SMPTE

MOTION IMAGING JOURNAL

Covering Emerging Technologies for the Global Media Community

UHD (2160/HDR) and HD (1080/BT.709) Simultaneous Workflows at the CBC

By Pierre Hugues Routhier

Introdução:

À medida que as fronteiras entre a televisão, o *streaming* e as redes sociais diminuem, as organizações de transmissão tradicionais se deparam com fluxos de trabalho cada vez mais complexos e uma lista, em constante crescimento, de opções para captura, pós-produção e entrega. Reconhecer proativamente que um aumento nas opções leva a um aumento explosivo nos custos e no tempo, devido à complexidade da cadeia, é dever de todos os gestores empenhados, sempre pela melhoria do processo. O artigo desta edição dá uma visão sistêmica de como unir diferentes times em prol de um *workflow* único do UHD/IP até o além. Boa leitura!

Tom Jones Moreira

Abstract

As video content evolves from traditional television into more complex and varied forms of media, more and more of our content creators capture in higher resolution, frame rate, and dynamic range instead of traditional high-definition (HD) to give themselves creative latitude in post-production. Generating ultrahigh-definition (UHD) content comes with a lot of challenges for a traditional HD infrastructure. For example, which version(s) should be archived: The source, ungraded asset, the HD version, the social media version(s) like square and vertical, or a combination of these? If content is shot in UHD, for example, and then cropped in post, should the final, cropped HD product be the one we archive, or the UHD original? If we archive the original, how do we store information about the cropping parameters? If color and dynamic range are wider in the source than the final render (e.g., SLog3/SGamut3), which 3D look-up table is the right one to use? Should 3D LUT also be archived with the source asset? How do we link a specific asset to a specific 3D LUT for posterity? Also, public broadcasters have legal obligations with regard to content preservation, but these do not necessarily specify the format(s) for archival purposes, so we need to balance obligations, cost, and future value of assets in the process of designing such workflows.

Keywords

3D look-up table (LUT), 4K, broadcasting, high-dynamic range (HDR), post-production, SLog3, ultra high-definition (UHD)

Introduction

In 2018, the Canadian Broadcasting Corporation (CBC)/Société Radio-Canada performed a series of tests in their television studios to validate the feasibility of producing ultrahigh-definition (UHD) and high-dynamic range (HDR) content. Following these tests, a series of post-mortem sessions were held with the technical and creative crews who attended. The consensus was that bringing production up to UHD was not a significant challenge, and therefore the requirements for the *Nouvelle Maison de Radio-Canada*, a new, all-IP broadcast center under construction in Montreal, included the capability for end-to-end UHD production, from capture, mix, post-production, and playout to archives.

In 2018, the Canadian Broadcasting Corporation (CBC)/Société Radio-Canada performed a series of tests in their television studios to validate the feasibility of producing ultrahigh-definition (UHD) and high-dynamic range (HDR) content. Following these tests, a series of post-mortem sessions were held with the technical and creative crews who attended. The consensus was that bringing production up to UHD was not a significant challenge, and therefore the requirements for the *Nouvelle Maison de Radio-Canada*, a new, all-IP broadcast center under construction in Montreal, included the capability for end-to-end UHD production, from capture, mix, post-production, and playout to archives.

For HDR, however, the conclusion was twofold. From a technical standpoint, not all content would benefit, and it would entail significant changes in art direction, lighting techniques, and a lot more light in the studios. From a commercial standpoint, the Canadian distribution market has no appetite for HDR yet, so it became difficult to justify the extra training and cost to deliver HDR content at this point in time. It was, therefore, decided that studio equipment would be HDR-compatible (which came with the UHD requirements anyhow), but would keep producing content in International Telecommunication Union-Radiocommunication (ITU-R) Recommendation BT.709 for the time being. That did not mean, however, that the Corporation could not keep moving forward with HDR and raising the technical expertise in-house for such

productions. In fact, some other CBC divisions outside the studio umbrella were already capturing UHD and HDR content using file-based workflows, but in a wide variety of formats and mainly for HD and standard-dynamic range (SDR) delivery. The desire to lay down the foundations for a near future where HDR is produced by the CBC under industry delivery specifications led to the mandate of "cleaning-up" capture and post-production standards beyond HD.

