

Principles of tracking

Tracking is a process of gathering information from previously filmed footage. It has become one of the key visual effects techniques of the present. Tracking does not produce finished shots but numerical information that can be used to match virtual camera and animated objects and characters with live-action camera. Tracking markers are usually used to help the camera tracking process. (Kerlow 2004: 375, 377) Before starting the tracking process, it is good to preview the shot in real time. Making notes of good tracking targets and when they go out of frame or are occluded can save time and trouble in the actual tracking process. Any parallax in the scene will disqualify otherwise attracting tracking targets. (Wright 2002: 183)

Many commercial tracking tools offer only solution for 2D tracking. 2D tracking tool is in fact a keyframe-generating assistant that analyzes images and calculates pixel shift on defined areas of interest. Visual effects artist could do all this manually but tracking tools are used to speed up the workflow. (Davies 2005: 196) 3D tracking resembles photogrammetry, a technique that can extract three dimensional models from two or more still images of a subject. Usually this is used to extract a depth map.(Kerlow 2004: 383)

In 3D camera tracking suitable points from the image are chosen as feature points either manually or automatically. The position of the feature points in the image plane

16
is measured and combined with the focal length of the camera. 3D tracking software analyses the footage frame by frame and estimates camera parameters from the paths of the chosen feature points. After analysing the acquired footage, the software can produce information usable in creation of a scene in 3D modelling software with virtual camera and the feature points placed in the 3D space.

Using gathered data in 3D modelling software

After the tracked data is exported from the tracking software, it is brought to the chosen 3D modelling software as a script. From the gathered data a virtual camera is created with the correct focal point and all the movement of the original camera used for acquiring the live-action footage. Also, all the feature points that were tracked are represented in the created 3D scene as a point cloud relative to the camera. Although the feature points are not in the correct position in the 3D scene, they can be used to determine the correct position and movement for the 3D models. (Ferguson & Heron) It is important to confirm that the 3D scene is in the correct aspect ratio before any animation is done. Easiest way to confirm this is to import the live-action footage as an image sequence to the scene and set it as a background image. After this you can easily see if the point cloud matches the background image from the virtual cameras perspective. When the point cloud matches the background image the 3D model can be imported to the scene. Using the point cloud as a reference, the model can be placed at the correct distance from the virtual camera.

What is Matchmoving, how does it work?

Matchmoving is a technique that is used to input computer generated imagery into a live footage. The result of this skill is a realistic illusion; ensuring viewers believe the CG element is an actual component of the live footage. Typically, a matchmover's purpose is to match the live footage camera movement into a virtual camera so that the perspective view of the CGI is equivalent to the live footage when merged together. Matchmoving or 'Motion-Tracking' has two varieties; the first being 2Dimensional. This is where you would track the camera movement of 2D footage to input special effects like motion blur or adding other 2D features e.g. changing the image of a billboard on the 2D footage.² This Matchmoving is sufficient when the special effect and the footage do not alter much in perspective. Software such as After Effects contains tools that allow users to select a feature of the footage to track and follow to create a path, which is the camera movement. The special effect is inserted and follows this path. The 3dimensional class of the technique is where you would recreate the 3dimensional space of the live footage and then track the camera movement of the footage using the footage and the 3dimensional space to insert the CG element, with the same perspective, into the live footage. Below is an example of the typical 3dimensional Matchmoving process, giving an idea of how CGI is inserted:

To load an image sequence

1 Select File > Load Sequence or click the Load Sequence icon in the toolbar.

The Load Sequence window appears.

Lesson 1: Tracking an image sequence | 7

2 In the Load Sequence window, browse to the location of the MatchMover Tutorials/Basic directory, and select lion000.jpg, which is the first image in the sequence.

3 Click Details.

Information about the image sequence, such as its length as well as preview images, are displayed in the Details section of the Load Sequence window.

The video you use for this lesson is in the interlaced PAL format, which means you need to set MatchMover sequence options for this type of video.

4 In the Sequence Options section, select Lower field first from the Interlace drop-down list.

5 In the Camera Settings section, select 25 fps from the Frame Rate drop-down list.

8 | Chapter 1 Getting Started

Leave Motion set to Free because in this example, the camera does not follow any of the motion types specified in the list.

