

Multi-Device Storyboards for Cinematic Narratives in VR

Rorik Henrikson* Bruno De Araujo* Fanny Chevalier†* Karan Singh* Ravin Balakrishnan*

*Department of Computer Science, University of Toronto

†Inria

{rorik | brar | karan | ravin} @dgp.toronto.edu

fanny.chevalier@inria.fr




Figure 1: Our multi-device system supporting the planning of virtual reality stories: a storyboard artist sketches the virtual environment of a scene moment using a lightweight tablet-based interface, while the director experiences the scene within VR. Additional dynamic overhead (top-right) and panoramic views (bottom-right) offer different perspectives on the scene to further support planning.

ABSTRACT

Virtual Reality (VR) narratives have the unprecedented potential to connect with an audience through *presence*, placing viewers within the narrative. The onset of consumer VR has resulted in an explosion of interest in immersive storytelling. Planning narratives for VR, however, is a grand challenge due to its unique affordances, its evolving cinematic vocabulary, and most importantly the lack of supporting tools to explore the creative process in VR.

In this paper, we distill key considerations with the planning process for VR stories, collected through a formative study conducted with film industry professionals. Based on these insights we propose a workflow, specific to the needs of professionals creating storyboards for VR film, and present a multi-device (tablet and head-mounted display) storyboard tool supporting this workflow. We discuss our design and report on feedback received from interviews following demonstration of our tool to VR film professionals.

Author Keywords

Virtual Reality; storyboard; sketching; 3D; movie.

ACM Classification Keywords

H5.2 [User interfaces]: Graphical user interfaces; Virtual Reality; H5.m [Miscellaneous]: Stereoscopic display

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UIST 2016, October 16-19, 2016, Tokyo, Japan
© 2016 ACM. ISBN 978-1-4503-4189-9/16/10...\$15.00
DOI: <http://dx.doi.org/10.1145/2984511.2984539>

INTRODUCTION

Head-mounted displays (HMDs) have evolved over half a century from research prototypes [29] to consumer products [36, 37, 38]. The entry of VR technology into homes has caused a surge of interest in the medium among developers, filmmakers and storytellers. Major film festivals, studios, and technology companies [12, 21, 30, 39] have created teams specifically targeted at *VR stories*: where live-action, or animated narratives occur in fully immersive environments, where the viewer experience is uniquely intimate.

VR is characterized by a quality known as *presence* – the feeling of actually being on location in a story rather than experiencing it from the outside [27]. This new medium provides a unique opportunity for viewer engagement and the exploration of novel storytelling, but requires the development of new cinematic constructs and film language [24].

Creating a movie for VR is not as simple as taking a regular script and going through the well-developed film making process. The unique properties of VR require directors to consider concepts such as presence and peripheral vision, and use them effectively. Directors also must address the challenge of guiding an audience through a narrative, while leaving them free in a fully immersive environment, to look or move in any direction, and even trigger events within the environment. Traditional cinematic principles of cuing and staging can help solve this problem, but need to evolve with the immersive use of spatial visual and auditory cues.

We learned through our multiple interviews that, no one currently knows how to properly plan for immersive narratives. Traditional storyboards and planning tools are shallow and restrictive given the full extent of the environment that needs to be discussed. Communicating ideas between individuals is further impeded by the experiential quality of VR. Film teams thus rapidly model, collect, and assemble

assets that are viewable in VR, simply to form the basis for team discussions, so everyone can understand and plan for the unique challenges being faced by the project.

To help address this unique set of problems, we sought to explore a workflow specific to the needs of professionals creating storyboards for VR and the related challenges; to enable teams to conduct initial explorations and collaborate at/near the speed currently possible with traditional film.

To facilitate this goal, we conducted a series of semi-structured interviews with a variety of film professionals to fathom the design and technical challenges. The interviews gave insight on how experts regard the space of VR narratives relative to traditional filmmaking, and the processes and tools to plan for immersive experiences, in use today.

We used this information to design a two-device interactive storyboard system for tablet and VR, that leverages both the speed and skills of a storyboard artist, and allows a director to experience the results in VR simultaneously.

We address a number of issues in a working prototype of our collaborative multi-device storyboard design: ensuring homogeneity between the different viewing environments; supporting quick authoring, fast navigation, and a proper overview of the fully charted story. This prototype was presented to VR film professionals and feedback was elicited on different aspects of the design, its ability to address the needs of planning narratives in VR, and to assess if the proposed workflow was worthy of further exploration.

Our contribution is thus a principled survey of the state of the art in cinematic VR planning, and the first multi-device system that supports a storyboard workflow for VR film.

FORMATIVE STUDY

To properly understand the differences between traditional and VR film creation and current planning processes, we talked to a variety of professionals, who each had at least 10 years of experience working with film, television, and VR. We distill here the result of semi-structured interviews with three producers, four directors, one executive creative director, one technical director and one storyboard artist.

Stories in Virtual Reality

Unlike movies that are projected on a silver screen, VR films can intimately engage audiences. These films are often referred to as *stories* in the industry, reflecting a looser terminology used by professionals. We will henceforth refer to these VR *stories* as VRSs. Also, as with our interviewees, we refer to “story” or “narrative” as the content being told, while the “VR experience” is what the audience consumes.