Legacy Capture Standards Versus Emerging Formats

A thorough analysis of file-based content generated at more than 50 CBC facilities showed the scope of the challenge we faced. As there were no official file-base format guidelines up to that point (officially, everything was still captured from BT.709 HD SDI sources), we found more than 100 different combinations of color gamut, dynamic range, spatial resolution, frame rate, codecs, and containers.

The wide variety of formats meant that it was difficult to share content across productions, as shots would look different and require heavy post-production. It also posed a significant challenge to our content preservation teams, especially since not all formats are clearly defined with standard metadata, and not every operator used the same exposure levels.

The lack of standards for formats beyond HD meant that a significant budget had to be allowed for the conversion of source assets, which, sometimes, also meant interpretation. For example, in certain cases, the shots were captured in 4K with the intention of cropping into the source image for the final product (a form of pan-and-scan, basically). However, the information about how the final image should be cropped was not available in the metadata, so archivists would have to make creative decisions: either to archive as-is (meaning the source cannot be used in production without post-production), scale down (thus losing resolution), or pan and scan. Since our archivists are specialized in content management, not post-production, nonconforming shots are added to the post-production workload and the costs charged to the production unit.

By the end of 2019, the backlog became significant enough to justify the design of new workflows, incorporating both UHD/HDR and HD deliveries without burdening the post-production department.

Codec, Resolution, and Color Volume Selection

As a vast majority of file-based cameras at the CBC is provided by one manufacturer, *Sony*, it was decided to use the SLog3/SGamut3 color volume for HDR. Beyond manufacturer prevalence, the following factors led to the decision.

- For dynamic range, SLog3 is compatible with both SMPTE ST 2084 (Perceptual Quantizer) and BT.2100-1 hybrid log gamma (HLG) formats (**Fig. 1**). ST 2084 is a more common electro-optical transfer function (EOTF) for file-based production workflows, while HLG is more common in live broadcast. Having a source format that can convert to either is preferable for CBC, as the source assets may be used in the future in either case (e.g., prerecorded TV program or live sporting event). To provide a quality image in all formats, a minimum of 10-bit coding is required, as internal CBC tests have shown significant banding artifacts at 8 bits.
- For color gamut, SGamut3 (**Fig. 2**) is wider than display-referred standards ITU-R BT.709 (HD) and BT.2020 (UHD). It is the closest gamut to the camera's native Bayer filters and therefore provides more room for adjustment in post-production, which is why it was chosen. Production teams used SGamut3.cine in the past, but since it is destructive to the source image and smaller in size (it was designed to emulate film print characteristics), it is more appropriate for Digital Cinema Initiatives (DCI) workflows in DCI-P3 than it is for HDR in BT.2020.
- For efficiency, the SLog3/SGamut3 color volume is much smaller in size, at 10-bit per color, than RAW formats like SonyRAW, which use a scene-linear, 16-bit per color format. With the volume of content generated by a national public broadcaster like the

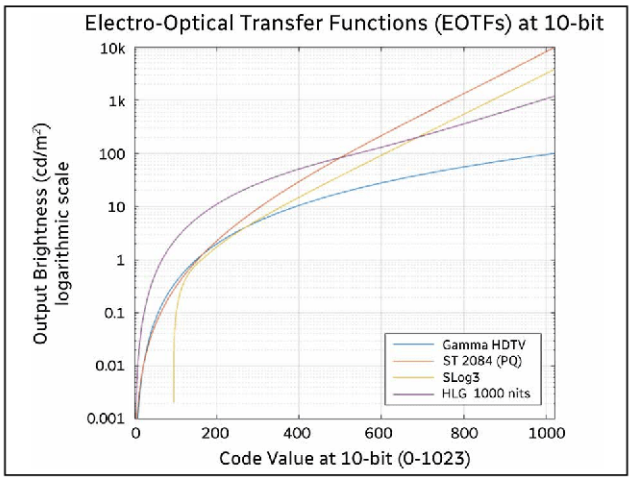

FIGURE 1. ST 2084 (PQ), SLog3, and HLG EOTF at 10-bit.

