Chapter 3

Spacetime Analysis

Many optical triangulation scanners operate by sweeping the illumination over an object while imaging the reflected light over time. The previous chapter clearly demonstrates that analyzing the imaged pulses at each time step leads to systematic range errors. In this chapter, we describe a new method for optical triangulation called *spacetime analysis*. We show that by analyzing the time evolution of the imaged pulses, this new method can reduce or eliminate range errors associated with traditional methods.

We begin with an intuitive description of how the spacetime analysis works (Section 3.1), followed by a rigorous derivation of its properties under a set of assumptions such as linear scanning and orthographic sensing (Section 3.2). We then describe how the method generalizes when weakening or removing these assumptions (Section 3.3). In Section 3.4, we consider the role of laser speckle in spacetime analysis and show that it still imposes a limit on triangulation accuracy. Finally, we put optical triangulation scanners into a signal processing context and draw conclusions about how to improve them (Section 3.5).

3.1 Geometric intuition

Figure 3.1 illustrates the principle of spacetime analysis for a laser triangulation scanner with Gaussian illuminant and orthographic sensor as it translates across the edge of an object. As the scanner steps to the right, the sensor images a smaller and smaller portion of the laser cross-section. By time t_3, the sensor no longer images the center of the illuminant, and

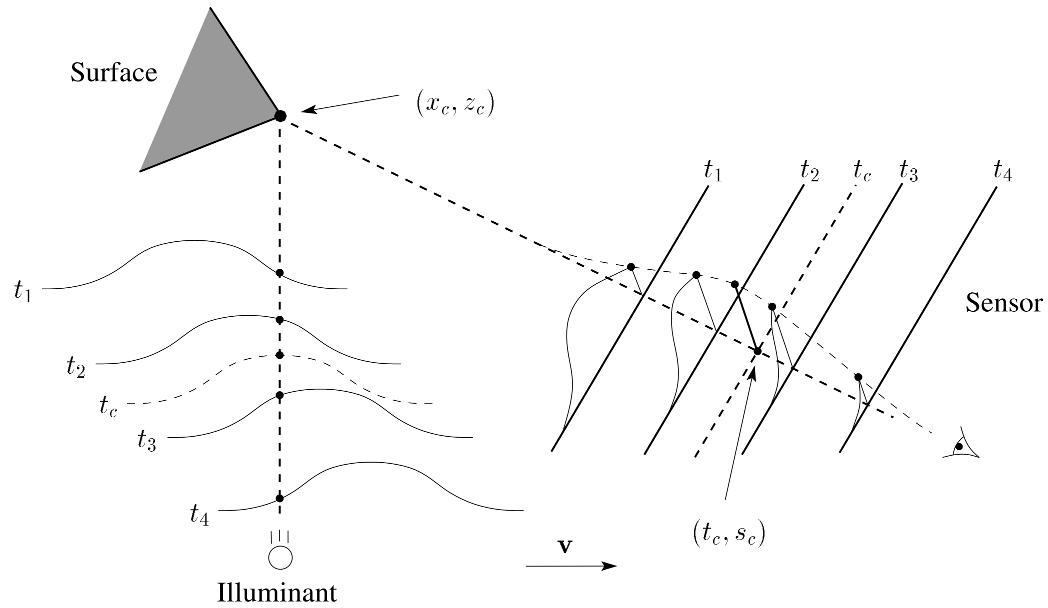

Figure 3.1: Spacetime mapping of a Gaussian illuminant. The Gaussians at left represent the profile of the illuminant at times t_1 through t_4. The solid Gaussians at right represent the images of the illuminant at these same times. The dashed Gaussian represents the time evolution of these images for a single line of sight, corresponding to a single point on the targeted surface. Although the solid Gaussians at right may be distorted or clipped by depth discontinuities or reflectance changes in the surface, the dashed Gaussian is always the same shape (unless it is missing entirely), and it is centered at a location that corresponds exactly to a range point. These properties of the dashed Gaussian are the keys to the spacetime analysis.

conventional methods of range estimation fail. However, if we look along the lines of sight from the corner to the laser and from the corner to the sensor, we see that the profile of the laser is being imaged *over time* onto the sensor (indicated by the dotted Gaussian envelope). Thus, we can find the coordinates of the corner point (x_c, z_c) by searching for the mean of a Gaussian along a constant line of sight through the sensor images. We can express the coordinates of this mean as a time and a position on the sensor, where the time is in general between sensor frames and the position is between sensor pixels. The position on the sensor indicates a depth, and the time indicates the lateral position of the center of the illuminant. In the example of Figure 3.1, we find that the spacetime Gaussian corresponding to the exact corner has its mean at position s_c on the sensor at a time t_c between t_2 and t_3 during the scan. We establish the corner's position by triangulating the center of the illuminant at time t_c

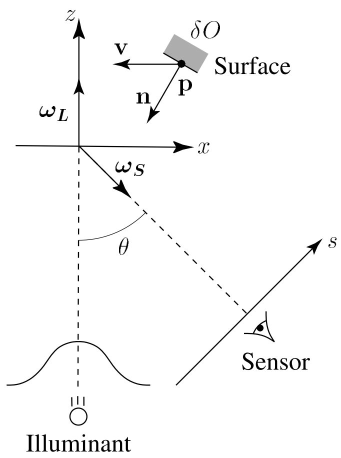

Figure 3.2: Triangulation scanner coordinate system.

with the sensor line of sight corresponding to the sensor coordinate s_c and the lateral sensor position at t_c.