Relationship of Story and Audience

At a high level, we formulate a film experience using three components: the *audience*, the *camera*, and the *narrative* (or story). The narrative is a collection of moments presented in some predetermined order. These moments may be temporally sequential, or non-linear, and they may be told within a continuous environment, or jump from one loca-

tion to another. The audience is a group of one or more individuals who experience the story being told. The camera is a mechanism to control this experience: framing shots, camera movement, stereoscopy, and field-of-view (FOV) are all variables in aligning the audience with the narrative.

With traditional film, the director is usually in complete control of the narrative and camera, with the audience generally experiencing the story from outside the movie’s environment. In VR, this arrangement is not as clear. With a VRS, the narrative remains in the director’s control, but the audience is part of the story; the narrative occurring all around them. The notion of a camera, is also less definite. The camera is loosely the viewer’s FOV (typically frameless) that the audience controls, though a director controls intrinsic camera parameters e.g. stereoscopy and lens filters (color, shading, graininess), and can subtly or explicitly guide or control audience gaze throughout the narrative.

Directors are actively experimenting with various configurations of our *audience-camera-narrative* model in VR:

virtual movie theatres; immersive environments where the viewer can look anywhere, but the narrative is focused in one primary place; environments where multiple narratives happen simultaneously in different places; a primary “stage” that slows down and eventually stops as the viewer looks away [41]; camera concepts where audiences attach themselves at any time to any character to follow that character’s story arc [4]; and the ability to teleport anywhere, to follow a story as much or as little as desired.

These techniques allow for different levels of interaction. We propose that a narrative shifts from a VRS to a game when a viewer’s interactions change the outcome of the story (e.g. a character dies as a result of a viewer’s choice). In this paper, we focus on VRSs, where the viewer control is limited to their consumption of the story, not its outcome.

Making VR Stories: Key Considerations

Though many of the concepts and approaches used for creating a VR experience are similar to conventional cinema, others are more akin to live theatre and plays. This interplay results in many new concepts that one must consider that are unique to authoring these experiences.

Presence – One of the biggest differences from traditional film is a concept that is referred to as presence [27]. Presence is the feeling of being inside a scene rather than looking at one from the outside. For instance, looking at a family vacation photo may cause you to remember your experience of being there, but is unlikely to invoke a similar feeling in an outside observer. VRSs evoke a sense of being there, that is not captured by 2D/stereoscopic imagery.

Active experience – Movies have generally been passive, static experiences for the audience, allowing one to sit down and be taken on a journey as the story is told. VR experiences allow audience participation, where viewers can actively look and potentially move around to explore the immersive setting, while following a narrative.

No framing – With traditional film, the view is defined by a border within which a director can layout or frame a shot. As a result, with traditional film, a significant amount of attention and effort is used to plan and design a shot. VR typically has no border and no guarantee where a user will look. Framing is thus harder, even conceptually, which changes how a director approaches the planning of a shot.

Blocking – With VR, the space all around the audience can be used to stage a film. For this reason, the concept of blocking, the process of deciding what will be where, where objects and people will move, and the composition of a scene as a whole within the set changes. “How much of the film will happen in front of the user?”, “What will happen behind the user?”, and “How should one guide the viewer’s attention from one area to another?” are just some of the considerations needed while planning the blocking for VR.

Peripheral vision – Since VR creates the sense of actually being on location, the director needs to think about the viewers’ lines of sight. In the past, things that would have been off-camera can now be seen peripherally.

Audience attention – Keeping the audience focused on the narrative is an added directorial burden in VR. Sound, staging cues, lighting and movement can all be used to draw the user’s attention; however, as an interviewee noted, these actions must be used sparingly, as audiences quickly desensitized and will ignore these techniques if used too often.

Limited Transitions – With traditional movies, directors rely on cuts, fade in/out, wipes and other shot transitions to keep the story engaging and to move the narrative from one moment to the next. In VR, transitions are more jarring and potentially disorienting, as they involuntarily “teleport” the audience. Interviewees indicated that VR directors have generally found that limited use of fades are the best way to transition between scenes. Due to this effect, many stories are told in one environment, effectively eliminating the need for transitions. Live theater has a similar constraint, not due to disorientation, but the inability to quickly or easily change physical sets or move the audience around.

Screen Format – A director has a wide choice of spatial staging formats in VR, from a virtual screen to a completely immersive environment.

Stitching – Current fully immersive VR video capture rigs mount multiple cameras omni-directionally, to immersively capture a live scene. Stitching these multiple camera images together often creates faint distortions and discontinuities along the seams. While future imagery may be seamless, current directors need to plan their shots to minimize viewer focus in the vicinity of seams.

Optics – A director’s choice to present their movie mono- or stereoscopically has a different dimension in VR. Stereoscopy is the effect created from each eye being shown a different image; one horizontally offset from the other, creating an illusion of depth through stereopsis. Leaving the



Figure 2: Previs (left) is used to plan for movie shots (right).
Image from Wee Biz Studios.

picture monoscopic, a director relies on depth cues like retinal size, overlap, and aerial perspective [20].

Duration – For the near future, VR films will be fairly short in length, both due to the complexity of authoring content and viewer fatigue from sustained viewing in VR.

Planning for VR Stories

Due to the multitude of decisions and technical challenges creating engaging experiences, directors need to plan a VRS in VR itself. This ensures that everyone works towards the same vision of the film and allows the film crew to organize a scene’s logistics. For instance, lighting technicians can figure out light types and placement, stunt coordinators can discuss choreography and rig setup, and problems can be anticipated before anyone is on set, preventing costly mistakes. The various unique considerations for VR, however, render traditional planning tools such as storyboards and previsualization (Figure 2), ineffective in VR.