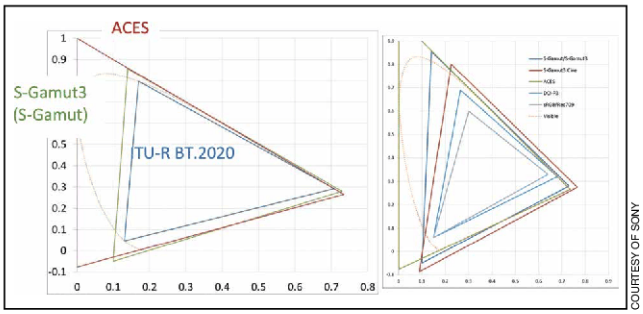

FIGURE 2. SGamut3 versus BT.2020 (left) and BT.709 (right) (courtesy of Sony).

Input reflection	0% Black (IRE0%)		18% Grey (IRE20%)		90% White (IRE100%)	
Output	IRE	CV	IRE	CV	IRE	CV
S-Log3	3.5%	95	41%	420	61%	598

FIGURE 3. CV at different exposure levels—Sony SLog3 (courtesy of Sony).

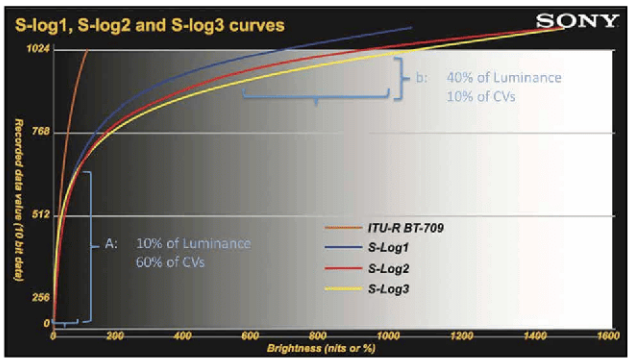

FIGURE 4. Dynamic range versus CVs—log formats (curves courtesy of Sony).

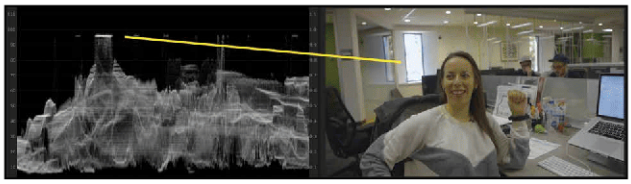

FIGURE 5. Over-exposed SLog3 image (before 3D LUT).

CBC, RAW workflows are not considered for normal production.

If a production wishes to use another color volume than SLog3/SGamut3 or BT.709/Gamma 2.4, then it is the production's responsibility to convert their source assets to BT.709 prior to delivery.

Issues With Nondisplay-Referred Capture

Producing content using formats that are not meant to be viewed as-is and thus require transformations or objective tools to be evaluated, comes with its load of practical issues. Here are the most significant ones we ran into at the CBC and their effects.

Exposure Specifications

One of the major challenges productions face when using formats that are not display-referred is that it is difficult to subjectively determine whether the image is properly exposed or not. One can get a false sense of security by relying on the subjective image presented on a camera's viewfinder, or an external display, which is usually the result of a 3D look-up-table (LUT) applied to the source image, and not what the camera actually captures. Each format has its own exposure specification (**Fig. 3**), which assigns specific code values (CVs) to relative exposure levels.

The failure to follow these specifications will not necessarily lead to an image that is unusable. However, it will prevent the automatic use of standardized

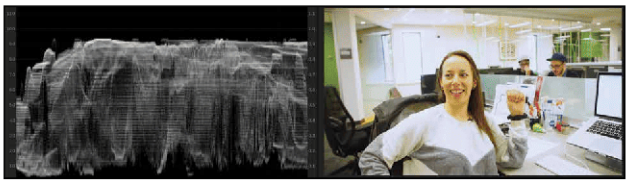

FIGURE 6. Applying technical 3D LUT on an over-exposed image, leading to over-exposed BT.709 final image.

