CS348C – Winter 1997

Brian Curless
Stanford University

Range Acquisition
Part II: Optical Methods

Outline:

1. Introduction
2. Illumination
3. Imaging radar
4. Triangulation
5. Moire
6. Active stereo
7. Active depth–from–defocus

Optical range acquisition

Strengths:

- non–contact
- safe
- can be inexpensive
- fast (many points per second)

Limitations:

- Can only acquire visible portions of the surface
- Sensitive to surface properties:
  - transparency
  - shininess
  - abrupt changes in color
  - darkness (zero reflection)
- Can be confused by interreflections

Illumination

Why are lasers a good idea?

- Compact (laser diodes)
- Low power
- Single wavelength means the ability to discriminate laser from other lighting
- Single wavelength means no chromatic aberration
- Can be held in tight focus over long distance

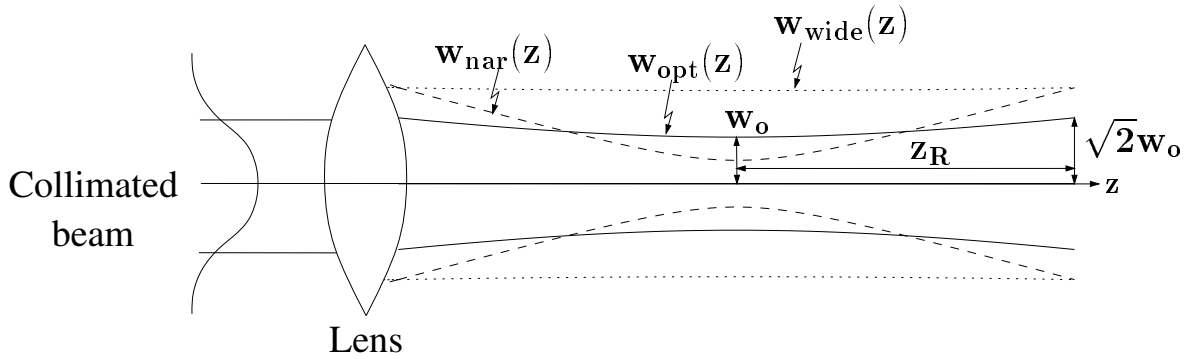

$$z_r = \frac{\pi w_o^2}{\lambda}$$

z_r = Rayleigh range (distance where $w = \sqrt{2}w_o$)
w_o = beam waist (narrowest laser width)
λ = wavelength of laser

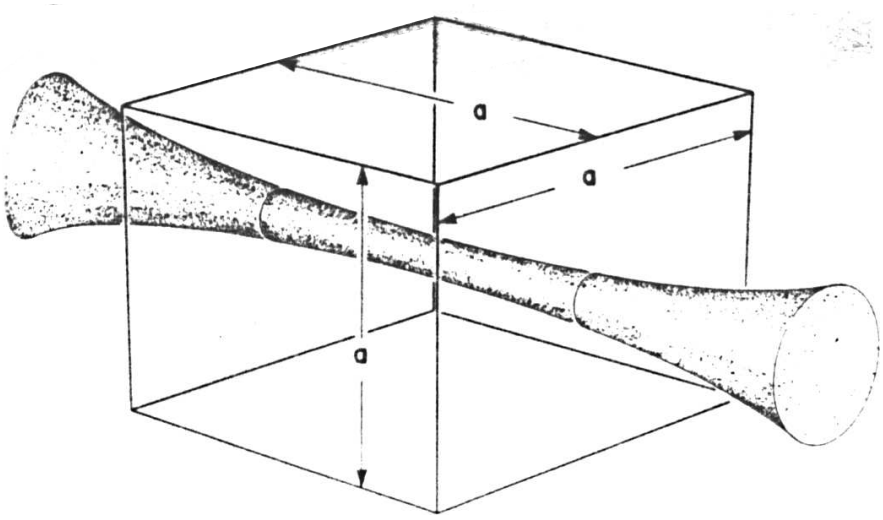

[Figure of 3D beam from Rioux 1987]

Limitations of lasers:

- Eye safety concerns
- Laser speckle adds noise – narrowing the aperture increases the noise

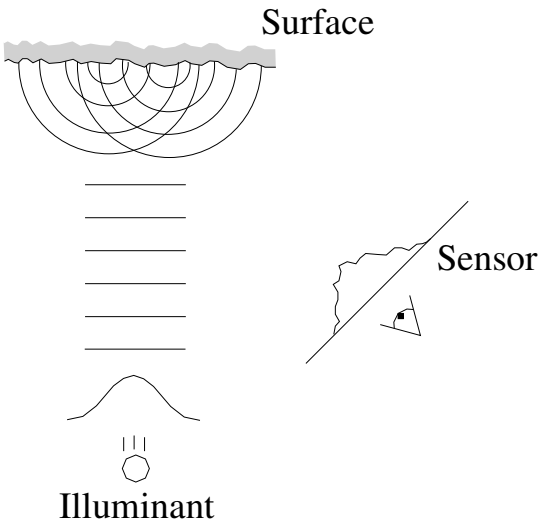

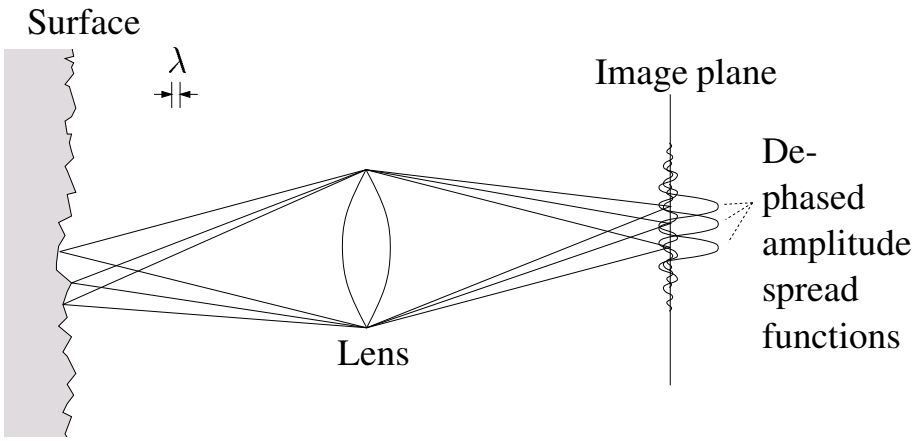

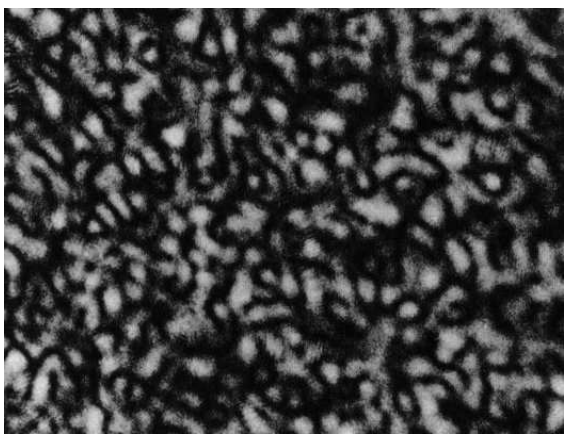

[Image of speckle from Goodman 1984]

Imaging radar: Time of flight

Principle:

A pulse of light is emitted, and the time of the reflected pulse ie measured:

c t = 2 r = roundtrip distance

Typical scanning configuration is two rotating mirrors to derive a range image.

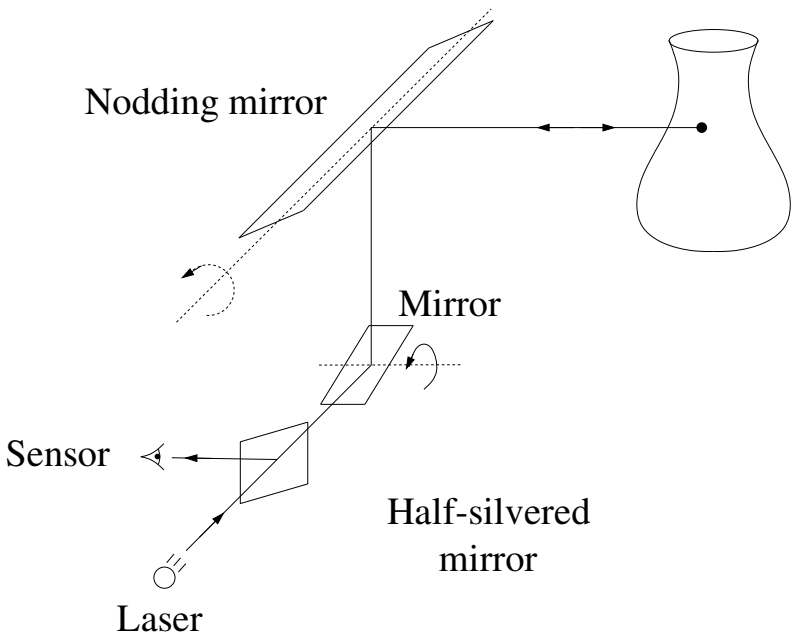

[The following statistics are for the *Cyra scanner*.]

Working volume:
50–100 meters range and about 60 degree field of view

Accuracy:
Targeted to achieve < 1 cm accuracy => 1 part in 10,000 or better.

Notes:
- 1 cm accuracy requires resolution down to ~10 picoseconds.
- Limitation is in detecting absolute time of flight – does not scale with smaller working volumes.
- show video...