3.2 A complete derivation

For a more rigorous analysis, we consider the time evolution of the irradiance from a translating differential surface element, δO, as recorded at the sensor. We refer the reader to Figure 3.2 and Table 3.1 for a description of coordinate systems and symbols. Note that in contrast to the previous section, the surface element is translating instead of the illuminant-sensor assembly. The element has a normal $\mathbf{n}$ and an initial position $\mathbf{p_o}$ and is translating with velocity $\mathbf{v}$, so that:

$$\mathbf{p}(t) = \mathbf{p_o} + t\mathbf{v}$$

Our objective is to compute the coordinates $\mathbf{p_o} = (x_o, z_o)$ given the temporal irradiance variations on the sensor. For simplicity, we assume that $\mathbf{v} = (-v, 0)$. The illuminant we consider is a laser with a unidirectional Gaussian radiance profile:

$$L_L(\mathbf{x}, \boldsymbol{\omega}) = I_L \exp\left[\frac{-2x^2}{w^2}\right] \delta(\boldsymbol{\omega} - \boldsymbol{\omega_L}) \tag{3.1}$$

δO	Differential surface element
$\mathbf{p}(t)$	Position of surface element
$\mathbf{p_o}$	Position of surface element (x_o, z_o) at time $t = 0$
$\mathbf{n}$	Normal of surface element
$\mathbf{v}$	Velocity of whole object
$\boldsymbol{\omega_L}$	Direction of illuminant
$\boldsymbol{\omega_S}$	View direction

Table 3.1: Symbol definitions for spacetime analysis.

where I_L is the power of the beam at its center, w is the e^{-2} half width, and the delta function indicates the directionality of the beam. In accordance with Figure 3.2, we will assume that $\boldsymbol{\omega_L} = (0, 1)$. The differential reflected radiance from the element in the direction of the sensor is simply:

$$dL(\mathbf{p}(t), \boldsymbol{\omega_S}) = f_r(\boldsymbol{\omega}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}|L_L(\mathbf{p}(t), \boldsymbol{\omega})d\omega \tag{3.2}$$

where $f_r(\boldsymbol{\omega}, \boldsymbol{\omega_S})$ is the bidirectional reflection distribution function (BRDF) of the point $\mathbf{p_o}$, evaluated at the incident direction, $\boldsymbol{\omega}$, and outgoing sensor direction, $\boldsymbol{\omega_S}$. Note that we assume that the BRDF is defined in a global sense for the surface element, rather than relative to its normal. Further, we treat incoming directions as pointing toward the surface rather than away from the surface, which is consistent with the description of the radiance field in terms of position and direction.

Substituting Equation 3.1 into Equation 3.2 and integrating over all incident directions, we find the total radiance reflected from the element to the sensor to be:

$$L(\mathbf{p}(t), \boldsymbol{\omega_S}) = f_r(\boldsymbol{\omega_L}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}|I_L \exp\left[\frac{-2(x_o - vt)^2}{w^2}\right] \tag{3.3}$$

Projecting the point $\mathbf{p}(t)$ onto the sensor, we find:

$$s = (x_o - vt)\cos\theta - z_o \sin\theta \tag{3.4}$$

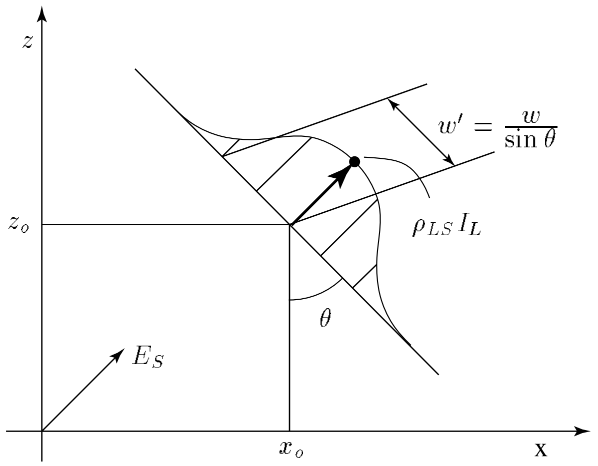

Figure 3.3: Spacetime image of the Gaussian illuminant. After scaling the sensor scanlines to map to depth and lateral displacement, we see that a point is imaged to a tilted Gaussian in the spacetime image. The center of the Gaussian maps to the coordinates of the point, the amplitude depends on the reflectance coefficient and beam strength, and the width is fixed by the scanner geometry.

where s is the position on the sensor and θ is the angle between the sensor and laser directions. We combine Equations 3.3-3.4 to give us an equation for the irradiance observed at the sensor as a function of time and position on the sensor:

$$E_S(t, s) = f_r(\boldsymbol{\omega_L}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}| I_L \exp\left[\frac{-2(x_o - vt)^2}{w^2}\right] \delta(s - (x_o - vt)\cos\theta - z_o \sin\theta)$$

To simplify this expression, we condense the light reflection terms into one measure:

$$\rho_{LS} = f_r(\boldsymbol{\omega_L}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}|$$

which we will refer to as the reflectance coefficient of point $\mathbf{p}$. We also note that $x = vt$ is a measure of the relative x-displacement of the point during a scan, and $z = s/\sin\theta$ is the relation between sensor coordinates and depth values along the center of the illuminant. Making these substitutions we have:

$$E_S(x, z) = \rho_{LS} I_L \exp\left[\frac{-2(x - x_o)^2}{w^2}\right] \delta((x - x_o)\cos\theta + (z - z_o)\sin\theta) \qquad (3.5)$$

