Dolby® 3D for Glasses-Free TV and Devices

Overview

Three-dimensional television (3D TV) offers tremendous market potential, but up to now, the annoyance of the required glasses has limited consumer viewing and slowed adoption. The real future of consumer 3D as an everyday experience that can truly tap into the market lies with glasses-free viewing. While the industry as a whole acknowledges this, glasses-free technology to date has had its own problems, generally suffering from restricted viewing positions and distracting visual artifacts.

Dolby® 3D changes that. It is a collection of processing and coding techniques brought to market through a collaboration between Dolby and Philips. The companies recently presented the Dolby 3D suite of innovative technologies to the public, garnering industry awards and accolades.

3D content is consumed from a variety of sources such as packaged media, broadcasts, and over the top (OTT) streaming. It is viewed on a variety of devices, increasingly including display devices equipped with advanced optical lenses that enable glasses-free viewing. When the Dolby 3D suite of technologies is implemented in devices along the path that 3D content takes to glasses-free 3D displays, Dolby ensures that the viewer's experience will surpass any previous 3D TV experience. The underlying architectures of these devices must accommodate the complexities of Dolby 3D processing specifications, which are key to ensuring that 3D perception consistently meets the expectations of sophisticated consumers.

But what is this underlying technology? How is it feasible in architectures prevalent today? When will all the pieces of the ecosystem come together? This white paper answers these questions in depth.

.

Table of Contents

Introduction

Glasses-based 3D adoption has been hampered by a number of factors, including cost of acquisition, new production grammar and techniques compared with 2D, compatible broadcast formats, increased bandwidth, and the performance of the first generations of 3D displays. Yet even taking these factors into account, the number-one cause for 3D's slow uptake outside the movie theatre is the lack of consumer displays that deliver a comfortable viewing experience. Investment, training, negotiations, and infrastructure can overcome the earlier shortcomings; but the lack of a practical, comfortable, and economical 3D consumer display platform for mass adoption has been at the crux of the issue. Specifically, the annoyance of having to wear glasses fundamentally defeats the likelihood of mass acceptance of 3D in the home.

The majority of consumers prefer to experience 3D in much the same way as they do 2D today:

- Without eyewear
- Without limitations in viewing angle or viewing distance
- Without visual discomfort
- Without a compromise in resolution
- Without any implications for viewing of 2D content

Fortunately for the 3D future, a number of companies have developed optics (such as optical sheets) that can deliver 3D from disparate images. Although a majority of these optics are meant for signage applications, a few are of high-enough quality to be used in a new generation of consumer glasses-free 3D TVs.

This white paper focuses on the specifics of how Dolby 3D processes content to a glasses-free panel. First, though, it offers a primer on the optics that make the experience possible.

Lenticular and Barrier Optics

Glasses-based 3D displays emit disparate images to each eye by interleaving images on the panel surface either temporally (with active glasses) or spatially (with passive polarized glasses). Glasses-free 3D displays, also known as autostereoscopic 3D (AS3D) displays, must always spatially interleave views onto the panel's red-green-blue (RGB) subpixel structure, using an optical splitter to emit certain subpixels in one direction and others in another direction.

Two well-known splitting methods exist:

- Barrier optics: a pinhole-parallax barrier prevents light from each subpixel from reaching one of the eyes (from a given viewpoint).
- Lenticular optics: a lens over a set of subpixels projects light in a certain direction.

Either technique can produce multiple views, rather than just two, allowing the viewer to move around freely without seeing annoying artifacts. The lenticular approach currently performs better for large-screen displays because it does not drastically reduce the light output.

Inevitably, spatially interleaving multiple views yields a resolution that is reduced from the native resolution of the display. However, this reduction is often mistakenly assumed to be 1/N for an N-view display. It is in fact a complicated factor of N that yields a perceptual resolution that is much closer to the input-content resolution. Also note that a two-view lenticular screen offers a higher perceived resolution and, when combined with the use of a head-tracking camera within the panel, can work well for a single viewer. However, head tracking has yet to be perfected and can cause occasional visual discomfort, both for the viewer whose head movements are being tracked and for viewers who may be watching nearby.

