3D video fragments

dynamic point samples for real-time free-viewpoint video

Report**Author(s):**

Würmlin, Stephan; Lamboray, Edouard; Gross, Markus H.

Publication date:

2003

Permanent link:

<https://doi.org/10.3929/ethz-a-006665881>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

CS technical report 397

3D Video Fragments: Dynamic Point Samples for Real-time Free-Viewpoint Video

Stephan Würmlin, Edouard Lamboray, Markus H. Gross

Computer Science Department
ETH Zurich, Switzerland
e-mail: {wuermli, lamboray, grossm}@inf.ethz.ch
<http://graphics.ethz.ch/>

3D Video Fragments: Dynamic Point Samples for Real-time Free-Viewpoint Video

Stephan Würmlin

Edouard Lamboray

Markus Gross

ETH Zürich
{wuermlin, lamboray, grossm}@inf.ethz.ch

Abstract

We present *3D video fragments*, a dynamic point sample framework for real-time 3D video. By generalizing 2D video pixels towards 3D irregular point samples we combine the simplicity of conventional 2D video processing with the power of more complex polygonal representations for 3D video. We propose a differential update scheme exploiting the spatio-temporal coherence of the video streams of multiple cameras. Updates are issued by operators such as inserts and deletes accounting for changes in the input video images. The operators from multiple cameras are processed, merged into a 3D video stream and transmitted to a remote site. We also introduce a novel concept for camera control which dynamically selects the set of relevant cameras for reconstruction. Moreover, it adapts to the processing load and rendering platform. Our framework is generic in the sense that it works with any real-time 3D reconstruction method which extracts depth from images. The video renderer displays 3D videos using an efficient point based splatting scheme and makes use of state-of-the-art vertex and pixel processing hardware for real-time visual processing.

Keywords: 3D video, distributed graphics systems, inter-frame coding, point based graphics, real-time graphics.

1 Introduction

Over the years, telepresence has become increasingly important in many applications including computer supported collaborative work (CSCW) and entertainment. While we dispose of commercial solutions for 2D teleconferencing in combination with CSCW, it is only in recent years that 3D video processing has been considered as a means to enhance the degree of immersion and visual realism of telepresence technology. The most comprehensive program dealing with 3D telepresence is the National Tele-Immersion Initiative (<http://www.advanced.org/teleimmersion.html>). Such 3D video processing poses a major technical challenge and is thus gaining interest in the computer graphics and computer vision communities. Here, a lot of research has been dedicated in particular to the extraction and reconstruction of real objects. The *representation* of 3D video streams, however, a fundamental prerequisite for efficient processing, has less intensively been investigated. In fact, most representations for 3D video streams are tailored for off-line postprocessing and, hence, share various limitations that makes them less practicable for advanced real-time 3D video processing.

Our work is devoted to the efficient representation, control and encoding of 3D video streams, facilitating sophisticated 3D rendering and visual effects. By introducing the concept of 3D video fragments, we generalize 2D video pixels towards irregular spatio-temporal point samples. Conceptually, each video fragment is a point sample with a set of attributes, like position, normal and color, which can be dynamically updated. A point-based object representation on the remote site stores all active fragments and can be accessed for efficient rendering. Our differential update stream inserts, deletes or updates fragments on-the-fly in real-time. It exploits the spatio-temporal coherence of individual 2D video streams by inter-frame prediction of input changes in image space. Our prediction does not require expensive calculations like texture

Figure 1: Examples for 3D video fragment objects. a) Point rendered view. b) Deferred operation: sine wave.

motion fields or 3D scene flows. While being conceptually lean and simple the presented approach effectively cuts down the number of expensive 3D shape computations. As opposed to mesh based representations, 3D video fragments provide a one-to-one mapping between points and associated color and normal attributes avoiding interpolation and alignment artifacts. In particular the lack of local connectivity makes 3D video fragments much more efficient for updating, coarse-to-fine sampling, progressive streaming, and compression.

Another benefit of retaining an underlying point based representation is graphics rendering. Since update operators of the 3D video stream dynamically change the representation, we have to carry out all necessary computations for rendering on-the-fly. Our 3D video renderer supports a wide range of visual effects, like explosions and warps, altering position and color attributes of individual fragments. To preserve data consistency we defer such operations to the final rendering stage and employ programmable hardware. By using a feedback loop which confines the number of active cameras we dynamically control the acquisition process and scale smoothly from view-dependence to view-independence. Moreover, virtual viewpoint- and resolution-driven sampling allows smooth transitions between a subset of the reference cameras and adapts to bandwidth or processing bottlenecks. The method features efficient rendering from arbitrary spatio-temporal positions and supports multiple viewers.

Our 3D video pipeline is designed and optimized for real-time applications. It performs fully automatically and does not need human intervention.

1.1 Related Work

Given the design philosophy of our 3D video fragments pipeline the work that is most related to ours includes real-time 3D acquisition, the MPEG-4 standard, and point sample rendering.

