Camille—Virtual Production Case Studies Considerations for Working With LED Walls for Independent Filmmakers

By Kathryn Brillhart

Abstract

Camille is a short research and development (R&D) film created with the intent to share learnings with other filmmakers. By design, the film references Terrence Malick's "Days of Heaven" and paintings by Andrew Wyeth. The director was interested in exploring ways to create naturalistic images and verité camera work using light-emitting diode (LED) volumes to prove that it is possible to capture supporting visual effects in-camera that add value to a project (even on an independent budget). Key reasons for using LED wall technology were: 1) no time or budget to travel crew to a real location; 2) no budget for permits, fire crew, and safety to execute a fire scene; and 3) to create extreme weather systems that could be art directed and scheduled around actors' availability. This article outlines key case studies that demonstrate how the creative team incorporated realtime workflows, a new department and display technology into the filmmaking process, revealing both successes and pivots in the process.

Keywords

2D plates, camera tracking, final pixel, ICVFX, LED wall, lighting in studios' low ceilings, photorealistic, realtime, VAD, virtual production, visualization, volume, volumetric clouds

Introduction

Virtual production is a broad term that encompasses several ways of incorporating realtime workflows into the production process. The content and director's creative process drive WHAT virtual production IS, which makes its implementation unique for every project.

At a minimum, the term virtual production (VP) encompasses realtime visualization, virtual scouting, Virtual Art Department (VAD), performance capture, virtual camera, simulcam, object tracking, and light-emitting diode (LED) wall/volume workflows. Many of these techniques can be used independently of LED walls or in different combinations with each other. When planning LED workflows, however, several of these techniques are essential to the success of the project.

LED workflows have gained popularity in recent years because they improve upon existing visual effects workflows such as using greenscreen/bluescreen or the use of trans lights, digital matte paintings, and projection techniques.

> **Because the ideal location for the story would require travel and a larger budget than existed, the filmmaker chose to use LED wall workflows so that they could bring the expansive field location to Los Angeles, control the weather and time of day, and schedule the shoot over two days.**

Sam Nicholson, ASC, who founded Stargate Studios, describes their early applications of LED wall workflows in 2001. "Our last experience like this was right after 9/11 when nobody would fly. We converted shows like "ER" and "Grey's Anatomy" over to virtual production. "Grey's Anatomy" didn't have to go to Seattle anymore, "CSI" didn't have to go to New York, and "ER" didn't have to go to Chicago. We created a virtual backlot, which meant capturing numerous locations and creating a footage library. Producers can peruse the library and reverse-engineer a scene in London or Paris or Moscow. It changes the business model away from green screen because no producer wants to deal with a thousand shots coming out of post. Editors don't like it, directors don't like it, and actors don't like acting on green screens. They'd much rather sit on the steps of the Capitol Building than stare at a laser dot."[1] As LED workflows began to be adopted by narrative feature directors, Alfonso Curón is sited by virtual effects (VFX) supervisor, Tim Webber, in 2013 as adopting an LED Lightbox methodology[3] to

Digital Object Identifier 10.5594/JMI.2023.3263888
Date of publication: 9 May 2023

achieve zero-G in the film, *Gravity* because it was better suited to his directing style.[2] The comment is in reference to Ron Howard's "vomit comet" used in *Apollo 13*. The Lightbox was used as a way for actors to have an accurate visual reference for what they saw in space and the interactive light on the subjects was able to be captured in camera by the DP.[3] Both these circumstances, in turn, improved overall performance during the shoot and integration of the elements in post with the visual effects team.

Since the application of the LED workflows and technology on *Gravity*,[4] films such as *Interstellar*[5] also paved the way for more extensive and updated use cases in both feature and episodic formats such as *The Mandalorian*,[6] *The Midnight Sky*,[7] *The Matrix Resurrections*,[8] *Red Notice*,[9] *Star Trek Discovery*,[10] *West World Season 03 & 04*,[11, 12] *Black Adam*,[13] *Avatar: The Way of the Water*,[14] and more….

By design, *Camille* references Terrence Malick's "Days of Heaven"[15] and paintings by Andrew Wyeth. The director was interested in exploring ways to create naturalistic images and verité camera work using LED volumes to prove that it is possible to capture supporting visual effects in-camera that add value to a project (even on an independent budget).[17-21] [Link to Visual Reference]

Visual references are important because it would typically be possible to capture these images in a camera without visual effects. Because the ideal location for the story would require travel and a larger budget than existed, the filmmaker chose to use LED wall workflows so that they could bring the expansive field location to Los Angeles, control the weather and time of day, and schedule the shoot over two days.

Limitation: No time or budget to travel crew to a real location.

Solution: LED workflows made it possible to move the shooting location to a comfortable stage in Los Angeles for the cast and crew. This also simplified scheduling around talent availability.

Limitation: No budget for permits, fire crew, and safety to execute a fire scene.

Solution: LED workflows made it possible to control the weather, time of day, and fire with the existing craft roles hired for the shoot. No additional safety crew is needed. Could shoot night scenes during the normal workday.

Limitation: Unpredictable weather. Would need to schedule shoots around weather systems, travel to those locations, and hire crew and actors for a much longer period of time.

Solution: LED workflows made it possible to create extreme weather systems that could be art directed and scheduled around actors' availability.

Camille is a short R&D film created with the intent to share learnings with other filmmakers and push the boundaries of LED wall techniques in a small space with low ceilings. *Camille's* on-set engineering and stage ops were executed by Orbital Studios, who, at the time, were operating within Fonco Studios' main stage. (Orbital Studios has since relocated to its own facility in downtown Los Angeles.) The team at Orbital provided funding, certain crew roles, and donated services to support the film. The LED wall and camera tracking system were set up as a test stage that *Camille* and other projects could use to provide user experience feedback to the stage engineers for tool development.

One of the tools that was in development for the project was a virtual set extension toolkit that would allow the camera to "see tracked content" outside the edges of the LED wall. Theoretically, this would give the director the ability to capture extremely wide shots in a small space. The story and environment were originally designed for this potential capability and were modified when the tools could not be implemented during the prep and shoot schedule.

In 2021, during the production of the film, there were very few independent filmmakers who had access to the technology and even fewer that were documenting their processes. Several stages were forming in Los Angeles and few had the client feedback or engineering support to build some of the tools stages advertise as "standards" today (i.e., virtual set extensions, virtual zoom, accurate virtual depth of field, and so on). Additionally, most visual samples using LED wall workflows were showcasing advancements in vehicle work, stories that took place in outer space, or had access to extremely large stages with studio funding that could be used toward "enhancing" in-camera shots in post. Many articles, panel sessions, and live talks focus on tools that are available to independent filmmakers, but do not necessarily synthesize how teams are formed in relation to the creative process.

Virtual Production Is Its Own Department

Over the past 20 years, certain techniques such as visualization, performance capture, and volumetric capture have been considered part of the visual effects budget. It has only been in recent years, as virtual production workflows have become accessible to a wider variety of teams, that breaking out certain techniques into a separate department budget has become necessary to track the ownership of certain responsibilities, staffing, and equipment.

If substantial realtime workflows are necessary to incorporate into a project, virtual production should be its own department. LED workflows require a hardware and software ecosystem that work best when teams are able to engineer tools for specific directors, as aspects of the workflow take the form of interactive games (i.e., virtual scouting, Vcam, and so on). In the same way, a DP might assemble a custom camera, lens,

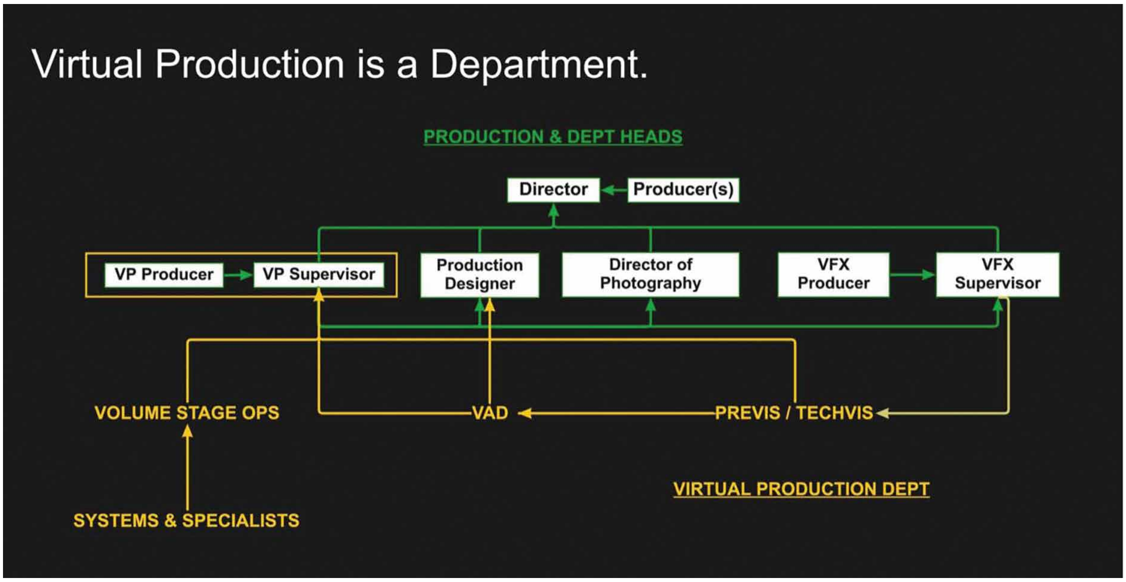

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 1. Organizational chart showing how the virtual production department fits into an LED wall-focused production.

and equipment package for a film, the virtual production VP supervisor takes into account the director's creative vision, desired user experience, and budget and customizes a solution that best fits the story (**Fig. 1**).