With lenticular optics, each subpixel (R, G, or B) is projected in a certain direction. The direction depends on the position of the subpixel with respect to the orientation of the lenses that make up the lenticular sheet. Pixel components that shine in the same direction together form a single view. All views together form the viewing cone. Beyond the cone, views are repeated. Generally, the primary or center cone has the best picture quality and 3D effect, but to achieve viewer comfort, the emphasis is on maintaining a reasonable picture quality and 3D effect throughout the viewing area. Important characteristics of lenticular optics include slant angle, pitch, cone size, and number of views per cone (see Figure 1). The lenticular display platform with peak performance characteristics as of now uses as many as 28 fractional views. Any optical stack (that is, any screen and lens) intended for glasses-free 3D might also be paired with a solution to restore the panel's native 2D resolution via supplemental switchable optical layers. However, Dolby 3D provides processing elements that allow a 4K display to render full HD 2D content with excellent quality.

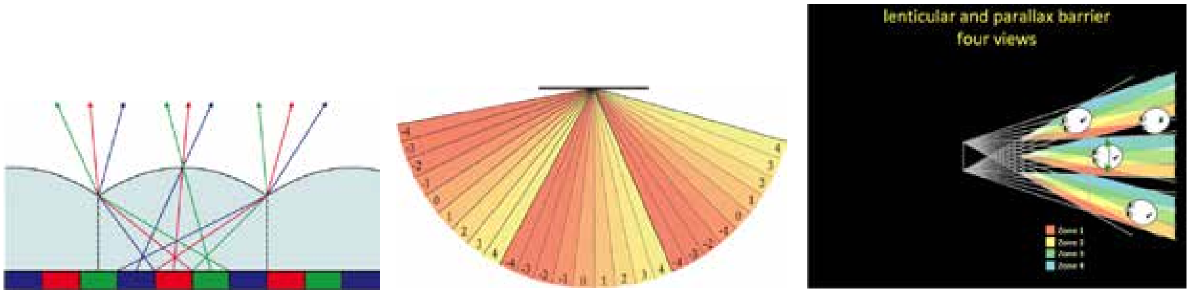

Figure 1 Characteristics of Lenticular Optical Stacks

As stated earlier, the optics chosen for the panel are just as crucial as the content processing and conversion achieved via Dolby 3D. Well-designed optics can meet the goal of increasing the viewers' comfort, making the overall display system very suitable for the consumer market. To meet the preferred viewing criteria, Dolby 3D processing is optimized to match the characteristics of the 3D optics.

Trends That Favor Comfortable 3D

In addition to the advent of optics and processing that enable comfortable glasses-free 3D, other important forces of change help make 3D a desirable feature for a wider range of consumers.

For example, panel manufacturers have stretched manufacturing capabilities to extend spatial resolution to 4K × 2K and temporal resolution beyond 120 Hz. Even without any current infrastructure for broadcast content in 4K, the use case for glasses-free 3D on a 4K panel makes a compelling combination for both consumers and display manufacturers. Multiview 3D through lenticular optics naturally suffers a reduction in spatial resolution (as many views spatially share the RGB pixels of the panel), so increasing the spatial resolution of the panel has a very positive effect on the perceived resolution of glasses-free 3D. In fact, demonstrations of Dolby 3D on a 4K prototype panel have garnered accolades for a 3D experience equivalent to 1080 pixel (1080p) HD resolution per eye.

The high temporal resolution of panels also allows for 3D to be rendered more comfortably. Recent films (for example, *The Hobbit: An Unexpected Journey*) are being captured at 48 frames per second (fps) rather than 24 fps because the creative community acknowledges the limitations of 24 fps for 3D. At the consumer level, the integration of good frame-rate conversion and good multiview autoconversion is critical to creating a pleasing fatigue-free experience.

The initial desire of filmmakers to demonstrate the technical capabilities of 3D by use of heavy negative parallax (also known as "pop out") has been displaced by the desire to use 3D as a tool to improve storytelling, a change evident in more recent 3D productions such as *Hugo*. This more subtle use of 3D not only improves viewer comfort when using glasses-based 3D displays, but also closely matches the capabilities of autostereoscopic displays in the foreseeable future.