Concepts for 3D Video Acquisition. There is a variety of methods for reconstruction of 3D video sequences. We distinguish between methods requiring off-line *postprocessing* and *real-time* methods. Examples of postprocessing algorithms include the work of Moezzi et al. [1996], who propose a batch-oriented computation

of 3D video sequences. Voxel representations are frequently derived by volume carving methods [Szeliski 1993]. While these methods can provide for point sampled representations they are not performing in real-time. An appealing approach for utilizing spatio-temporal coherence for 3D video is the work of Vedula et al. [2002] which computes a 3D scene flow for spatio-temporal view interpolation. It produces impressing results but the lack of real-time performance makes it impractical for our purposes. Carceroni and Kutulakos [2001] present a dynamic surfel sampling representation and algorithm for estimation of 3D motion and dynamic appearance. However, they use a volumetric reconstruction for a small working volume and do not demonstrate real-time performance either. Würmlin et al. [2002] present a 3D video recorder which stores a spatio-temporal representation in which users can freely navigate.

As opposed to post-processing approaches real-time methods are much more demanding with regard to computational efficiency. Matusik et al. [2000] present an image-based 3D acquisition system which calculates the visual hull [Laurentini 1994] of an object. It is build on epipolar geometry and outputs a view-dependent LDI representation. Their system neither exploits spatio-temporal coherence, nor is it scalable in the number of cameras. Similarly, polyhedral visual hulls [Matusik et al. 2001] are based on epipolar geometry and provide view-independent rendering through a mesh and texture representation. It shares the same limitations and furthermore introduces interpolation artifacts due to improper alignment of geometry and texture, a common drawback of mesh based methods. Kanade et al. [1997] and Narayanan et al. [1998] employ a triangular texture-mapped mesh representation. A similar approach was presented by Mulligan and Daniilidis [2000] utilizing trinocular stereo depth maps from overlapping triples of cameras. It also features the aforementioned limitations of mesh based techniques. Pollard and Hayes [1998] utilize depth map representations for novel view synthesis by morphing live video streams. This representation can suffer from inconsistencies between different views.

Special-purpose hardware solutions for real-time depth estimation from video images, such as 3DV Systems' *ZCam*TM (<http://www.3dvsystems.com>), and Tyxz's *DeapSea* chips (<http://www.tyxz.com>) have recently become available. They can be seen as complementary to our work since they solve the problem of real-time 3D reconstruction and can be incorporated into our framework.

3D Video Standards. Even though the MPEG-4 committee is actively deliberating future 3D video standards, no standard for dynamic, free view-point 3D video objects has yet been defined [ISO/IEC 2002]. The MPEG-4 multiple auxiliary components can encode depth maps and disparity information. But these are not complete 3D representations and possible shortcomings and artifacts due to DCT encoding and unrelated texture motion fields and depth or disparity motion fields still need to be investigated. Our concept for real-time 3D video fits well into the system architecture proposed by the MPEG committee including acquisition, representation and display stages with back-channel transmission of viewpoint selection. Additionally, we support all types of interactivity, i.e., interaction at the encoder side and interaction with or without all data available at the decoder side.

Point Sample Rendering. In recent years points have experienced a renaissance as a graphics primitive. While there are various methods for fast and high quality rendering of point sampled geometry at our disposal, to date, none of them can efficiently cope with dynamically changing objects or scenes. For instance, the surfel system [Pfister et al. 2000] samples an object with three orthogonal LDI's, building a so-called LDC-tree. It is well suited for

progressive rendering but has to be rebuilt once the object changes. Surface splatting [Zwicker et al. 2001] extends the surfel system with a high-quality interactive software renderer which is based on a screen space formulation of the elliptical weighted average (EWA) filter adapted for irregular point samples. While Ren et al. [2002] present a hardware-accelerated extension based on multi-pass rendering they all share the aforementioned limitations imposed by pre-processing and setup. Qsplat [Rusinkiewicz and Levoy 2000] is a progressive point sample system for representation and display of very large geometry. They represent static objects by a multi-resolution hierarchy of point samples based on bounding spheres. As with the surfel system they rely on extensive pre-processing for splat size and shape estimation making it impracticable for our needs. As we will discuss, our dynamic 3D video engine performs all computations for high quality point sample rendering on-the-fly in real-time.

1.2 Conceptual Overview

Figure 1 depicts a conceptual overview of the 3D video fragments pipeline. We acquire images from multiple calibrated video cameras. The images are processed to segment foreground from background. By means of dynamic camera control (Section 3) we determine a set of active cameras from which we generate 3D point samples as well as a set of supporting cameras delivering additional data to improve the 3D reconstruction. Using inter-frame prediction in image space we generate a stream of differential operators (Section 2) which dynamically update point sample attributes including position or color. We thus avoid to recompute the full 3D representation in each frame. The dynamic point samples are rendered by an efficient point splatting scheme (Section 4) and are composited with a virtual scene. In a final stage we apply deferred operations like visual effects by running them directly on the graphics hardware.

Figure 1: Conceptual components of the 3D video fragments pipeline

2 Differential 3D Fragment Operators

Our concept of video fragments exploits the spatio-temporal inter-frame coherence of multiple input streams by using a differential update scheme for dynamic point samples. The basic primitives of this scheme are *3D fragment operators*, each of which is derived from a 2D pixel operator as illustrated in Figure 2.

We distinguish between three different types of operators:

- **INSERT** adds new 3D video fragments into the representation after they have become visible in one of the input cameras. Insert operators are streamed coarse-to-fine as discussed in Section 3.
- **DELETE** removes fragments from the representation once they vanish from the view of the input camera.

Figure 2: Relationship between 2D pixel operators and 3D fragment operators.