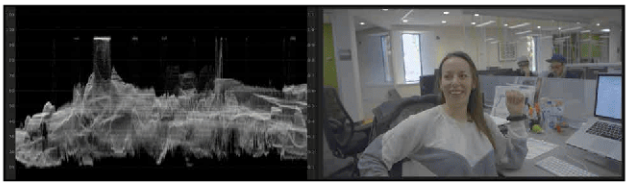

FIGURE 7. Exposure correction of source to match SLog3 levels prior to 3D LUT.

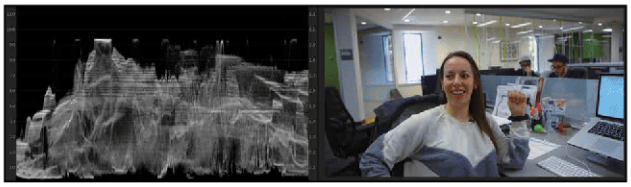

FIGURE 8. Final BT.709 image after applying standard 3D LUT to the exposure-corrected source—notice the unrecoverable, clipped highlight in the back window.

FIGURE 9. Under-exposed section of a SLog3 image (left) and the resulting BT.709 output (right).

FIGURE 10. Exposure-corrected SLog3 image (left) and the resulting, noisy BT.709 output (right).

transformations or 3D LUT, as these rely on strictly following exposure specifications.

Contrary to scene-linear formats, which allocate equal range to all CVs, scene-referred log formats allocate a wider range to high-exposure CVs and a narrower range to low-exposure CVs (**Fig. 4**). It is an efficient form of compression, as it follows the human eye's natural, higher sensitivity in the darks. However, it becomes an issue in post-production when images are under-exposed or over-exposed, as manual exposure compensation must be applied before converting the source assets to display-referred BT.709 and

FIGURE 11. UHD frame for HD crop—notice the presence of rigging, making the image unusable "as is."

BT.2100, with the risk of encountering clipping, banding, and dark noise artifacts in the final image.

Over-Exposure

Over-exposing the image leads to unrecoverable, clipped highlights. As **Fig. 5** illustrates, the window in the back of the image (right) is over-exposed, which can also be seen on the waveform monitor (left), where highlights are at 95%, instead of the specified 61% (**Fig. 3**). The neutral grey of the desk is around 50%, instead of the specified 41%. As stated in *Exposure Specifications* section, an automatic conversion process cannot be used, as it will deliver an over-exposed image (**Fig. 6**). We, therefore, need to add a manual exposure adjustment step to the process (**Fig. 7**), to bring the neutral grey down to spec, before the final conversion (**Fig. 8**). Unfortunately, in this example, over-exposure meant that detail in the background window was clipped at capture and cannot be recovered.

Under-Exposure

Under-exposed images (**Fig. 9**) allocate too many CVs to low-light areas of the image, which are prone to exhibiting dark noise. When correcting exposure on such images (**Fig. 10, left**), the noise is amplified and will likely be present in the final display-referred render (**Fig. 10, right**), whether it is in SDR (BT.709) or HDR (BT.2100).

As we have just seen, over- or under-exposing a camera in nonlinear formats may lead to image quality issues in post-production and prevents us from using automatic conversion processes, leading to an increase in cost and delivery time. It is, therefore, important to understand a camera's latitude and the exposure specifications of the recording format to achieve optimal as well as cost-effective images. Sampling through our previous production files, it became obvious that there was a lack of understanding of these principles among our content-creating community. This has led us to design and implement a training program that combines theory and practical exercises, performed with the cameras our storytellers use on a day-to-day basis.

FIGURE 12. Tighter square (1:1) crop.

FIGURE 13. Wider HD (16:9) crop.