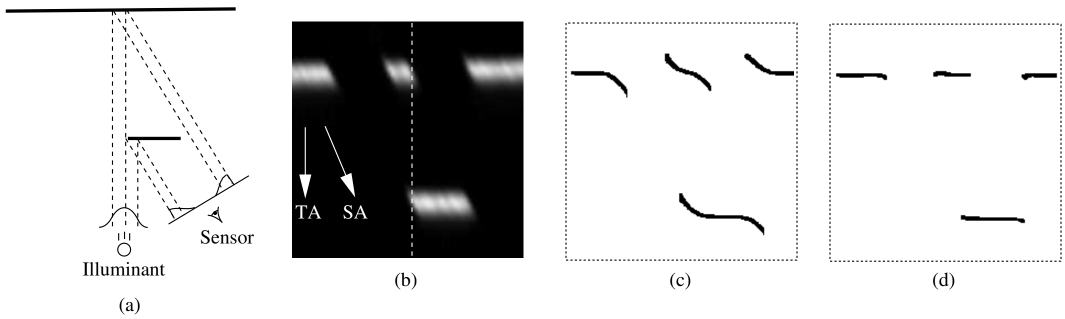

Figure 3.4: From geometry to spacetime image to range data. (a) The original geometry. (b) The resulting spacetime image. TA indicates the direction of traditional analysis, while SA is the direction of the spacetime analysis. The dotted line corresponds to the scanline generated at the instant shown in (a). (c) Range data after traditional mean analysis. (d) Range data after spacetime analysis.

This equation describes a Gaussian running along a *tilted line* through the spacetime sensor plane or *spacetime image* as depicted in Figure 3.3. We define the spacetime image to be the image whose columns are filled with sensor scanlines that evolve over time. Through the substitutions above, position within a column of this image represents displacement in depth, and position within a row represents time or displacement in lateral position. From Figure 3.3, we see that the tilt angle is θ with respect to the z-axis, and the width of the Gaussian along the line is:

$$w' = \frac{w}{\sin \theta}$$

The peak value of the Gaussian is $\rho_{LS} I_L$, and its mean along the line is located at (x_o, z_o), the exact location of the range point. Note that the angle of the line and the width of the Gaussian are solely determined by the fixed parameters of the scanner, *not* the position, orientation, or BRDF of the surface element. Thus, extraction of range points should proceed by computing low order statistics along tilted lines through the spacetime image, rather than along columns (scanlines) as in the conventional method. Figure 3.4 illustrates this distinction. Further, we establish the position of the surface element independent of the orientation and BRDF of the element and, assuming no interreflections, independent of any other nearby surface elements. The decoupling of range determination from local shape and reflectance is *complete*. As a side effect, the peak of the Gaussian yields the irradiance at the sensor due to the point. Thus, we automatically obtain an intensity image precisely registered to the range image, information which can assist in a task such as object segmentation.

3.3 Generalizing the geometry

We can easily generalize the previous results to other scanner geometries under the following conditions:

- The illumination direction is constant over the path of each range point.
- The sensor is orthographic.
- The motion is purely translational.
- The illumination and sensing directions are different.
- The illumination and translation directions are different

The first three conditions ensure that the reflectance coefficient, $\rho_{LS} = f_r(\boldsymbol{\omega_L}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}|$, is constant, the fourth condition guarantees that actual triangulation can occur, and the fifth condition ensures that the illumination is scanning *across* points in the scene. Note that the illumination need only be directional; coherent or incoherent light of any pattern is acceptable, though a narrow pattern will avoid potential depth of field problems for the sensor optics. In addition, the translational motion need not be of constant speed, only constant direction. We can correct for known variations in speed by applying a suitable warp to the spacetime image.

We can weaken the need for orthography if we assume that the BRDF does not vary appreciably over the sensor's viewing directions. The range of viewing directions depends on the focal length of the lens system and the range of positions a single point can occupy while being illuminated during a scan, which in turn depends on the width of the illuminant. In the general perspective case, the spacetime analysis proceeds along curves rather than straight lines. If the image of a point traverses a small displacement on the sensor relative to the focal length, then assumption of local orthography is reasonable. In addition, general perspective can lead to changes in occlusion relationships during a scan that cause an illuminated point to be visible for a finite time and then become invisible while still illuminated. This effect results in a partial image of the occluded point's spacetime Gaussian and will result in a range error for this point. These points will be easy to identify: they will be near step discontinuities in the resulting range image. When acquiring multiple range images, these points can be eliminated in favor of unoccluded acquisitions of the same portions of the surface. Thus, while perspective sensors can create some complications in applying the spacetime analysis, typical lens configurations will yield reasonable results.

We can also weaken the requirement for pure translation, as long as the rotational component of motion represents a small change in direction with respect to the $f_r(\boldsymbol{\omega_L}, \boldsymbol{\omega_S})|\mathbf{n} \cdot \boldsymbol{\omega_L}|$ product. This constraint is reasonable for motion trajectories with large radii of curvature relative to the beam width. As in the perspective case, the spacetime analysis proceeds in general along curves, and changes in occlusion relationships can lead to erroneous range data. Nonetheless, for moderate rotational trajectories, the spacetime analysis will lead to improved range data.

3.4 Influence of laser speckle

In the previous chapter, we developed the equations that are needed to analyze the influence of laser speckle on range determination. The results of statistical analysis show that laser speckle places a limit on the attainable accuracy when using traditional triangulation. Does spacetime analysis eliminate this limitation? One encouraging observation is the fact that the imaging equation (2.4) is space-invariant, so that moving the object field corresponds to moving the image field. In this way, speckle behaves as though it adheres to the surface, much like a noisy reflectance variation. This observation would seem to indicate that the spacetime analysis would correct for speckle noise. However, when we move the object with respect to a non-uniform illuminant (such as a laser beam with Gaussian cross-section), then the object field is *not* simply translating. As a result, the speckle does not completely adhere to the surface, and spacetime analysis faces the same limitation as traditional triangulation: laser speckle imposes a fundamental limit on accuracy. In this section, we will develop some of the equations necessary to demonstrate this result.