The roles in yellow represent the virtual production department for LED workflows and how it hooks into the existing relationships between other department heads and the director in green. The department head roles are not arranged in a hierarchy, but in a way that helps organize the team relationships with the virtual production subsets in yellow.

Relationship Between the VP Supervisor and Director

The director and VP supervisor should work closely together in the same way they would work with a DP, production designer, or VFX supervisor. The VP supervisor reports to the director. This individual supports the director's vision and ensures a positive "user experience" with the technology and crew from prep to production. The more informed the director is about the strengths and limitations of the technology, the better decisions that person can make under pressure on set to prioritize getting the shots needed.

On *Camille*, the crew was primarily volunteer-based, including the VAD. With only 14 weeks allocated for both prep and production, the VP team needed to be small and nimble. The director (who had experience with these workflows) assumed the role of VP supervisor and guided the visualization process for the DP and production designer both with virtual VAD meetings and with on-set testing. Having the director engaged in the visualization process streamlined content validation, prelights, and camera testing because they could approve creative, experiment with the technology over time, and strategize how to navigate the limitations of the tools to get the best results during the shoot.

Table 1. Virtual production roles for LED wall workflows.

Role	Description	Reports To:
VP Supervisor	Head of Department	Director
VP Producer	Bidding & Budget	VP Supe
VAD Lead	–	VP Supe
VAD Artists	–	VAD Lead
Visualization Lead	–	VP Supe
Visualization Artists	–	Visualization Lead
Volume Stage Supe	–	VP Supe
Volume Stage Prod Manager	–	VP Supe & Producer
Virtual Gaffer	–	Volume Stage Supe
Content Operators	–	Volume Stage Supe
VCam Technicians	–	Volume Stage Supe
Unreal Engineer	–	Volume Stage Supe
Engineering Lead	–	Volume Stage Supe
LED Technicians	–	Volume Stage Supe
Video Routing Tech	–	Volume Stage Supe
Systems Admin / IT	–	Engineering Lead
Color Science Specialist		VP Supe

**Some roles will report to different people during prep versus production.*

If the scope of work and complexity of the workflows increased, the roles on *Camille* would have created a bottleneck in the creative process (**Table 1** as a reference for suggested roles for LED wall workflows).

Virtual Production Team Structure for Projects That Require LED Wall Workflows

The film *Camille* was a low-budget, independent R&D film that did not operate with the formality of larger sets. In addition to the VP supervisor, there were three VAD artists and two unreal operators on stage. The VP supervisor generated visualization, organized test shoots on stage, and communicated technical plans to each department.

Table 1 shows a more robust list of the types of roles to expect on a project with formal roles.

Inherently R&D, LED VProd workflows require a deep knowledge base in cinematography, visual effects, and engineering. The team's skill sets often need to overlap in ways that are essential to agile workflows and may contradict the teamwork style of traditional craft roles or visual effects "waterfall" workflows. That means that within one department, there is an intersection of at least three different communication and management styles working together. Engineering in the context of LED wall workflows can be both physical and virtual. In a physical sense, it may be necessary to hire an engineer that can 3D print custom parts on set to improve camera tracking mounts or modify physical components. The other necessary tools will need to be developed by a programming team that works in C or C++ on the game engine itself or developing new systems that are not possible to realize with the various scripting systems that game engines utilize. For example, if the director would like the depth of field to shift accurately in a virtual environment when the camera operator racks focus, this tool will need to be built. Many stages that offer "end-to-end" LED services (i.e., visualization through production) are developing basic toolkits that creatives can use without billing additional tool development to a project. However, customization is often necessary to support unique workflow requirements.

The virtual production department requires leadership that understands these nuances, anticipates how the differences can cause friction on set, and clears a path for the artists involved, the department heads they are collaborating with, and the director. The flow map in **Fig. 2** shows another way to see how the roles in the LED Team Org Chart fit in the phases of prep and production. This flow map assumes that the LED stage/ecosystem is built and functioning during the story phase of the project. If the team were to build a custom stage and hire a team, the Volume Stage Ops and systems and specialists would need to be hired during the story phase.

It is important to note that the VP supervisor and producer should join a project during the story or development phase to help establish the scope of work needed for content on the LED walls. For this department, the time between story development and the final approval of designs made by a production designer is the "prep" period for the VP supervisor. It is a critical time to build a road map for the team and lock in team hires. As soon as the designs are approved to pass to a VAD, the virtual production department is considered "in Production." **Figure 3** shows very general phases for when each team inside a virtual production department should be brought onto a project.

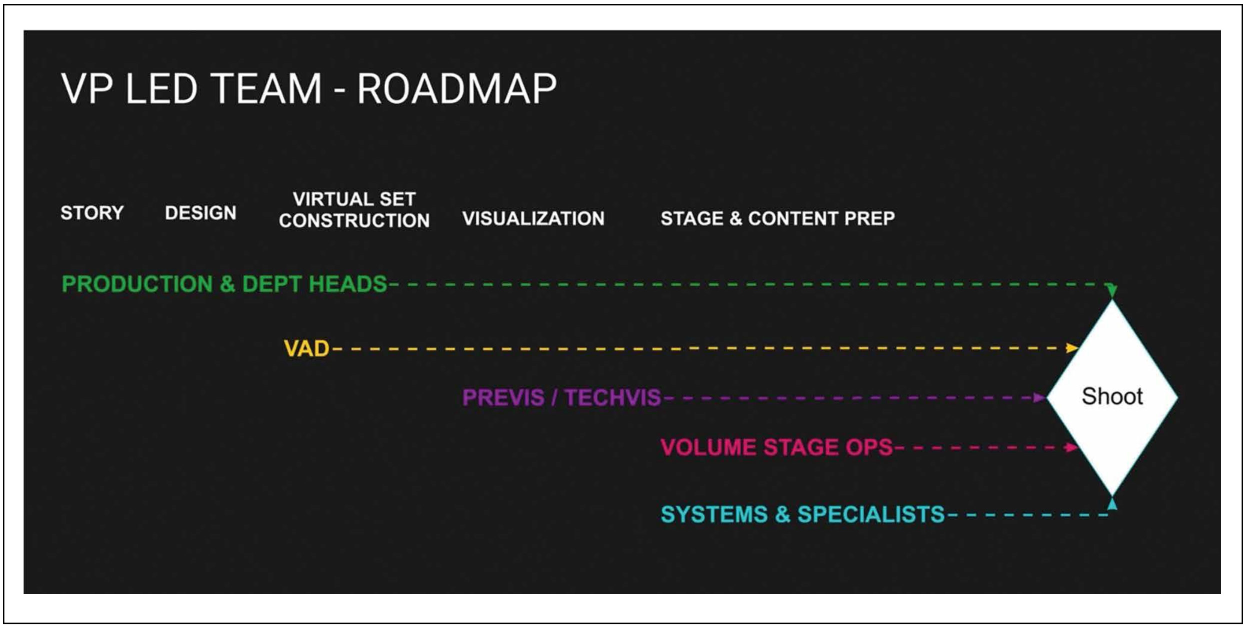

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 2. Chart showing general timing for members of a virtual production team to be involved during the preproduction process. The VP supervisor is hired during the story or development phase of the project and is responsible for hiring and managing the rest of the team.