- UPDATE corrects appearance and geometry attributes of fragments that are already part of the representation, but whose attributes have changed with respect to prior frames.

The time sequence of these operators creates a differential fragment operator stream that updates a 3D video data structure on a remote site. An INSERT operator results from the reprojection of a pixel with color attributes from image space back into three-dimensional object space. Any real time 3D reconstruction method which extracts depth and normals from images can be employed for this purpose. Note that the point primitives feature a one-to-one mapping between depth and color/texture samples. The depth value is stored in a *depth cache*. This structure accelerates the DELETE operator which performs a lookup in the depth cache and can thus be carried out very efficiently.

UPDATE operators are generated by all pixels which have been inserted into previous frames and which are still foreground pixels. They can be divided into three categories: The detection of *color changes* is performed during inter-frame prediction and leads to an UPDATECOL operator. UPDATEPOS operators take care of *geometry changes* and are analyzed on spatially coherent clusters of pixels in image space. We also use the depth cache for this purpose. We define independent blocks of points according to a pre-defined grid. Typically, for a 640×480 resolution, a block comprises 16×16 pixels and for each frame new depth values are calculated for the four grid corners only. If the differences to the previous depths exceed a threshold, we recompute 3D information for the entire block of points. Thus, our scheme proposes an efficient solution to the problem of un-correlated texture and depth motion fields. Note that position and color updates can be combined to an UPDATEPOSCOL operator. All other candidate pixels for updates remain *unchanged* and no further processing is necessary. The 3D operators and associated data can be summarized as follows:

$$\left\{ \begin{array}{ll} \text{INSERT:} & (\bar{\mathbf{p}}, \mathbf{c}) \rightarrow (\dot{\mathbf{P}}, \mathbf{c}, \mathbf{n}) \\ \text{DELETE:} & (\bar{\mathbf{p}}) \rightarrow (\dot{\mathbf{P}}) \\ \text{UPDATECOL:} & (\bar{\mathbf{p}}, \mathbf{c}) \rightarrow (\dot{\mathbf{P}}, \mathbf{c}) \\ \text{UPDATEPOS:} & (\bar{\mathbf{p}}) \rightarrow (\dot{\mathbf{P}}_{old}, \dot{\mathbf{P}}_{new}, \mathbf{n}_{new}) \\ \text{UPDATEPOSCOL:} & (\bar{\mathbf{p}}, \mathbf{c}) \rightarrow (\dot{\mathbf{P}}_{old}, \dot{\mathbf{P}}_{new}, \mathbf{n}_{new}, \mathbf{c}) \end{array} \right. \quad (1)$$

where $\bar{\mathbf{p}}$ are the coordinates of a pixel, $\mathbf{c}$ its color, $\dot{\mathbf{P}}$ the respective 3D coordinates of the point sample, and $\mathbf{n}$ its surface normal. The encoding of the 3D operators will be explained in Section 5.3.

We propose an image space inter-frame prediction mechanism which derives the 3D fragment operators from the original video images. We define two functions for pixel classification: A foreground-background segmentation defines a Boolean function $fg(\bar{\mathbf{p}}, t)$ returning TRUE if the pixel $\bar{\mathbf{p}}$ is in the foreground at frame t . A second function $cd(\bar{\mathbf{p}}, t, t')$ returns TRUE if the color difference of a pixel $\bar{\mathbf{p}}$ exceeds a certain threshold in between the time instants t and t' . These two functions allow us to assign to each pixel one of the following five classes using simple Boolean operations,

$$\left\{ \begin{array}{ll} fg(\bar{\mathbf{p}}, t) \wedge fg(\bar{\mathbf{p}}, t-1) \wedge \neg cd(\bar{\mathbf{p}}, t, t-1): & \text{color unchanged} \\ fg(\bar{\mathbf{p}}, t) \wedge fg(\bar{\mathbf{p}}, t-1) \wedge cd(\bar{\mathbf{p}}, t, t-1): & \text{color changed} \\ fg(\bar{\mathbf{p}}, t) \wedge \neg fg(\bar{\mathbf{p}}, t-1): & \text{new} \\ \neg fg(\bar{\mathbf{p}}, t) \wedge fg(\bar{\mathbf{p}}, t-1): & \text{expired} \\ \neg fg(\bar{\mathbf{p}}, t) \wedge \neg fg(\bar{\mathbf{p}}, t-1): & \text{background} \end{array} \right. \quad (2)$$

where $\wedge$ denotes the Boolean AND and $\neg$ the Boolean NOT operator. Figure 3 illustrates the image acquisition processing and depicts the five possible pixel states.

Figure 3: 2D pixel operators. a) Silhouette at time $t-1$. b) Silhouette at time t . c) Pixel classification. Green indicates *new pixels*, red *expired pixels*, blue *color changed pixels*, white *color unchanged pixels*, and black background.

Finally, a *new* pixel invokes an INSERT operator, an *expired* pixel a DELETE operator and a *color change* an UPDATECOL operator. As previously described, *unchanged* and *color changed* pixels can nonetheless lead to an UPDATEPOS operator.