Imaging radar: Amplitude Modulation (AM)

Principle:

A laser diode is driven at frequency:

$$f_{AM} = c / \lambda_{AM}$$

where c is the speed of light and λ_{AM} is the wavelength of the modulated signal. We can compute the range based on the phase difference between the outgoing and incoming signal:

$$r(\Delta\varphi) = 1/2 \; \lambda_{AM} \; (\Delta\varphi +/- 2\pi n)/2\pi$$

roundtrip correction

wavelength of signal

fraction of interval

Note the phase ambiguity due to the $+/- 2\pi n$. This translates to a range ambiguity:

$$r_{ambig} = n \; \lambda_{AM} / 2$$

Typical scanning configuration is two rotating mirrors to derive a range image.

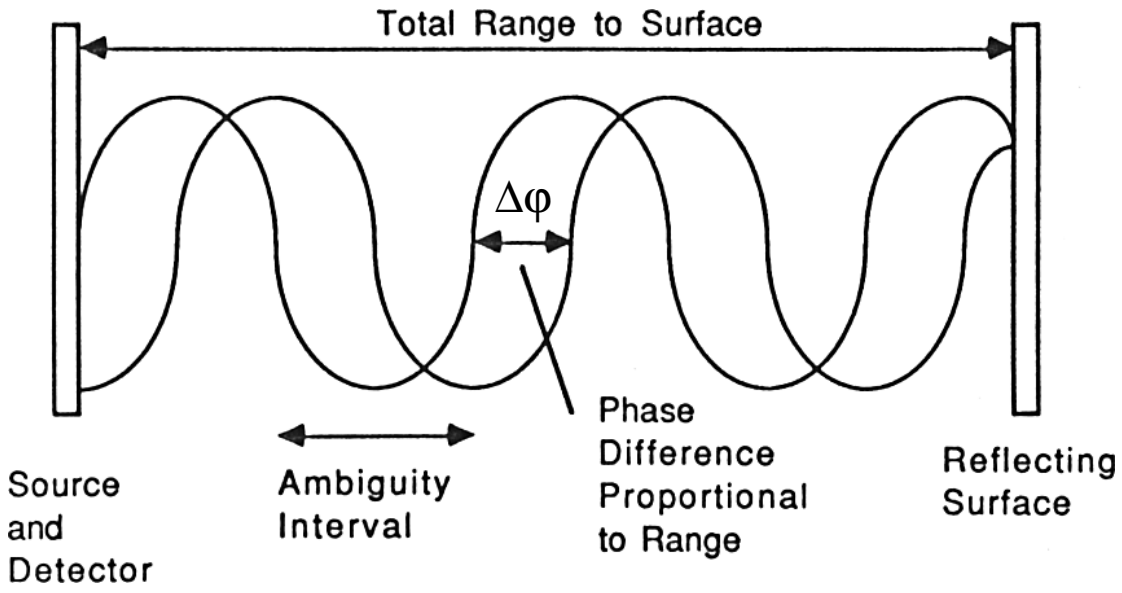

[Figure from Besl 1989]

Imaging radar: AM (cont'd)

[The following statistics are for the *Perceptron scanner* as reported by R. Pito at U. Penn.]

Working volume:
0.6 m

Accuracy:
1 mm => 1 part in 600

Notes:
– How can you resolve the phase ambiguity problem?

Optical triangulation

Principle:

A narrow beam of light strikes a surface, and the reflection is imaged by an off−axis sensor. The "center" of the imaged reflection approximately maps to the center of the illumination. The intersection of these two lines of sight correspond to a range point.

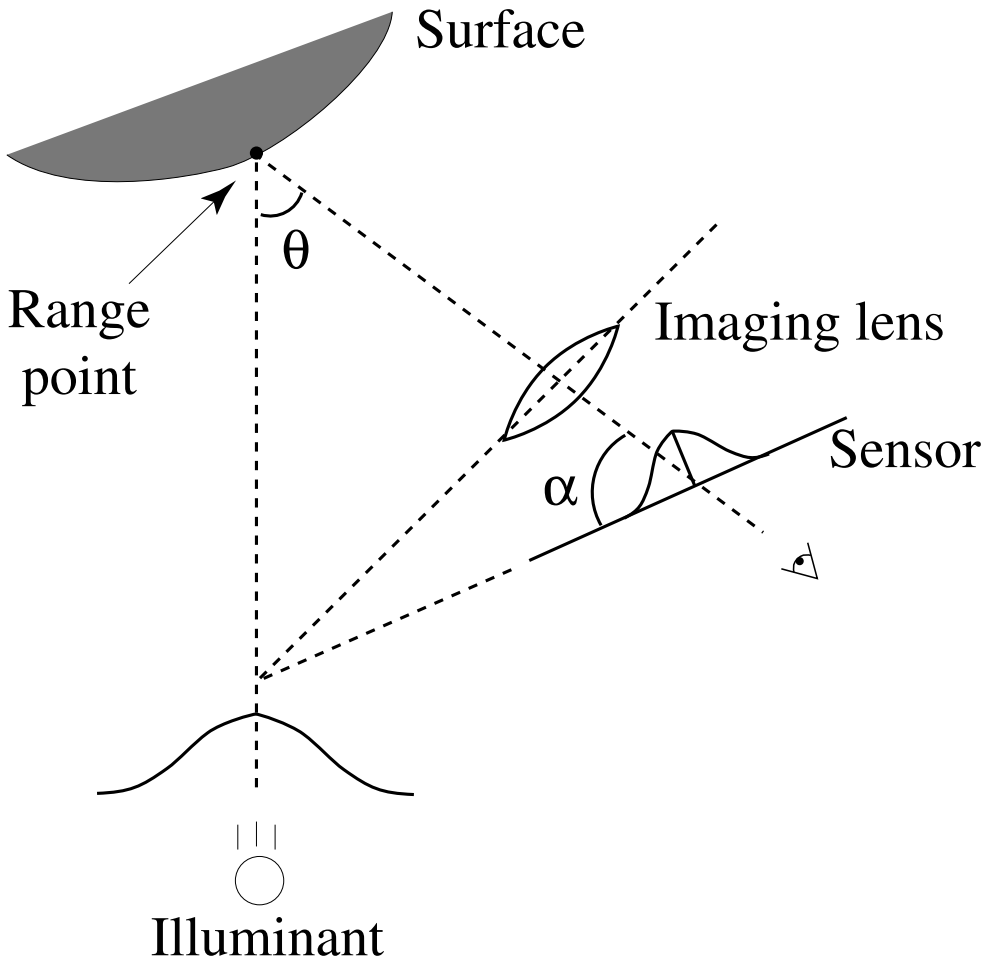

A lens maps a line to a line and a plane to a plane. If the line or plane is tilted, then the image plane is also tilted. The image tilt is related to the object tilt through the *Scheimpflug condition*:

$$\tan\alpha = 1/M \tan\theta$$

where M is the on−axis magnification (d_o/d_i for a thin lens).

Optical triangulation: scanning configurations

To scan a scene with a triangulation scanner, one approach is to keep the sensor stationary and sweep the laser across the scene. This approach has several problems:

- Loss of resolution due to defocus
- Field of view varies widely
- Resolution varies widely