To play the image sequence

1 Select Sequence > Play or click the play button which is in Play Sequence Toolbar at the bottom of the Workspace.

Note that when you play the beginning of the sequence it is slightly jerky. This is because the sequence is loading into your computer's RAM. After the images are in the cache, the sequence plays smoothly. The size of the current cache is displayed in the bottom-right corner of MatchMover user interface.

Since the image sequence is only 50 frames, it plays back too quickly to closely examine. Changing the play mode to PingPong will continuously loop the sequence back and forth, making it easier to check the sequence motion.

Tracking an image sequence |

2 Select Sequence > Play Mode > PingPong or click the PingPong button in the Play Sequence Toolbar.

3 To stop the playback, click the Stop button or press Esc.

TIP To play the scene manually, you can press Ctrl, click anywhere in the image, and drag the pointer to the left or to the right.

4 To go back to the first frame, press Ctrl+Home (Windows) or Command+Home (Mac) or press .

Running the automatic tracker

You can use the automatic tracking feature to have MatchMover select the best location of place track points in the image sequence.

NOTE You can bypass this section of the lesson by opening Basic_Calibrated.mmf. This file can be found in the Tutorials/Basic directory. You can then proceed to

To run the automatic tracker

1 Select 2D Tracking > Automatic Tracking or click the Run the Automatic Tracking icon .

The Automatic Tracking window appears listing the steps in the automatic matchmoving process.

Getting Started

2 Click Run.

Colored indicators display beside the option name to show you the status of each step.

After MatchMover has computed the camera path, a collection of 3D points are displayed as cross-hairs in the Workspace. The points are also represented by colored icons in the Project window in the Point Track > Auto Tracks folders.

Tracking an image sequence |

The colors indicate the quality of the tracking, where green indicates good, yellow fair, and red bad tracking. Some grayed tracks have not been reconstructed, which means that MatchMover considered them unfit to calculate.

Checking the tracking results

To check the results of the automatic tracking you can:

- View the track quality by looking at the cross-hairs superimposed on the Workspace, or by view the color icons in the Auto Tracks folder. Most of your reconstructed tracks should be green, a few yellow, and none red.
- View the reconstructed 3D scene, and check to ensure that it is consistent with the real scene.
- Render a video sequence where the 3D objects are superimposed on the image sequence and check to ensure that the integration is seamless.

To check the tracking results in 3D space

1 Click the 3D space icon that is located in the top left of the workspace

view.

MatchMover switches to 3D mode. The tracks are now displayed as 3D cones rather than cross-hairs.

2 Use the following Navigation tools to examine the tracking:

- Click-hold the Dolly icon to move backward and forward in the space.
- Click-hold the Orbit icon to view the space in 360 degrees.

When orbiting the camera, the 3D View automatically orbits around the selected 3D object, point cloud, or camera.

When you navigate in the 3D scene, you can recognize the different elements of the scene, such as the lamppost, statue and wall, as well as the camera path.

3 To center the view on a track, select it, then select View > Look At.

4 To view the Workspace through the computed camera with the image in the background, select View > Lock on Camera or click the icon.

When viewing the Workspace, you can use of the location of the 3D cones to judge the potential orientation of any elements that you might want to introduce into the scene.

5 If the 3D cones appear too big or small, you can change their size by selecting the Project tab in the Parameters window. Change the 3D Helper size value from 1 to 0.5.

Rendering the tracked image sequence

After you have created and checked the quality of the 3D tracking in the Workspace, you can render the scene in a video format to further check its quality. You can also export the tracked sequence to a file format that you can open in your favorite 3D animation or compositing program.

To render the tracked sequence

1 Select 3D Scene > Render Setup.

The Render Setup window appears.

2 In the Render Setup window, click Setup/Browse and do the following:

- Type a file name for your rendered sequence.
- From the File type list, select an output file format.
- Click Save.