Last but not least, the production cost of shooting live 3D, especially sports, is being addressed by the industry. While there are production differences between 2D and 3D shoots, opportunities to share resources exist and are being exploited.

Dolby 3D Impact

So what exactly is Dolby 3D? What part does Dolby play in the ecosystem? Which components does Dolby supply, and which components does Dolby enable others to build?

Today consumers can enjoy 3D content from existing packaged media, broadcast, and OTT sources. This content is not created with glasses-free displays in mind, so each display must convert the left-eye and right-eye input into a multiview output suitable for that display. For this purpose, Dolby 3D specifies multiview autoconversion components that reside in the system on a chip (SoC) integrated circuits within the display.

New 3D service providers can take advantage of the versatile Dolby 3D format, which heightens the experience on newer devices without hindering existing devices and services. Real-time autoconversion requires implementation trade-offs that limit its capability in scenes that are traditionally difficult (for example, fast motion). Dolby 3D autoconversion provides the ideal trade-offs to ensure that the autoconversion artifacts are reduced to a point where viewer comfort is maintained, even in difficult scenes. However, the overall experience can be greatly improved if depth maps are created and transmitted to the display. Such information is almost always generated during the content-creation process and can greatly aid in ensuring that the creative intent is rendered as desired on any display device. Newer launches of 3D service by broadcast, OTT, or emerging packaged-media delivery can benefit from a format that ensures an even more consistent, compelling, and comfortable 3D experience on any device (see Figure 2).

Figure 2 Impact of Dolby 3D from Capture to Consumer

It is important to recognize that Dolby 3D is delivery-mechanism agnostic. Dolby 3D is compatible with any 3D optics and any 3D display panel. However, in order to optimize the viewing experience, Dolby has some “best practice” guidelines.

The next section of this white paper covers Dolby 3D components that enable real-time autoconversion on display devices. It then introduces the Dolby 3D format and considers the practical advantages at various points from creation to exhibition. Finally, the paper will put these together to discuss how both would coexist on the display device.

Multiview Autoconversion

The task of real-time autoconversion for multiview glasses-free 3D displays is simply described as ingesting a stereo pair of images and producing N images that are spatially packed into an RGB frame suited to the panel and optics being targeted. Table 1 shows the three options to generate the necessary views in real time.

Table 1 Options for Generating Multiple Views from a Stereo Pair

Option	Pros and Cons	Pattern
Repeat available stereo views.	• No new view generation required. • High crosstalk. • Strong inversion. • Not comfortable.	L R L R L R L R L L L L L L R R R R
Estimate depth for L view and render new views around it.	• A single depth map yields consistent behavior. • Comfortable and stable behavior. • Amenable to real-time (RT) implementation.	$L_{-4} L_{-3} L_{-2} L_{-1} L\ L_{1} L_{2} L_{3} L_{4}$
Estimate depth for L and R views; render views in between and around both.	• Increased motion parallax. • Inconsistency between L and R can be very disturbing. • Presently impractical for RT implementation.	$L_{-1} L\ L_{1} L_{2}\ L\ R_{-2} R_{-1} R R_{1}$

The first option, repeating available stereo views, is clearly inadequate. The third option, estimating depth for left and right views and then rendering views in between and around both, has some advantages but requires at least twice the complexity of the second option. The second option, estimating depth for left view and rendering new views around it, has benefited from a decade of research into efficiently stabilizing and producing consistent multiview images on a variety of optics. Only one depth map is created, and the renderer creates views around one eye, typically the left eye.

Stages of Autoconversion Processing

The simplest division of any autoconversion is into two stages:

1. Extract the depth map.
2. Generate N-views using the depth map.

Depth extraction can be subdivided into two steps:

1. Obtain the disparity between the left and right views.
2. Translate that disparity into depth, and enhance it for use in rendering.

The rendering can be further subdivided into two steps:

1. Render N-views.
2. Interleave pixels from specific views and compose a single frame of packed RGB pixels that contain elements from both views, spatially positioned to match the characteristics of the optical sheet.

The majority of the complex intellectual property at the heart of a comfortable and enjoyable 3D viewing experience lies in the depth-enhancement stage and the subsequent rendering stage. Full details of these critical stages are provided to licensees of Dolby 3D.