One major difference between VP supervisor and other department head roles is that it is the only role that directly assists the director, production designer, DP, and VFX supervisor (and other department heads) with prep for their department's needs on set, as well. Virtual set construction, visualization, and virtual scouting are only a few ways the VP team supports other departments. What may not be visible to everyone, during prep with production, is that the VP supervisor is also assembling their Stage Ops team and supporting systems and specialists. On projects of different scales, this will take on different forms (i.e., short film, episodic, feature). However, early content testing, playback system testing, color workflow testing, and machine room testing are necessary before formally organizing prelights and test days for the director, DP, and other crew. The scale (or scope) of a project is determined by the content (which aspects of the script the director would like to apply these techniques) and the director's working style (how the techniques should be applied to have the best user experience).

Production must be aware of and budget for prelight and test days before the final shoot days. They are necessary for each department to be prepared to work with technology as well as the new VProd department crew.

Case Study #01—Exploring "Final Pixel" Computer Generated (CG) Environments Created in Realtime Engine

Final pixel is a subjective term that implies that the content being captured on an LED wall is of the highest quality possible. Whether the goal for the images is stylistic or photorealistic, the implication is that when viewed through the camera's lens, they are approved by the director on set as "final" quality to suggest they will not need to be replaced in post. Depending on the creative needs, these environments could be 3D, 2D, 2.5D, or a combination of these. One goal for the film, *Camille*, was to capture supporting "final pixel" in camera visual effects by using an LED wall to create visible set extensions.

How Long Does It Take VAD to Create a Final Pixel Environment?

Independent projects often have limited funding and resources. When budgeting for 3D environments with dynamic lighting and/or FX, allow VAD a 16-week minimum time frame to deliver final pixel quality. Each environment should be bid individually. This estimate may vary depending on the complexity of the creative, artist seniority, and crew size. The team's showreel should be referenced to understand what result 16 weeks of work would produce and whether more or less time may be needed. Additionally, as technology improves over time, this may also factor into a VAD team's estimate.

For *Camille*, the scope of work was contained by focusing on the construction of one hero environment that could be repurposed.

A three-person VAD team built the initial 3D wheatfield environment using free assets from Unreal Engine's marketplace in less than one day and increased photorealism by integrating 2D photographic elements. This did not include final lighting or scene optimization.

VAD had a 12-week construction schedule for layout, lighting, and FX integration for 4–5 variations of the environment that would showcase different times of day, weather, and fire. Each variation required a unique approach to construction and an approval process with the director. Through a series of tests, it was determined that 2.5D and 2D solutions for the storm system and fire would help the VAD team reach their visual quality goals and complete the work on schedule.

Validating "Final Pixel" Quality

To determine whether an environment had reached the approved "final pixel" quality, the director and VP supervisor would review each environment remotely on a color-calibrated computer screen with the VAD lead at Rotu Entertainment. If the director approved the overall look, the VP supervisor would prep the content for a visual test day on the LED wall to include physical lighting and set dressing. The VP supervisor would work with other departments and the stage manager to coordinate tests. Each test was also a prelight opportunity for the DP. The tests required a camera and lens kit with similar specs to the camera that would be used during the shoot. Any differences in sensor or lens relationships were noted and their mathematical comparisons were recorded for reference. Viewing the content and physical set dressing with basic lighting setups gave the VP supervisor more visual information to help make the decision that the virtual sets were of the appropriate visual quality for the final shoot. The tests were recorded for the VP supervisor to analyze. Only specific tests were shared with the director, VAD, and other department heads to help communicate certain changes that needed to be made. This process worked because *Camille* consisted of a small team with less formality than a feature or episodic set (**Fig. 3**).

It is important to note that "final pixel" quality expectations should not be set by a VAD team or Stage Ops team independently from each other. For example, if a VAD team promises a photo real result at 60 frames/sec during the shoot (a goal that has not yet been achieved), this "look" and frame rate should be tested on the desired final LED stage before establishing a scope of work on the project. To achieve that result is to assume the computer processors, playback system, network, and display meet the requirements for this promise. Alternatively, a Stage Ops team should not make the same guarantee independent of a VAD team. They would not necessarily be able to accurately speak

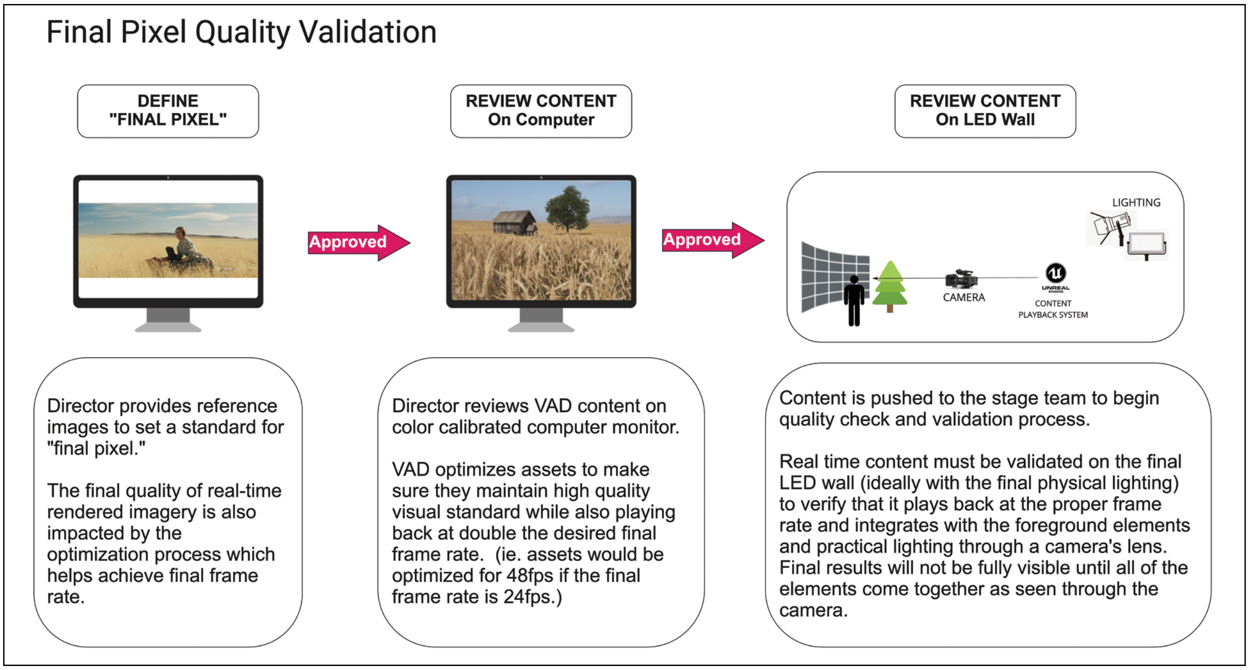

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 3. Content quality validation workflow diagram.

to the craftsmanship of the content or the scope of the asset optimization process needed. The VP Supe can help manage these expectations with the director and teams within the VProd department.

A major challenge early on was helping the director, production designer, and VAD team see how the 3D wheatgrass assets translated to the large LED wall display. On a computer screen, the 3D wheatgrass selected by the production designer and approved by the director looked like the correct scale. However, when the same environment was displayed on the LED wall, the scale of the wheatgrass, although "accurate" inside the game engine, looked very large and disproportionate in person and through a camera's lens in relation to real, living grasses sourced by the production designer.

During a formal test on stage, the production designer auditioned seven species of grass in front of the wheatfield content displayed on the LED wall. The VAD team provided two types of preapproved virtual grasses that the team could try to match with the physical options: a wheatfield with 3D-modeled grasses and a wheatfield populated with 2D card grasses.

The choice between each of these virtual grass options became an important crossroads as each required a different asset workflow, would affect the frame rate of the scene differently, require different methods for wind simulation, and integrate with the physical grass differently based on the physical lighting setup (**Fig. 4**).

3D Versus 2D Solutions for Selling Environments as "Real"

Whether a background is 2D or 3D, a wider shot framing and deeper depth of field require more detail for the environment to read as "real" to one's eye.

Integrating photographic elements, projected textures, and/or matte paintings are ways to enhance a 3D world if VAD has limited time.

Lighting choices also significantly impacted how virtual and physical elements integrated when viewed through a camera.

For scenes that required bright daylight, a mix between 3D-modeled grasses and 2D card grasses was used. 2D card grasses were placed in hero areas near the foreground of the environment because their scale and appearance were a closer match to the physical grasses selected. Although 2D card grass was a closer match, it did not read as photoreal to the eye when foreground elements were visible to the camera and required additional physical elements (**Fig. 5**, left image).