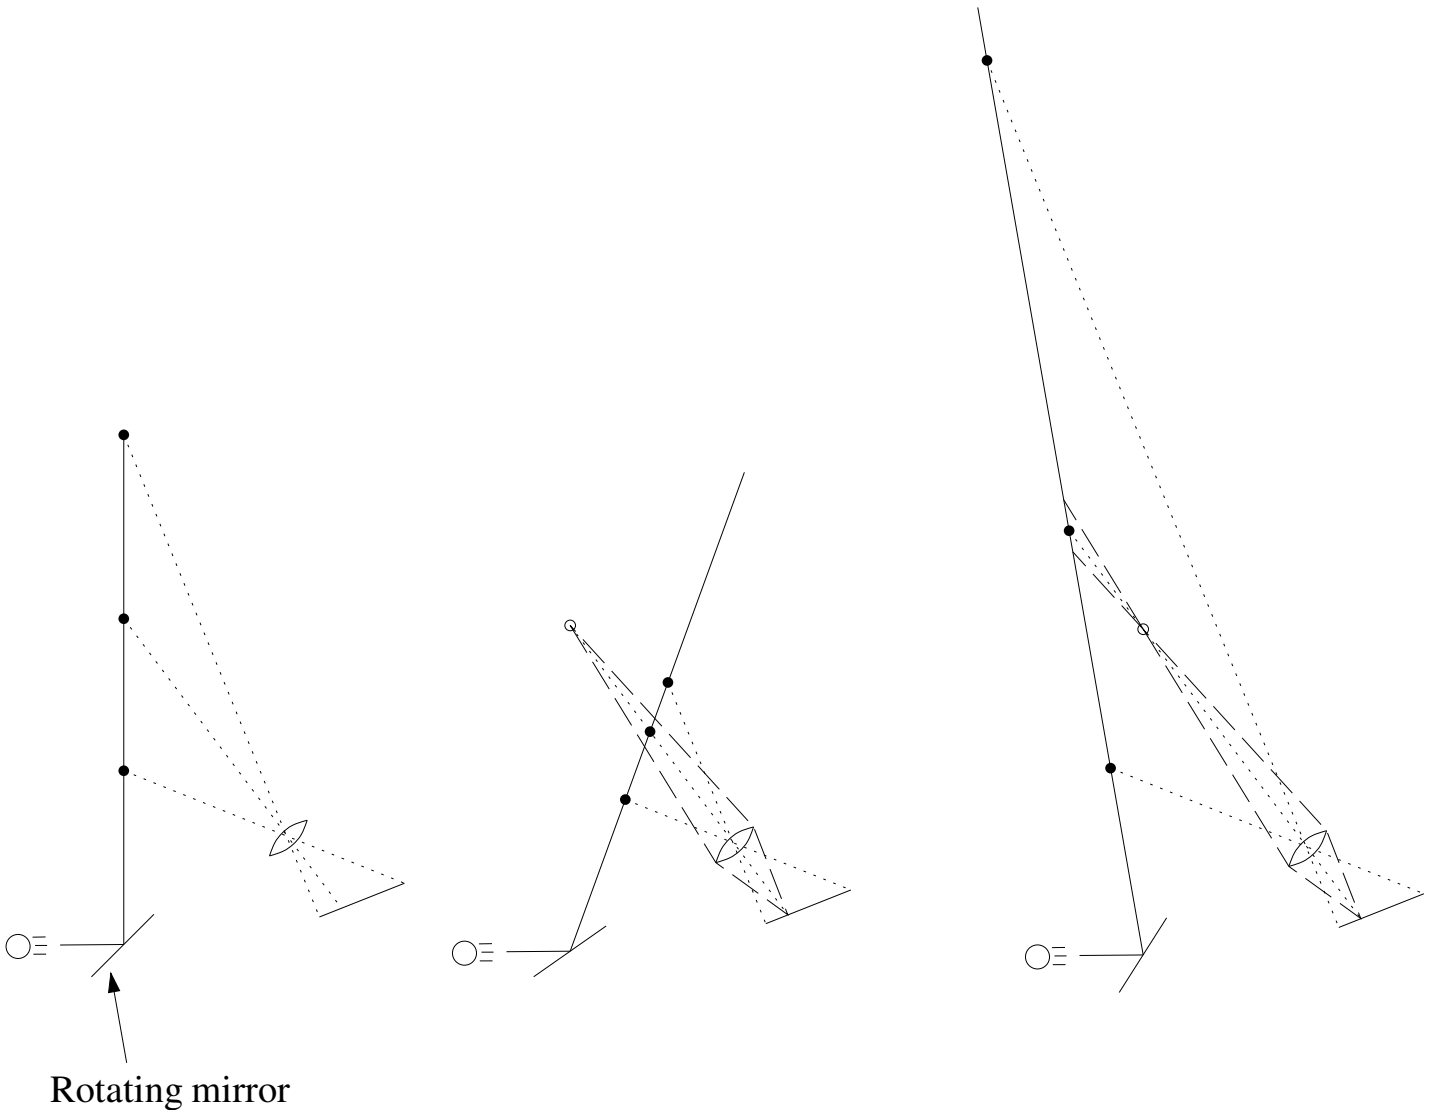

Instead of moving the laser independently of the sensor, we can move the laser and sensor as a unit. This will maintain focus and more consistent resolution.

Scanning is achieved by translating and/or rotating the triangulation unit.

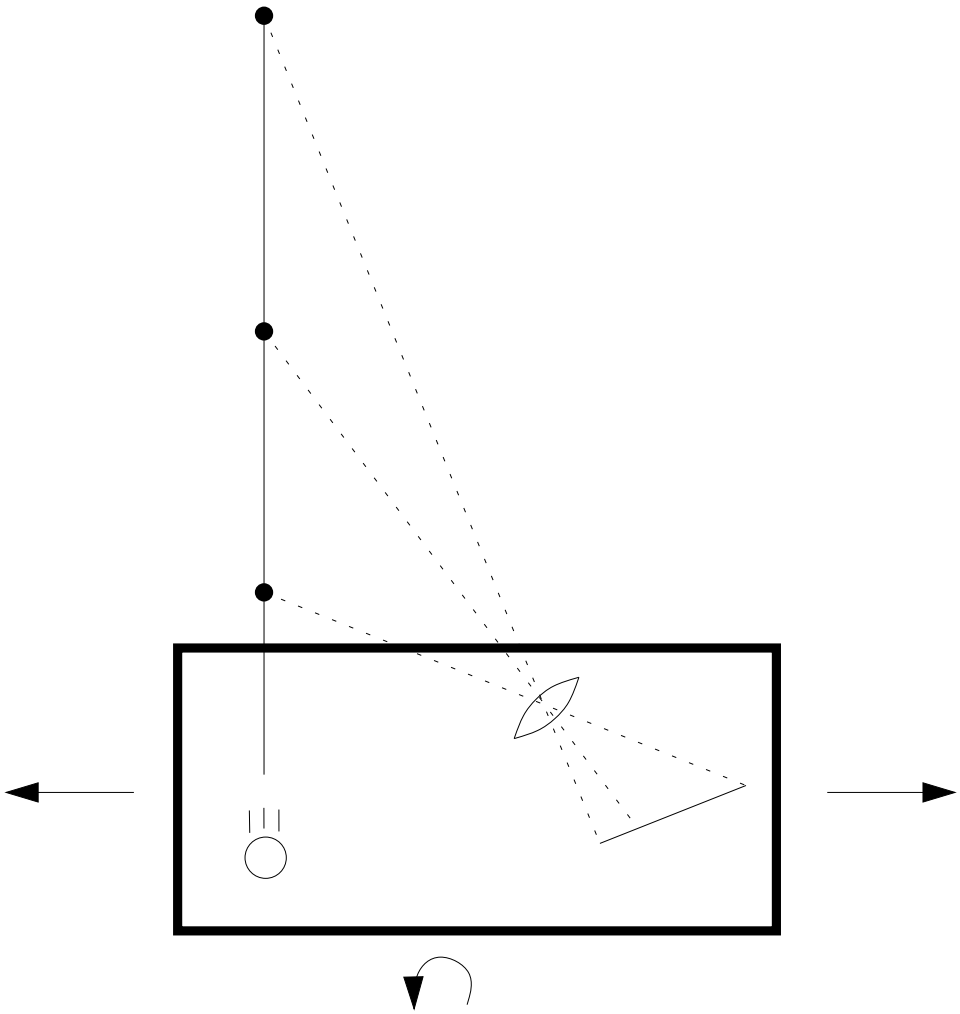

A novel design has been created (and patented) by Marc Rioux at the NRC in Canada. This design scans the laser and the sensor *simultaneously* over a scene.

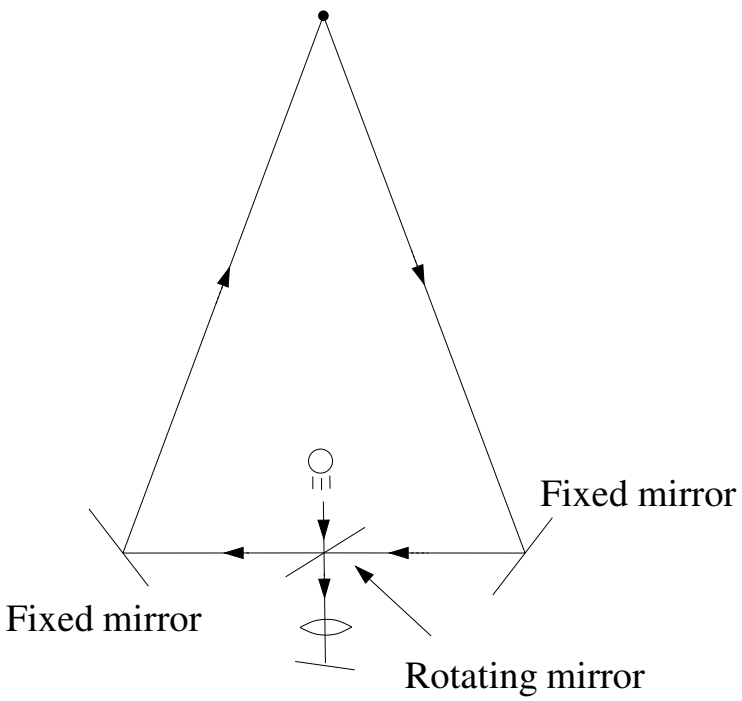

Rioux scanner

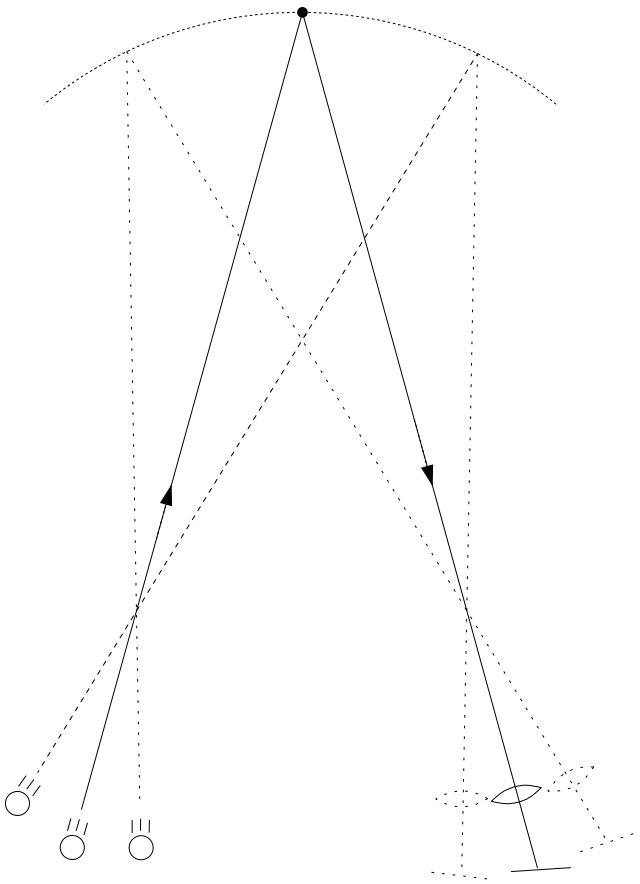

After unfolding the optics

Extension into 3D can be achieved by spreading the light beam into a plane of line to make a stripe scanner. Alternatively, a nodding mirror can be added to a flying spot scanner.