As Figure 3 shows, the main precepts of Dolby 3D autoconversion are to:

- Create temporally coherent depth maps
- Incorporate methods that yield significantly better sharpness
- Dynamically adapt to the parallax intended for each scene of the content
- Successfully mitigate inversion artifacts in the cone boundary zone

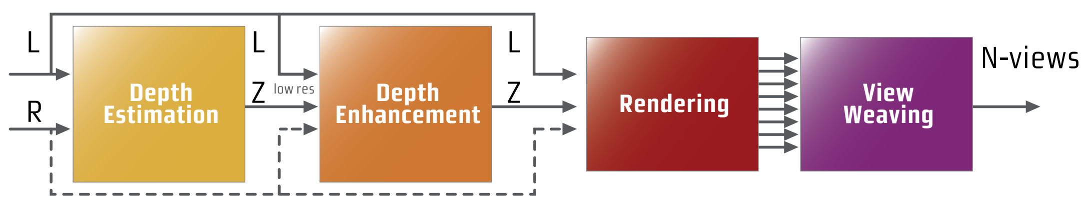

Figure 3 High-Level Data Flow for Stereo to Multiview Autoconversion via Dolby 3D

Dolby 3D autoconversion also exposes several controls that allow an OEM or the end user to modify viewing preferences, including defining the preferred viewing distance, the desired gain, and other parameters.

Depth Estimation

To produce depth maps from a stereo pair, we begin with coarse-level disparity estimation. Assume a stereo pair of images or videos captured at a certain distance from each other. This distance is generally called the camera baseline. Objects located at different distances from this baseline are visible in the stereo images at horizontally shifted positions, typically called disparities. A disparity estimator provides estimations of these disparities, and therefore also of the relative depths of the objects (see Figure 4).

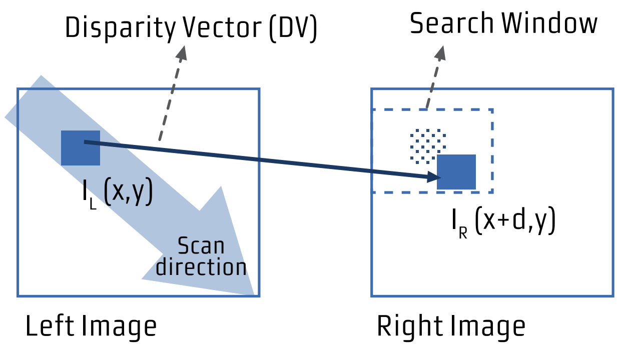

Figure 4 Disparity Estimation

This task is very similar to motion estimation between subsequent frames in a video sequence, apart from the fact that only horizontal motions are expected for our stereo capture. We can therefore use an off-the-shelf motion estimator based on block-matching algorithms. A block-based motion estimator tries to find a corresponding location for each block of pixels in the other image. Instead of doing a full search for the best matching window over the entire image, it evaluates a small number of candidate vectors and keeps the best match.

Depth Enhancement

Quality depth maps translate into quality renderings. Enhancing coarse-level disparity estimates into consistent, temporally stable, and accurate full-resolution depth maps plays no small part in creating a comfortable viewing experience. This is the goal of the depth-enhancement block.

Here are the key criteria:

- Ensure depth consistency, especially of high-contrast objects.
- Ensure that depth information remains stable over time.
- Ensure that coarse-level disparity maps are accurately scaled to full resolution.
- Dynamically map the artistic intent into the optical stack (that is, the screen and lens) in the best achievable manner.
- Ameliorate issues associated with noisy depth maps.

The emphasis of this block is on spatial and temporal stability of the depth map. This is where many autoconversion algorithms break down and where the Dolby 3D approach shows significant improvements. Implementing these improvements would normally consume exponentially larger resources and throughput, but the Dolby team has taken steps to moderate the increases.

Multiview Rendering

The other half of the complexity of autoconversion lies in the creation of the views that lie in between (interpolated) and on either side (extrapolated) of the input stereo pair. Creating information that does not exist requires a deep understanding of the issues, but clean depth maps, clever view-synthesis algorithms, and forgiving optics go a long way toward creating an excellent glasses-free 3D TV experience.