For scenes that required overcast light, 2D card grasses were used and appeared more photorealistic to the camera due to a combination of virtual and physical lighting changes, as well as lower overall contrast. To one's eye, the 2D card grass did not appear photoreal. However, when viewed through a camera with a curated physical set dressing and applying the lens' depth of field, the virtual and physical merged as one seemingly photoreal image. **Figure 5** is a frame taken from early test footage showing how diffused lighting and lower contrast helped blend the physical and virtual grasses (**Fig. 5**, right image).

Scenes that required visible, dark clouds portraying a large storm moving in required more art direction and custom world construction. Like the grasses, this introduced conversations about the use of volumetric or card-based cloud systems and wind simulation.

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 4. Foreground elements in a 3D environment will appear much larger in scale on an LED wall.

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 5. Although 2D card grass was a closer visual match to the physical grass, it did not read as photoreal to the eye when foreground elements were visible to the camera and required additional physical elements to improve integration (left). Grass tests were performed in several lighting setups before selecting either virtual or physical grass species (right).

Each of these choices impacts not only how the world is constructed, but also the game engine's ability to render them at the proper frame rate in realtime. For more exploration of this topic, refer to the section titled "Case Study #02."

For in-depth discoveries about 3D versus 2D assets for the fire use case, see the section titled "Case Study #05."

Replacing 3D wheatgrass with 2D card grasses improved scene optimization, reduced scale issues in the foreground of the virtual world, and matched the physical grasses available to the production designer. Prelight and test days are essential for lighting, production design, and VAD content to be integrated before the director review.

Strengths and Limitations of Marketplace Assets

Strengths: Free or purchased assets can accelerate environment layout which, in turn, expedite decisions on which areas of the world need higher fidelity textures. **Limitations:** Many assets are not created for cinematic use and are intended for use in games that have lower quality requirements. Purchased assets often need to be conformed to the project quality and/or optimization standard to avoid issues later in the pipeline. Confirm assets are compatible with the game engine version in use.

Case Study #02—Explore Extreme Weather and Creating Photorealistic Skies

In addition to subtle weather cues, *Camille's* script also required a large, distant weather system that approaches the farm creating a ticking time clock for the characters who must take shelter for safety.

Art directing clouds is not a simple task, whether one is working in postvisual effects or in-camera VFX. In both scenarios, the cost of rendering must be considered, as well as the ability to achieve a proper playback frame rate in realtime workflows.

An essential step in the process was determining which cloud workflows were available to the filmmakers, what could be accomplished on their timeline, and which technique would not significantly drop our frame rate during playback on set.

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 6. *Camille's* sky script shows the progression of weather from a bright sunny day with dappled clouds, overcast, dark storm, to fire.

Because the landscape had been optimized by featuring 2D card grasses in the overcast environment, the team could afford to spend more realtime rendered milliseconds toward clouds and lighting (**Fig. 6**).

Volumetric Clouds Versus 2D Animated Card Clouds

"Modeling and animating complex volumetric natural phenomena, such as clouds, is a difficult task. They are amorphous phenomena that elude traditional modeling techniques with their ever-changing volume-filling microstructures."[22] In the spring of 2021, during the production of *Camille*, the VAD team was researching both 3D and 2D workflows to achieve the moving, dark cumulus clouds provided in the sky script.

Scenes with bright daylight were most simple to make because only distant static clouds would be visible. For overcast scenes, the director wanted to see slow-moving clouds. For example, if the performance warranted a long take, the director wanted to be able to see the passage of clouds in the sky over time (in realtime) and also visible if scrubbing through the clip during playback. It was important to the director and DP that the movement of the clouds helped tie together the virtual and physical lighting.

For the bright sunny look, the VAD team used Unreal Engine's SunSky Actor in combination with a skybox. This would allow the director to change the angle of the sun and produce realistic patterns of sunlight and shadow wherever it was positioned. Skyboxes are textures that can be used to display distant objects and environments in a realtime scene. In this instance, a skybox was used to create distant clouds in the sky to match the provided reference.

Because the overcast skies required moving clouds, the VAD team applied a more advanced solution. Using the same skybox principle, they were able to apply a high-dynamic-range timelapse recording of a real sky instead of applying a static texture. The Hemisphere Plugin, which can be downloaded in the UE Marketplace, was designed for Unreal Engine (4.24–4.27, 5.0) and made it possible to select prerecorded, photoreal skies that could be adjusted in Sequencer. The VAD team could further manipulate the footage by altering the directional movement of the clouds, the cloud speed, and the time of day, however, these modifications were limited by the footage itself.

The scenes that required strong winds and large, dark cumulus clouds needed to have more components: movement, light interaction, apparent depth, and density. These requirements added a significant level of complexity to their design requiring the VP supervisor and VAD team to come up with a custom solution that combined 3D and 2D approaches.

Because the environment would be displayed on an LED wall on a very small stage, the VAD team had more control when combining 3D and 2D elements because they could be art directed from a single perspective. This meant there would be minimal visible parallax and the camera would not need to pass through the clouds in any way.

For the dark storm scenes, the VAD team had a library of preconstructed volumetric skies that could be used as a base and then further manipulated. The director could change the direction and speed of the cloud movement, and the time of day. In **Fig. 7**, there are three iterations of the dark storm clouds.

Both versions 01 and 02 are examples of volumetric skies that the VAD team had access to that could be used as base layers for the storm and clouds. Version 01 was not approved by the director because it was too monochromatic and had low contrast. The director wanted the clouds and lighting designed in a way that

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 7. Images show improvements in the art direction of the storm cloud system. Version 03 shows the use of a volumetric sky from version 02 with added 2D photographic elements to enhance the storm by making it appear more dangerous, dark, and looming lower to the ground.

created more depth and visible distance. Version 02 met those basic artistic requirements, but did not convey how close the dark, ominous storm was to the physical set and the distance it traveled to get there.

The director wanted the storm clouds to appear as photorealistic as possible. The photographic reference provided specified backlit cumulus clouds with varying top edges rimmed with light and dark shadows toward the bottom of the cloud.

Due to the limited project timeline and artist support for developing custom solutions in 3D, the director decided to use version 02 of the realtime volumetric clouds as a base layer and create a 2D photographic layer that could be composited onto the 3D scene. An example of the final composite is visible in version 03 (**Fig. 7**). The VP supervisor took a high-resolution photographic image of a storm cloud that fits the director's description and composited pieces of the photograph over different areas of the foreground sky. The VAD team was able to take the VP supervisor's photoshop file and apply the final composite in Unreal Engine.

This method ensured that the volumetric elements retained movement toward the camera. The volumetric clouds could interact with the lighting in the virtual environment including the sun source, resulting in visible shadows interacting with the grass. The 2D photographic images darkened the storm and made it appear closer to the characters in the foreground. Their sharp detail and contrast made the sky appear more photoreal to both one's eye and through the camera once the depth of field (and other settings) were applied. This is an example of a 2.5D approach (**Fig. 8**).

When evaluating which methodologies for cloud design are the best fit for a creative project, it is important to consider the quality needed, how difficult they are to make, and whether they meet the project's realtime requirements.

Volumetric objects, such as clouds, fog, or smoke, are typically created in software such as Houdini where the physics and look can be designed either as a static element or one that changes over time. A benefit of volumetric objects is seeing parallax as the camera moves horizontally or vertically in a scene. Elements that change over time are more complex. One methodology is to add computational fluid dynamics (CFD) to the volumes or mesh fluid simulations. These "are dependent on accurate physical conditions to simulate cloud development, which can be very time-consuming."[23] To use these volumetric objects in a game engine, they need to be encoded as "volume textures" or Voxel DataBase files (VDB) files. Whether an object is custom-made or downloaded from an existing source, it will need to be encoded for the game engine.

Smoke, fog, and clouds are unique objects because they do not have surfaces, they are volumes. To render these types of objects, a technique called ray marching must be used. With ray marching in the game engine, "the [light] ray's condition and its interaction with the scene are sampled at uniform steps along its path. If, for example, the ray were passing through the fog, each

FIGURE 8. View through the camera's lens—A 2.5D approach was used to create dark storm clouds.

step along its march would return the amount by which it has illuminated an area of fog. The cumulative effect of all of these samples can be accounted for in the color of the final rendered pixel."[24] Ray marching is computationally very intensive and in a game engine, which affects the frame rate at which imagery can render in realtime. Sometimes, this methodology is used until the director approves the final look. At that time the elements are "baked" (prerendered) and dynamic lighting is removed from the project to increase the frame rate. This is a form of optimization within a 3D workflow. This option would yield high visual quality, require an additional software pipeline with a prerendered output that would then need to be integrated into the Unreal Engine environment, and slow down the frame rate below 24 frames/sec. This methodology was not attainable in *Camille's* timeline or budget and was not necessary as the camera would not need to pass through the clouds.