3 Dynamic System Adaptation

Many real-time 3D video systems are employed for point-to-point communication. In such cases, the 3D video representation can be optimized for a single view point. Multi-point connections, however, require truly view-independent 3D video. In addition, 3D video systems can suffer from performance bottlenecks at all pipeline stages. Some performance issues can be locally solved, for instance by lowering the input resolution, or by utilizing hierarchical rendering. However, only the combined consideration of application, network and 3D video processing state leads to an effective handling of critical bandwidth and 3D processing bottlenecks. In the point-to-point setting the current virtual viewpoint allows one to optimize the 3D video computations by confining the set of relevant cameras. As a matter of fact, reducing the number of involved cameras or the resolution of the reconstructed 3D video object implicitly decreases the required networking bandwidth. Furthermore, the acquisition frame rate can be adapted dynamically.

The aforementioned issues suggest a concept for dynamic system adaptation for the 3D video system, which will be described in this section.

3.1 Active Camera Control

We devise a concept for dynamic control of *active* cameras which allows for smooth transitions between subsets of reference cameras and efficiently reduces the number of involved cameras for 3D reconstruction. Furthermore, increasing the number of so-called *texture active* cameras enables a smooth transition from a view-dependent to a view-independent representation for 3D video.

A texture active camera is a reference camera applying the intra-frame prediction scheme as explained in Section 2. Each pixel classified as foreground in such a camera frame contributes color or texture samples to the set of 3D points in the 3D representation. Additionally, each camera might provide auxiliary information for the employed 3D reconstruction algorithm. We call the state of these cameras *reconstruction active*. Note that a camera can be both texture and reconstruction active. The state of a camera which does not provide data at all is called *silent*. Figure 6 illustrates the dynamic control of active cameras.

Figure 4: Illustration of the dynamic camera control. Green cameras are texture active, red cameras are reconstruction active, and yellow cameras are both texture and reconstruction active. Uncolored cameras are silent. a) for one active camera. b) for three active cameras

In order to select the k -closest cameras for the desired viewpoint as texture active cameras, we compare the angles between all camera look-at vectors and the desired viewing vector. Choosing the k -closest views minimizes artifacts arising from occlusions in the reference views. Experimentally, we found that for our target objects, i.e. humans, $k = 3$ performs well. The selection of reconstruction active cameras has to be computed for all texture active cameras and is dependent on the employed 3D reconstruction method. Since our prototype system uses a shape-from-silhouette algorithm, each reconstruction active camera provides silhouette contours. The set of candidate cameras is chosen by two simple rules. First, the angles between a texture active camera and its associated reconstruction active cameras have to be smaller than some threshold, 100 degrees in our setting. The candidate set is thus confined to cameras lying in approximately the same hemisphere. Second, the angle must not be smaller than 20 degrees. Although this is hardly the case in our setup, cameras which are too close to each other only provide in-significant differences in their silhouette information. Optionally, we reduce the candidate set to a maximum size. We compute the angle between all candidate camera pairs and subsequently discard one camera of the closest pair. In our case, this scheme leads to an optimally smooth coverage of silhouettes for every texture active camera. The set of texture active cameras needs to be computed on-the-fly accounting for viewpoint changes. The map of corresponding texture and reconstruction active cameras can be pre-computed at system start-up time.

Although we did not investigate 3D reconstruction techniques other than shape-from-silhouette, our active camera approach is versatile and can be employed with other algorithms. A multi-view stereo algorithm for example requires from each reconstruction active camera the texture of the whole frame in combination with some features. In this case, the k -nearest cameras would be selected such to enable correspondence calculations.

Overall, the dynamic camera control allows to actively trade-off 3D reconstruction performance versus 3D video quality. In our shape-from-silhouette setting, the quality of the 3D reconstruction is improved by a growing number of reconstruction active cameras, but so increases the processing time.

3.2 Texture Activity Levels

A second strategy for dynamic system adaptation involves the number of reconstructed fragments. We define a *texture activity level* A_i for each camera i to determine the number of pixels fed into the 3D video pipeline. Initial levels for k texture active cameras are derived from the weight formulas for Unstructured Lumigraph Rendering [Buehler et al. 2001, Vlasic et al. 2003]:

$$r_i = \frac{\cos \theta_i - \cos \theta_{k+1}}{1 - \cos \theta_i}, \quad w_i = \frac{r_i}{\sum_{j=1}^k r_j} \quad (3)$$

Here, r_i represent the relative weights of the closest k views. r_i is calculated from the cosine of the angles between the desired view and each texture active camera. w_i are the normalized weights which sum up to one. The texture activity level allows for smooth transitions between cameras and enforces epipole consistency. In addition, texture activity levels are scaled with a system load penalty $penalty_{load}$ regarding the reconstruction process. The penalty takes into account not only the current frame load but also the activity levels of prior frames. Finally, the resolution of the virtual view is taken into account with a factor ρ leading to the following equation:

$$A_i = s_{max} \cdot w_i \cdot \rho - penalty_{load}, \quad \text{with } \rho = \frac{res_{target}}{res_{camera}} \quad (4)$$

Note that this equation is recomputed in each frame for each texture active camera. The maximum number of sampling levels s_{max} discretizes A_i to a linear sampling pattern in the camera image, allowing for coarse-to-fine sampling. All negative values of A_i are clamped to zero. Figure 5a illustrates the linear pixel sampling pattern which is basically a multi-grid sampling and Figure 5b shows activity levels for a set of cameras for a given virtual viewpoint. Currently, we use NTSC-cameras which leads to $38^{\circ}400$ pixels per sampling level.