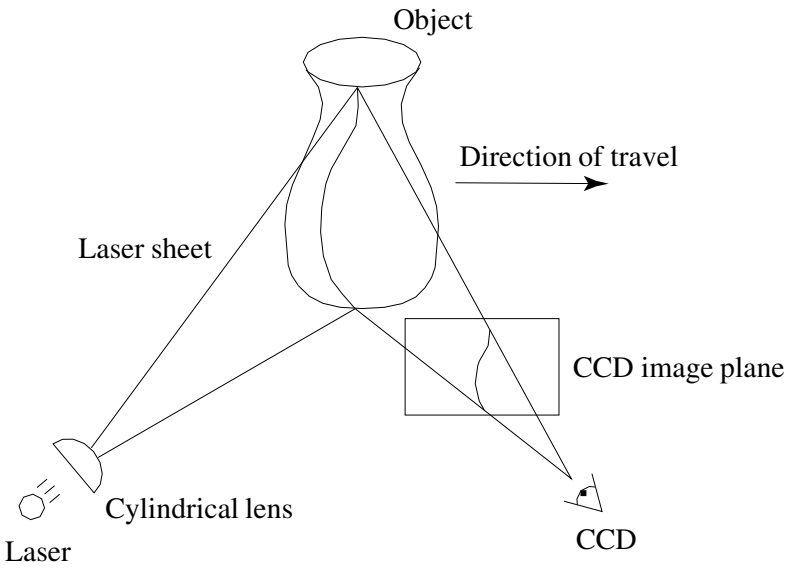

E.g., Cyberware scanner

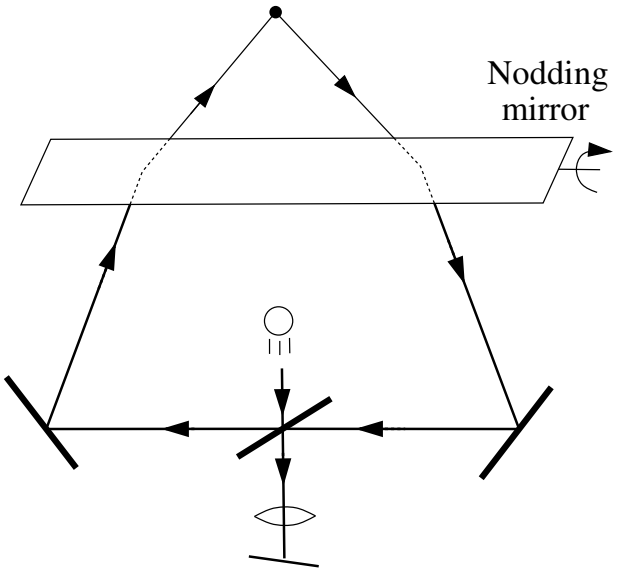

Rioux scanner

Cyberware specs for MS3030

Working volume:
30 cm

Triangulation angle:
30 degrees

Accuracy:
0.3 mm => 1 part in 1000

Speed:
15,000 pts/sec

(show video)

Hymarc (Rioux) scanner specs:

Working volume:
8 cm

Triangulation angle:
~15 degrees

Accuracy:
50 microns => 1 part in 1500

Speed:
10,0000 pts/sec

Multi−spot and multi−stripe triangulation

To achieve faster acquisition, some scanners use multiple spots or stripes. Such systems have the same limitations as described with moving illuminant and stationary sensor.

Still, they are attractive because of the potential for acquiring many range points in a short period of time.

There is, however, one additional problem: ambiguity.

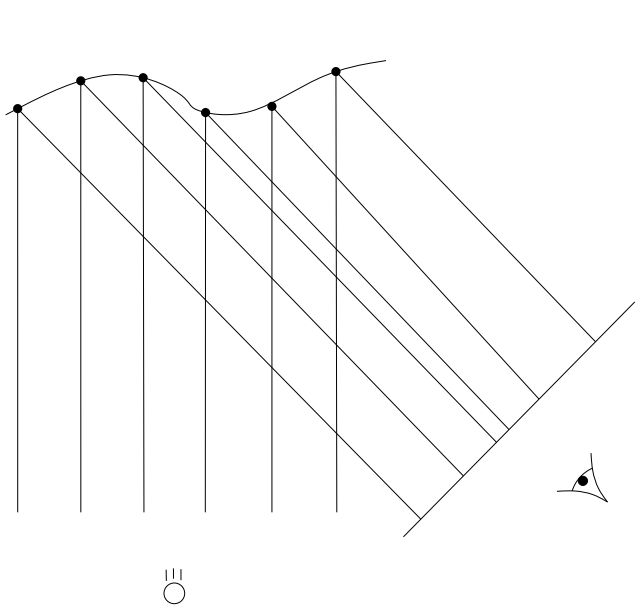

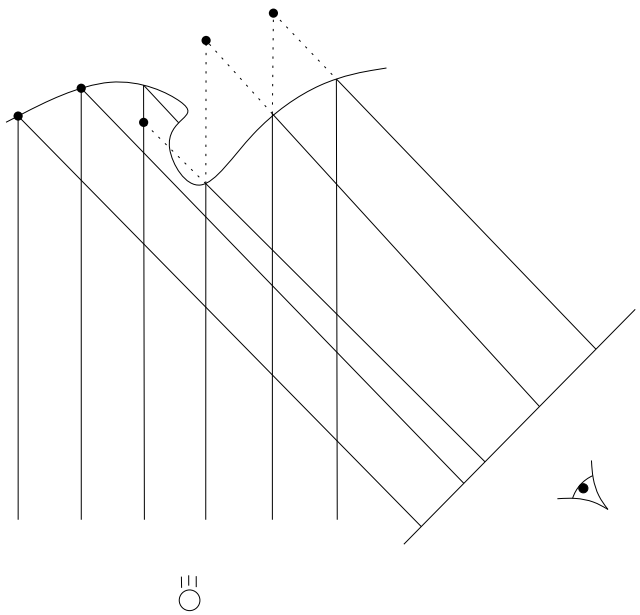

How can we resolve this ambiguity?

Binary coded multiple illumination

One approach to solving the ambiguity problem without projecting N stripes separately, is to project a binary pattern of light and dark regions. This method, illustrated below, resolves the ambiguity with $\log_2 N$ steps.

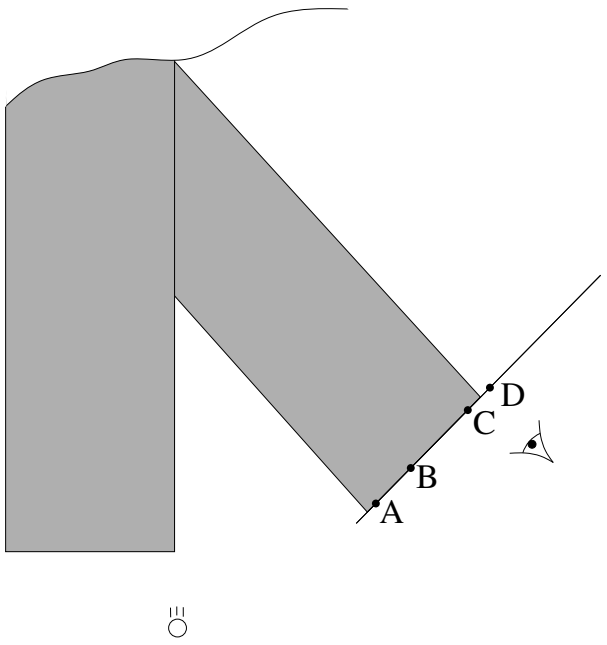

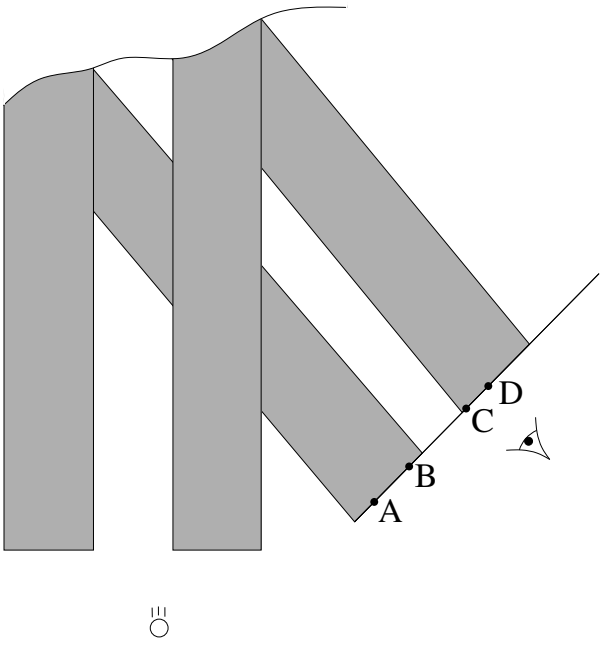

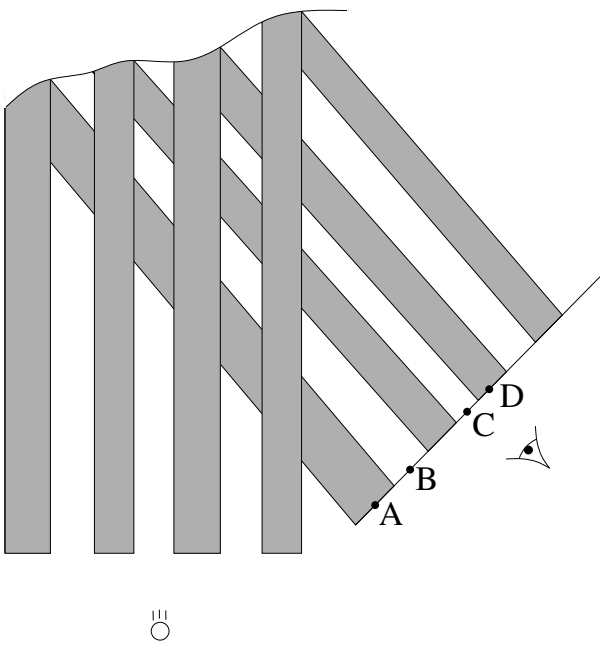

Binary codes:

A = 1 1 1
B = 1 1 0
C = 1 0 0
D = 0 1 1

Notes on triangulation

Scalable to very small working volumes.