The use of animated 2D card-based clouds (flipbooks) and "shader hacks" were implemented on *Camille* because they improved optimization and frame rate in realtime rendered scenes and produced high-quality results for most cloud types depending on the source images used. Although card-based solutions can deliver extremely high visual quality as they tend to use photographic sources, their limitations should be noted.

1) The camera cannot travel through cards in a way where they appear to be surrounded by the cloud (the camera would "pop" through the card).
2) Composition is limited in the game engine and would need to be manipulated in the source image that has been applied to the card as a texture.
3) Dynamic lighting is limited to information available in the source image(s).

A shader-based methodology might be as simple as creating a 3D sphere in a game engine, applying a Skybox material, and lighting the scene using a Sun-Sky Actor. In this scenario, a high-resolution photograph would produce the highest quality. The dynamic light would be limited to the sun angle captured in the photograph.

Shading is part of the rendering pipeline that gives each sample on a surface a value. Shaders can not only modify the color of an object, but they can also change the shape of an object. When applied to a CG object, a shader can also be referred to as a material. "Surface shaders deal primarily with color. Displacement/bump shaders alter the appearance of the geometry of the surface. They provide a way for the lighting artist to add visual detail to a CG object without adding complexity and weight to the model... Adding texture to an object is analogous to painting an object in the physical world. The added complication is understanding how to apply a flat, 2D image (or texture) to a 3D object. This process of transforming the 2D image space to 3D object space is known as mapping."[24]

Cloud maps are a slightly more complex 2D shader methodology because the 2D projected parcels can be stacked to create a unit 3D simulation cube. Each stacked plane represents a sampling point on the cast ray, for which the same aggregate texture is used (**Fig. 9**). "Cloud map CM is a two-dimensional bounded planar texture consisting of the spatial density values for the cloud mass at a given moment. The idea behind its use is to employ a pregenerated 2D texture of noise that guides cloud coverage and shapes in the sky. Furthermore, instead of the density values, CM stores spatially consistent and continuous noise values, which are then used to generate cloud densities...One possibility could be to use several such maps from successive timestamps

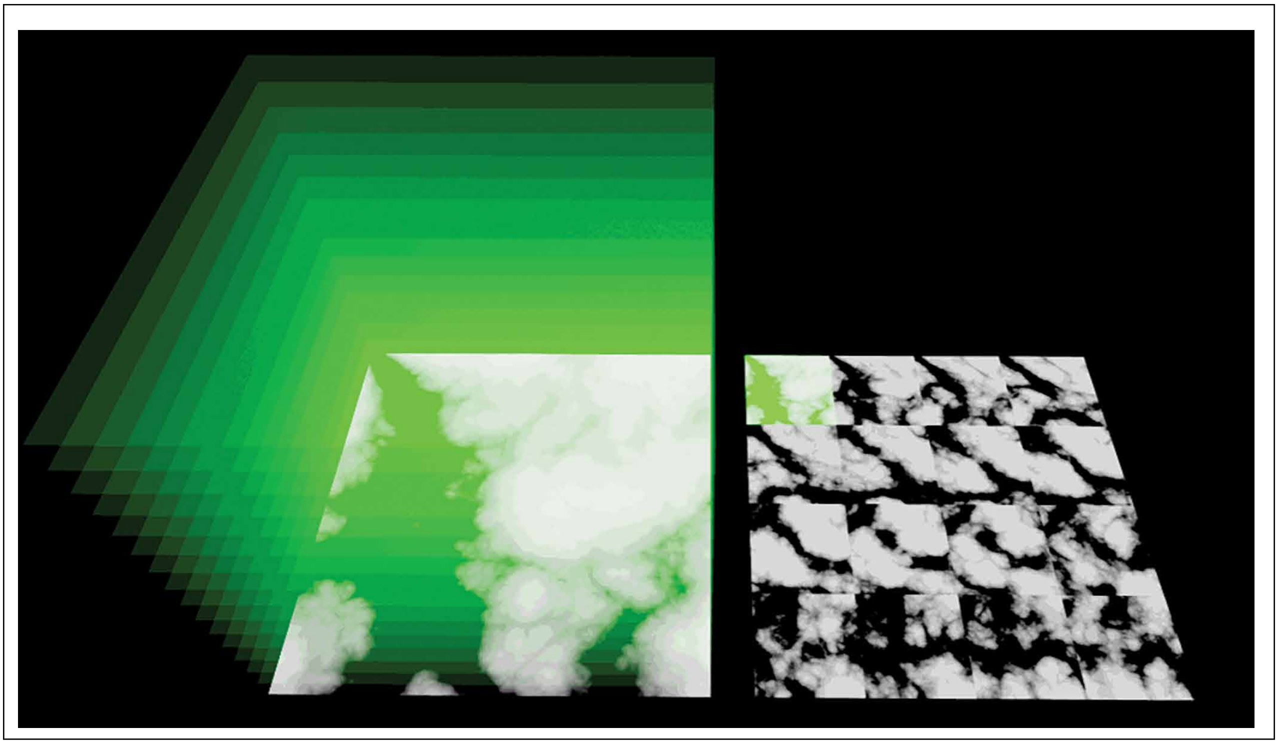
IMAGE COURTESY OF RYAN BRUCKS.

FIGURE 9. On the left, we can see all of the frames stacked vertically, representing the encoded position of each frame. The texture content is only displayed for a single frame at a time to help visualize this. On the right, you can see the 2D frames laid out. The frame that is highlighted on the right is the same as the frame displayed on the 3D slice on the left, to show how the top and bottom of the cube form the extent of the volume texture (and also the first and last frame of the "flipbook").

to create a continuous animation. The presence of a cloud map allows one to produce realistic cloud shapes but itself is insufficient to simulate the underlying cloud physics."[23] Today, blueprints exist that use cloud map concepts for plugins that allow for intuitive cloud painting[24] and are available to artists.[25] One way to add movement to 2D card workflows in a game engine is by using flow maps. Flow maps encode directional information into a card-based workflow.

The volumetric skies available to *Camille's* VAD team were available on the Unreal Engine marketplace. They consisted of a collection of materials and cloud volumes that replicated a VDB-like actor workflow for highly customizable volumetric skies. Adding individual 3D volumetric clouds and/or 2D sprites are other ways to enhance a scene.

Once a volumetric object is encoded into a volume texture (VDB file) for use in a game engine, the same "flipbook" principle is applied, however, to a 3D volume. "A flipbook is simply a collection of 2D frames laid out in a series. Flipbooks are actually very similar to volume textures: the only 'difference' is that for a flipbook, we consider **time** to be the 3D and for a volume texture we consider another **spatial dimension** to be the 3rd dimension, typically the *Z*-axis. This distinction is subjective as they are both just another dimension in the data. This means that we can treat volume textures just like flipbooks by replacing time with a spatial axis."[26] If time also needed to be visually communicated using this methodology, it would be considered 4D and custom engineering would need to be factored into the design process.

As technology advances, adopting machine learning and artificial intelligence to improve materials and texture animation may continue to be ways to improve the illusion of "realistic 3D cloud animation without introducing too much computational overhead."[27]

Wind Simulation

The scenes leading up to and during the storm required increasing wind speeds, which needed to be represented in both the physical and virtual worlds. Choices made in regard to the construction of the environment from clouds to grass impacted how wind simulation would be achieved and interpreted in the game engine.

A major difference between the 3D and 2D grasses was that they required different processes for emulating breeze and wind, which impacted the final production design. The 3D grass would interact with physics-based wind simulation in the Unreal Engine and would appear to sway and bend as individual stalks. The card grasses used animation to emulate a breeze affecting large sections of grass at one time. In a game engine, both types of breeze emulation read as "real" to the viewer on a computer screen. However, the differences were significantly noticeable when displayed on the LED wall.

For example, the animated card grass movement displayed at a larger scale did not translate as wind to the eye. The movement looked like ocean waves on the large LED wall. What looked like wind on a computer screen appeared to be an exaggerated up-and-down movement across large areas of grass.

The issue with the 3D grass was not the appearance of the wind simulation. It was how "heavy" the scene became when they were densely populated in the scene. Populating with individual 3D assets that also interact with physics-based wind is taxing to the computer processor and GPU. The VAD team was not able to optimize this workflow in the short project timeline provided.