3 From the Resolution list, select 50%.

4 In the Options section, turn on Antialiasing.

5 Click Render.

The Render window appears, displaying your rendered sequence with 3D cones composited on your original image sequence. With perfect tracking, the 3D cones should appear as if they were part of the scene. Otherwise, you would notice slight discrepancies between the motions of the real scene and the 3D elements.

Lesson 1: Tracking an image sequence | 15

Now that you have automatically tracked a PAL video sequence and verified it by rendering it, you can export the results to a file that can be opened in your favorite animation or compositing software.

6 Select File > Export or click the icon.

The Export window appears.

7 In the Export window, do the following:

- Type a name for your exported file
- From the File type list, select a format for you exported file.

For example, to open your tracked image sequence in Maya, select Maya (*.ma).

- Click Save.

The image sequence is exported to the specified file format for use in your 3D animation or compositing program.

Introduction In recent years, the process we call “Matchmoving” (also referred to as “3-D Camera Tracking”) has come to play an increasingly crucial role at the foremost edge of the digital graphics world. Matchmoving is made possible by analytical processes applied to original motion picture sequences, by which we define and recreate authentic 3D motion data for all secondary visual elements. Our matchmoving system, 3D-Equalizer (“3DE”), is an elegant, powerful, easy-to-operate solution for synthetically reconstructing critical motion and structural data from practically any existing quantifiable image elements. Exploring all the possibilities of your 3DE postproduction system is guaranteed to trigger a widening range of exciting creative opportunities, and also to maximize workflow efficiency. 3DE, now in its 3rd generation, is a stand-alone tool which automatically resolves the extensive challenges and specific problems of matchmoving, based on its ingenious scientific calculation core and supported by an intelligent, highly adaptive user interface. As confirmed by our system users, the industry’s foremost image creators, 3DE V3 is designed and tested from the ground up to be the most powerful and efficient matchmoving tool you will find anywhere. Most matchmoving scenarios involve a combination, or more precisely, a composition of real elements and virtual elements based on time-dependent spatial relationships between your cameras and the objects being recorded by these cameras. 3DE is your principal ally throughout this mathematically critical process. 3DE automates complex time-space-motion computation, then seamlessly transfers all the motion data essential to further processing stages for your modeling, animation and compositing modules via specially-created export filters. This manual describes in detail all of 3DE’s V3 stand-alone functions, controls and characteristics as well as 3DE system workflow in relation to other applications. Please contact us with any questions or suggestions regarding further developments. Matchmoving is increasingly used in the fields of commercial, industrial and feature film production, as well as for visualization in science and architecture, in video reenactments and in virtual studios. The basic situations of 3DE usage are categorized as follows: a) Virtual (or “CG”) objects placed in a real visual environment. In architectural previsualization, CG is often used to integrate partial or entire structural illustrations into real-life photographs or video sequences. Using related methods, CG prosthetics and “virtual make-up” can be applied to alter the features of actors or artifacts in a motion picture scene. 4 3D-Equalizer Version 3 b) Real objects placed in a virtual (or CG) environment. An example of this would be a real actor moving among computer-generated scenery. In TV studio situations, presentations often involve “virtual sets” in a process normally requiring lots of expensive hardware to accomplish - however not if it can be done in postproduction, because now the entire motion reconstruction process can be worked out quickly using 3D-Equalizer. c) Real objects placed within real environments. Hybrid images combining different real elements do not work convincingly if scales and perspectives from all the various elements are too divergent. Remember “The Incredible Shrinking Man”? You won’t have to, when you’re using 3DE! Your real environment is recorded in the conventional way, then all motion information is reconstructed in 3DE and the data fed into a motion control rig which shoots other real objects to be composited into the primary real environment. d) Stabilizing creative optical illusions. In “frozen moment” shots (also variously known as “bullet time” or “slice of life”) which are based on a geometric combination of simultaneously triggered still cameras, 3DE is used for calculating the exact position and orientation of each individual camera and a point-based geometric model reconstruction of the set. Your motion data and the point model exported from 3DE define and compensate for all and any random fluctuations of position and orientation of the group of cameras, so that your shot will then be flawlessly composited.