Figure 5: Texture activity levels. a) Linear image pixel sampling pattern. b) Activity levels for three texture active cameras as seen from the virtual viewpoint.

4 Dynamic Point Processing and Rendering

The final stage of our 3D video pipeline constitutes the processing and rendering of the 3D video fragments. All necessary computations for rendering must be performed on-the-fly and in real-time. In particular, the size and shape of the splat kernels for high quality rendering must be estimated dynamically for each point sample. For that purpose we propose a new data structure for 3D video rendering. We organize the point samples for processing on a per-camera basis similar to a depth image. However, instead of storing a depth value per pixel we store references to the respective point attributes. We name this representation *3D fragment image*. The

Figure 6: 3D fragment images. Fragments can be referenced from an image space representation for point processing, but rendered efficiently from a flat OpenGL vertex array.

point attributes themselves are organized in an OpenGL vertex array which can be directly transferred to graphics memory. With this representation we combine efficient insert, update and delete operations with efficient processing for rendering. Figure 6 illustrates the data structure.

In addition, the 3D video renderer supports compositing with a virtual scene by Z-buffering. Finally, we also support deferred operations, such as 3D visual effects, which are applicable to the real-time 3D video stream without destroying the consistency of the data structure.

4.1 Local Density Estimation

For static objects, local point sample densities can either be estimated in a pre-processing step [Rusinkiewicz and Levoy 2000] or during the acquisition procedure [Zwicker et al. 2001][2002]. In our approach however, the acquisition process leads to irregular 3D point sampling patterns. Hence, we cannot estimate the local point sample density during acquisition. Moreover, in a dynamic real-time system, it is not economical to maintain an advanced spatial search structure supporting point density estimations like nearest-neighbor searches [Pauly and Gross 2001]. Instead, we propose to estimate the local point sample density for each point based on incremental nearest-neighbor search in the 3D fragment cache. The resulting neighbors are only approximations of the real neighbors, but they prove to be sufficiently close for local sampling density estimation. Our algorithm, which considers only two neighbors, uses the following heuristics. First, it calculates the nearest-neighbor N_1 of a given point in the 3D fragment cache. Then we search for a second neighbor N_{60} , forming an angle of at least 60 degrees with N_1 . Our neighbor search needs approximately four more iterations for finding N_{60} .

4.2 Point Sample Rendering

As in Ren et al. [2002], we render the point samples as polygonal splats with a semi-transparent alpha texture using a two-pass algorithm. In the first pass, opaque polygons are rendered for each point sample and thus visibility splatting is performed [Pfister et al. 2000]. The second pass renders the splat polygons with an alpha texture. The splats are multiplied with the color of the point sample and accumulated in each pixel. A depth test with the Z-buffer from the first pass resolves visibility problems during rasterization. This ensures correct blending between the splats.

The neighbors N_1 and N_{60} can be used for computing polygon vertices of our splat in object space. The splat lies in the plane which is spanned by the coordinates of the point p and its normal n . We now distinguish between circular and elliptical splat shapes. In the first case, all side lengths of the polygon are twice the distance to the second neighbor N_{60} , which corresponds also to the diameter of the enclosing circle. For elliptical shapes, we determine the minor axis by projecting the first neighbor N_1 onto the tangential plane. The length of the minor axis is determined by the

Figure 7: Polygon setup for elliptical splat rendering. r_1 and r_{60} denote the distances from the point sample P to N_1 and N_{60} , respectively. c_1 to c_4 denote the calculated polygon vertices.

distance to N_1 . The major axis is computed as the cross product of the minor axis and the normal. Its length is the distance to N_{60} . Figure 7 illustrates the polygon setup for elliptical splats.

The alpha texture of the polygon is a discrete unit Gaussian function, stretched and scaled according to the polygon vertices deploying texture mapping hardware. The vertex positions of the polygon can be entirely calculated in the programmable vertex processor of contemporary graphics engines.

A current limitation of this rendering concept is the lack of per-pixel normalization in the frame buffer after splatting. As this feature is not yet available in today's graphics hardware and our dynamically changing models do not allow a pre-processing approach, it remains unsolved. However, in our experiments, we never experienced severe visible artifacts due to intensity variations.

4.3 Deferred Operations

We implemented a framework for deferred operations on all attributes of the 3D video fragments. Since vertex programs only modify the color and position attribute of the point samples during rendering, we do not destroy the consistency of the representation and of the differential update mechanism. Note that fragments can only be processed independently, and, because of the consistency constraints, fragments cannot be created or deleted in the vertex programs. Furthermore, temporal effects across multiple frames cannot be implemented by storing intermediate results, because the 3D operator stream modifies the representation asynchronously. Nevertheless, we designed a large number of visual effects from procedural warping to explosions and beaming. Periodic functions can be employed to devise effects such as *ripple*, *pulsate*, or *sine waves*. In the latter, we displace the fragment along its normal based on the sine to its distance to the origin in object space. For explosions, a fragment's position is modified along its normal according to its velocity. As illustrated in the accompanying video the explosion operation can take gravity into account. Note that all deferred operations are performed on-the-fly during rendering without any pre-processing.

5 Prototype System

We built a prototype real-time 3D video fragments system together with a coding scheme which will be explained in the following subsections.