In reference to a game engine project hitting a 60 frames/sec goal where the amount of time, one is allowed to resolve the entire update loop and draw and the entire frame is 16.6 ms—"For the GPU, a typical distribution is that the environments are allotted around 9 ms to draw, less than 0.5 ms—totaling 15.5 ms. This leaves a touch of headroom for the remaining operations to be complete. In an ideal world, the developer will aim closer to a low 16 ms to allow for headroom so that spikes in complexity will not compromise the frame rate. Culling mechanisms such as the frame rate can be used to reduce the number of objects that need to be drawn and the level of detail required".[24] *Camille's* target frame rate was 24 frames/sec, which meant the VAD team technically had 41.667 ms to work with in terms of resolving the update loop and draw. As a metric, the team worked toward optimizing the scene for playback at 48 frames/sec (20.834 ms) to account for frame rate loss that could occur due to NDisplay, Frustum, and camera tracking integration.

Instead of using physics-based wind simulation or card animation, the creative team decided to rely on virtual and physical lighting techniques as well as physical wind machines that would affect the foreground grass to sell breeze and wind. The ability to control the timing and speed of the wind on a studio stage helped keep production within its two-day shooting schedule.

How the Art Direction of the Weather and Clouds Impacted Virtual and Physical Lighting Choices

Creating believable weather systems and clouds was a collaboration between the VAD, the production designer, and DP. During early tests on stage, it became evident that believable wind simulation and lighting would need to be carefully intertwined in both the virtual and physical worlds.

The VAD team was able to use volumetric skies, which allowed for cloud movement, interactive light, and shadow movement. By seeing shadows form and move across the landscape over time, it made the illusion of clouds moving in the wind come to life. There was just enough interaction to sell the effect of wind over a field. During prelight, the gaffer was able to match the subtle moving shadows on the actors in the foreground as well, helping tie the physical and virtual lighting together. He used a large black flag to manually create more or less negative fill during each take by obstructing or revealing light on the subject.

Case Study #03—Testing Lighting Options on a Small Stage With Low Ceilings—Hard Daylight, Overcast, and Dark Storm

Stage Dimensions and Technical Setup

The stage was relatively small with dimensions of 50-ft long by 34-ft wide with a 15-ft high ceiling. The LED wall was 31-ft wide and 9-ft 10-in. tall with a 1.9-mm pixel pitch. An OptiTrack camera tracking system was configured for the 25 × 25 ft area in front of the LED wall. The LED panels were provided by Planar (Planar CLI 1.9 mm) and the processors were provided by Colorlight. The Planar CLI 1.9-mm LED panels formed a wall that was 20 panels wide ×6 panels high, with a total screen resolution of 5,120 × 1,536. The playback system on *Camille* was Unreal Engine 4.26, which supported texture resolutions up to 8,192 × 8,192 and a maximum final output of about 4k mapped to the LED walls via NDisplay plugin. The computer processors' maximum resolution playback per node was also 4K, 3,280 × 2,160. For 2D video playback the team used Quicktime Player to display 4K footage. Quicktime Player would not be ideal for most productions. This was used as a cost-free, quick workaround on set as there was little time to address issues with the other playback system.

There are a few factors that impact the resolution of the content when played back in an LED wall ecosystem. The panel cabinet resolution, the total resolution of the wall, the maximum resolution output of the content playback software, the maximum resolution output of computer the processors, and in some cases the GPU (graphics card) all impact the final resolution of the images that are displayed on the LED wall. Ensuring that all cables connecting each hardware component support 4K or higher will also impact whether the proper output is visible.

In the center of the stage was a large support pole that determined the furthest distance the camera could be from the wall and which camera angles could be used (**Fig. 10**).

Measuring from the front, center of the LED wall, there were about 25 ft in front of the wall to stage actors and set design and an additional 12 ft behind the center pole on either side for crew. As the stage was not configured with in-camera set extensions at the time, the ability to capture wide shots in the camera was limited by the height and width of the LED wall. Any extreme wide shots would become postvisual effects shots.

With these tight physical parameters, it was ideal to have panels that had a 1.9-mm pixel pitch, so the camera and actors could play in focus closer to the LED wall without the camera seeing moiré (see **Fig. 11**).

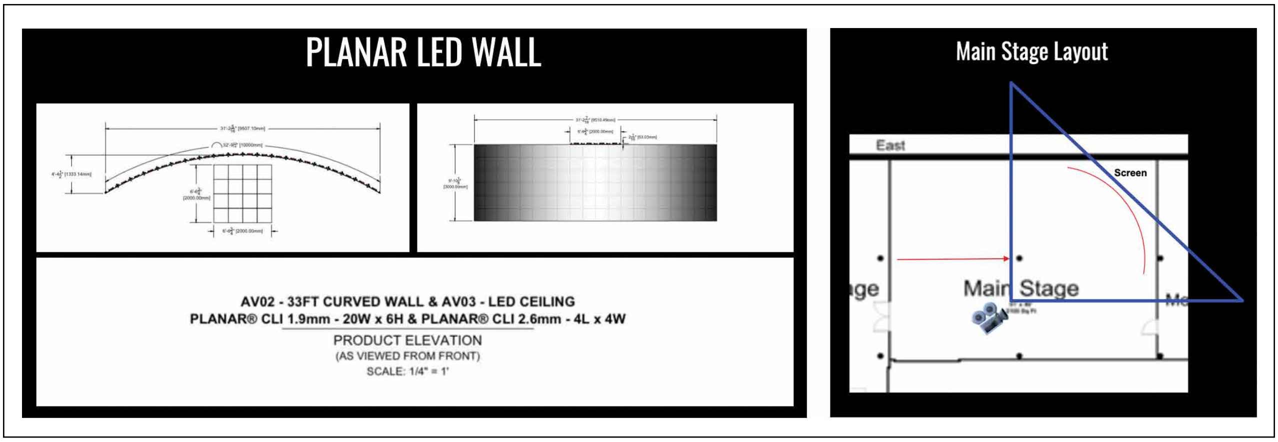

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 10. LED wall dimensions (left). Main stage layout showing the LED wall screen placement in relation to the support pole in the center of the room (right). The blue triangle shows the main area for the actors and set. The red arrow identifies the pole in the center of the room, which limited camera placement in the space.

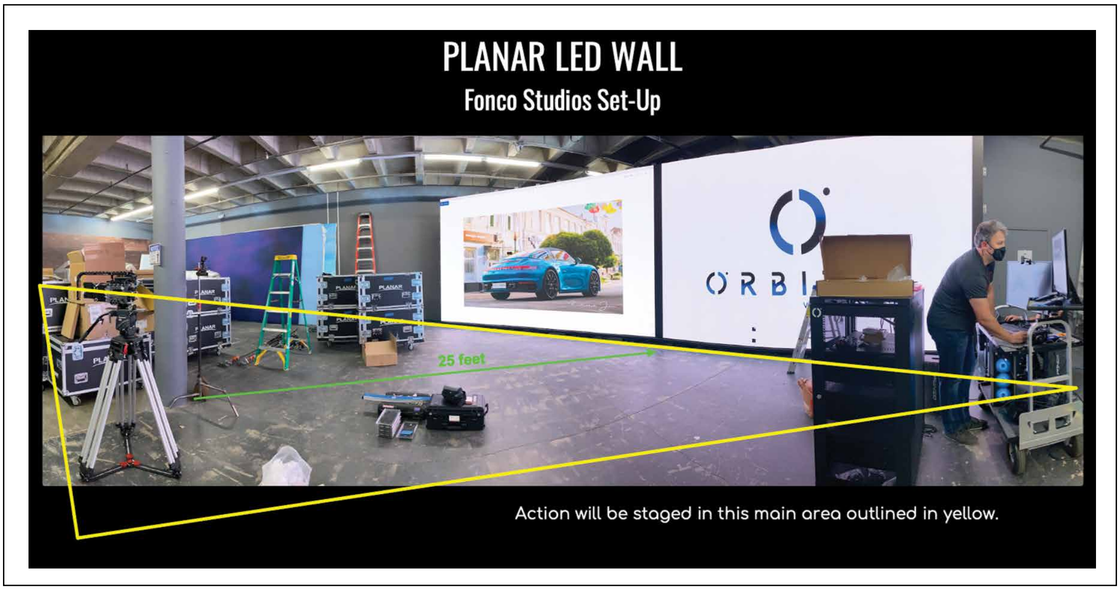

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 11. View of the area in front of the LED wall that was not obstructed by the center pole.

Lighting Considerations for a Small Stage

To maximize the space and ceiling height the gaffer used smaller lighting fixtures that were LED with RGBW color spectrum. To avoid issues with fire sprinklers, fixtures that produced minimal heat were ideal (**Fig. 12**).