Conventions 3DE's graphical user interface is designed in accordance with the SGI Indigo Magic Style Guide. All 3DE console colours mentioned in this manual are selections from the Indigo Magic desktop colour scheme. When you wish to enter something in a text field, just place the mouse pointer over the field and enter your text. Removing the mouse pointer from the text field then confirms the entry.

Terminology Within 3DE and this manual certain specific expressions are used, most of which will be familiar to everyone experienced in modelling or compositing. In the following section we will review some of these expressions.

5 Introduction Basic geometric expressions

When a scene is recorded by a camera, for each moment in time camera characteristics are defined by position and orientation. When we fix coordinates in space, we can represent this position by three values (x,y,z) . The orientation of the camera, essentially the direction it looks in and the rotation around this view axis, can be represented in many different ways, most familiar of which the reader will be with angular representation (rx,ry,rz) . The explicit choice of this representation is not so important here. In a similar sense we can characterize a recorded object by its position and orientation, provided we associate an origin and some coordinate axis to this object. Let us, however, consider only the case of a moving camera in the following.

Cameras and Images Recording something with a film or video camera is a process which evolves in time, it has a beginning and an end. Hence, it is clear, what we understand by a camera motion path of a sequence: it is the set of all camera positions/orientations for a given shot. We can easily visualize such a camera motion path. Camera motion path For simplicity, we will represent a camera by a pyramid, which is also the representation within our software (e.g. Orientation Window) and also within many animation packages. This corresponds to a very simple pinhole camera model. Until we begin to deal with non-linear effects like lens distortion, this model is sufficient for our purposes. Camera pyramid grip this grip this real object

6 3D-Equalizer Version 3

The geometric data of a camera relevant to 3DE (except for lens distortion) are contained in the shape of this pyramid: the focal length, filmback width and filmback height. Another important parameter is the pixel aspect, but this is more associated with the digitizing process than to the camera optics itself. The camera geometry and digitizing process is explained in the section "Camera Adjustment", chapter "Selected Topics". The precise definition of our lens distortion model is explained in the appendix (see also tutorial "Lens Distortion"). Let us assume we want to attach some CG-content to the scenery so as to get the correct perspective throughout the shot. The CG-content is rendered by a virtual camera, as you already know from your rendering software. If it is defined correctly according to the original footage, that is, if the virtual camera performs exactly the same movement as the real camera, the perspective of CG objects and real objects will always be correct. Therefore the real camera movement must be known or calculated from the footage, and this is what 3DE does.

Sequence and reference frames

The basic input for 3DE consists of image files, the so-called frames. These frames can be divided into two classes, namely sequence frames and reference frames. A (real) sequence is a set of frames, recorded with a camera, which fulfill the following condition: the image content varies more or less continuously, which reflects the continuity of time. In other words, neither the camera positions and orientations nor the position and orientation of the recorded objects "jump" from one frame to the next frame. In contrast to a sequence, reference frames are usually single, isolated recordings of the real scenery or some real objects. They could e.g. consist of photographs. The task which 3DE performs is to calculate the motion path of a camera/object for a sequence. The purpose of reference frames is to provide additional information to 3DE. Usually, in reference frames, the image content does not exhibit any kind of "continuity" like a sequence would.