Storyboards

Storyboards, traditionally, are hand drawn sketches similar to comic strips that outline the key moments of a shot for a scene in a film [13]. Typically, they are small, quickly drawn rough sketches made to capture the essence of a moment in the story without the need for any underlying infrastructure. They allow a director to step through a movie to discuss concepts and moments, their relationship in the narrative, and to plan a compelling story. These boards are generally created very quickly by one or two artists. Storyboards are often distributed on set to aid organization.

Though great for discussing shots, traditional storyboards have disadvantages with VR. First, one cannot convey presence with a small 2D sketch. Also, since VR is mostly frameless, framing can distract from understanding. Lines of sight on a 2D picture have an abrupt border, where our visual acuity gradually fades at the periphery in VR. VR planning further entails blocking a 3D environment around the audience, instead of a limited field-of-view 2D plane.

Previs

Previsualization, or previs, is the technique of creating (usually) a 3D computer model of a film set, populated with virtual cameras and characters to explore and refine narrative breakdown and plan shots (Figure 2). While directors can freely explore camera angles and transitions, previs is expensive, time-consuming, and requires the creation of many digital assets early in the filmmaking process. Conceptually, previs can be used in other film contexts [35].




Figure 3: Examples of co-located / synchronous (a, b) and distant / asynchronous (c, d) workflows between the storyboard artist (in a white shirt) and the director (in a dark shirt) using our multi-device approach.

Drawing and Painting in VR

There are some very compelling 3D tools for drawing and painting directly in VR, like TiltBrush [40], where users paint in the air using 3D controllers, and Hyve3D [9], where drawing on a spatially tracked tablet is mapped to virtual planes in a cave-like environment. Storyboards for early planning and blocking, however, should leverage existing skills of storyboard artist and support quick sketches of appropriate roughness, that provide a good spatial proxy for scene elements at the right level of abstraction.

Neither of the approaches are well suited to storyboards: in-air interfaces like TiltBrush are too slow for prolonged use when creating many storyboard panels, and have no natural abstraction for blocking. A planar sketching interface like Hyve3D provides such an abstraction, but controlling arbitrary planes in space can be overkill and thus cumbersome.

Further, while stacked planar sketching is ideal for traditional and stereo storyboards [14], a more natural spatial extension to VR would be drawing on concentric cylinders or spheres around the audience.

Current VR Planning Strategies

We learned from professionals that there are currently no standards for planning VR projects, as there is no sufficient process for discussing evolving VR cinematic concepts. Consequently, groups cobble together 3D visual mockups (using images, projected 2D storyboard panels, 3D objects, toys, cardboard cut-outs, previs and other proxies), so presence in VR can be experienced and discussed.

For example, one group we interviewed uses a workflow that combines Mettler (an Adobe After Effects plugin for 360/VR) and Adobe Photoshop. Each asset is sketched in Photoshop and imported and manipulated onto planes in After Effects. The overall scene mock-up is exported to be later viewed using a Samsung Gear VR headset. Any iterations require the entire process to be painfully repeated, which obviously affects productivity. We thus set out to provide a solution of the same quality and efficiency as traditional storyboards, that leverage domain expertise (mimics pen and paper), is easy to edit by an artist or director, and provides real-time feedback in VR, which are all design requirements critical to effective storyboarding [14].

TOWARD A SYSTEM FOR VR STORY PLANNING

There are a number of questions and issues that need to be addressed in the design of a VR storyboard system:

How does one allow for the speed and familiarity of drawing on a tablet/paper in an immersive environment? What should a VR storyboard look like? How does one abstract an immersive environment that surrounds the audience into canvases for sketching? How does one collaborate and communicate in a virtual space, when visually separated from other discussants? How do we ensure storyboard sketches provide sufficient volume to perceive depth and spatial relationships between objects in a scene?

We started by asking our interviewees if a VR storyboard would be useful. The executive creative director commented “to [him], storyboarding in VR is the first step of tackling a big problem when creating content for the medium, and that is, all parties involved have something different in their imagination. Similar to a film, storyboarding in VR serves the purposes of getting everyone on the same page, gathering around the creative vision. Having a VR storyboard [would allow] for all parties involved to understand how their roles play into the larger picture.” It was also stated that “Using a traditional storyboarding method would require exponentially more drawings/work/time, and even then there would be no sense of space or scale.”

We propose a planning tool based on a networked multi-device system, where one can simultaneously exploit the sketching abilities of a tablet, and immersive viewing by wearing an HMD (Figure 1). Through a “best of both worlds” approach, an artist can draw on a tablet, working with panoramic sketch planes, and themselves, or collaboratively with a director, view the results instantly mirrored in VR to experience the scene from within.

We illustrate different possible workflows enabled by our multi-device approach in Figure 3. If the director (in a dark shirt) and artist (in a white shirt) are working synchronously in the same location (Figure 3-a, b), they could both sit around the tablet and have a discussion similarly to what is done for traditional film, or both view it in VR. Alternately, the director can immerse themselves in VR, while the artist sketches on the tablet live.

In an asynchronous scenario, the artist could start by sketching and viewing on a tablet as is typical, or drawing on the tablet and viewing in VR. The director, working asynchronously, could at any time view the immersive VRS panels with the HMD and with the tablet add comments and notes for the artist to subsequently address (Figure 3-c, d).

While the physical tablet is not visible in VR, with the tactile feedback of the physical tablet, and the instant visual feedback of the sketch lines and stylus in VR, sketching is not a problem for digital artists used to drawing on digitizing tablets that are physically separate from the display.