5.1 Physical Setup

We use a prototype acquisition system with sixteen Point Grey Research Dragonfly digital cameras, equipped with 640×480 CCD imaging sensors. Fifteen cameras surround the acquisition stage on a circular shape and one camera captures the scene from the top. We use CS-mount lenses with focal lengths between 2.8mm and 6mm. The sixteen cameras are calibrated by a fully automatic procedure based on self-calibration [Pollefeys et al. 1999]. For estimating the radial lens distortion, we employ the Caltech camera calibration toolbox and we use the Open Computer

Vision Library for correction. Our test environment is an immersive telepresence installation which allows for simultaneous acquisition and projection. For projection, switchable glasses are between the cameras and the object producing slightly less saturated textures from most cameras [Gross et al. 2003]. We cope with this problem by adjusting the colors of individual cameras in a correction procedure. Although our cameras are capable of capturing synchronized images at 15 frames per second, the test environment limits us to 5 or 9 frames per second. For most of the tests we triggered the cameras with 5 frames per second. The physical setup of the acquisition environment is depicted in Figure 8.

Figure 8: Physical setup of our prototype immersive telepresence installation allowing simultaneous acquisition and projection.

Our processing system can be decomposed into three major parts. A set of camera nodes runs the acquisition clients, which do all image space processing steps in parallel, and finally transmit their data to a reconstruction node. The reconstruction node, in our case a dual processor machine with two AMD AthlonMP 2400+ CPUs, runs the multi-threaded 3D point processing which transforms the 2D pixel operations into 3D fragment operations. The rendering node receives the 3D fragment operations from the reconstruction node and maintains the data structure representing the 3D video object. This data structure is finally rendered to the screen. We currently use an 1.8 Ghz Pentium4 machine equipped with an NVIDIA GeForce4 Ti200 graphics accelerator. The rendering node transmits feedback information, i.e., the current viewpoint, to the reconstruction node. The reconstruction node controls the acquisition process at the camera nodes accordingly. All nodes are currently interconnected in a Fast-Ethernet local area network at 100 Mbit/s. In a typical application, the rendering node could also be connected via a wide area network.

5.2 3D Processing

For 3D video fragment processing we use a 3D reconstruction built upon the image based visual hulls method [Matusik et al. 2000]. However, instead of calculating depth and normal in a desired view, our 3D video fragments approach calculates 3D information for the camera views. Our current implementation would be able to deal with up to 85k INSERT or UPDATEPOS operations or more than 800k UPDATECOL and DELETE operations per second. A caching scheme ensures that the computation time for the costly INSERT and UPDATEPOS operations decreases logarithmically with the number of processed operations. The raw performance is sufficient for processing objects with less than 30k points. In a typical 3D video sequence, processed at 5 frames per second and containing in between 15k and 25k points, the position of 6% of the 3D video fragments is updated in each frame, whereas more than 10% might get a new color.

The operation scheduling at the reconstruction node is organized as follows: The contour data can simply be handed over to the visual hull reconstruction module. DELETE and UPDATE operations are immediately applied to the corresponding points. INSERT operations however, require a prescribed set of contours, which is derived from the active camera control. Furthermore, an efficient 3D video fragment processing requires that all DELETE operations from one camera node are executed before the INSERT operations of the same. The camera nodes support this operation flow by first transmitting contour data, then DELETE and UPDATE operations and, finally, INSERT operations. Note that the INSERT operations are generated in the order prescribed by the sampling strategy of the input image. On the reconstruction node, the operation scheduler will only forward INSERT operations to the visual hull unit if no other type of data is at hand.

Furthermore, every camera node, even the contour inactive cameras, transmit at least an empty set of contours for every frame. This strategy allows the reconstruction node to check if all cameras are still synchronized. The acknowledgement message of contour data contains the new state information for the corresponding camera node. The reconstruction node detects a frame switch while receiving contour data of a new frame. At that point in time, the reconstruction node triggers state computations, i.e. recomputes the sets of reconstruction and texture active cameras for the following frames.

The 3D video fragment operations are transmitted in the same order in which they are generated. But the relative ordering of operations from the same camera node is guaranteed. This property is sufficient for a consistent 3D data representation. Figure 9 depicts an example of a differential 3D video stream.

Figure 9: Differential 3D video fragments streaming pipeline.

The organization of our 3D video pipeline leads to a system inherent latency of 3 to 5 frames. Thus, the absolute latency is directly depending from the camera acquisition rate. The processing bottleneck in our pipeline are currently the camera nodes which can only handle up to 9 frames per second at full NTSC-resolution.

5.3 Streaming and Compression

The 3D video fragments pipeline requires a distributed consistent data representation. Each camera node shares with the reconstruction node a coherent representation of its differentially updated input image. The incremental updates of the rendering data structure also require a consistent data representation between the reconstruction and rendering nodes. Hence, all network links must support lossless and in-order data transmission. A TCP byte stream fulfills these requirements, however, TCP is not very well suited for real-time systems because of its built-in flow control mechanisms. Thus, we implemented an appropriate scheme for reliable data transmission based on the connectionless and unreliable UDP protocol and on explicit positive and negative acknowl-

edgements. An application with multiple renderers can be implemented by multicasting the 3D video fragments stream, using a similar technique as the Reliable Multicast Protocol RMP in the source-ordered reliability level [Whetten et al. 1994]. The implementation of our communication layer is based on the TAO/ACE framework (<http://www.cs.wustl.edu/~schmidt/TAO.html>).