Not easily scalable to large working volumes – baseline requirements become prohibitive.

Smaller triangulation angles mean less self–occlusion, but smaller angles mean trading off depth resolution at the back of the field of view.

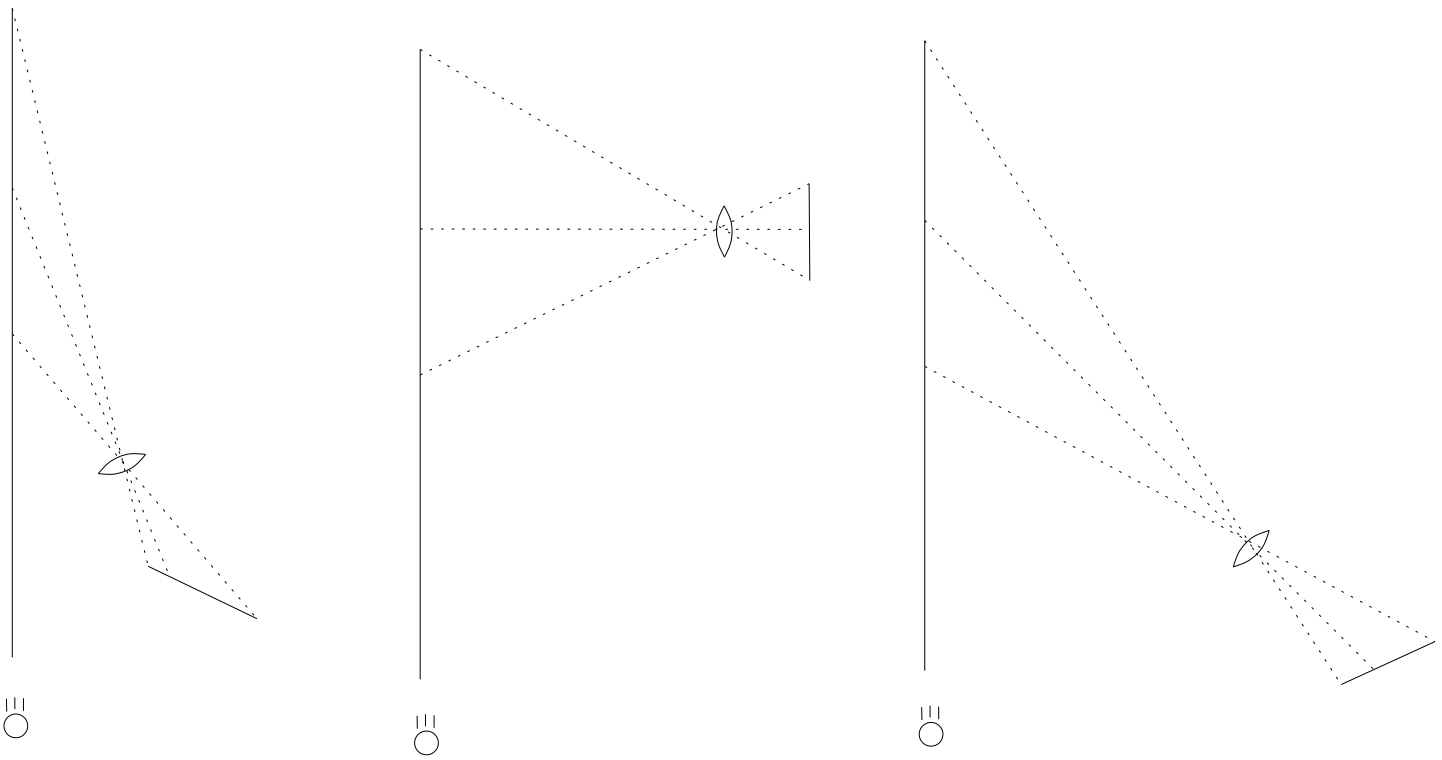

Triangulation Errors

Finding the center of the imaged light is the tricky part. If light reflects off of a planar, diffuse surface, then results are accurate (up to speckle noise).

If the surface exhibits variations in reflectance and shape, then the accuracy is ultimately limited by the width of the laser sheet... unless *spacetime analysis* is applicable.