In synthesizing views from a different viewpoint (camera baseline), one inevitably reveals objects that were occluded from the adjacent view. The cleverness in view synthesis comes from approaches to recreate or discover this information as accurately as possible, albeit without drastically taxing the resources of real-time autoconversion. The Dolby 3D approach to handling occlusion in the view-synthesis stage suitably achieves this balance.

In addition to those key steps, a number of small adjustments to the rendering are made just prior to matching the output to the display and optical sheet. Various filters and feedback loops are employed to optimize the final sharpness and depth impression, and to largely remove the cone-boundary inversion artifacts.

As a last step prior to driving the display, the views must be spatially weaved into a single RGB frame that is sized appropriately for the panel. In view weaving, Dolby 3D uses the characteristics of both the target panel and the target optical sheet to ideally tune the resulting picture.

The Dolby 3D Format

To meet the needs of the widest subscriber base for 3D, broadcasters initially adopted what is known as a frame-compatible (FC) approach to 3D transmissions. By packing the views of the left and right eye into a single full HD frame, even legacy set-top boxes (STBs) could pass a 3D signal through to any 3D display. The manufacturers of glasses-based displays that those broadcasters were targeting had more pressing issues than this loss of resolution. However, over time, manufacturers refined these displays and corrected issues with glasses, polarizers, retarders, graphical user interfaces (GUIs), HDMI® connections, de-judder, packing signaling, prefilters, and so on. In spite of this, today's advanced glasses-based 3D displays still render a broadcaster's 3D signal in a way that lacks the crispness of its full HD counterparts such as Blu-ray™. The same signal is also ill suited to target a newer generation of 3D devices coming into the mainstream. Even more importantly, the creative artist cannot yet convey the full artistic intent from one universal format onto the plethora of 3D devices.

Rationale Behind the Dolby 3D Format

There are three main methods of delivering 3D content to consumers:

- Broadcast
- OTT streaming
- Packaged media

These delivery methods have varying requirements, but they share these core goals:

- Improving the resolution of 3D content
- Ensuring that 3D content is rendered consistently, compellingly, and comfortably on a wide range of devices
- Ensuring that existing devices, services, and infrastructure are not affected
- Ensuring that the format can be adopted by the widest array of devices
- Providing information to the display device to reconstruct the creative intent in the best achievable manner
- Allowing bit-rate targets to be met

Of the three delivery methods, OTT ecosystems have a unique set of additional requirements. OTT aggregators launched 3D services that leveraged broadcast and packaged-media formats for the broadest user base. However, these formats don’t suitably accommodate the full requirements preferred by the aggregators and content delivery networks in their ecosystems. Neither do they ensure, as with broadcast, that the exhibition of the artist's work will be achieved in the best possible manner on any device. To an aggregator, the most important traits of a 3D format include the ability to do the following:

- Ensure the highest-quality rendering on any device, from mobile devices to cinematic displays
- Maintain good 3D at very low bit rates
- Choose a format that permits ease of implementation on specialized architectures, such as gaming devices
- Meet or exceed (infrastructure permitting) the resolution guarantees of competitive sources of content, including broadcast and packaged media
- Allow adaptation to fluctuations in the nondeterministic delivery infrastructure
- Leverage the aggregator’s knowledge and relationship with each customer by accommodating the customer’s viewing preferences into a customized stream for delivery
- Prevent an unwarranted explosion in inventory

Dolby 3D is attuned to these additional constraints. It is expected that delivery via OTT to newer 3D devices can be rolled out sooner than delivery via broadcast or packaged media. In fact, a conceptual prototype of the Dolby 3D format has already been demonstrated in the Vudu™ streaming ecosystem in North America.

The following section discusses three main attributes of the Dolby 3D format.

Frame-Compatible Resolution Enhancement

Early glasses-free displays did not require high-resolution input images because, at the time, issues that had little to do with resolution dominated the artifacts of the rendered 3D image. Now that the issues of optics, substrates between optics, autoconversion algorithms, panel resolution, and sheet alignment have been addressed, the resolution of the source signal from which other views are synthesized makes a very noticeable difference. It is clear that glasses-free displays with native 4K × 2K resolution would benefit from source content that is at least full HD 1080p per eye.