A 3D video fragment is defined by a position, a surface normal vector and a color. For splat footprint estimation issues (Section 4.1), the renderer further needs knowledge about the camera identifier and the image coordinates of the original 2D pixel. The geometry reconstruction is computed in float precision, but the resulting 3D position can accurately be quantized using 27 bits. This position encoding scheme leads in our acquisition stage to a spatial resolution of approximately $6 \times 4 \times 6 \text{ mm}^3$. The remaining 5 bits of a 4 byte word can be used to encode the camera identifier. We encode the surface normal vector by quantizing the two angles describing the spherical coordinates of a unit length vector. We implemented a real-time surface normal encoder, which does not require any trigonometric computations on-the-fly, and which uses variable precision of up to 16 bits as suggested by [Deering 1995, Botsch et al. 2002]. Colors are encoded in RGB 5:6:5 format. At the reconstruction node, color information and 2D pixel coordinates are simply copied into the corresponding 3D video fragment.

Since all 3D fragment operators are transmitted over the same communication channel from the reconstruction node to the renderer, we need to encode the operation type explicitly. For simplicity, we use one prefix byte to each operation, and encode the operation type within. For UPDATE and DELETE operations, it is necessary to reference the corresponding 3D video fragment. We exploit the feature that the combination of quantized position and camera identifier references every single primitive. The renderer maintains the 3D video fragments in a hash table. Thus, each primitive can efficiently be accessed by its reference key. Figure 10 shows the byte layout for all possible attributes of a 3D fragment operator.

Figure 10: Byte layout of all possible attributes of a 3D fragment operator.

Figure 11 shows the bandwidth required by a typical sequence of differential 3D video, generated from 5 contour active and 3 texture active cameras at 5 frames per second. The average bandwidth in this 30 second sample sequence is 1.2 Megabit per second. The bandwidth is strongly correlated to the movements of the reconstructed person and to the changes of active cameras, which are related to the changes of the virtual viewpoint. The peaks in the sequence are mainly due to switches between active cameras. On average, INSERT and DELETE operators contribute to 25% and 12% of the bandwidth respectively, the remaining bandwidth is consumed by UPDATE operators. The compression performance of our system is difficult to analyze, since, up to our knowledge, no previous work was done on streaming and compressing dynamically

Figure 11: Bitrate of a differential 3D video stream. The contributions from the different 3D operators are highlighted.

changing geometric data. Relating our work to 2D video compression, the bandwidth required by a 3D video fragment stream is in the same order of magnitude than the bandwidth required by two MPEG-2 streams. At 5 frames per second, MPEG-2 recommendations ask for three intra-coded frames per second, which leads to a bitrate of approximately 1 Megabit per second.

Furthermore, entropy coding can be applied to the 3D video fragments stream. Our experiments show that an additional compression ratio of 2:1 can be expected.

6 Results

We have tested our 3D fragments pipeline in an immersive telepresence system with different persons. Figure 12 shows some snapshots of 3D video sequences with corresponding depth maps. Especially in Figure 12b the depths from the involved texture active cameras are visible. The cameras have to be shuttered at approximately 5 ms for simultaneous acquisition and projection of the prototype system. Thus, a high camera gain has to be set which leads to oversaturations in the region of the head.

Figure 1b illustrates the sine wave operation on the 3D video object from Figure 13a. Figure 14 shows an explosion of a 3D video object. The accompanying video shows example 3D video sequences with arbitrary virtual viewpoints and more effects. For our 3D video objects of less than 30k point samples we never experienced performance problems on the point rendering engine, even with enabled deferred operations.

7 Conclusions and Future Work

The 3D video fragments framework proposes a real-time 3D video pipeline which generalizes the 2D video pixels towards 3D irregular point samples. A differential update scheme exploits the inter-frame coherence of consecutive frames and a dynamic camera control scheme continuously reconfigures the 3D video process for optimal performance, according to the feedback from application and pipeline stages. The 3D video sequences are rendered using efficient point based splatting schemes and state-of-the-art vertex and pixel processing hardware.

In our current implementation, blending between several cameras might still lead to discontinuities. Moreover, texture artifacts due to improper color alignment between cameras are still visible. Furthermore, the requirement of a coherent distributed data representation is a severe constraint for the networking layer and an error-resilient representation for our 3D video pipeline needs to be investigated. In the future, we also plan to utilize the concepts of 3D video fragments for 3D video recording and try to integrate high-quality re-shading which requires smoothing of the poorly reconstructed normals.

Figure 12: 3D video fragment objects: a+c) Point rendered views. b+d) Depth maps from the respective views.

Figure 13: 3D video fragment objects. a) Point rendered view corresponding to the sine operation in Figure 1b. b) Another view from the 3D video sequence of Figure 1a.

Figure 14: Deferred Operations: a) Point rendered view. b) View from (a) exploded with gravity taking into account.

References

BOTSCH, M., WIRATANAYA, A., AND KOBELT, L. 2002. Efficient high quality rendering of point sampled geometry. In *Proceedings of the Thirteenth Eurographics Workshop on Rendering*, pages 53–64.

BUEHLER, C., BOSSE, M., MCMILLAN, L., GORTLER, S. J., AND COHEN, M. F. 2001. Unstructured Lumigraph Rendering. In *SIGGRAPH 2001 Conference Proceedings*, ACM Siggraph Annual Conference Series, pages 425–432.

