*To be presented at the **ACM SIGGRAPH 2002** Sketches & Applications Program*

An Inexpensive 3D Camera

Askold Strat
Symbol Technologies and
SUNY at Stony Brook

Manuel M. Oliveira
Department of Computer Science
SUNY at Stony Brook

Abstract

We describe the implementation of a portable 3D-camera prototype based on a consumer grade digital camera and an inexpensive laser raster generator. Such hand-held device can be used to capture smooth shapes by acquiring one or more images. Its output can be either a 3D wireframe or a textured model.

CR Categories: I.4.1 (Image Processing and Computer Vision): Digitization and Image Capture; B.0 (Hardware):
Keywords: 3D camera, 3D shape acquisition

1 Description

Shape acquisition has become an essential component of computer graphics systems that require modeling and visualization of objects and scenes from the real world. Most of the existing devices for capturing 3D shape have high cost and limited portability. We are building an inexpensive 3D camera using consumer grade technology to support the acquisition of smooth shapes.

Our 3D camera consists of three modules: a digital camera, a laser raster generator and some software. Although, in principle, any digital cameral can be used, our current prototype uses a monochrome VGA CCD device. The laser raster generator, in turn, consists of a laser module, a high frequency oscillating mirror (x-mirror) and a low frequency oscillating mirror (y-mirror). Pictures showing the raster generator and the high frequency element are shown in Figure 1.

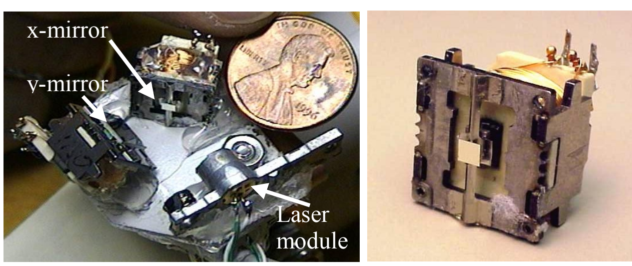

Figure 1. Raster generator (left). High frequency element (right)

The laser raster generator projects a number of parallel lines, which are deformed by the objects' surfaces (Figure 2 - left). The geometry is then recovered from the image using triangulation (Figure 2 – right). A synchronized circuit turns the laser on and off (to eliminate the horizontal retrace), depending on the movement direction of the x-mirror. A PC sound card produces the waveforms necessary to drive the two mirrors. The left and right channels of the audio output are fed to drivers of the x-mirror (2KHz) and of the y-mirror (5Hz), respectively. The faster mirror is driven with a sinusoidal waveform that matches its mechanical resonating frequency. Driving the element at resonance is necessary to maintain lower power consumption, which in our current prototype is below 0.5 W. Unlike the fast element, the y-mirror should not operate at resonance, since at resonance its movement would be sinusoidal. In order to obtain laser planes whose angles are equally spaced, the angular velocity of the y-mirror has to be constant throughout the entire range of motion. This is not possible since the mirror and the magnet make a large mass that needs to be decelerated/accelerated when the mirror changes direction. Experimentally, we have found that the mirror's angular deflection is proportional to the current through the coil for angles up to +-17 degrees. Thus, linear motion can be obtained away from both ends of the motion, where velocity is uniform. The camera's shutter is synchronized with the motion of the y-mirror using the mirror drive signal to trigger the shutter. The exposure time is controlled independently in the camera.

Once an image containing the raster pattern (*e.g.*, Figure 2 – left) has been acquired, it is sent to the computer (the current prototype uses a serial interface), where it is processed. The 3D-shape reconstruction can be obtained using two different approaches. If all raster lines are visible in the image, the angular value associated with each line is immediately available and reconstruction is straightforward. In order to guarantee that all raster lines are visible, one needs to constrain the working range of the device, possibly using a back plane at the maximum acceptable range. This approach reduces the hardware requirements considerably, but the software has to make more intelligent decisions to disambiguate the lines that correspond to different projection angles. For unconstrained working ranges, two images are taken in a row from the exact same position. One of the images contain only one (or a few) raster line(s) with angular value(s) known a priori, from which the angles of all remaining lines are obtained. The use of more than one line is required if the scene contains several objects at different heights.

Figure 2 (left) shows a stripe light pattern produced by the raster generator and projected onto a person's face. The image to its right shows a wire-frame reconstruction of the visible lines from a different viewpoint.

The 3D-shape acquisition system described here is subject to the same limitations as other laser-based shape acquisition approaches. Especially challenging are surface discontinuities, which can make the lines that have been projected at different angles to appear as if they are continuous. Image segmentation is currently being considered to deal with this situation. By treating continuous areas in the image as separate entities, the entire image can be reassembled from the reconstructed entities. Once the geometry has been recovered, it will be possible to use the original texture on the imaged objects and estimate the position of the light source. We are currently working on these features.

Figure 2. Laser raster projected onto a person's face (left). 3D reconstruction of the lines (right).