Resolution, Aspect Ratio, and Framing

As is the case throughout the industry, the CBC is reaching out to an audience that consumes content in more and more diversified ways, beyond the traditional television ecosystem. It is not uncommon today that an image will be generated in one format and then delivered in two or three other formats. For example, a team may capture an important event in UHD and then deliver shots for a feature in UHD, an edited piece for the evening news in HD and a social media story in 9:16 1080 vertical video. It is also possible that the team's short-term goal is only one of those formats and that later, another unit within the Corporation wants to reuse the assets with another format in mind. This may entail a reframing of the source, similar in process to the traditional "pan and scan" conversion process.

When ingesting assets into our systems, it is, therefore, important to ensure that the source is usable in multiple resolutions, aspect ratios, and framing types. As we have discussed in the Introduction, our backlog includes shots in multiple formats, and there are currently no standards

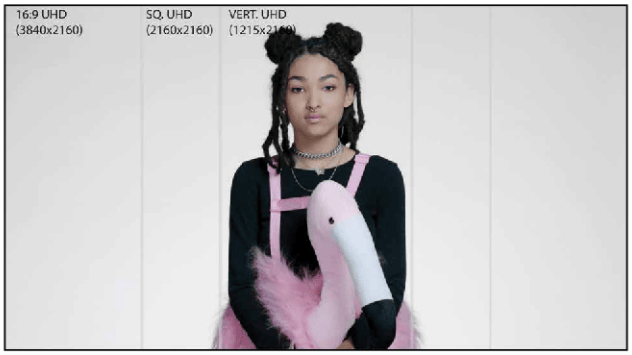

FIGURE 14. Framing guide in UHD for 16:9, square, and vertical deliveries.

in source asset metadata to indicate the desired final framing for very high-resolution captures.

For example, take an image (**Fig. 11**) shot in UHD for a social media and HD television production. From the image, we can clearly see that the creative intent was to frame within the source, as there are wires and fixtures within the frame.

The source does not indicate, however, the final creative intent. It could have been, for example, intended for a 1:1 square video (**Fig. 12**) or a 16:9 HD video (**Fig. 13**). Furthermore, the determination of how tight the final deliverable should be framed is subjective. One solution would be to archive the entire source frame; however, this entails that all future users of this asset will need to go through a mandatory, manual post-production step (if only to remove the unwanted gear in frame), adding cost and time.

Instead of trying to manage aspect ratios, we opted for a multi-aspect ratio strategy that ensures that future shots can be used either as they are, manually pan-and-scanned, or cropped automatically, by introducing aspect ratio guidelines (**Fig. 14**). It is the production team's responsibility to ensure that the entire frame is usable and that the guides are respected.

By ensuring that the framing guidelines are followed, we can archive the source image and then we can automatically crop the source to the desired aspect ratios, saving a significant amount of time in post-production.

In the future, we are looking into adding to our delivery workflow an option to automatically crop the source to the delivery's aspect ratio so that users can order any delivery format by simply checking the desired option at the time they request an asset. In the meantime, at least, the assets will be usable as they are, or cropped on the left and right sides without requiring time-consuming manual pan-and-scan.

Automated Workflows

Now that we have standardized color gamut, dynamic range, exposure levels, and aspect ratios, we can implement automated workflows and benefit from significant time and cost savings, while preserving the creative freedom that our content creators desire. The two major aspects of our automated workflows are the use of a standardized 3D LUT for conversion of SLog3/SGamut3 to BT.709, and the automatic transcoding and uploading of assets to the different groups that may need them, in their formats of predilection.

3D LUT Selection

For the conversion of SLog3/SGamut3 assets to BT.709, we have decided to use a "technical" conversion LUT, which corresponds to the exact conversion of SGamut3 primaries to BT.709 according to the *Sony* specifications, and the Tone Mapping of SLog3 to 2.4 Gamma using the same specifications. However, we added an extra knee at the end of the highlights to CV 960 in 10-bit, to ensure that the signal, when converted, meets the specifications of ITU-R BT.2100-2 and BT.709 for narrow range, which is in use at the CBC, without harsh clipping.