To enable storyboard artists to sketch in a fast and familiar method on a canvas surrounding the viewer, and capture the well-understood cinematic concept of scene depth [14], we divide the space into a set of concentric cylinders around the viewer. Artists draw on unrolled cylinders shown as flat panoramic panels – a surface that is easier for artists to mentally unroll than a wrapped sphere. These panels can be rolled back into cylinders for immersive displays (Figure 5). While this depiction is best suited to a 360° view, it is nonetheless able to quickly depict 3D scenes that can be arbitrarily navigated and viewed for VR cinematic discussions.




Figure 5: Concentric cylinders are unrolled into panoramic views for sketching.

WORKING PROTOTYPE

We present a solution to enable initial assessments of our proposed workflow. Our implementation combines a tablet-based sketching tool (Microsoft Surface Pro 2.0) to author VR storyboards, and a synchronized HMD (Oculus Rift DK2) to experience the boards immersively.

Many design innovations of our dual system were based on supporting synchronous use of the tablet for sketching, and the HMD for experiencing the VRS panels from within. Given the lack of ability to currently storyboard for VR, rather than trying to establish specific input techniques, the goal of this first prototype was to explore and assess a workflow that enables teams to conduct early explorations, collaborate at/near the speeds currently possible for traditional film with a VR system, and to see whether storyboards are even useful at this early stage for VR.

We first describe the authoring of a storyboard on the tablet. Then we describe how the artist or the director experiences the result in VR, and can annotate using the tablet. Finally, we discuss how our dual system supports collaborations between the artist and director (Figure 3).




Figure 4: The tablet sketching view, integrating director's notes and indicators.

Authoring VR stories on the tablet

The primary storyboarding activity consists of sketching, and viewing a series of moments (*panels*) that make up a narrative. The tablet presents a manifestation of the unrolled concentric cylinders, enabling the artist to sketch on the 360° environment of the VRS panels.

Sketching View

Like traditional storyboard tools, our system includes a sketch canvas, where the artist can draw content for a panel. As in Storeoboard [14], she can add strokes at different view depths (i.e., concentric cylinders), by navigating a slider bar, to support the director's spatial vision. The bar is augmented with thumbnails to show the placement of the strokes that are present in the environment (Figure 4).

To maximize the drawing area, rather than presenting the entire panel unwrapped, the tablet shows a smaller portion of the current panel. Swiping left or right, the artist can access the rest of the drawing space. The canvas smoothly wraps around as the artist reaches the edge of the panorama.

As the artist develops the narrative, she can add new panels to describe new moments in the story. Tapping on the edge of the canvas, the artist can easily access the previous (left) or the next (right) panel of the VRS.

Technically, one of the main challenges of drawing on a cylindrical panorama is to provide a drawing surface that accounts for the canvas' curvature when projecting on the drawing surface. We discard such correction, i.e. we provide a flat canvas, since in practice, artists draw at sufficient viewing depths that the distortion becomes negligible.

Overhead View

An important aspect of storyboarding in VR is to visualize overall layout of the immersive scene. Drawing inspiration from a technique currently used by VR directors, we aid the quick placement and blocking of objects via an overhead view (Figure 6), which allows the artist (and the entire film crew), to see all sketch strokes for a VRS panel, in a target view (right) with a preview window (left).




Figure 6: Overhead view of the sketch surface, showing preview window (left) and overhead target (right).

The overhead target is divided into 30° segments with concentric circles to highlight given depths. Colored segments drawn over the circle convey the radial locations of scene content, allowing the artist to identify crowded and empty spaces, as well as distances between sketch elements in the virtual environment. The current view direction (seen in elevation in the top-left preview window) is represented by a gray wedge, similar to radar, and can be manipulated by rotating the wedge around, to update the preview window.

Flipping back and forth between panels in this view allows the artist (and director) to see how audience attention is being guided over the course of the narrative, understand the flow of action, as well as spot abrupt changes in depth and location, which can be disorienting for the audience.

Simple interactions allow the artist to quickly edit the spatial composition of the panel. She can adjust depth and location of strokes by lasso selecting with the pen (in either the overhead target or the preview window) and dragging the selected lines with a finger. We thus leverage pen and touch interaction for different activities, supporting artistic workflow with minimal interruption [15].

Panoramic Storyboard View

As in traditional storyboards, we provide an overview of the VRS, showing the sequence of panels as panoramic thumbnails, under which the artist can add notes (Figure 7). From this view, the artist can also quickly jump between panels for further edits in the sketching view.

To heed the VRS planning of audience attention, a concern unique to VR storyboards, we highlight the primary area of focus for each panel (green box and gradient in Figure 7). Panoramas are aligned with respect to the view angle (see angular scale under thumbnails), allowing the artist to understand and experiment with different focus transitions along subsequent panels, and maximize the use of the stage.

With its sketching, overhead, and panoramic storyboard views, the tablet provides the artist with a rich set of perspectives of the VRS, supporting expressive authoring and providing assets for discussion between the director and crew. Transitions between the views is achieved through a simple pinch gesture, without disrupting the creative flow.




Figure 7: Panoramic storyboard view showing the focus bounding box and focus gradient.

Exploring the virtual environment with the HMD

At any time during the process, the director, the artist, or any other collaborator can decide to experience how the VRS looks within the virtual environment, using the HMD. Design decisions for VR were largely based around a philosophy that the tablet draws, and the HMD views content.