Edges

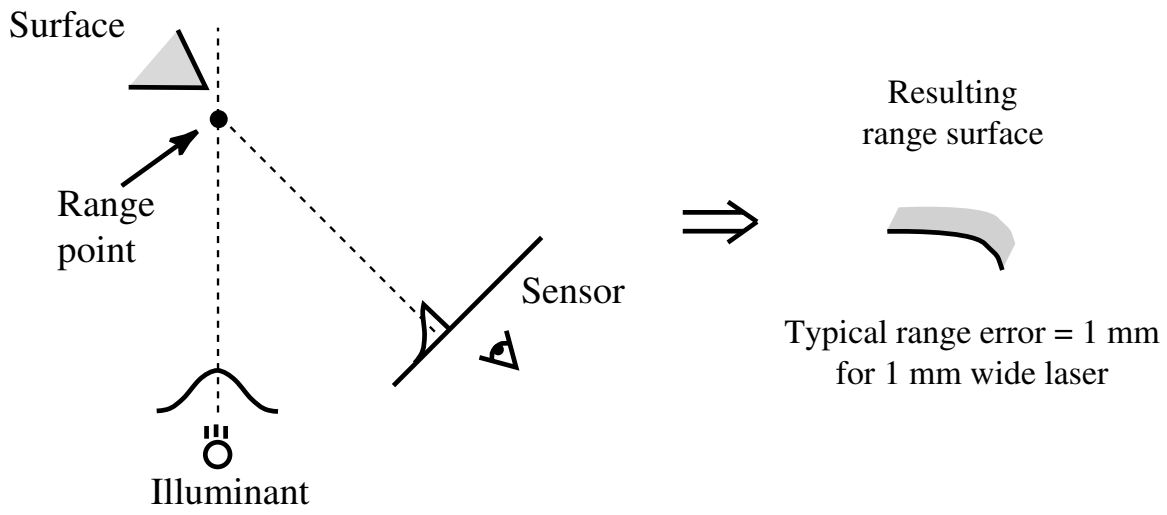

Shape variations

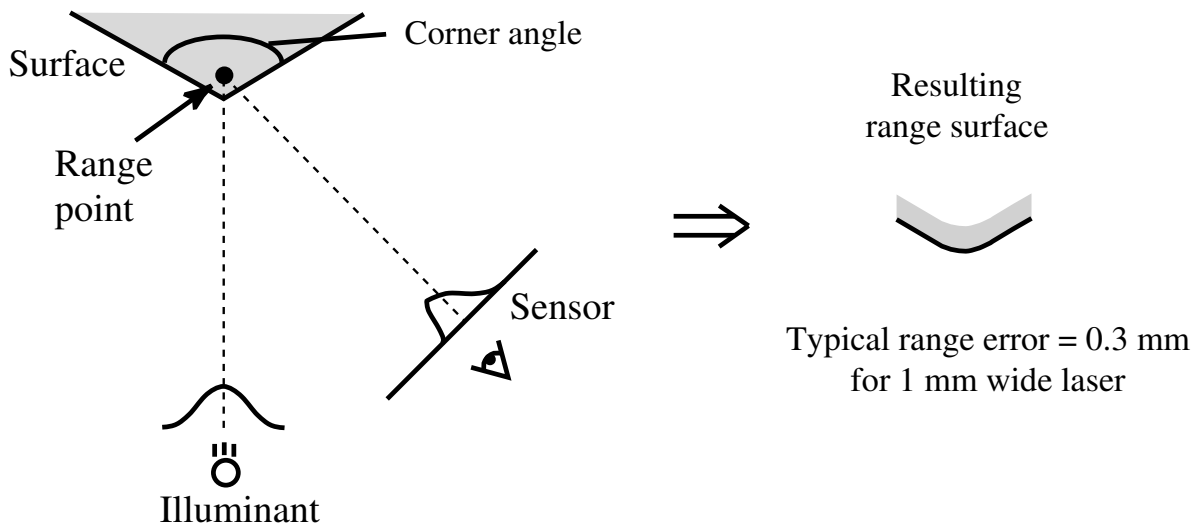

Triangulation Errors

Reflectance variations

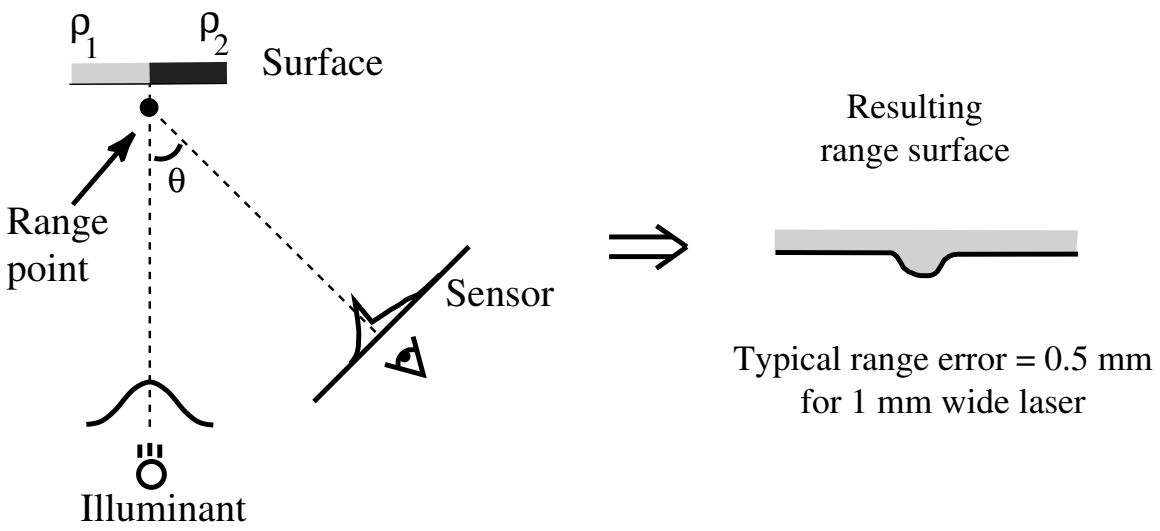

Laser speckle

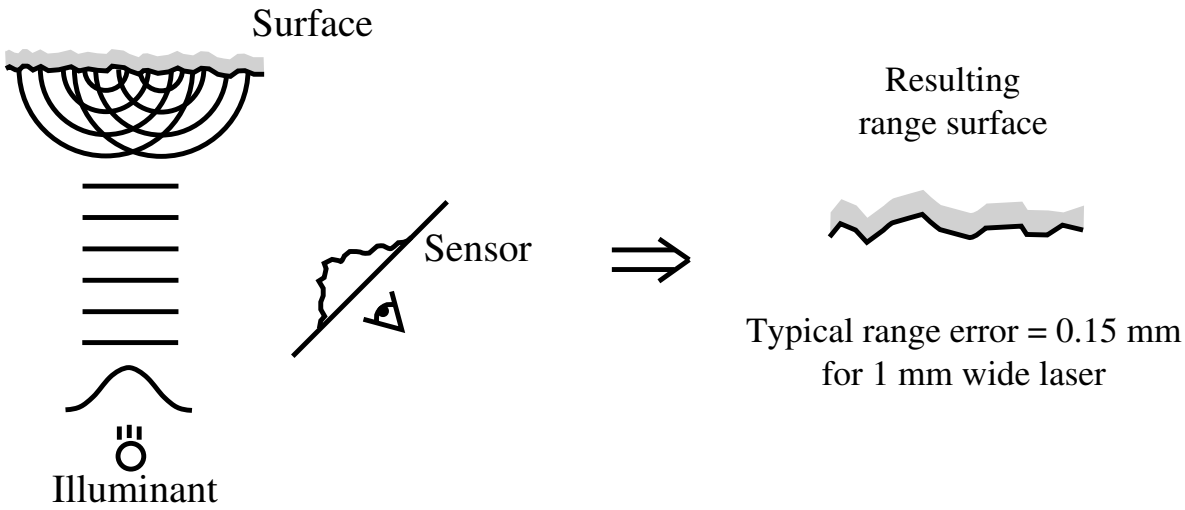

Errors and corrections through spacetime analysis

Reflectance correction

Reflectance card

Traditional analysis

Spacetime analysis

Edge curl reduction

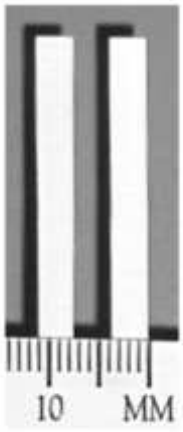

Two thin strips

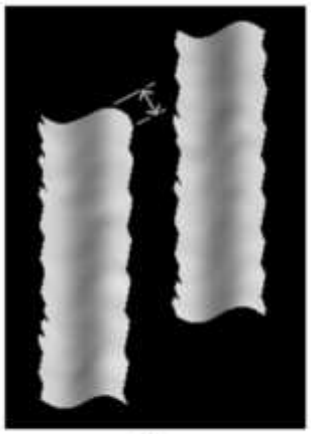

Traditional analysis

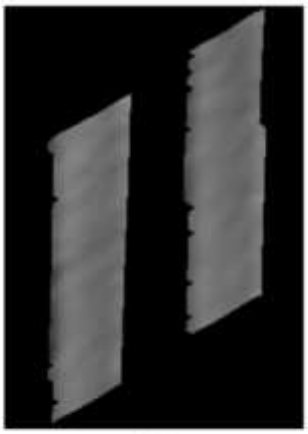

Spacetime analysis

Improved shape extraction

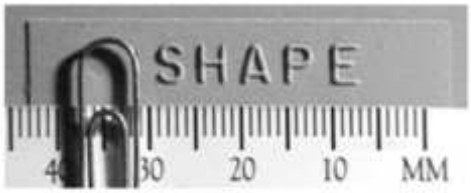

Shape ribbon

Traditional analysis

Spacetime analysis

Moire

Principle:

Illuminate a surface at an angle with a periodic grating such as a sinusoid. Image the reflection from a different angle. The reflected image is of the form:

$$I(x, y) = a(x, y) \cos[\, 2\pi f_o x + \phi(x, y)\,]$$

where:

a(x, y) is the amount of light reflected from the illuminant to the sensor modulated by the illumination function.

f_o is the period of the image of the grating when reflecting from a flat surface

$\phi(x, y)$ is the phase difference imparted by variations in the shape of the surface

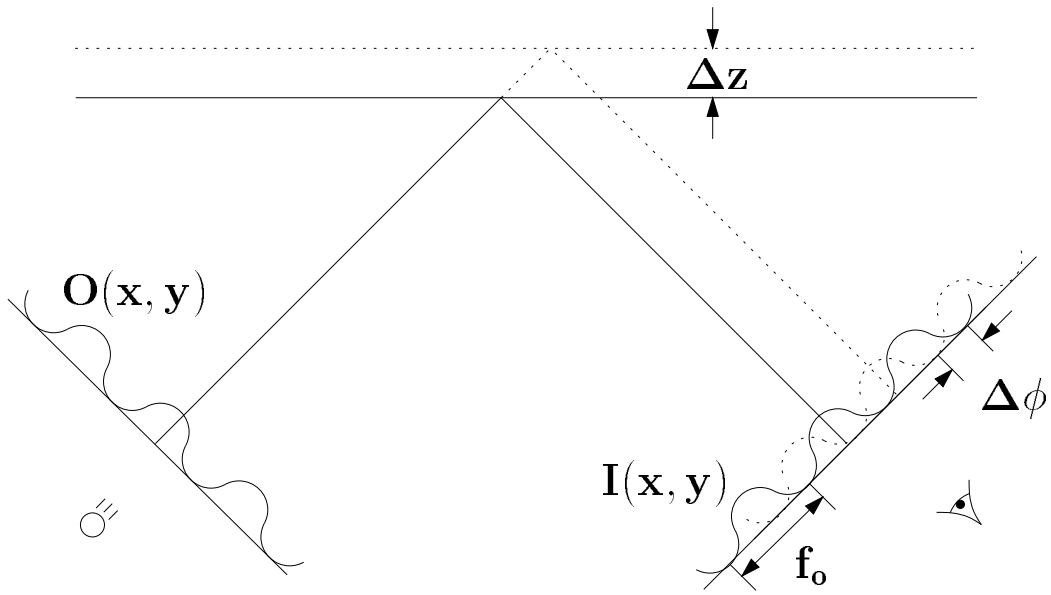

It is easy to show that:

$\Delta\phi(x, y) \sim \Delta z(x, y)$

Quadrature Multiplicative Moire

To recover the shape, we can "demodulate" the signal; i.e., extract the phase. To do so, we multiply by a sinusoid of the same frequency:

$$
\begin{aligned}
& I(x, y) \cos[\, 2\pi f_o x \,] \\
& \quad = a(x, y) \cos[\, 2\pi f_o x + \phi(x, y) \,] \cos[\, 2\pi f_o x \,] \\
& \quad = 1/2\ a(x, y) \{\cos[\, \phi(x, y) \,] + \cos[\, 4\pi f_o x + \phi(x, y)]\}
\end{aligned}
$$

Assuming a(x,y) and $\phi(x, y)$ are slowly varying, we can low−pass filter (LPF) the image:

$$M_1(x, y) = \mathrm{LPF}\{\, I(x, y) \cos[\, 2\pi f_o x \,] \,\} = 1/2\ a(x, y) \cos[\, \phi(x, y) \,]$$

Next, we multiply by a phase shifted sinusoid:

$$
\begin{aligned}
& I(x, y) \sin[\, 2\pi f_o x \,] \\
& \quad = a(x, y) \cos[\, 2\pi f_o x + \phi(x, y) \,] \cos[\, 2\pi f_o x + \pi/2 \,] \\
& \quad = 1/2\ a(x, y) \{\sin[\, \phi(x, y) \,] - \sin[\, 4\pi f_o x + \phi(x, y)]\}
\end{aligned}
$$

After filtering:

$$M_2(x, y) = \mathrm{LPF}\{\, I(x, y) \cos[\, 2\pi f_o x + \pi/2 \,] \,\} = 1/2\ a(x, y) \sin[\, \phi(x, y) \,]$$

Thus, we have

$$\tan[\, \phi(x, y) \,] = M_1(x, y) / M_2(x, y)$$

The phase corresponds to the range values to be extracted. In order for the phase term not to be filtered out, it cannot change rapidly. This requires that the surface be fairly smooth.