For the spacetime analysis, we must consider a moving surface and follow a single point on the surface to observe the image of the Gaussian mapped over time. Conversely, we can consider a moving illuminant, and follow a single line of sight from a stationary sensor. These two cases are equivalent as far as the resulting speckle patterns are concerned, but the latter yields a simpler analysis.

We can describe the field at the surface of the object in terms of the moving illuminant:

$$\mathbf{U}_o(x_o, t) = G_o(x_o - v_o t)\exp[i\phi_o(x_o)]$$

where the Gaussian is moving with velocity v_o. Mapping this to the image plane:

$$\mathbf{U}_i(x_i, t) = \{G_g(x_i - v_i t)\exp[i\phi_g(x_i)]\} * h(x_i)$$

where G_g and ϕ_g are the images of the Gaussian and phase functions given by geometric optics, and v_i is the velocity of the imaged Gaussian. To simplify this relation, we can first expand the convolution:

$$\mathbf{U}_i(x_i,t) = \int_{-\infty}^{\infty} G_g(x_i - v_i t - \alpha)\exp[i\phi_g(x_i - \alpha)]h(\alpha)d\alpha$$

Next, we consider a single line of sight, $x_i = 0$, so that:

$$\mathbf{U}_i(0,t) = \int_{-\infty}^{\infty} G_g(-v_i t - \alpha)\exp[i\phi_g(-\alpha)]h(\alpha)d\alpha$$

For a symmetric amplitude spread function, and after some straightforward manipulations, this relation becomes:

$$\mathbf{U}_i(0,t) = G_g(\eta) * \{\exp[i\phi_g(\eta)]h(\eta)\} \mid_{\eta=v_i t} \tag{3.6}$$

As described in the previous chapter, the intensity is the square magnitude of the field, but in this case, the mean we seek is with respect to *time*:

$$c_t = \frac{\int_{-\infty}^{\infty} t I_i(0,t)dt}{\int_{-\infty}^{\infty} I_i(0,t)dt} \tag{3.7}$$

The error can be solved as the variance of the mean as described in the previous chapter:

$$\sigma_{c_t} = [\langle c_t^2 \rangle - \langle c_t \rangle^2]^{\frac{1}{2}} = \langle c_t^2 \rangle^{\frac{1}{2}} \tag{3.8}$$

Figure 3.5 illustrates the influence of laser speckle when using spacetime triangulation. For traditional analysis, the uncertainty in the position of the mean on a sensor scanline corresponds directly to a depth error along the center of the illuminant. For spacetime triangulation, the error in the mean maps to an error in determining *when* the laser sheet's center intersects a line of sight at a range point. The range error is distributed along the sensor's line of sight, and obeys the relation:

$$\sigma_r = \frac{v_o \sigma_{c_t}}{\sin\theta} \tag{3.9}$$

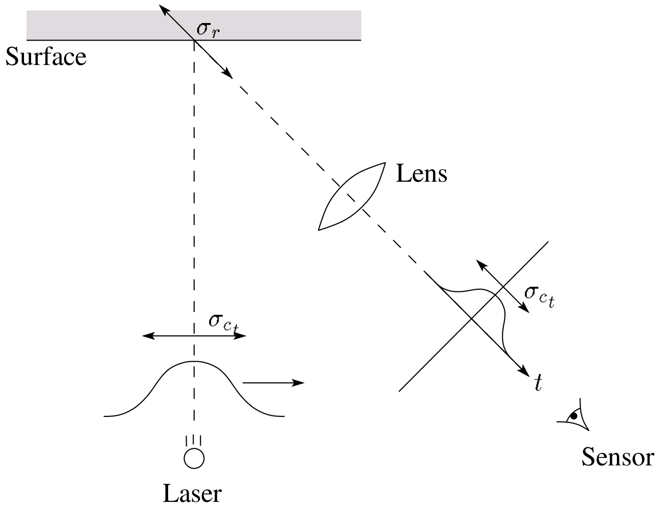

Figure 3.5: Influence of laser speckle on spacetime triangulation. When performing the spacetime analysis, the line of sight to the point is fixed, but the time at which the center of the laser passes over the range point varies with range. Laser speckle introduces uncertainty, σ_{c_t} as to when the laser is centered over the range point. The effect is range uncertainty, σ_r, along the line of sight from the sensor.

As in the case for traditional analysis, computing an analytical solution using Equations 3.6-3.9 has proven to be extremely difficult. Instead, we have performed numerical simulations that lead us to the same result as we found for traditional analysis (Equation 2.11):

$$\sigma_r \approx \frac{1}{2\sin\theta}\sqrt{\frac{w\lambda d_o}{2a}} \tag{3.10}$$

Thus, the error in spacetime triangulation under coherent illumination is independent of sensor resolution and increases with the narrowing of the aperture and the widening of the laser sheet. In addition, numerical analysis reveals that, just as in the case for traditional triangulation, smoother surfaces lead to reductions in range error.