reference frame reference frame sequence (frames) camera motion path 7 Introduction As mentioned, the frames represent the basic input for 3DE, but the 3d-calculations of 3DE are not based on evaluating the entire image, only some very restricted subsets of information. Again, let us consider the situation in which a camera is recording scenery. The user identifies this scenery with a point group. Such a point group consists of points, objects which essentially represent a 3d-position in space, but which are equipped with a lot of other properties in the context of 3DE. Usually the user designates striking points in the scenery which can easily be recognized in the images. These striking points happen to be in the scenery "naturally" or are placed intentionally (e.g. markers). In the course of the sequence and also in the reference frames, the points are identified by the user within the frames. For sequences this process is referred to as motion tracking. The expression "Tracking" makes sense, because, as we mentioned above, the image content varies continuously. The result is a so-called tracking curve, a two-dimensional curve (e.g. expressed in pixel coordinates). In practice, a tracking curve often will not represent the motion of the point during the entire sequence. On the contrary, the curve will be truncated when a point in the real scenery is not yet (or no longer) visible, or covered by obstacles. However, for simplicity, we can assume here that the points are tracked during the entire sequence. Camera point group and object point groups When a camera is recording a single object, 3DE calculates the geometric relationship between camera and object. We will illustrate in the following how camera movement and object movement relate to each other. In order to do this, we consider a circle-shaped camera motion around a cube. For the calculation kernel of 3DE this situation is completely equivalent to a fixed camera and a rotating cube. Circle-shaped camera motion vs. rotating object On the other hand, imagine a camera fixed to a tripod, rotating around its vertical axis, and a fixed cube. This situation is equivalent to a totally fixed camera and a moving cube. It would also be possible to interpret both camera and cube as moving objects. 8 3D-Equalizer Version 3 rotating camera vs. translating object In practice, it may happen that both types of information are required: the absolute movement of the camera with respect to the scenery and the absolute movement of the object with respect to the scenery. As an example, imagine a street with a moving car. The camera as well as the car are moving with respect to the street. The task is to replace the real car by some CG-car moving on the street. You want to add realistic lighting (e.g. from the sun) and realistic shadowcasting on the street. So, it would be very useful to know the absolute movement of the car with respect to the street/scenery. Object pointgroup: car - Camera pointgroup: street In order to solve this type of problem, 3DE provides methods for calculating several point groups within one project. The two important expressions in this context are "camera group" and "object group", which we define below. In order to calculate the motion of the camera with respect to the street one would choose some tracking points on the street. These points form the so-called camera group. The car, however, is a moving object, and we want to determine its (absolute) movement with respect to the street. Again, one chooses a few tracking points on the car, and these points form a point 9 Introduction group called the object group. There can be at most one camera group, but there could be in principle an arbitrary number of object point groups. As you see, object groups and camera groups now clearly distinguish between moving cameras and moving objects. It would make no sense to talk about a street moving with respect to a fixed camera. Deviation When 3DE calculates a camera motion path (or object motion path) by evaluating the tracking curves, the positions of the points in three-dimensional space are also calculated. In a proper arrangement of camera positions and point positions the points in space lie on the rays from the camera position through the image plane as in figure

“Sequence and reference frames”. In practice however, there will always be a slight discrepancy between the rays and the points. This discrepancy can be projected back to the image plane and is usually expressed in pixel. We call it the deviation in the context of 3DE and in this manual. So, the deviation is a number associated to a frame and a point group (or to a frame and a point). In 3DE’s Status Window you can examine the deviation for an entire sequence and a point group as a curve. The precise definition of the deviation is not so straightforward, and we shall omit the details here. It may be sufficient to say that it expresses the average difference between the “is”-position of the 3dpoints projected onto the screen from their “should-be”-position given by the tracking curves of the respective points.

Calculation methods Let us imagine a situation in which the motion of camera is to be calculated by 3DE. In the current version, 3DE distinguishes two situations: (1) the camera is translating and rotating in three-dimensional space, (2) the camera is rotating but its position is fixed, e.g. when the camera is fixed to a tripod. The attribute “fixed or not” only makes sense for a camera group, and it is associated to a sequence (Frames Window::Modify...). The method for processing this kind of sequence is called the fixed camera position mode (FCPM), and is enabled by activating the “fixed camera position constraint”.