Example: Shadow Moire

We can place a grating near a surface and then illuminate it. When viewing it from a different angle, the grating is the demodulating signal and creates an interference pattern. Instant Moire!

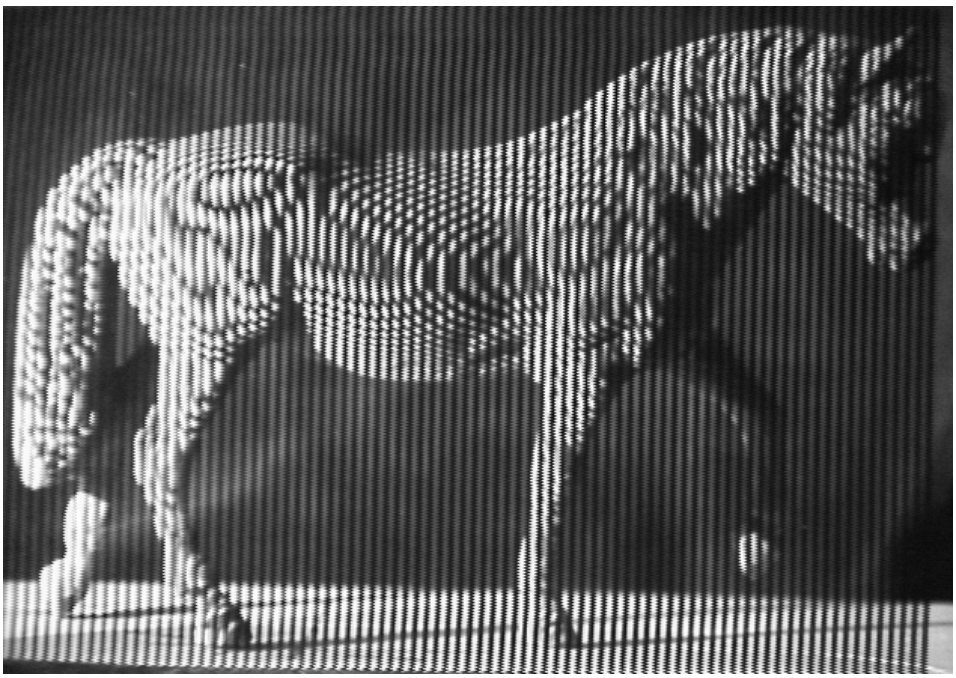

Shadow Moire

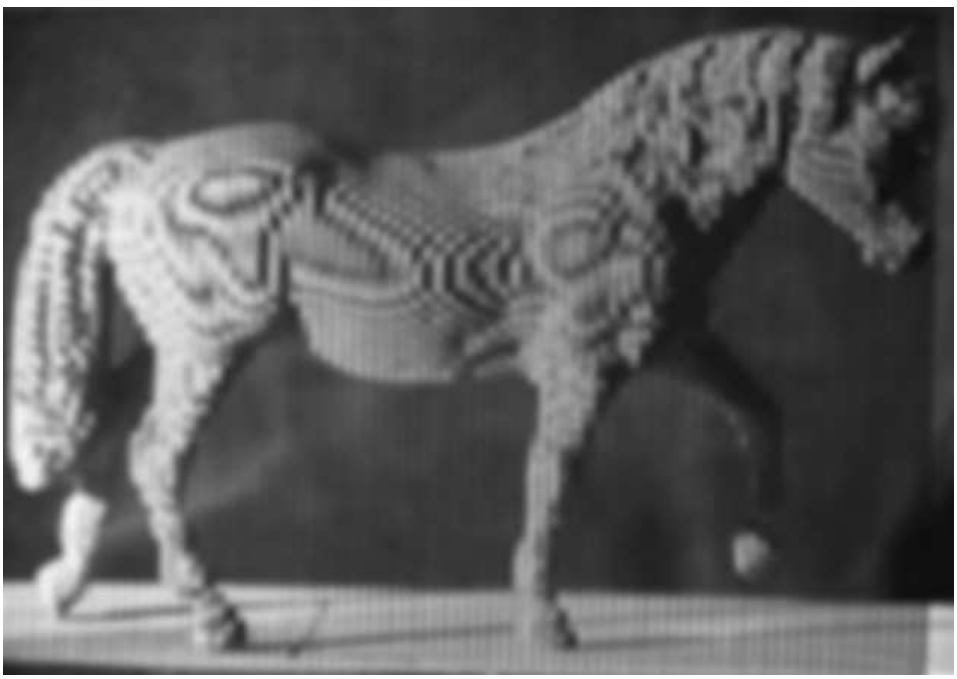

Filtered to extract the phase term

Active stereo

Principle:

Stereo methods use two offset cameras and match features in one image to features in another image. The lines of sight of matching features meet at a point on the surface.

Active stereo uses projected illumination to avoid the problem of finding textures.

Distinguishing accurate matches from false matches is not always robust.

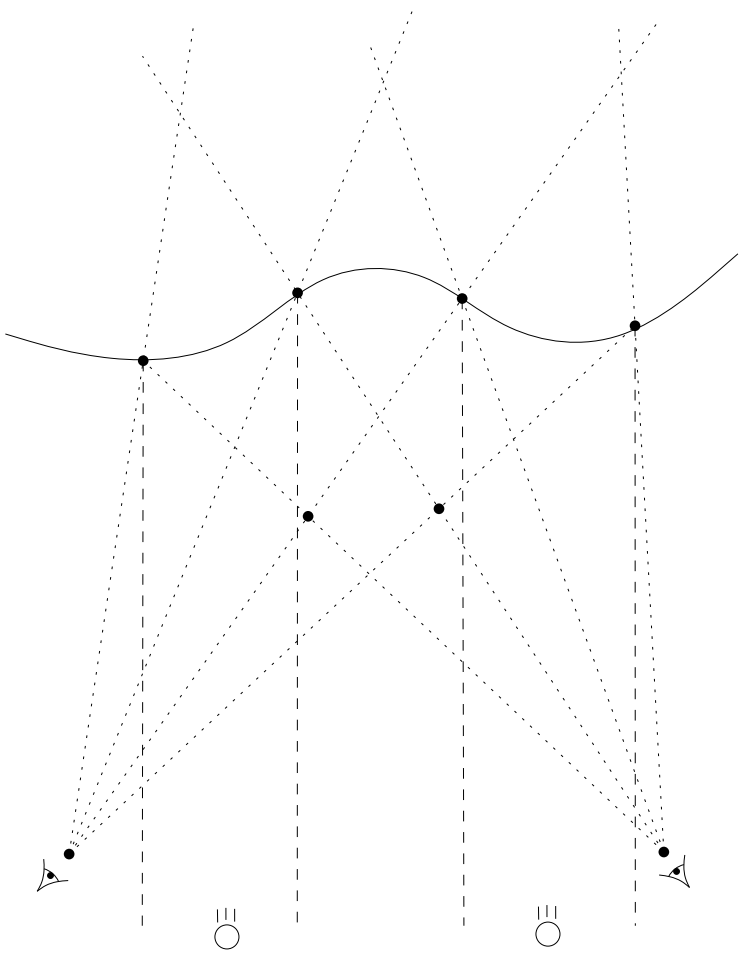

Note that for a stripe projected texture (shown here) we can limit matches to those sharing the same gradient directions on the image plane; i.e., don't match light−dark edges with dark−light edges.

Active multi–baseline stereo

Using multiple cameras helps reduce the likelihood of false matches.

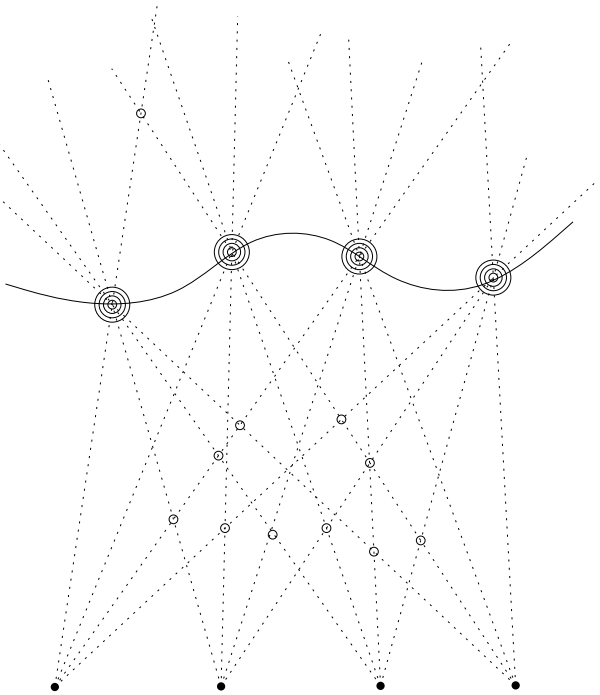

All possible matches

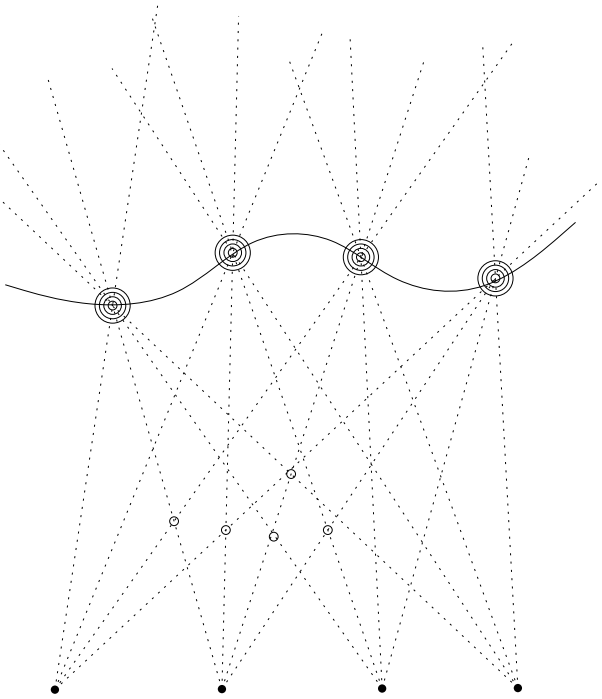

Matches against one camera

A new product has been released by Visual Interface, started by Jon Webb from CMU.