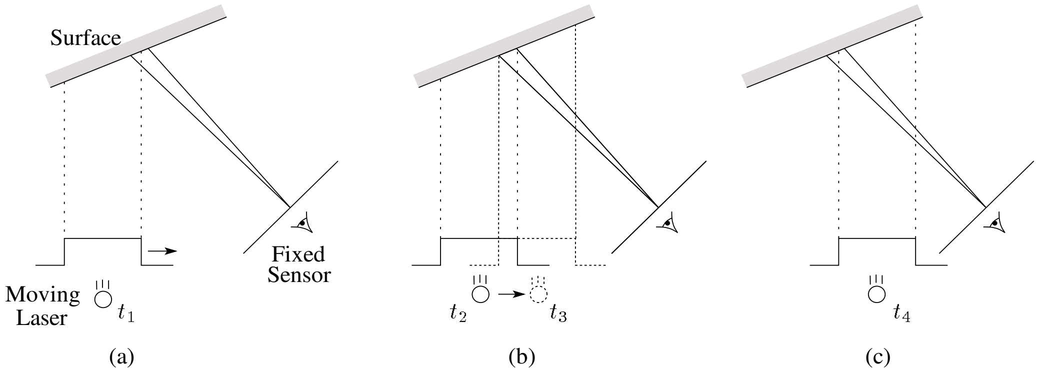

Figure 3.6: Speckle contributions at the sensor due to a moving coherent illuminant with square cross-section. The sensor and surface are stationary while the illuminant sweeps across the surface. A single pixel maps to a single point on the surface, but a neighborhood of points on the surface influence the pixel intensity as a result of diffraction at the aperture. This neighborhood is indicated by the region between the lines drawn from the sensor to the surface. (a) As the illumination translates to the right (shown here at time t_1), only part of the neighboring points, mostly to the left of the range point, are illuminated. Only these points will contribute to the observed, speckle-corrupted intensity at the sensor. (b) At time t_2, the illumination uniformly covers the neighborhood and continues to do so until time t_3. During this time, the pixel intensity will not vary as the illumination moves. (c) As the left edge of the illumination crosses into the neighborhood near the range point (shown here at time t_4), then only the the points mostly to the *right* of the range point will affect the pixel intensity. Thus, we expect the speckle noise in (a) and (c) to be uncorrelated.

The impact of laser speckle on *traditional* triangulation is not surprising, given the obvious corruption of the imaged illuminant. The effect of speckle on *spacetime* triangulation, however, is harder to visualize. To gain a more intuitive grasp, consider the example in Figure 3.6, which illustrates the didactic example of a stationary object and sensor and a moving illuminant with square irradiance profile. A single pixel on the sensor maps to a single point on the surface, but due to diffraction at the aperture, neighboring points on the surface influence the intensity measured at this pixel. When the sensor approaches the range point from the left, then surface points to the left of the range point affect the pixel intensity. When the illumination is completely covering the neighborhood of the range point, then translations of the illuminant do not affect the measured intensity. As the illumination passes the range point to the right, the surface points to the *right* of the range point affect the pixel intensity. Thus, when the illuminant is mostly to the left or right of the range point, the pixel intensity is determined by *different* points on the surface.

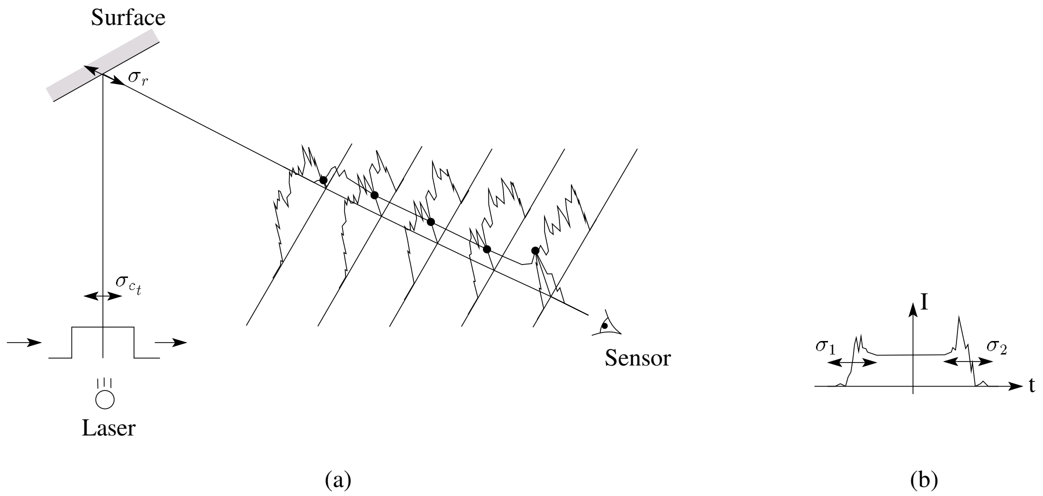

Figure 3.7: Spacetime speckle error for a coherent illuminant with square cross-section. (a) As the sensor and illuminant sweep over the surface, the sensor images noisy square waves. For regions near the center of the square wave, the speckle behaves as though it adheres to the surface as the scanning progresses. Near the edges, however, the speckle influence reveals itself. For the range point being tracked, the amplitude in the first image depends on the properties of the left portion of the surface, while the amplitude in the last image depends on the right portion of the surface. In (b) we see the resulting spacetime image of the range point. The central regions of the pulse are very flat, unlike the individual images of the square pulse at each time step in (a). However, the speckle influences the edges of the pulse, resulting in uncertainty about the mean of the pulse.

Figure 3.7 illustrates how this distinction leads to errors in spacetime triangulation, this time including the motion of the sensor along with the illuminant. As the center of the square illuminant moves over the object, the speckled image of the surface does indeed move with the surface, because the illuminant is fairly constant near the center. This explains why the spacetime pulse is very flat in the center (Figure 3.7b). However, near the edges of the pulse, the illumination is not at all constant, and speckle noise appears. The result is an uncertainty in determining the location of these edges. The noise at the edges has different properties for the left edge versus the right edge as described above, meaning that the uncertainty in determining the position of the two edges is fairly uncorrelated. Since resolving the edges is the only way of determining the center of the spacetime pulse, spacetime triangulation is limited in accuracy by laser speckle.

3.5 A Signal Processing Perspective

With insights provided by the previous section, we can now study optical range triangulation in a signal processing context. This signal processing point of view aids in visualizing how the ideal spacetime image is formed and how it is modified during acquisition and reconstruction. Further, analysis of the Fourier spectrum of the spacetime image leads to insights about how we might improve range extraction.