Workflow In the figure below you see the typical workflow of a 3DE project for a simple project. In the tutorials the single steps are carried out in detail. As mentioned, the basic input when setting up a project consists of frames, namely reference frames and/or a sequence. It is possible to build a project that only consists of reference frames, e.g. if you want to use 3DE in the original 10 3D-Equalizer Version 3 sense of photogrammetry, i.e. surveying a three-dimensional object by means of recorded images from this object. However, in practice many projects contain only a sequence with, in some situations, reference frames being added. The next step is to create a point group. The expression “point group” is new in version 3. Since in contrast to version 3 of 3DE, version 2 handled only a single point group it has now become necessary to distinguish different point groups by name, to interpret them as objects. In the tutorials you can find examples of projects with more than one point group. The points are then tracked in the sequence (and placed in the reference frames if possible). What we simply called “track points” in our diagram is in practice an interactive procedure in which several functions of the software can be involved: choosing the most precise method, placing key-frames, tracking forward and backward, using image controls when the image material has low contrast, etc. We have now gathered the basic information about the object or the scenery. However, we have also to specify in what way, the real world is “projected” into the image files, i.e. we need to specify the camera geometry and the digitizing parameters. The camera geometry consists of the focal length and the filmback width and height (or for video cameras: CCD chip size). There are many ways to express these data, e.g. in terms of horizontal and vertical aperture angle or, also, horizontal angle and film aspect ratio. It is up to the user to choose the most convenient way. The digitizing parameters are essentially the field of view width and height and the pixel aspect. All the parameters mentioned here are in some way dependent on one another, and this dependency structure defines the special appearance of our Camera Adjustment Window. Usually, it is not necessary to enter all parameters because, as mentioned, some of the information is redundant and unknown parameters can be calculated by 3DE. Apart from the basic geometry of the camera it may be necessary to compensate for non-linear effects of the camera optics. These effects are usually referred to as “lens distortion”. Please read the tutorial “Lens Distortion” for more detailed information. Often, some of the camera parameters are unknown or only approximately known, others are known or can be guessed. The camera parameters which are unknown are calculated

by the software. This also includes a calculation of the points in space and the camera motion path. In the diagram we have simply summed up this procedure in the field "Calculate points and motion path". When all 3D calculations are done, you could, in principle, export all data to the animation software (or to a compositing system that is able to deal with 3d-objects). However, it might be necessary to render a preview within 3DE and then improve the quality of the calculation. The Status Window displays a curve which expresses the deviation of a point group for a sequence, resolved in frames. It is also possible to examine this curve for each point, and these curves can be used to detect errors in the motion tracking. This is the working cycle "check deviation curve -> improve motion tracking". In some situations (see tutorials) it may also be necessary to use the postfilter in order to smooth out the motion path. The working cycle 11 Introduction usually is "tweak postfilter->generate preview". Some users prefer to use the final rendering output in high resolution to check the result. In this case the procedure reads "tweak postfilter- >export data->render and compose->tweak again" or similar. Workflow There are projects (particularly when not enough points are available because the image material does not allow you to track a sufficient number of points) in which it may even be necessary to edit the camera/object motion path. We consider this function optional, and it should be the last step (apart from rendering a preview movie) before exporting the data to the Reference Frames Sequence Frames Define pointgroup and points Track points Adjust field of view Specify camera parameters Calculate points and motion path Generate preview movie Tweak postfilters Check deviation curve Improve motion tracking Edit motion path Export data Setup project 12 3D-Equalizer Version 3 animation package. Everything else such as rendering and compositing is done outside 3DE. If your sequence exhibits lens distortion, you may use the external tool WarpDistort in order to distort the rendered images or to remove distortion from the real images. WarpDistort reads the project file of 3DE and constructs an image operator according to the specified distortion parameters. Differences between 3DE V2 and V3 The project files are upwardly compatible, i.e. version 2 files are understood by version 3, but not vice versa. The most important difference between version 2 and version 3 is that the latter one has "more of everything". In other words, it is possible to handle more than one point group (as mentioned, the expression point group did not make sense in version 2), more than one sequence and more than one set of reference frames, associated to different cameras. Version 3 is able to deal with so-called "infinite points", points which are too far away to be triangulated properly, and which are therefore declared to be located at "infinity". These points are particularly important when the foreground of the recorded scenery is changing quickly, making it difficult to find a sufficient number of points with a sufficient overlap in time. The lens distortion model of version 3 consists of five parameters and enables users to handle anamorphic lenses, while the single-parameter model of version 2 was insufficient for this type of lens. Finally, in version 3 there are a limited number of constraints which allow the user to force the camera to move only in certain limits. 13 Introduction 14 Tutorial In the following chapter the main features of 3D-Equalizer are presented by means of several examples. The 3D-Equalizer project files that belong to these examples can be found in the directory "/projects/". Tutorial 1: Introductory project "ball_bluebox1" The matchmoving process, determining 3D motion data out of a 2D image sequence can be summed up as follows. The "Environments concept" section explains the efficient handling of 3D-Equalizer's user interface. In the sections "Creating a new project" and "Creating a new sequence" we describe how to set up a project, including the acquisition of the sequence frames. The sections "Creating a new pointgroup" and "Creating a new point" describe how to enter points which are going