HMD as an External Observer

An external observer (the director, or any other collaborator) can experience the panels as the audience would, and view sketched content wrapped around the cylinder in 3D, while the artist continues to work on the tablet (Figure 3-b).

To facilitate discussions around content, all changes that the artist makes on the tablet are immediately reflected in the VR environment. Meanwhile, the observer can look around at will without affecting the artist's environment, and may even decide to experience the scene from a slightly different viewing angle, through navigation with the keyboard. Similarly, the artist can freely work in any of the views (sketching, overhead, panoramic) without affecting the VR environment. Only when the artist changes the panel on the tablet, will the observer be teleported into the new scene.

HMD as a Display of the Tablet

A single user can immerse themselves in the VRS, while using the tablet, e.g. the artist can look at their creation in the HMD, and keep editing the content on the tablet (Figure 3-c, left); later, the director can look at the VRS immersively and make annotations on the tablet (Figure 3-c, d, right).

While wearing the VR headset, the surrounding physical environment is invisible to the immersed user. Thus, the tablet and keyboard are less accessible. We rely on “within arm’s reach” proprioception [22] and haptic cues for interactions. In practice, drawing on the tablet is not hindered, as the canvas and pen are represented by a frame widget and cursor in VR (Figure 9), which is similar to artists working with 2D digitizing tablets. In this setup, we reduce the tablet interface to pen interaction for sketching, and touch to control the view, negating the need to remove the HMD. Finger dragging up and down changes the current working depth on the sketch surface, while panning rotates the view, i.e. the VR environment is rotated to the position shown on the tablet; but the user can still look left and right. While large




Figure 8: Diorama storyboard view on the VR system.

rotations performed via the tablet can cause vestibular discomfort [1], it is preferable to using the HMD gaze to control the tablet view. As the tablet is used for sketching, it should be insensitive to ambient HMD head movements. Still, free navigation using the keyboard remains available, and hotkeys allow easy access to tools and the undo action.

Diorama Storyboard View

Our interviews indicated that professionals were unsure what a storyboard for VR would look like. To initiate the discussion on this topic, we introduced a diorama storyboard (Figure 8) that acts as the VR counterpart of the panoramic storyboard view on the tablet. Panels are represented as a series of 3D *miniatures* within cylindrical discs with white backgrounds that contrast the strokes. Viewed from a 3rd person perspective, these dioramas surround the user, rather than being organized traditionally in a flat grid. The panels can be viewed from different angles and proximities with keyboard navigation. Like the panoramic board, the user can jump between panels by looking at the desired diorama and with a key switch to sketching for that panel.

Support for Collaboration

Collaboration is an important design requirement of our dual system, as artists and directors often work together on storyboards, and may discuss them with other film crew. To enable the workflows illustrated in Figure 3, we synchronize the VRS panels via a wireless LAN using UDP messages, supporting consistency between devices, increasing awareness between users, and facilitating discussions.

Consistency Across Devices

To establish common landmarks for all users, we ensure consistency between the flat tablet's view and that of the HMD. First, the tablet and VR system display the same FOV (~90°), which occupies only half the screen given the same aspect ratio (freeing the other half for other purposes). Further, to help the director and artist mutually understand the VRS space, both interfaces contain a ground plane represented by a radial grid, and a degree scale wrapped around the environment for orientation (Figure 4 and 9).

Though both devices show grid lines with the same perceptual spacing, they are not exactly the same. The HMD shows the 3D scene in accurate linear perspective. The tablet on the other hand presents the sketch cylinders unrolled into rectangles, ideal for sketching. Therefore, the ground

plane projects as straight lines rather than radial lines on the tablet. We forego this accurate depiction for one that provides a better perception and interactive control. We draw the ground plane grid to perceptually match the grid in VR: with the concentric circular arcs as horizontal lines (that show the depth of the sketch cylinder), and the radial lines near vertical, but converging inward. Note that, two sliver regions outside the outermost radial lines are outside the VR field-of-view. We further show a dark horizontal line that moves along the ground plane to help the artist distinguish strokes above and below the ground plane (Figure 4).

Collaboration Feedback and Additional Features

Users in disparate visual spaces introduce a communication challenge, which we address using overlays indicating the current view, or gaze, of the other user on both devices.

While sketching on the tablet, the artist may notice a red horizontal bar along the bottom of the screen that moves to track the director's gaze in VR (Figure 4). With this mechanism, the artist can quickly pan the sketch surface to match that of the director during discussions. A similar VR field-of-view indicator appears in the overhead view target on the tablet, as a pie wedge. Similarly, at the depth of the current sketch cylinder, an outline of the current view area of the tablet is shown in 3D on the HMD. This render overlay not only helps by showing what the tablet sees, but also defines the current 3D canvas for sketching while immersed.

Alternately, the tablet's view can be temporarily coupled to that of the HMD, sparing the artist the need to actively track where the director is looking, when the latter momentarily leads the discussion. The reverse is not advised, as an external controlling source may cause nausea on the HMD.

Finally, we support director's notes as a dedicated canvas (Figure 4) that we display in purple, distinct from content on the tablet and in VR (which can be toggled as needed).

Ultimately, we support the production of VRSs by artists following traditional storyboard language, and enables VR explorations which were not possible before. Integrated in a cohesive dual system, these enable new collaborative workflows between the artist, the director and overall film crew.