The Dolby 3D team has devised a remedy that requires neither a significant reworking of broadcast infrastructure nor a complex re-spin of SoCs found in 3D-capable STBs. It would not alienate the customers of operators around the world that have already begun 3D broadcasts that leverage the frame-compatible (FC) approach. FC packing sacrifices the spatial resolution by low-pass filtering and subsampling, usually by half for each view, while maintaining compatibility with existing codecs and currently deployed infrastructures. Both FC side-by-side (SbS) packing and FC top-and-bottom (TaB) packing yield a significant loss in sharpness as compared to the original.

Dolby 3D addresses this issue with Orthogonally Multiplexed Frame-Compatible Full Resolution (referred to as OM-FCFR), which is a scalable extension of the FC Advanced Video Coding (AVC) base layer to full resolution with the addition of an enhancement layer. The base layer carries the FC signal; thus, legacy decoders can still decode a pair of half-resolution views and perform the proper upsampling on the base-layer pictures. The enhancement layer carries the orthogonal high frequencies lost in the base layer. The base layer and enhancement layer use orthogonal FC formats. As an example, when the base layer is SbS and the enhancement layer is TaB, OM-FCFR retains full resolution by preserving the vertical high frequencies intact in the base layer and the horizontal high frequencies intact in the enhancement layer.

Figure 5 shows a frame from a *Newsroom* clip with its original spectral components, the spectra of what is transmitted today, and the spectra that can be recovered via a backward-compatible enhancement.

(a) Full resolution (b) FC-SbS (c) OM-FCFR

Figure 5 Spectrum Visualization of *Newsroom* in OM-FCFR

OM-FCFR is built on the MVC Stereo High Profile of the AVC specification. Since OM-FCFR's base layer and enhancement layer have different FC formats, the current MVC architecture cannot explore the correlation between these two layers directly. Dolby 3D extends MVC by introducing an additional processing element inside the MVC architecture to better explore the characteristics between the two layers. The additional element, the Reference Processing Unit (RPU), first preprocesses a base-layer reference picture before it is used for the prediction of its corresponding enhancement-layer signal. The preprocessing step includes the downsampling and upsampling filters for the format conversion from the base-layer FC format to the enhancement-layer FC format as well as for better prediction accuracy for the enhancement layer. The simplified diagram of OM-FCFR is shown in Figure 6.

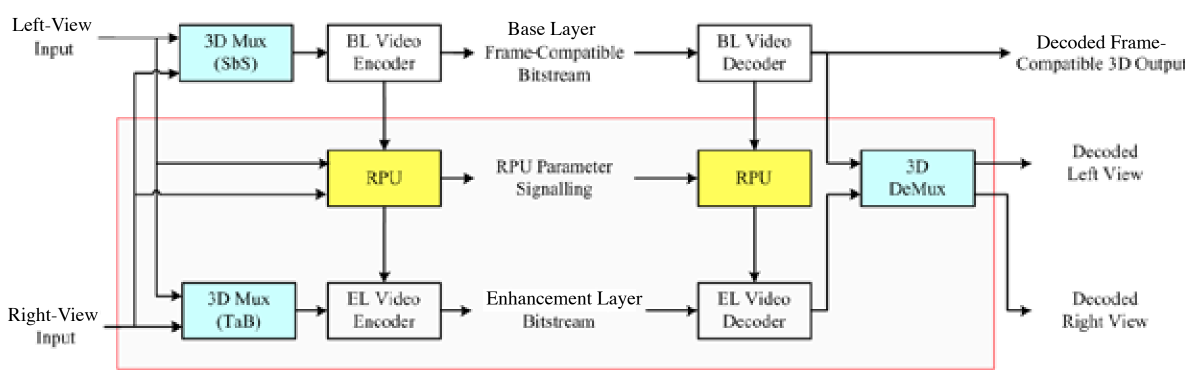

Note: BL = base layer; EL = enhancement layer; RPU = Reference Processing Unit; SbS = side by side; TaB = top and bottom.