3.5.1 Ideal triangulation impulse response

In examining the response of a scanner to a differential element in Section 3.2, we have discovered the impulse response of a translating triangulation scanner with Gaussian illuminant and orthographic sensor. The key is to re-write Equation 3.5, as a convolution:

$$E_S(x,z) = f(x,z) * h_{ST}(x,z) \tag{3.11}$$

where h_{ST} is the spacetime impulse response,

$$h_{ST}(x,z) = I_L \exp\left[\frac{-2x^2}{w^2}\right] \delta(z\sin\theta - x\cos\theta) \tag{3.12}$$

and $f(x,z)$ is a point in space,

$$f(x,z) = \rho_{LS}(x_o)\delta(x-x_o)\delta(z-z_o) \tag{3.13}$$

In general, $f(x,z)$ describes the piecewise continuous set of points on the surface of the object that are visible to both the laser and the CCD. This set of points can be represented as a function, $z = r(x)$, with corresponding reflectance variations $\rho_{LS} = \rho_{LS}(x)$, so that we arrive at:

$$f(x,z) = \rho_{LS}(x)\delta(z - r(x)) \tag{3.14}$$

Figure 3.8 illustrates how the spacetime image is derived from a two dimensional scene.

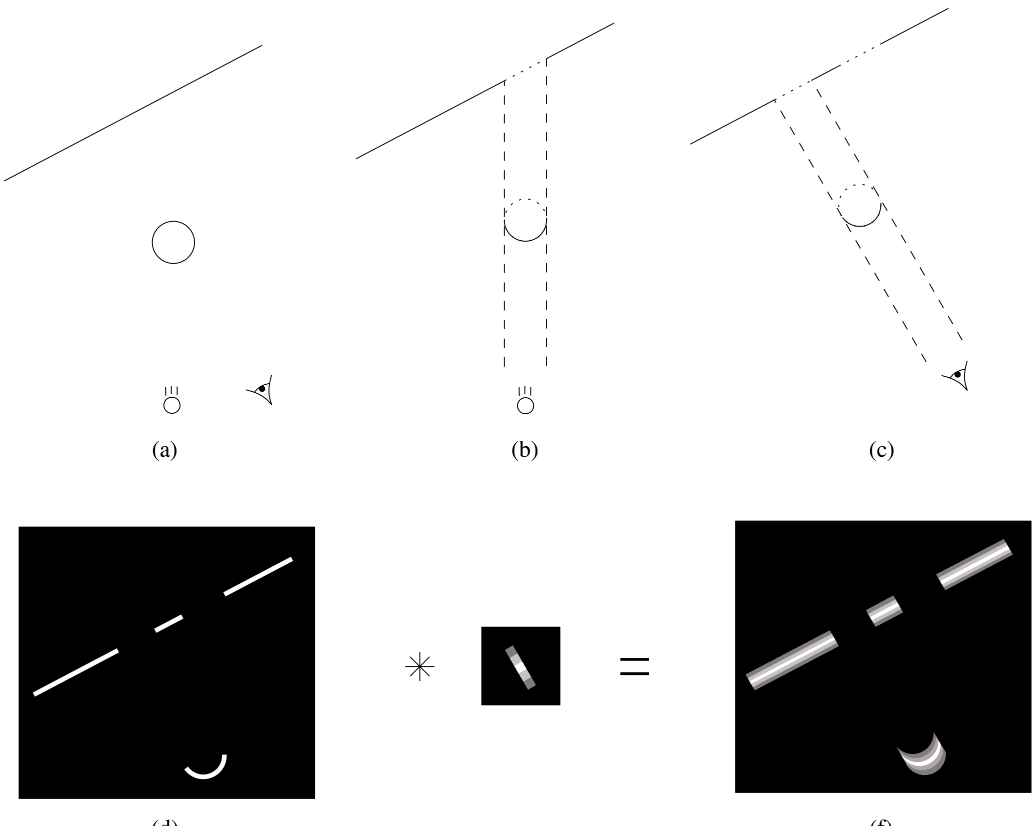

(a) (b) (c)

(d) (e) (f)

Figure 3.8: Formation of the spacetime image. The scene in (a) is being scanned horizontally. First we determine the portions of the surface visible to the laser (b) and to the sensor (c). The laser and sensor lines of sight that clip the surface are shown dashed, while the removed portions of the surface are shown dotted. The resulting surface (d) is convolved with the spacetime Gaussian (e) to create the spacetime image (f). For the purposes of this illustration, the contribution of $\rho_{LS}(x)$ has been omitted. The result of $\rho_{LS}(x)$ would be to modulate the amplitude of the spacetime Gaussian as it is being convolved across the surface.

3.5.2 Filtering, noise, sampling, and reconstruction

The function $E_S(x, z)$ represents the ideal spacetime image, and if we could acquire it perfectly, then we would be able to compute the *exact* shape of the object. In the real world, however, this image is filtered, corrupted by noise, sampled, and ultimately reconstructed before being analyzed for shape extraction.

The filtering can be decomposed into two components: spatial and temporal. Spatial blurring occurs at the sensor due to the imaging optics and the fact that each sensor cell has a finite area over which it gathers photons and accumulates a charge. We call this filter $h_{pix}(z)$, because it spans sensor pixels. The other filtering component operates temporally; as the object moves relative to the scanner, the sensor captures photons over a fixed interval of time before sampling the result and starting over. The motion is effectively blurred during this interval. We call this temporal filter $h_{frame}(x)$, where "frame" refers to the time interval and x is the representative of the time component (remember, $x = vt$).

Several noise sources also affect the spacetime image. As described in the previous section, coherent illumination leads to speckle noise that corrupts the spacetime Gaussian. In addition, the sensor will have some inherent noisiness caused by a variety of factors (see, e.g., [Janesick et al. 1987]). Finally, during the sampling step, the pixel values will be quantized for representation in the computer, introducing quantization noise. We will represent all of these noise sources with one function, $\eta()$.