USER FEEDBACK

We conducted interviews with professionals from the VR cinematic industry to determine whether our dual system approach was worth pursuing. Our interviews focused on high-level feedback on the dual system to assess the workflow and its ability to support the creative process rather than focusing on particular UI components.

We solicited feedback from a producer and an executive creative director from an Emmy winning VR team (*p1*, *p2*), a director/producer from a leading VR company (*p3*), and a 3D Lumière award winning director of stereoscopic and VR films (*p4*). We invited these professionals to our lab to experiment with our system following a ten-minute demonstration of its functionality. They experimented with each




Figure 9. Examples of VRS panels created with our system, as seen in the HMD.
Note, in the leftmost example, the gray frame widget that conveys the current view of the tablet.

context freely without any time constraint (i.e. using the tablet and HMD simultaneously, just the tablet then looking at the results in VR, etc.). We concluded the sessions with a semi-structured interview, where we elicited feedback on the system as a whole, as well as exploring specific points.

General Feedback

The general response we received regarding the system was all positive, eliciting comments like “*I loved it... hugely valuable*” (p1) and “*from a prep side, it’s a huge step forward*” (p4). When asked about these responses, there was a consensus that the coupling of the devices and the ability to see the drawings simultaneously on the tablet and in VR was the biggest contributing factor to this feeling, “*the simultaneous view in the [HMD] is amazing...a hit hands down*” (p2).

Dual System Approach

We were interested in user response to the need for a tablet and a HMD. We found that the users were happy with this setup and viewed the flexibility this introduced as a benefit. One participant noted that the two-system approach allows artists to “*intuitively block things in*” and that when they put on the headset, “*it’s exciting to actually see [the scene]*” (p2).

Participants also acknowledged that very few people can remain in a VR environment for a long period of time without feeling slightly ill. For a storyboard artist, to be immersed all day is impractical. It was felt that our system was a “*good way to preview without using the [HMD]*” (p1), and that it was actually positive that “*[storyboard artists] don’t have to be in [the VR environment] the whole time*” (p2).

We inquired about possible problems caused by the disparity in the way viewers perceive the scene (i.e. panoramic panel on the tablet vs 3D interactive space in VR). None of the interviewees were concerned about this. It was felt that “*people will start to learn the relationship between what it looks like in panorama and what it looks like in a [HMD]. Having an experienced artist using your tool, you can get pretty far just in panorama view*” (p1). Further, it was stated that “*once [the storyboard artists] understands the concept they can just draw, [they] don’t need to be in there.*” (p4)

Authoring

One goal with our system was to allow for quick storyboard creation despite the complexity of VR. Participants were able to quickly generate simple content during the interviews. Figure 9 shows some examples of their creations.

We inquired about what users liked about our system workflow and what could be improved. General comments were positive. Abilities such as “*being able to toggle back and forth [between the panels] is awesome*” (p1). It was felt, however, that there were a number of features we could add to improve on the existing system. The thumbnail slider bar was found to be a nuisance, and that it “*got in the way*” (p2). The ground plane grid was found useful for a sense of spatial dimension, and similar scale feedback in the vertical dimension was suggested, as was a feature to toggle to-scale silhouettes of an “*apple, person and dinosaur*” (p2).

General comments about the UI included suggestions that combine interactive control of the overhead and sketch surface view. It was also noted that with complex scenes, it is difficult to isolate items in the overhead view as they often overlapped. Therefore, a “*function to explode out the lines that are overlapping*” (p1) would be a useful addition.

VR Storyboards

When starting this project, we learned that, currently there are no such things as VR storyboards. Consequently, we created both the panoramic storyboard view and the diorama storyboard view. Users seemed to prefer the panorama storyboard view: “*[I] can go scene one, scene two, scene three, and just see the subtle differences*” (p2) and “*you can see your primary point of focus*” (p1).

Several of the professionals noted that though they did not have much use for the dioramas, they could see them being useful to other departments. “*I could see [dioramas] being more useful in the future... if you end up moving cameras*” (p2); “*I think it would be useful for other departments*” (p4).

Collaboration

The potential for a collaborative workflow with our system was recognized by users as a significant contribution. It was felt our system “*allowed for creative collaboration where [the director and artist] were able to be in the same space and bounce things off of one another*” (P4). One user stated that showing the location of the other user was “*helpful ‘big time’...when you’re in the same space, it’s pretty intuitive*” (p1). However, it was also noted by this user that “*if you were over the phone... [it would be useful] if there were some way for a person with the headset on to stare at something*” (p1) to explicitly highlight object(s) of discussion.

Final thoughts

The users who tried our system were very excited about its potential. Being able to try shots without having to set-up cameras for test footage, and the ability to quickly and collaboratively create and discuss ideas, were all factors that contributed to their enthusiasm. One can often measure the success of an implementation by a user's willingness to use the system, and when we asked interviewees for final thoughts, we were told that, *"if you had a [consumer] version ready tomorrow we'd use it" (p1)*.

RELATED WORK

Telling stories in a virtual environment is not a new concept. In early work, Randy Pausch et al. [24] identified the need for a VR storytelling lexicon (that is currently evolving), and the need for rapid prototyping in VR. Our talks with current filmmakers suggest that this is still very much an open problem. Below we focus on existing tools that enable 3D sketching, and drawing in VR.