Figure 6 Simplified Diagram of OM-FCFR

The final step of OM-FCFR decoding entails reconstructing the full-resolution frame from its FC parts.

Note that frame compatibility can also be defined as being service compatible with 2D transmission. This formulation can be useful for certain use cases, including OTT services. The Dolby 3D format accommodates approaches that employ service compatibility to reduce inventory for OTT aggregators.

Carriage of Depth Maps

The exhibition of 3D on display devices can be drastically improved with consistent and correct depth maps. These depth maps should be created with care and transmitted to the device in order to do the following:

- Achieve the best rendering on glasses-free devices and avoid the inevitable artifacts that ensue from display-side depth-extraction algorithms, which have to make compromises for the sake of real-time implementation
- Take advantage of clean depth maps that are already part of the movie content-creation process
- Allow unconstrained real-time autoconversion algorithms to manifest in live content-creation workflows
- Allow creative artists to convey the intent of their 3D scene consistently to the display device, so that it can accurately adjust to its parameters
- Remove the need to remaster content to each display device on which the content can be viewed
- Reduce the complexity of the display architecture
- Achieve the best possible accommodation of depth preferences on any device while preserving the relationships between objects in the scene

The Dolby 3D format accommodates the carriage of depth for backward-compatible broadcast as well as OTT use cases. Both texture and depth are available in the base layer. The resolution of both texture and depth are improved by way of an enhancement layer.

While the accuracy of depth maps is left to the creative process, the precision required for a given display can be variable—at least for the time being—as glasses-free displays become even more discerning. To take into account the dual and competing requirements of backward compatibility and high quality, Dolby 3D incorporates a trade-off for lower-resolution depth maps in the base layer while also providing the means of upgrading the depth information by way of an enhancement layer.

The Dolby 3D team is creating production tools to enable the production and carriage of depth and metadata information. Fortunately, the creation of 3D content forces artists to create depth maps as a by-product of the creative process. Depth metadata is also often part of the cinematographer’s note sheets. Quite a bit of 3D content originates from 2D capture and is postprocessed into 3D via the synthesis of expertly created depth maps. Thus, expertly created depth information for encapsulation into the Dolby 3D format appears to be readily available. Dolby has recently partnered with leading tool providers to create a complete solution.

Details of the configuration of texture and depth channels for each use case are presented in the Dolby 3D Format Specification. Contact your Dolby representative for more information.

Carriage of Metadata

Aside from the images and depth maps, the format carries additional information to:

- Accurately carry the artist's creative intent, irrespective of the 3D glasses-free display design
- Handle picture elements such as graphics, subtitles, logos, downstream onscreen-display overlays, and so on
- Provide additional aids to improve the rendering of 3D content

Metadata is carried within the stream, aligned with the image and depth information.

The Integrated Dolby 3D Device

As Figure 7 shows, the overall system diagram of a Dolby 3D display device contains both the format-decoder component and the autoconversion and format-renderer component. The renderer part of the autoconversion block is capable of ingesting a variety of metadata encapsulated into the data stream as specified in the format. Most display devices will need to accommodate both autoconversion and format decoding. Some mobile devices that cannot ingest current stereoscopic broadcasts or packaged media could implement the format decoder and renderer alone. Figure 7 shows the decoder component separate from the autoconversion component.