When reproducing daylight in a small space, tools that shaped light enhanced its texture, and intensity gave the DP more flexibility. The use of negative fill, rim lighting, egg crate diffusion, as well as parabolic reflectors and/or fixtures helped control light on subjects and keep light from spilling onto the LED wall's surface.

Overcast lighting was a softer variation of the bright, evenly lit environment. The Arri L5-C Fresnel sources were replaced with Arri S30 Sky Panels and these were also directed into white bounce cards and the ceiling to fill the volume with light. The gaffer was able to create depth on the small set by using negative fill on subjects and emulating moving shadows motivated by volumetric clouds visible in the content.

During the darkest moments of the storm, when the subjects enter the shed set piece, the gaffer relied on the LED wall as the key light for the subjects and used Arri S30 Sky Panels bounced off the ceiling to evenly light areas further from the wall (**Fig. 13**).

Case Study #04—Explore FG Set Piece Integration With Virtual Content on LED Wall

The use of foreground set dressing helps to blend the virtual world with the physical. In *Camille*, the use of natural elements and partial structures was key to repurposing virtual backdrops, including a large three-walled shed set piece and a wooden fence.

The director and DP preplanned where the camera positions would be located inside the shed to reduce the scope of the physical build. The production

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 12. Gaffer used a mixture of Arri S30 Sky Panels, Arri L5-C Fresnel, Hudson Spider Parabolic LED light, large white bounce, and black flags to control light on a small stage with low ceilings.

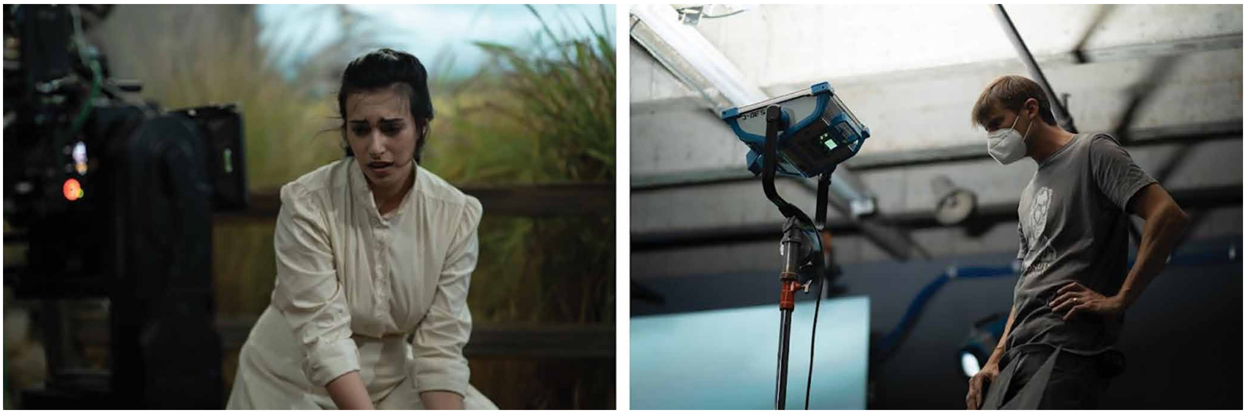

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 13. Gaffer used the cement ceiling as a bounce for Arri S30 Sky Panel to help create an even overcast lighting setup.

designer maximized the small stage by designing a set piece that hid the large support pole in the center of the room, which, for other scenes, was a challenge to work around and gave the camera several angles to choose from. The shed walls allowed the crew to extend the set 6–12 ft further from the LED wall than in previously photographed scenes.

Adding physical foreground elements and set pieces draws attention away from the background and enhances the photorealism of a scene in the camera.

Planning the camera positions during preproduction can help reduce the cost of physical set builds (**Figs. 14** and **15**).

Case Study #05—Testing In-Camera Capture of Low-Light, Darkness, Silhouettes, and Fire FX in Realtime

Producing and shooting "real fire" is challenging, seasonal, and maybe a safety issue for film crews. Night shoots are challenging to schedule as they can wear out a cast and crew if they are scheduled too early in a production schedule.

Camille tested both 3D and 2D options for re-creating fire on an LED wall to avoid night shoots, additional safety costs, permitting, and improving working conditions for the crew (**Fig. 16**).

In 12 weeks, the VAD was able to produce three iterations of a 3D wheatfield at night with a burning line of fire. Although each version improved in realism, the team did not have enough time to get the visuals to final pixel quality. Details such as the spawning, growth, and looping of the fire and smoke were going to take additional time whether they were designed in an Unreal Engine or on software such as Houdini. A Houdini pipeline may have improved the visual quality of the rendered fire FX, however, it would require additional software, technical artists, art direction, rendering, and integration with Unreal Engine, which time or budget would not permit on the project.

Additionally, each iteration of the 3D realtime fire was not able to be optimized to playback at a minimum of 24 frames/sec. The slow frame rate was visible to the camera and reduced the realism of the fire. The highest

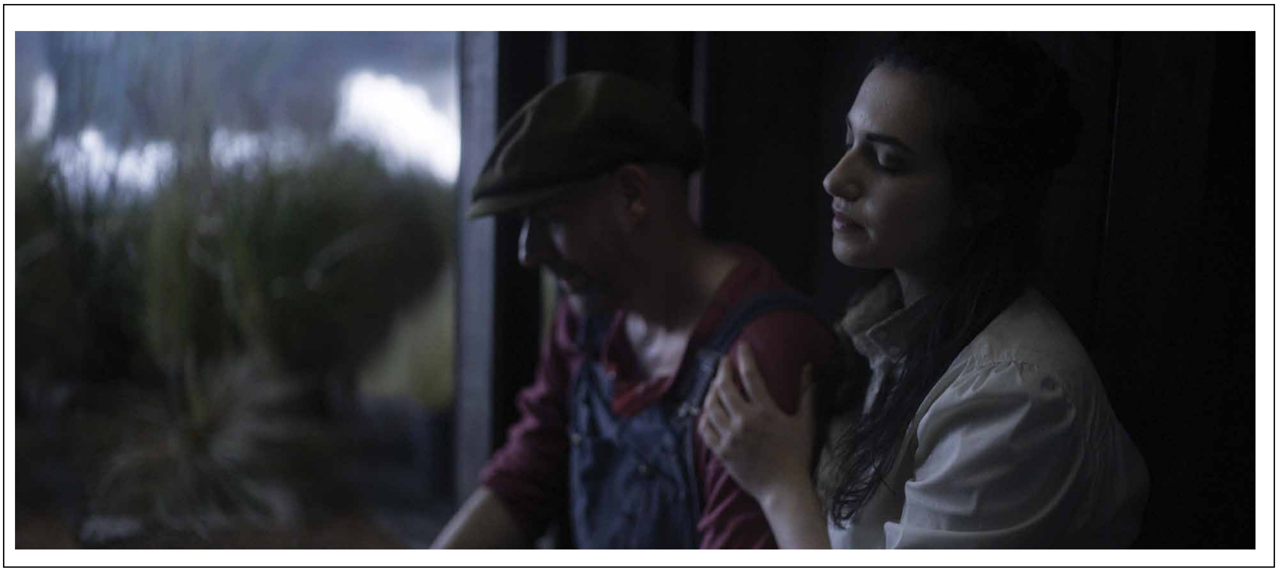

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 14. View through the camera lens—in-camera visual effects are integrated with the physical set.

