at the score, even though there was also a score reader shouting out cue numbers.

An alternative model would be to have the score reader responsible for cue timing, triggering the cues directly via a second USB keypad, while the director focuses more on the shot selection and overall 'look'. This is something we would like to experiment with further.

FURTHER POTENTIAL USE-CASES

It could be argued that what we had demonstrated was just a more complex way of achieving something that could, and is, already done manually now. While that's possibly true in this specific test case, given that we had hired in an experienced multicamera director and we 'only' had 4 PTZ cameras – we believe there are significant opportunities to scale this approach in a number of interesting directions.

We identified the following potential use-cases:

1. Increasing the visual complexity of a performance capture without a corresponding increase in cost. So adding more cameras to give more shots and angles without having to add more operators and increase costs.
2. Decreasing the technical skill level required to achieve a usable result with a small team. An example of this would be for a radio visualisation of a live music performance, something that is beyond the scope of a typical PTZ type setup in radio studios.
3. Splitting out the pre-vizualisation and preparation tasks from the operation tasks. This would maximise the value from the more specialist staff to choose shots and angles, without then requiring them to be onsite during the actual recording. Someone else could 'play back' their timeline without needing the same specialist knowledge. This, combined with the cloud-syncing capability of the CuePilot could enable an interesting decoupled workflow.
4. Exploring non-traditional ways of presenting the performance to audiences that would not otherwise be possible. One example of this would be to create multiple 'mixer out' feeds at the same time, that are all under CuePilot control that together make a quad split 'montage' of angles all of which the audience see simultaneously.

CONCLUSION

Having adapted the CuePilot tool to this rather different use-case and having been able to use this experimental workflow for four full performances, we now believe that there is significant potential in using automation for orchestra live streaming.

This approach allows you to add more cameras/angles/cuts to give a more engaging viewing experience without multiplying up the cost.

The system can also be retro-fitted to existing PTZ camera installs.

There remain some constraints to the current implementation and it retains quite a few manual steps particularly in the preparation of cues and their timing.

We would be very interested in finding a way to capture timing directly from the score without needing to enter it manually and this is certainly an area for future development.

We believe this approach helps to achieve the goal of being able to live stream more performances in a more visually interesting way for the same budget.

REFERENCES

1. https://www.cuepilot.com

ACKNOWLEDGEMENTS

I would like to thank Johnny, Alan and Jenn and other colleagues at BBC Scotland for their contributions to this work. I would also like to thank Kriss Russman and Morag MacIntosh, along with Per and Jakob from CuePilot.