CARCERONI, R. AND KUTULAKOS, K. 2001. Multi-View scene capture by surfel sampling: From video streams to non-rigid 3D motion, shape & reflectance. In *Proceedings of the 7th International Conference on Computer Vision*, pages 60–67.

DEERING, M. F. 1995. Geometry compression. In Cook, R., editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 13–20. ACM SIGGRAPH, Addison Wesley, held in Los Angeles, California, 06–11 August 1995.

GROSS, M., WÜRLIN, S., NAEF, M., LAMBORAY, E., SPAGNO, C., KUNZ, A., KOLLER-MEIER, E., SVOBODA, T., VAN GOOL, L., LANG, S., STREHLKE, K., VANDE MOERE, A., AND STAADT, O. “blue-c: a spatially immersive display and 3D video portal for telepresence.” In *Proceedings of SIGGRAPH 2003*. ACM Press / ACM SIGGRAPH, 2003. To Appear.

ISO/IEC 2002. Draft requirements for 3DAV. In Smolic, A. and Yamashita, R., editors, *JTC1/SC29/WG11 N4795*. ISO/IEC.

KANADE, T., RANDEP, P., AND NARAYANAN, P. 1997. Virtualized reality: Constructing virtual worlds from real scenes. In *IEEE MultiMedia*, Vol.4, No.1, Jan.-Mar. 1997, pages 43–54.

LAURENTINI, A. 1994. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162.

MATUSIK, W., BUEHLER, C., AND MCMILLAN, L. 2001. Polyhedral visual hulls for real-time rendering. In *Proceedings of Twelfth Eurographics Workshop on Rendering*, pages 115–125.

MATUSIK, W., BUEHLER, C., RASKAR, R., GORTLER, S. J., AND MCMILLAN, L. 2000. Image-based visual hulls. In Akeley, K., editor, *Proceedings of SIGGRAPH 2000*, pages 369–374. ACM Press / ACM SIGGRAPH / New York.

MOEZZI, S., KATKERE, A., KURAMURA, D. Y., AND JAIN, R. 1996. Immersive video. In *Proceedings of the 1996 Virtual Reality Annual International Symposium*, pages 17–24. IEEE Computer Society Press.

MULLIGAN, J. AND DANIILIDIS, K. 2000. View-independent scene acquisition for tele-presence. In *Proceedings of the International Symposium on Augmented Reality*, pages 105–108.

NARAYANAN, P. J., RANDEP, P., AND KANADE, T. 1998. Constructing virtual worlds using dense stereo. In *Proceedings of the International Conference on Computer Vision ICCV 98*, pages 3–10.

PAULY, M. AND GROSS, M. 2001. Spectral processing of point-sampled geometry. In *Proceedings of SIGGRAPH 2001*, pages 379–386. ACM Press / ACM SIGGRAPH.

PFISTER, H., ZWICKER, M., VAN BAAR, J., AND GROSS, M. 2000. Surfels: Surface elements as rendering primitives. In Akeley, K., editor, *Proceedings of SIGGRAPH 2000*, pages 335–342. ACM Press / ACM SIGGRAPH / New York.

POLLARD, S. AND HAYES, S. 1998. View synthesis by edge transfer with application to the generation of immersive video objects. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, pages 91–98. ACM Press / ACM SIGGRAPH, New York.

POLLEFEYS, M., KOCH, R., AND GOOL, L. V. 1999. Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. *International Journal of Computer Vision*, 32(1):7–25.

REN, L., PFISTER, H., AND ZWICKER, M. 2002. Object space EWA surface splatting: A hardware accelerated approach to high quality point rendering. In *Proceedings of Eurographics 2002*, COMPUTER GRAPHICS Forum, Conference Issue, pages 461–470.

RUSINKIEWICZ, S. AND LEVOY, M. 2000. QSplat: A multiresolution point rendering system for large meshes. In *Proceedings of SIGGRAPH 2000*, pages 343–352. ACM Press / ACM SIGGRAPH, New York.

SZELISKI, R. 1993. Rapid octree construction from image sequences. *CVGIP: Image Understanding*, 58(1):23–32.

VEDULA, S., BAKER, S., AND KANADE, T. 2002. Spatio-temporal view interpolation. In *Proceedings of the Thirteenth Eurographics Workshop on Rendering*, pages 65–76.

VLASIC, D., PFISTER, H., MOLINOV, S., GRZESZCZUK, R., AND MATUSIK, W. 2003. Opacity light fields: Interactive rendering of surface light fields with view-dependent opacity. In *Proceedings of the 2003 Symposium on Interactive 3D Graphics*. to appear.

WHETTEN, B., MONTGOMERY, T., AND KAPLAN, S. M. 1994. A high performance totally ordered multicast protocol. In *Dagstuhl Seminar on Distributed Systems*, pages 33–57.

WÜRLIN, S., LAMBORAY, E., STAADT, O. G., AND GROSS, M. H. 2002. 3D video recorder. In *Proceedings of Pacific Graphics '02*, pages 325–334. IEEE Computer Society Press.

ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. 2001. Surface splatting. In *Proceedings of SIGGRAPH 2001*, pages 371–378. ACM Press / ACM SIGGRAPH, New York.