to be tracked in the sequence frames. Section "Tracking markers" describes the various methods of tracking the points on the screen. Afterwards, in section "Defining camera parameters", the camera geometry and motion path are reconstructed. This enables us to take a first look at the preview movie, which is generated as described in the sections "Editing dummy objects" and "Generating a preview movie". In order to get in contact with your preferred animation package the reconstructed points and the reconstructed camera path are exported. The various exportfilters are mainly described in the "Reference Manual". This completes the basic procedure of 3D-Equalizer. 15 Tutorial Getting started • Please start 3D-Equalizer by double-clicking on its icon in the icon catalog window (SGI) or by entering "3dequalizerV3" in a Unix shell window (Linux/SGI). After a few moments, the requester 3DEqualizer Software Licence opens. • Enter the name or IP address of the Flc server workstation (see installation) in the text field Licence Server in order to get a list of the currently available licences. Choose one of them

let's take care of the lens distortion parameter. Real world camera lens systems sometimes distort the filmed images in a non-linear manner. That means, the image content is "bent" towards the edges. • Activate the toggle button Overview Window::View::Show Distortion Grid. A red grid is now displayed over the current frame in the Overview Window. It visualizes the lens distortion of the camera linked to the current sequence. • Enter several different values, ranging from -0.05 to 0.2, into the Lens Distortion textfield and watch the Overview Window's distortion grid. A value of 0.0 yields no lens distortion effect. Values smaller than 0.0 results in a pillow-type, whereas values greater than 0.0 results in a barrel type distortion grid. Typical values of lens distortion for a video camera are 0.02 to 0.08. Good film camera lenses will often exhibit no visible lens distortion at all, so that the value may be set to zero. Distortion Grid A consumer video camera has been used to shoot the sequence. There are some minor lens distortion artefacts visible. • Enter a lens distortion value of 0.05 and set the respective option menu to Fixed. Finally, we have to define the type of camera lens used for recording the sequence. In principle, there are two different ways in to do that in 3D-Equalizer. You can specify a focal length, a horizontal or a vertical camera angle. 43 Tutorial Usually, information is available about the focal length of the camera lens. Unfortunately, this parameter depends highly on the size of the film gate (size of the negative). For instance a 10mm lens for a video camera is totally different from a 10mm lens for a 35mm film camera, because the CCD chip (equal to film gate) of a video camera is much smaller than the film gate of a 35mm film camera. If one knew the exact size of the film gate and the focal length, then the camera lens could be properly defined. In addition there is another problem. Images are often cropped during the digitizing process. That means, the digital image may represent only a fraction of the actual film gate size. In this case the respective focal length must be properly adapted to the "cropped" film gate size, which can be quite difficult. It is much more intuitive to define the camera lens by its horizontal or vertical camera angles. These parameters are not dependent on film gate size. For instance, the 45.0 ° horizontal camera angle of a video or a film camera will produce the same kind of image. A horizontal angle greater than 50.0 ° represents a wide angle camera lens, and a horizontal angle smaller than 30.0 ° represents a narrow angle (zoom) camera lens (see pictures below). Image shot by a narrow angle (zoom), a regular angle and a wide angle camera lens Unfortunately, information on the camera angle is not often available, and cropping the images causes the camera angle to be

changed, too. Therefore, it is often useful to let 3D-Equalizer establish the camera lens for you. In our example, we have no information about the camera lens. The image material appears to be taken by a regular lens.