6 cameras
1 light stripe pattern projector
1 color camera

Working volume:
1 meter

Accuracy (typical):
1 mm => 1 part in 1000

[Image from Visual Interface web pages – http://www.visint.com]

Active depth from defocus

Principle:

When an imaging device is placed in a plane other than the image plane corresponding to an object point, the image of that point will be blurred. The amount of blur corresponds to the distance to the object.

Because a point will be equally blurry when the imaging device is at certain positions on either side of the image plane, at least two images will be necessary to make an unambiguous conclusion.

To avoid the problem of finding blurry textures already on the object, an active depth from defocus system projects light onto the scene.

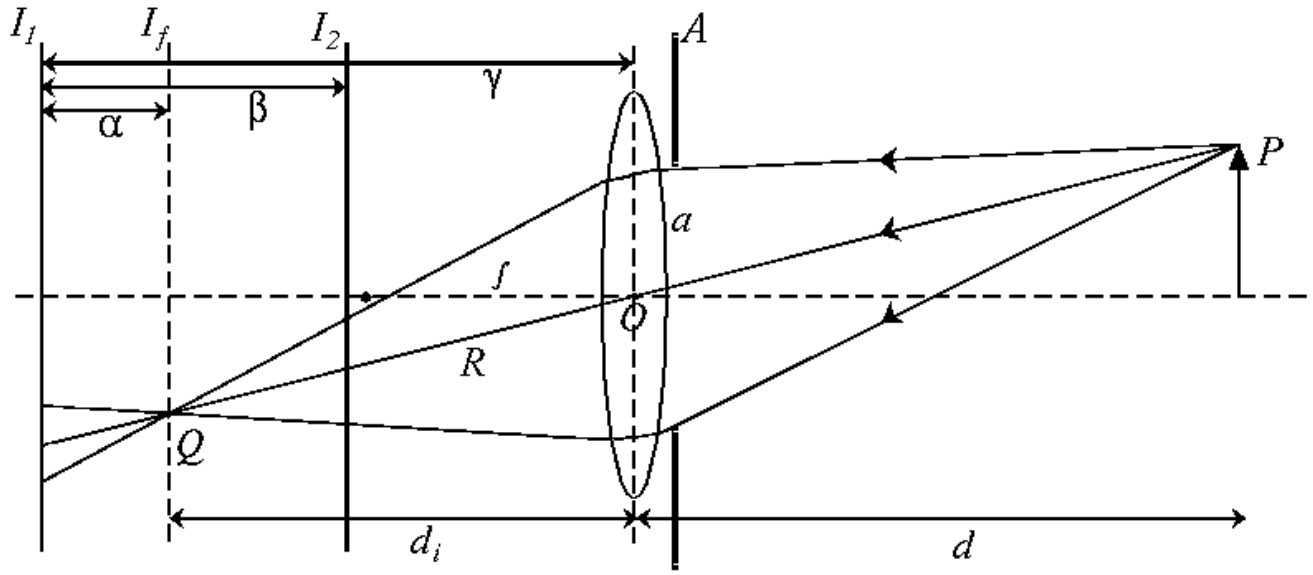

[Figure from Nayar 1995]

Example: Nayar's real–time system

Highlights:

- Telecentric optics for constant magnification between the image planes
- Checkerboard illumination pattern that follows the same optical path as the cameras – no occlusions
- Careful consideration of all the sampling and filtering issues
- Simple "focus operator" that returns the amount of defocus at each point
- Accuracy to 1 part in 300 over 30 cm
- Real–time!

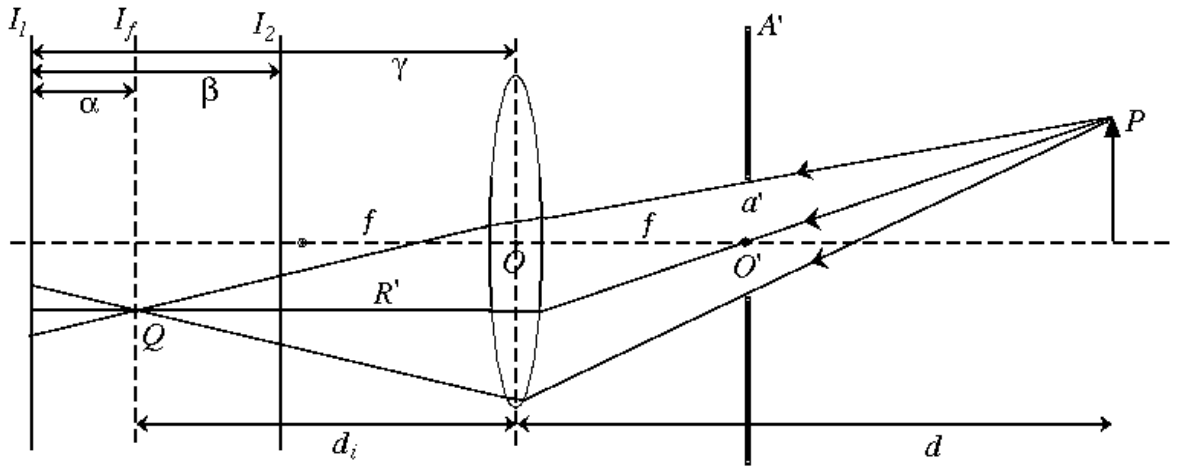

Telecentric optics

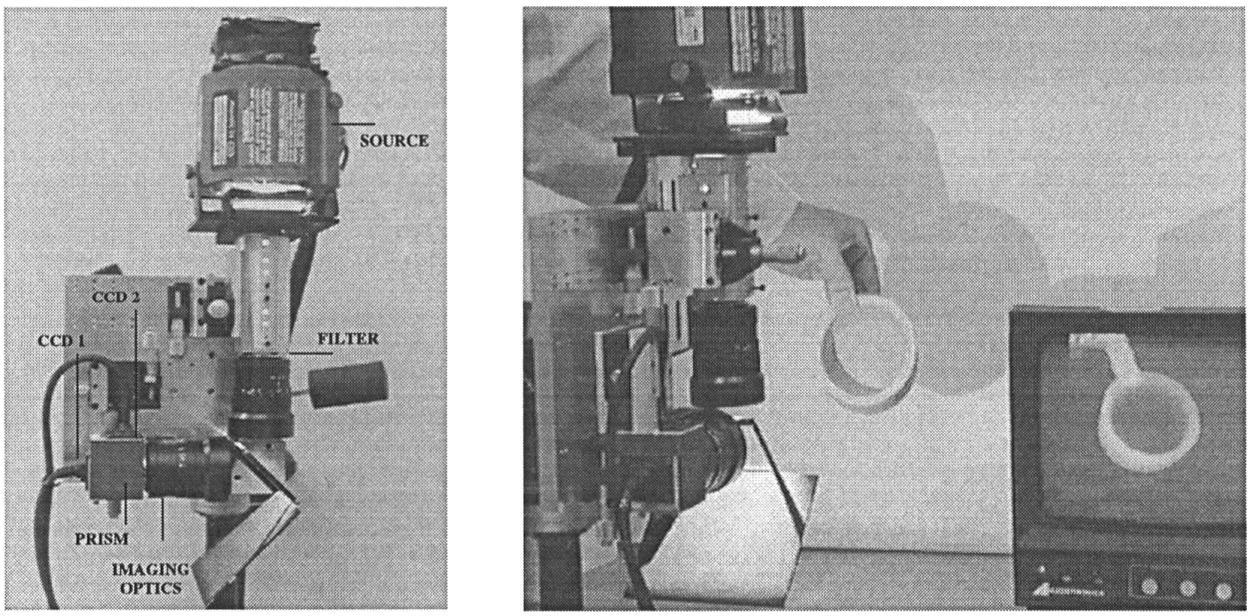

The hardware: 2 CCD's, 1 prism, 1 HS mirror, 2 identical lens systems

[Figures from Nayar 1995]

Bibliography

Besl, P. *Advances in Machine Vision. Chapter 1: Active Optical Range Imaging Sensors.* Sanz, J., editor, Springer−Verlag, New York, 1989.

Goodman, J. *Laser Speckle and Related Phenomena. Chapter 1: Statistical properties of laser speckle patterns*. Dainty, J., editor. Springer−Verlag, New York, 1984.

Nayar, S.K., Watanabe, M., and Noguchi, M. "Real−time focus range sensor", Fifth International Conference on Computer Vision (1995), pp. 995−1001.

Rioux, M., Bechthold, G., Taylor, D., and Duggan, M. "Design of a large depth of view three−dimensional camera for robot vision," *Optical Engineering* (1987), vol. 26, no. 12, pp. 1245−1250.