Putting the filtering and noise sources together, we arrive at an expression for the spacetime image prior to sampling and reconstruction:

$$E_S^{filt}(x, z) = \eta[E_S(x, z) * h_{pix}(z) * h_{frame}(x)] \tag{3.15}$$

This image is then sampled according to the spatial resolution of the sensor and the frame rate. We represent this sampling process as an impulse train, $i(\frac{x}{\tau_x}, \frac{z}{\tau_z})$, where τ_z is the sensor pixel separation projected along the laser sheet and τ_x is the distance the object moves relative to the scanner between frames. Once we have the sampled, filtered spacetime image, we can employ a reconstruction filter, $h_{rec}(x, z)$, to obtain an approximation to the original continuous image. This final image is then:

$$\tilde{E}_S(x,z) = \left[E_S^{filt}(x,z) \cdot i\left(\frac{x}{\tau_x}, \frac{z}{\tau_z}\right)\right] * h_{rec}(x,z) \tag{3.16}$$

One of the first observations that arises from these equations is the fact that the reconstructed spacetime image is a *continuous* function. Thus, this reconstruction yields a (piecewise) continuous surface, rather than a set of range samples. In the next chapter, we exploit this observation to extract as much range data as possible.

3.5.3 The spacetime spectrum

Taking the Fourier transforms of Equations 3.11 - 3.16 will provide several insights. For simplicity, we omit the influence of noise and arrive at the following relations:

$$\mathcal{F}(E_S) = F(s_x, s_z)H_{ST}(s_x, s_z) \tag{3.17}$$

$$\mathcal{F}(E_S^{filt}) = \mathcal{F}(E_S)H_{pix}(s_z)H_{frame}(s_x)] \tag{3.18}$$

$$\mathcal{F}(\tilde{E}_S) = [\mathcal{F}(E_S^{filt}) * \tau_x\tau_z i(\tau_x s_x, \tau_z s_z)]H_{rec}(s_x, s_z) \tag{3.19}$$

where $\mathcal{F}()$ is the Fourier transform operator, s_x and s_z are the spatial frequencies in x and z, and F, H_{ST}, H_{pix}, H_{frame}, and H_{rec} are the Fourier transforms of f, h_{ST}, h_{pix}, h_{frame}, and h_{rec}, respectively. Figure 3.9 shows the steps in creating, filtering, and sampling the spacetime image.

Several observations arise from thinking about the Fourier spectrum of the spacetime image. The least surprising is the fact that higher sampling rates permit (and are usually accompanied by) wider H_{pix} and H_{frame} filters, allowing for the acquisition of more of the spacetime spectrum, i.e., higher frequencies. In other words, a higher sampling rate leads to more captured detail.

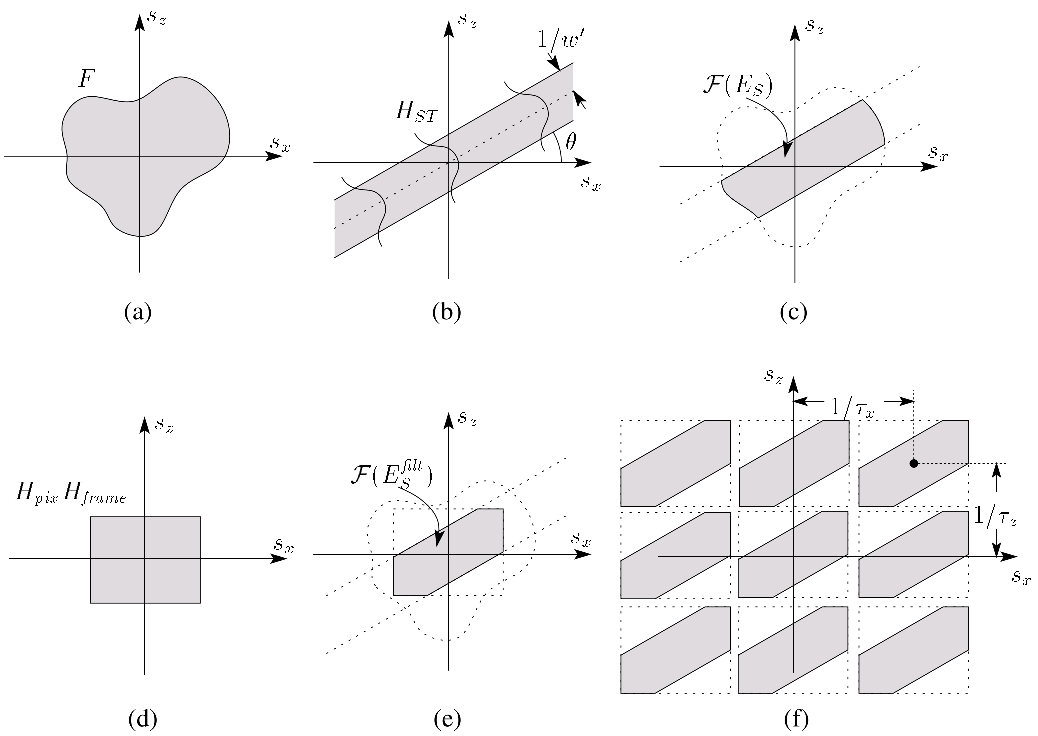

Figure 3.9: The spacetime spectrum: a visualization for Equations 3.17 - 3.19. The Fourier transform of the shape function (a) is multiplied by the spectrum of the spacetime Gaussian (b) to give the spectrum of the ideal spacetime image (c). Before sampling, this spectrum is multiplied by the pixel and frame filters (shown as idealized low-pass filters in (d)) to give the filtered spacetime spectrum without noise (e). The sampling process creates replicas of this filtered spectrum (f) spaced $1/\tau_x$ and $1/\tau_z$ apart in s_x and s_z respectively. The reconstruction step entails extracting the central replica of the spectrum after multiplying by H_{rec} (not shown).