Sketching 3D models

Since Sutherland's seminal sketchpad system [28] more than half a century ago, digital sketching of 3D objects has been an active area of research (see [23] for a survey). A large body of research deals with interpreting a collection of 2D strokes as a network of 3D curves. Tools like Analytic 3D Drawing [26], ILoveSketch [3], and True2Form [33], define geometric relationships between curves in 3D based on inferences of 2D strokes, using regularity constraints. These curves are often treated as cleanly drawn 2D strokes projected onto 3D geometry (like our sketch cylinders). Other research, illustrated by the organic shape modeler Teddy [17], builds 3D objects by extruding or inflating sketch strokes.

As an ideation-stage sketching tool, 3D objects are both overkill and distract from the primary goal of story layout. Gesture based drawing systems provide rules to build 3D geometry [8, 11, 34] or suggest plausible alternatives [16].

Our work is better served by systems that mimic the fluidity of pen-and-paper [9, 31], where no assumptions are made on the nature or semantics of sketched 2D strokes, save for projecting them onto concentric sketch cylinders that reflect view depth. Our abstraction of view depth is akin to 2.5D modeling [25], or techniques that conceptually subdivide space into volumetric layers [7].

Drawing in VR

Direct drawing in VR has been explored for over two decades [6], and has even been enhanced to provide techniques to paint in VR [18]. Numerous systems use a 2D plane as a mode of interaction for constructing in the 3D environment [2, 10] whether it is a virtual 2D plane, or a physical drawing tablet represented in the 3D space [5]. LaViola et al. [19] provide a summary of different techniques for drawing in VR.

When worrying about speed and accuracy, research has shown that sketching 3D objects in a virtual space is less accurate and more time consuming than traditional sketch-

ing [32]. Given the large volume of sketches needed for storyboarding, we thus choose tablet-based sketching over direct drawing in VR.

LIMITATIONS AND FUTURE WORK

While the response to our prototype was very positive, there were many insightful suggestions that would improve future versions of the system, such as the following:

Our system does not address the concept of stitch lines, which vary by camera-rig. Directors often film test footage to check blocking of elements. Showing stitch lines on storyboards would greatly benefit planning and layout in VR.

Our system represents object using sketch strokes. Though effective and lightweight, this approach lacks explicit mass and volume, which can be useful for planning. Combining our sketches with previs quality 3D assets is interesting future work, along with a richer set of input options while in VR. Different input techniques: bi-manual input, sketch language, augmented real 3D sketching, and visualization all provide other promising areas for future exploration.

Also, though this paper has focused on VR stories, we believe our dual system approach, and its planning features for 3D dynamic scenes would be useful in many other contexts such as gaming, education, architecture, and AR.

Finally, our prototype was designed with two collaborators in mind. While many of our concepts scale from two people to a full team, it is yet to be deployed or tested in such a setting. In fact, a deployment study to observe the tool actively being used is a good next step to determine how to improve the workflow, and where to focus next.

CONCLUSION

With the consumerization of VR headsets, there is an explosion of interest in creating content for this medium. While filmmakers wish to tell stories in VR, tools to support this endeavor are next to none. We thus contribute the first system targeted at storyboarding in VR.

We performed a formative study to assess the challenges and goals of VR storytellers. We studied the tools available to them today, and the limitations. We then proposed a workflow that pairs a sketch tablet and VR display, and allows users to sketch on stacked panoramic panels, corresponding to concentric cylindrical surfaces of increasing view depth. This setup leverages the domain expertise of storyboard artists and enables interactions in VR between the director and artist. We implemented this approach and elicited early feedback from professional VR storytellers on our prototype. This system was received very positively, and responses indicated that our approach provided the speed and functionality needed for early stage planning, and the artifacts to properly discuss VR films.

Acknowledgments – We are grateful to the professionals who participated in our interviews and evaluation for their time and valuable feedback. We also thank our colleagues at the DGP lab for discussions throughout the project.