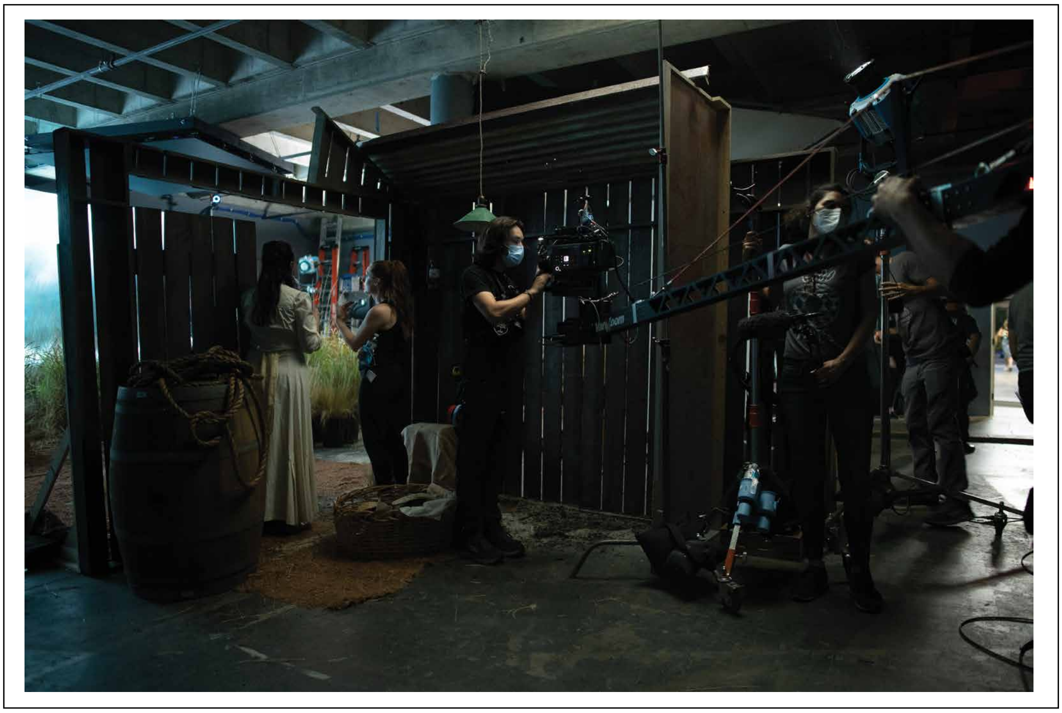

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 15. Shed set piece allowed the crew to maximize the studio space. The base of the camera crane, which is out of frame, is touching the back wall of the stage as a reference for scale.

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 16. First pass 3D fire does not run at the proper frame rate and requires further art direction (left). The second pass 3D fire and smoke require further art direction regarding the growth and decay of each effect, looping animation, and improved optimization (right).

IMAGE COURTESY OF KATHRYN BRILLHART.

FIGURE 17. View through the camera lens of final footage—the director chose a 2D 4k photographic plate of fire as a background element displayed on the LED wall. The actors are lit only by the LED wall.

frame rate that the 3D content would playback was about 7–12 frames/sec.

During prep, 2D plate playback tests with fire elements were also conducted on the wall using photographically captured 4K footage. The tests produced high-quality results and became the approved technique used during the shoot. Before planning a highly complex shoot, where the crew would actually set a field on fire at night, the director looked through existing high-resolution footage libraries to see if they could find photographically captured images that had the correct look and feel for the scene. The director was able to make a final selection from a footage library and saved travel, crew, equipment, and safety costs to be applied to other areas of the production. The fire scene proved to be the most cost-effective in the film, which might not have been the case had it been filmed on location.

The gaffer was able to eliminate light spills on the LED wall by using the footage playback as the key source light, which also improved the richness of the panels, especially in the blacks. The final captured footage that show the use of 2D photographic plates are illustrated in **Fig. 17**.

Conclusion

When planning to incorporate LED wall workflows into a project, whether a Hollywood feature or an independent project, it is important to identify the key problems that the technology will be used to solve to achieve the director's creative goals. Strategically applied, one can reduce travel costs for crew and talent, art direction and schedule weather systems, avoid night shoots, and reduce carbon footprint. A virtual production supervisor and producer can help manage these expectations if brought onto a project early during the story development phase of the project.

One of the most challenging aspects at the beginning of an LED wall workflow is agreeing on the scope of work including the in-camera visual effects requirements for the project. It is important to identify whether there are "final pixel" requirements and if so, select visual references to set a standard early on.

Each component of a realtime environment must share a total amount of time to resolve the entire update loop and draw based on the final playback frame rate requirements on an LED wall. Carefully assessing the construction of the world and how all of the elements are integrated will help inform how it can be art directed virtually in VAD and physically on an LED wall during a shoot. The hardware and software ecosystem can also impact these draw calls and may require further optimization once the content is finalized.

When selecting an LED stage for production, make sure that the stage team has a developed set of tools to improve the director and DP's user experience. If customization is foreseen, make sure to plan tool development into the project's budget and schedule well in advance of testing and production.

Although independent filmmakers will always need to be more scrappy in terms of developing a creative process, structuring a team, and accessing resources, there are now more accessible realtime virtual production workflows at low cost that can help improve the quality of a project.

Acknowledgments

This endeavor would not have been possible without the support of Orbital Studios, Fonco Studios, Rotu Entertainment, Epic Games Unreal Engine, and the cast & crew.

References

1. N. Kadner, *The Virtual Production Field Guide Volume 02,* M. Perkins and M. Bousquet, Eds., Epic Games, p. 96, Apr. 2021.
2. "Gravity: Vfx That's Anything But Down to Earth—Fxguide." Fxguide, Oct. 8, 2013. Accessed: Mar. 2023. [Online]. Available: https://www.fxguide.com/fxfeatured/gravity/
3. S.-F. Wong, "Gravity—The Light Box." YouTube, May 28, 2014. Accessed: Mar. 2023. [Online]. Available: https://www.youtube.com/watch?v=OXpKvFwDOF4.
4. A. Cuarón and J. Cuarón, *Gravity*. Directed by A. Cuarón, N/A, 2013, N/A.
5. J. Nolan and C. Nolan, *Interstellar*. Directed by C. Nolan, N/A, 2014, N/A.
6. J. Favreau, *The Mandalorian*. Disney, 2019.
7. G. Clooney, *The Midnight Sky*. Netflix, 2020.
8. L. Wachowski et al., *The Matrix Resurrections*. Directed by L. Wachowski, Warner Bros. Pictures, 2021, N/A.
9. R. M. Thurber, *Red Notice*. Directed by R. M. Thurber, Netflix, 2021, N/A.
10. B. Fuller and A. Kurtzman, *Star Trek: Discovery*, 2017.
11. J. Nolan and L. Joy, *Westworld Season 03*. HBO, 2020.
12. J. Nolan and L. Joy, *Westworld Season 04*. HBO, 2020.
13. J. Collet-Serra, *Black Adam*. Directed by J. Collet-Serra, Warner Bros., 2022, N/A.
14. J. Cameron, *Avatar: The Way of Water*. Disney, 2022.
15. T. Malick, *Days of Heaven,* T. Malick, Ed., Paramount Pictures, 1978, N/A.
16. J. Nolan and L. Joy, *Westworld Season 04*. HBO, 2022.
17. A. Wyeth, *Christina Olson* (Painting), 1956.
18. A. Wyeth, *Christina's World* (Painting), 1948.
19. A. Wyeth, *Fields of Broad Cove* (Painting), 1941.
20. A. Wyeth, *Oil Lamp* (Painting), 1945.
21. A. Wyeth, *The Stone Fence* (Painting), 1946.
22. J. Schpok et al., "A Real-Time Cloud Modeling, Rendering, and Animation System," in *Eurographics/SIGGRAPH Symposium on Computer Animation*, D. Breen and M. Lin, Eds., The Eurographics Association, 2003.
23. P. Goswami, "Interactive Animation of Single-Layer Cumulus Clouds Using Cloud Map," in *STAG: Smart Tools and Applications in Graphics*, M. Agus et al., Eds., The Eurographics Association, 2019.
24. A. Whitehurst et al., *The VES Handbook of Visual Effects*, J. Okun and S. Z., Third, Eds., Routledge, p. 595, 597, 601, 655, 2021.
25. A. Zhu, "All-in-One Cloud Painter." YouTube, Aug. 5, 2020. Accessed: Mar. 2023. [Online]. Available: https://www.youtube.com/watch?v=HKD7X_y7K1M.
26. R. Brucks , "Authoring Pseudo Volume Textures," in *Shader Bits: HLSL and UE4 Development Blog*, Oct. 18, 2016. Accessed: Mar. 2023. [Online]. Available: https://shaderbits.com/blog/authoring-pseudo-volume-textures
27. P. Goswami, A. Cheddad, F. Junede, and S. Asp, "Interactive Landscape–Scale Cloud Animation Using DCGAN," *Front. Comput. Sci.*, Mar. 2023, doi:10.3389/fcomp.2023.957920.

About the Author

Kathryn Brillhart is a cinematographer and director who leverages volumetric capture, visualization techniques, and supervising visual effects to enhance projects. For the past decade, she has helped to redefine best practices and standards in virtual production through her role on the Global Board of Directors for the Visual Effects Society. Brillhart is a Member of the Virtual Production Committee within the ASC Motion Imaging and Technology Council. Her recent work as a virtual production supervisor incorporates realtime game engine and virtual production techniques to achieve in-camera visual effects, including such projects as Warner Bros. *Black Adam (2022)* and *Amazon Studios' Fallout (2023)*.