- Set the option menu beside the horizontal angle textfield to Unknown. Check the Pixel Aspect textfield, and if necessary set its option menu to Fixed and enter a value of 1.0926. The camera is now prepared for the adjustment process.
- Press the Adjust button, in order to start the adjustment process. Several messages will appear on the screen, indicating the current status of the calculation. 3DEqualizer will try to calculate the parameters declared to be Unknown, then trim the parameters with setting Fine Adjust. You can terminate each level of calculation by pressing Cancel in the progress requester. The parameters yielding the best value up to that moment will then be used.

44 3D-Equalizer Version 3 Status Window 3D-Equalizer finishes this adjustment process and calculates a horizontal angle of about 44.88°. Now, please take a look at the Status Window. You see a green curve, which represents the reconstruction quality of each single frame. It expresses the deviation of the 2D motiontracked points from the reconstructed 3D points displayed in pixel. There is a dotted line with a black label "av:" and a number. This number represents the average deviation of the entire sequence. It should be around 0.35 pixel, which is already quite a good value. In order to learn more about adjusting a camera please take a look at the tutorial "Camera Adjustment". The project "ball_bluebox1.3de" found in the directory "/projects/", contains all steps we have taken up to now. Please feel free to review it at any time.

Error analysis using the Status Window In the following section we will present how you can identify tracking errors by means of the Status Window.

- Please open the project "ball_bluebox2.3de". You can find it in the directory "/projects/".
- Select the "Orientation" environment from the Environments menu, and take a look at the 3D Object Browser Window on the right side of the desktop. The 3D Object Browser Window contains a list of all 3D objects displayed in the Orientation Window. It is mainly used for selection purposes.

50 3D-Equalizer Version 3 3D Object Browser Window

- Click with the left mouse button on the Scene entry in the 3D Object Browser Window. The Scene entry becomes selected and the coordinate frame labeled "Scene" is drawn in redgreen-blue in the Orientation Window, which indicates that it is selected as well. This 3D object is called the "scene node". It represents a coordinate frame under which all other relevant 3D objects are grouped. Aligning the 3D data to the global coordinate frame means transforming the scene node!
- Select some other objects in the 3D Object Browser Window, and watch the Orientation Window's display. Now, let's take care of the Orientation Window's handling.

51 Tutorial Orientation Window

- Hold down the Alt-key and click-drag the left mouse button. The mouse cursor changes into a rotation symbol. Dragging the mouse, you are able to rotate the view around the displayed 3D objects.
- Click on the Top, Side and Front buttons found on the left border one after another, and watch the results in the Orientation Window. The 3D objects are shown from the top, from the right and from the front, respectively.
- Hold down the Alt-key and click-drag the middle mouse button. The mouse cursor changes into a pan symbol. Dragging the mouse, you are able to pan the view in a horizontal or vertical direction.
- Hold down the Alt-key and click-drag the right mouse button. If a menu should open up, please hold down the Shift- and Alt-key instead. The mouse cursor changes into a dolly symbol. Dragging the mouse, you are able to move into or out of the view.

52 3D-Equalizer Version 3

- Drag the vertical slider located on the right border of the Orientation Window. As you see, dragging this slider you are able to move into and out of the view, as well. Let's continue to modify an object.
- Select the scene node in the 3D Object Browser Window. Activate the scale toggle button, located on the left border, and click-drag the left mouse button in the

Orientation Window's display. As you see, by moving the mouse to the left and to the right, the scene node is scaled uniformly. Since all other objects are grouped under the scene node, they are scaled as well. • Activate the rotate toggle button, located on the left border, and click-drag the left mouse button in the Orientation Window's display. The scene node and all other relevant objects are rotated around its local x-axis. • Click-drag the middle or the right mouse button in the Orientation Window's display. The scene node is rotated around its local y-axis or its local z-axis, respectively. • Hold down the Shift-key and click-drag the left mouse button in the Orientation Window's display. The scene node is rotated around the x-axis of the global coordinate frame. • Please play around with the remaining translate toggle button and the Shift-key. Using these mouse functions, you can modify every 3D object in the Orientation Window. In addition, there is another method to modify attributes of 3D objects. • Double-click on the scene node entry of the 3D Object Browser Window, or simply click on its Modify button.