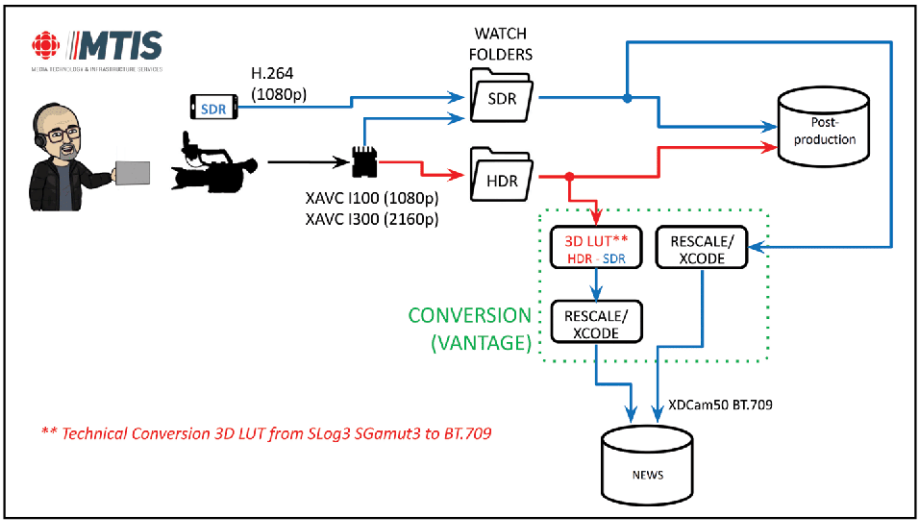

FIGURE 15. Automated HDR and SDR workflow for digital groups.

Strict adherence to the exposure specifications frees us from having to manage a plethora of corrective 3D LUTs and assigning those LUTs to specific assets in the asset management system. We do not restrict the use of "creative" 3D LUTs in post-production, but those are applied to properly exposed sources.

Automated Transcoding for Content Sharing

Figure 15 illustrates the process of automated transcoding and content distribution that we have designed and implemented for our digital media groups for French Services at Radio-Canada, our first test of this workflow.

- When content is captured following exposure specifications, the production team delivers files in one of the two available drop folders: SDR, which collects all sources captured in BT.709, and HDR, which collects all sources captured in SLog3/SGamut3.
- SDR files (blue connectors in **Fig. 15**) are copied to the post-production storage space allocated to the digital group, with an additional copy sent to a *Telestream Vantage* conversion process, programmed to rescale and transcode source assets to XDCam50, BT.709 HD assets. These assets are then delivered to the daily news feed partition on the news server, where content managers for the news organization determine whether the assets are meant to be used that week or sent to retention for eventual archiving.
- HDR files (red connectors in **Fig. 15**) are copied directly to the post-production storage space as well, as our digital groups have the skill set, budget, and time to perform additional post-production to deliver a polished product with a distinctive visual signature. For the news organization, however, a copy is sent to the *Vantage* conversion process, where the standardized 3D LUT (see the section titled "3D LUT Selection") converts the HDR source to a BT.709 narrow-range format and then rescales and transcodes it to a HD, XDCam50 BT.709 asset. As in "SDR files" above, the assets are now available to the news organization, without requiring any post-production.
- When HDR assets are converted to SDR, an email notification (not illustrated in **Fig. 15**) is sent to the asset manager from the digital group and the news organization. This allows them to ensure that the source was shot with the proper exposure specifications and that the LUT provided the desired results and to clear the shot for use.

Archival

Even though both the digital groups and the news organization may use the same source and deliver a final product in BT.709, the look will usually be significantly different, due to the heavy post-production performed by the digital groups. It is, therefore, worthwhile to archive both final products, as they are in fact different assets.

News Content Archival

The content management and preservation team determines, on a weekly basis, which assets meet our archival mandate and will flag assets for retention and then eventual archival. Those are the XDCam50 BT.709 HD assets transcoded from the *Vantage* process (**Fig. 15**).