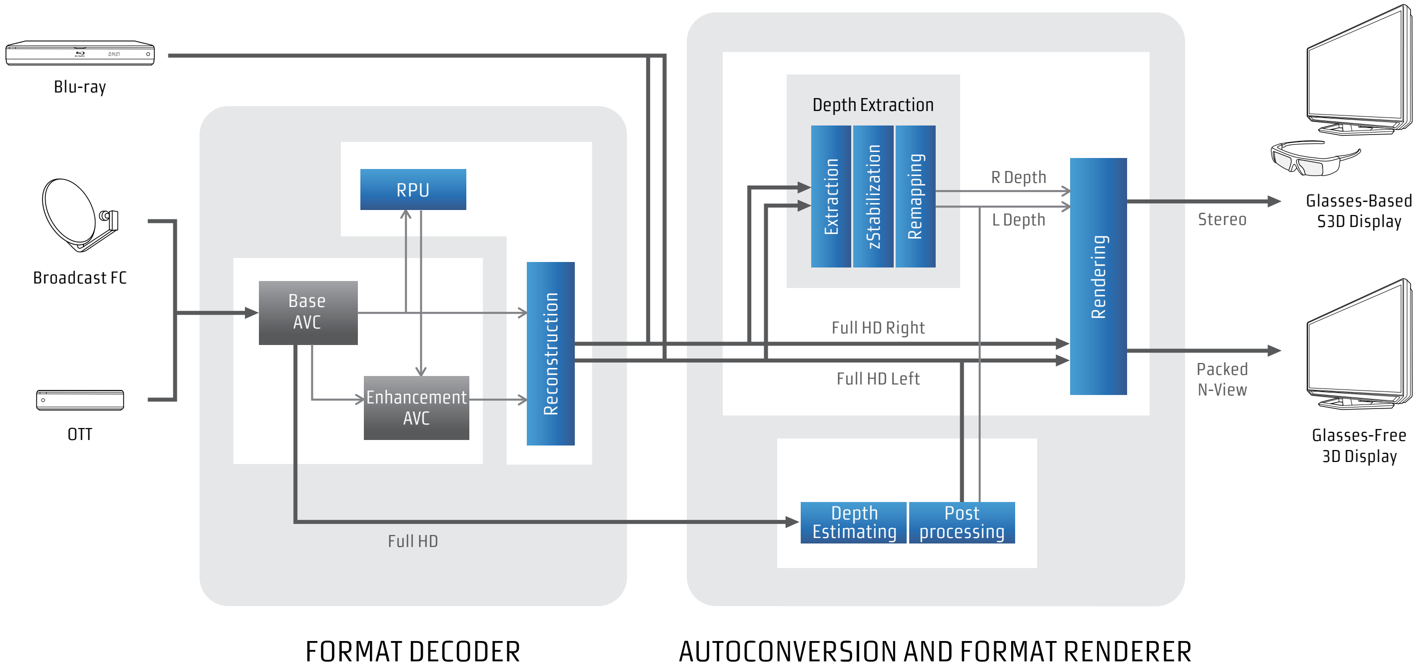

Figure 7 Overview of the Dolby 3D Decoder System

The autoconversion block can integrate additional features commonly found in consumer 3D displays, including 2D to 3D conversion and stereoscopic depth adjustment. These blocks share processing resources with the multiview autoconversion block and belong in 3D sets that are glasses-free or glasses-based, as shown.

The end-to-end flow of content from capture to exhibition on a consumer panel is shown in Figure 8. The Dolby 3D format is enabled at various stages in the ecosystem. Dolby 3D delivers the know-how to enable the creation of these components, including the end display device.

Figure 8 End-to-End Flow of Content

Summary

The combination of Dolby 3D format and Dolby 3D autoconversion enables an experience that is compelling, comfortable, and deployable in the emerging market for glasses-free 3D displays. The Dolby 3D team works with all of the partners in the chain, from content to display, to ensure that the resulting experience is consistent across all platforms. Dolby 3D licensing kits ensure that the amalgam of format and conversion technologies is delivered in a form that spurs quick integration on multiple platforms. The ensuing experience from a Dolby 3D capable device creates the optimal 3D viewing experience.

Abbreviations

Abbreviation	Full Term
4K	4K × 2K display technology
AS3D	autostereoscopic 3D
AVC	Advanced Video Coding (also known as H.264)
BL	base layer
EL	enhancement layer
FCFR	frame-compatible full resolution
FPS	frames per second
MVC	Multiview Video Coding
NRT	non-real time
OM	orthogonal multiplexing
OTT	over the top
RPU	Reference Processing Unit
RT	real time
S3D	stereoscopic 3D (typically referring to 3D with the aid of glasses)
SbS	side by side
SoC	system on a chip
mux	multiplexer
SVC	Scalable Video Coding
TaB	top and bottom

Dolby Laboratories, Inc.
100 Potrero Avenue, San Francisco, CA 94103-4813 USA T 415-558-0200 F 415-645-4000
dolby.com