REFERENCES

1. Akiduki, H., Nishiike, S., Watanabe, H., Matsuoka, K., Kubo, T. and Takeda, N., 2003. Visual-vestibular conflict induced by virtual reality in humans. *Neuroscience letters*, 340(3), 197-200.
2. Angus, I. & Sowizral, H. 1995. Embedding the 2D Interaction Metaphor in a Real 3D Virtual Environment. In *Proc. of SPIE – The International Society for Optical Engineering*, 282-293.
3. Bae, S.Y., Balakrishnan, R. & Singh, K. 2008. ILoveSketch: as-natural-as-possible sketching system for creating 3d curve models. In *Proc. UIST'08*, 151-160.
4. Bishop, B. 2015. Inside Industrial Light & Magic's secret Star Wars VR lab. Retrieved: January 2016 from <http://www.theverge.com/2015/8/13/9131805/ilm-ilmxlab-interview-virtual-reality-star-wars-movies>
5. Bowman, D., Wineman, J., Hodges, L. & Allison, D. 1998. Designing Animal Habitats Within an Immersive VE. *Computer Graphics & Applications*, 18(5):9-13.
6. Deering, M. 1995. HoloSketch: a virtual reality sketching/animation tool. *ACM Transactions on Computer-Human Interaction*, 2(3):220-238.
7. De Paoli, C. & Singh, K. 2015. SecondSkin: sketch-based construction of layered 3D models. *ACM Trans. Graph.* 34, 4, Article 126 (July 2015), 10 pages.
8. Ding, Z., Zhang, S., Peng, W., Ye, X. & Hu, H. 2005. PENCIL: a framework for expressing free-hand sketching in 3D. In *Proc. International conference on Advances in Natural Computation*, Vol. Part I, 835-838.
9. Dorta, T., Kinayoglu, G. & Hoffmann, M. 2014. Hyve-3D: a new embodied interface for immersive collaborative 3D sketching. In *ACM SIGGRAPH 2014 Studio*, 37.
10. Ens, B., Hincapié-Ramos, J. & Irani, P. 2014. Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In *Proc SUI '14*, 2-12.
11. Forsberg, A., LaViola, J., Markosian, L. & Zeleznik, R. 1997. Seamless Interaction in Virtual Reality. *Computer Graphics & Applications*, 17(6):6-9.
12. Frost, J. 2016. 360-degree VR Video invades Disney's theme parks. Retrieved April 2016 from <http://thedisneyblog.com/2016/01/21/360-vr-video-invades-disneys-theme-parks/>
13. Hart, J. 2007. *The Art of the Storyboard: A Filmmaker's Introduction* (2nd Ed.). Focal Press.
14. Henrikson, R., Araujo, B., Chevalier, F., Singh, K. & Balakrishnan, R. 2016. Storeoboard: Sketching Stereoscopic Storyboards. In *Proc. CHI'16*, 4587-4598.
15. Hinckley, K., Yatani, K., Pahud, M., Coddington, N., Rodenhouse, J., Wilson, A., Benko, H., and Buxton, B. 2010. Pen + touch = new tools. In *Proc UIST '10*, 27-36.
16. Igarashi, T. & Hughes, J. 2001. A suggestive interface for 3D drawing. In *Proc. UIST '01*, 173-181.
17. Igarashi, T., Matsuoka, S. & Tanaka, H. 1999. Teddy: a sketching interface for 3D freeform design. In *Proc. SIGGRAPH '99*, 409-416.
18. Keefe, D., Acevedo Feliz, D., Moscovich, T. Laidlaw, D. & LaViola, J. 2001. CavePainting: a fully immersive 3D artistic medium and interactive experience. In *Proc. Interactive 3D graphics (I3D '01)*, 85-93.
19. LaViola, J. & Keefe, D. 2011. 3D spatial interaction: applications for art, design, and science. In *ACM SIGGRAPH 2011 Courses*. Article 1, 75 pages.
20. Lipton, L. 1982. *Foundations of the Stereoscopic Cinema: A Study in Depth*. Van Nostrand Reinhold Inc.
21. McNally, V. 2014. Warner Brothers Announces Plans for a Virtual Reality Batman: The Animated Series Project. Retrieved April 2016 from <http://www.themarysue.com/virtual-reality-batman/>
22. Mine, M., Brooks, F., & Sequin, C. 1997. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proc. SIGGRAPH '97*, 19-26.
23. Olsen, L., Samavati, F., Costa Sousa, M. & Jorge, J. 2009. Sketch-based modeling: A survey. *Computers & Graphics*, 33(1), 85-103.
24. Pausch, R., Snoddy, J., Taylor, R., Watson, S. & Haseltine, E. 1996. Disney's Aladdin: first steps toward storytelling in virtual reality. In *Proc. SIGGRAPH '96*, 193-203.
25. Rivers, A., Igarashi, T. & Durand, F. 2010. 2.5D cartoon models. In *ACM SIGGRAPH 2010 papers (SIGGRAPH '10)*, Article 59, 7 pages.
26. Schmidt, R., Khan, A., Singh, K., & Kurtenbach, G. (2009). Analytic drawing of 3D scaffolds. In *ACM Trans. Graph.* 28(5) Article 149.
27. Slater, M., Usoh, M. & Steed, A. 1994. Depth of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 3(2):130-144.
28. Sutherland, I. E. 1963. Sketchpad: a man-machine graphical communication system. In *Proc. AFIPS '63*, 329-346.
29. Sutherland, I. E. 1968. A head-mounted three dimensional display. In *Proc. AFIPS '68*, 757-764.
30. Trenholm, R. 2016 Samsung to open VR movie studio in New York: 'We love stories'. Retrieved: April 2016 from <http://www.cnet.com/news/samsung-to-open-vr-movie-studio-in-new-york/>
31. Tsandilas, T., Grammatikou, M. & Huot, S. 2015. BriccoSketch: Mixing Paper and Computer Drawing Tools in Professional Illustration. In *Proc. ITS'15*, 127-136.
32. Wiese, E., Israel, J.H., Meyer, A. & Bongartz, S. 2010. Investigating the learnability of immersive free-hand sketching. In *Proc. Sketch-Based Interfaces and Modeling Symposium (SBIM '10)*, 135-142.
33. Xu, B., Chang, W., Sheffer, A., Bousseau, A., McCrae, J. & Singh, K. 2014. True2Form: 3D curve networks from 2D sketches via selective regularization. *ACM Trans. Graph.* 33, 4, Article 131 (July 2014), 13 pages.
34. Zeleznik, R., Herndon, K. & Hughes, J. 1996. SKETCH: an interface for sketching 3D scenes. In *Proc. SIGGRAPH '96*, 163-170.
35. Zwerman, S. & Okun, J. A. 2014. *The VES Handbook of Visual Effects: Industry Standard VFX Practices and Procedures*, Focal Press.
36. Google Cardboard. <https://www.google.com/get/cardboard/>
37. HTC Vive. <https://www.htcvive.com/ca/>
38. Oculus. <https://www.oculus.com/en-us/>
39. Oculus Story Studio. <https://storystudio.oculus.com/en-us/>
40. TiltBrush. <http://www.tiltbrush.com/>
41. Visionary VR. <http://visionaryvr.com/>