Post-Produced Content Archival

The digital groups flag their final, post-produced videos, so the content management and preservation teams can send them to retention and archival. As for the source assets that are captured in HDR, they are archived as-is, as long as they meet the exposure specifications and deliver a quality image when applying the standardized 3D LUT. If there were issues with exposure, the assets are flagged for post-production exposure compensation and will be sent to an archival post-production backlog, so the exposure levels can be corrected prior to archival.

Lessons Learned

We have been in this process for almost a year, now, and can already draw some lessons about the way we approached the problem, what went according to expectations and what did not. Hopefully, these lessons can be valuable to other organizations sharing similar challenges.

Address Alternate Formats Quickly

One of the main lessons of this project is that it is better to tackle the challenge of managing multiple file-based formats as soon as possible. For us, the alarm bell rang a little late. We have a three-year accumulated backlog that will need to be manually converted by a skilled colorist before we can archive the assets. Every day that an organization lets alternate formats proliferate means time and cost are going to accumulate down the road.

Limit Capture Standards

It is not a trivial thing to do, but we had to limit the number of creative options available to our user base. Selecting a single wide dynamic range and color gamut format has allowed us to automate the conversion process, with a single 3D LUT by format (BT.709 for current SDR broadcasts and BT.2100/HLG for future HDR broadcasts). Since every camera manufacturer has proprietary formats, our selection narrows the choice of equipment, but it is the price we decided to pay to maintain cost-effective and manageable processes.

Provide Training

It is essential that images be exposed accurately to match the format specifications. Achieving this goal may not be realistic without setting up a full hands-on training program and addressing all of the crew concerns for their given work environments. We provided one full day of training to each group working with the new formats; the feedback was excellent, and we were able to address the vast majority of issues and concerns in a real work environment.

Show Patience and Persistence

The transition from a certain *laisser-aller* to a more stringent set of specifications does not end after training is provided. We have had follow-up activities for a good six months afterwards, which has enabled us to address the issues of exceptional shooting conditions, lack of comprehension and sometimes, lack of will to change previous methods to the new process. In these instances, setting up follow-up meetings and mini-workshops to help crews transition properly has yielded more success than confrontation. Expect that the transition is not going to be instantaneous. Patience and persistence are key.

Foster Collaboration Between Stakeholders

Another important factor in our successful transition is the involvement of all key stakeholders, from the camera department, supervising technicians, journalists, videographers, archivists, editors, and so on. We ensured we got buy-in and critical input. Without collaboration, we could have designed a nice architecture on paper, but it is doubtful that it would have been implemented successfully.

Conclusion

As the frontiers between cinema, television, streaming, and social media continue to shrink, traditional broadcast organizations are faced with increasingly complex workflows and an ever-growing list of options for capture, post-production, and delivery. Proactively recognizing that an increase in options leads to an explosion in cost and time through the sheer increase in complexity leads to a streamlining of those processes and the possible options.

To provide a reasonable balance between the creative desires and the corporation's budgets, we have limited capture options to one color volume for traditional production (BT.709) and one for productions requiring more creative latitude and having post-production budgets (SLog3/SGamut3).

Strict observance of exposure specifications allowed us to automate cross-format conversion between SLog3/SGamut3 capture and either BT.709 or BT.2100 deliveries, using "technical" 3D LUT.

The implementation of this process required collaboration between all key players, a significant amount of training as well as several months of follow-up and patience, until all units consistently delivered assets that meet the specifications. It has been an intensive effort, but the cost savings by eliminating manual cross-conversion and color grading were worth the effort.

As a final observation, this project's experience leads us to believe that to successfully implement HDR production, it is better to gradually bring teams to proficiency by starting with file-based workflows, where we still have a chance to correct issues in post-production during the learning curve, than to go live from scratch and learn on-the-fly with a product that does not necessarily offer the best quality to our viewers from day one.

Acknowledgments

This project would never have been a success without the active participation of so many people at CBC/Radio-Canada. Special recognition goes to Maxime Labrie, supervising technician, digital groups, who organized the training sessions, Patrick Handfield, project coordinator, who designed the *Vantage* workflow, Claude Dubé, supervising technician, who delivered the training workshops, and Jonathan Dupras, Manager of the Architecture team, who gave us the vote of confidence to implement this new approach.

Glossary and References

3D LUT

A matrix of points determining the conversion values between two color volumes for R, G, and B values. Values between the matrix's points are interpolated, so the more points a matrix has, the more accurate the conversion will be.

UHD

Ultrahigh-definition: The general term used to describe video resolutions beyond HD. Commonly called *4K (UHD-1)* and *8K (UHD-2)*, the television resolutions are actually 3840 × 2160 (4K) and 7680 × 4320 (8K). May or may not include wider color gamut and HDR transfer functions.

HDR

High-dynamic range: In television, refers to a contrast ratio (or exposure "stops") that is superior to SDR, or standard dynamic range of HDTV. Different organizations have different criteria to determine what constitutes HDR in content as well as display capabilities. See the UHD Alliance, for example: https://alliance.experienceuhd.com/uhd-premium-features.

EOTF

Electro-optical transfer function: The formula(e) used to convert CVs into light at the display level.

HLG

Hybrid log gamma: One of the transfer functions that encodes optical information into CVs. Developed by the BBC and NHK as a relative system (e.g., it converts CVs to light relative to a screen's maximum luminance).

SMPTE ST 2084

One of the transfer functions that encodes optical information into CVs—also called *perceptual quantizer* (PQ), as the CVs are allocated based on human contrast perception. It is an absolute transfer function as specific CVs are equivalent to absolute light levels on the display, independent of the display's maximum luminance.

BT.709

The HDTV color volume uses the 709 color volume, with primaries and a gamma transfer function based on backwards compatibility with CRT televisions. Modern displays are capable of exhibiting.

BT.2020

Color space for next-generation television: Based on maximum color representation achievable with RGB lasers. Mostly an aspirational standard for home and mobile displays, as no device currently on the market can actually represent those colors on their screen.

DCI-P3

The color space used in digital cinema, as per the Digital Cinema Initiatives, L.L.C. specification. Offers a wider range of colors than BT.709 but less than BT.2020.

BT.2100

Color volume for next-generation television, which combines the BT.2020 color space and an HDR-compatible transfer function like ST 2084 or HLG. Can be used either in HD or in UHD.

SLog3/SGamut3

A Sony proprietary color volume for quantizing dynamic range (SLog3) and color gamut (SGamut3). A technical summary of the format can be found here: https://pro.sony/s3/cms-static-content/uploadfile/06/1237494271406.pdf

About the Author

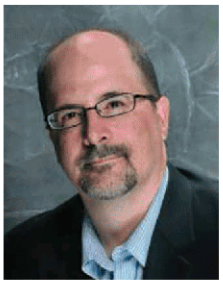

Pierre Hugues Routhier is a former aerospace engineer and a specialist in advanced imaging technologies that provides training to studios, manufacturers, and broadcasters worldwide. He has developed innovative solutions and workflows for studios and post-production facilities in all fields of advanced imaging: ultrahigh-resolution, high frame rate, VR/AR, HDR, and others. He is credited with more than a dozen patents related to advanced imaging and has been published in several industry papers. His latest book, *High Dynamic Range for Television and Motion Pictures: A Digital Troublemaker Guide*, has quickly become an industry reference on HDR for broadcast and motion pictures.

SMPTE Virtual Courses

Sharpen your skills in the latest digital media technologies

Last year, nearly 10,000 media professionals, technologists and engineers chose our courses to help them deepen their technical knowledge, with 97% already planning their next SMPTE class.

That's why we're constantly expanding our course offerings with classes on technologies including HDR, UHD and DCP, and special focus on transformative standards like ATSC 3.0 and ST 2110. Choose Instructor-led courses for personal attention and feedback, or start learning immediately with our flexible self-study option.

SMPTE

View the latest offerings and register today at smpte.org/virtual-courses