3.5.4 Widening the laser sheet

The Fourier spectrum also tells us that the spacetime Gaussian acts like a bandlimiting filter that runs along the sensor lines of sight, as indicated in Figure 3.9c. Note that this “filter” does not remove detail, as in the case of the pixel and frame filters. The resulting spacetime image is still ideal and contains all the information required to extract the original surface geometry precisely. Thus, by narrowing the spectrum of the spacetime Gaussian, we can restrict the spectrum of the spacetime image without loss of detail. The surprising conclusion: *widening* the illuminant can improve results.

Consider scanning a line (a plane in 3D), as depicted in Figure 3.10. The Fourier transform of this line is another line, perpendicular to the first. After superimposing the spectrum of the spacetime Gaussian, we arrive at the ideal spacetime spectrum, which is basically a line segment in the Fourier plane. The length of this line segment corresponds to the maximum frequency content of the spectrum and thus determines the sampling rate needed to acquire the surface. The parameters that determine the length of the Fourier line segment are the orientation of the original line, the width of the laser sheet, and the triangulation angle. For a fixed laser width and triangulation angle, the spectral line segment grows in length as the line is rotated to be at more of a grazing angle to the sensor. At a critical angle, the pixel and frame filters begin to clip the spectral line segment and degrade the quality of the reconstruction; worse, if the filters are not ideal (and, generally, they are not), aliasing will occur. By widening the laser sheet, however, this critical angle can be extended closer to the triangulation angle.

This notion of widening the laser sheet runs counter to the intuition that a narrower illuminant gives better localization of the point on the surface being acquired. This intuition, however, stems from thinking of triangulation in the traditional context. In the spacetime context, the wider beam does not degrade resolution by “blurring” neighboring points together. It is interesting to note that while the accuracy of the traditional method improves with a narrower laser sheet in principle, in practice the laser must have a minimum thickness to avoid range errors due to aliasing. For a surface of fixed orientation, narrowing the laser sheet also narrows the imaged Gaussian on the sensor, until eventually it is narrower than a pixel cell. At this point, the accuracy is limited by the width of a pixel projected onto the

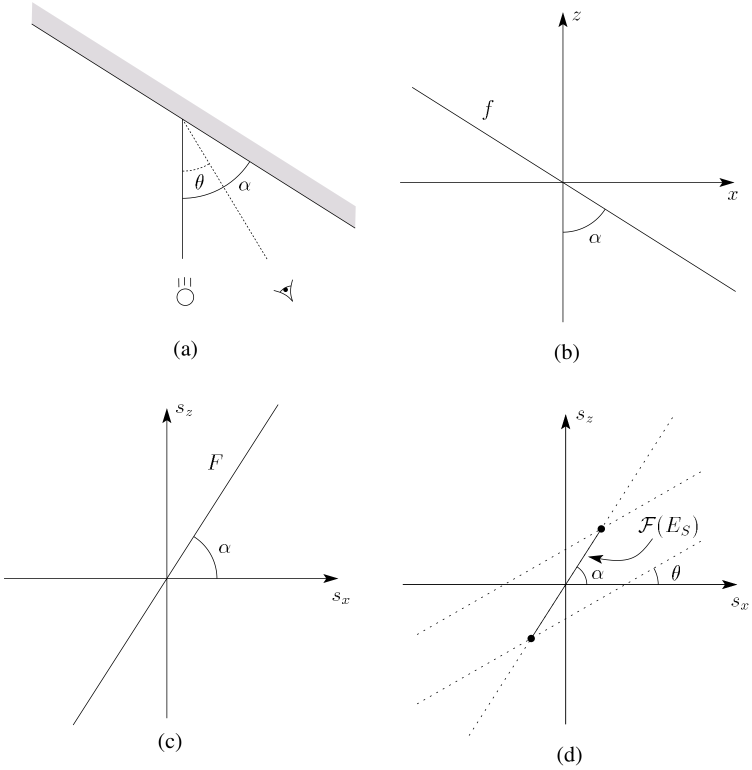

Figure 3.10: Spacetime spectrum for a line. (a) A line at angle α to the laser is scanned. (b) The function to be convolved with the spacetime Gaussian is a delta function line. (c) Its Fourier transform is another delta function line oriented perpendicular to the original. (d) After "filtering" with the spacetime Gaussian, the spacetime spectrum is essentially a line segment in the Fourier plane.

laser sheet, *not* the width of the laser.

In practice, the laser sheet cannot be widened arbitrarily. The primary limitation is the depth of field of the sensor. As the laser is widened, the sensor must record more of the object that is out of the plane of focus. The out of focus points will be blurred, and the improvement from the laser sheet will be countered. One solution might be to use a brighter beam and then stop down the sensor's aperture to increase the depth of field. However, for coherent illumination, the narrowed aperture leads to increased speckle noise which can also defeat the improved accuracy. Even incoherent imaging systems will approach a diffraction limit as the aperture is narrowed. In addition, as the laser sheet widens, then the effects of using a sensor with perspective projection become more prevalent, as noted in Section 3.3. Further, speckle analysis indicates that widening the laser sheet increases the influence due to speckle noise. Nevertheless, by balancing these trade-offs, it should be possible to realize a system with improved accuracy. This remains an area for future work.

3.5.5 Improving resolution

Fourier analysis of the spacetime image also reminds us of the importance of using good bandlimiting filters, H_{pix} and H_{frame}, to avoid aliasing artifacts. However, we do not have much control over these filters. The pixel filter is essentially the convolution of the blurring due to finite pixel cell size with any optical blurring that may occur. When the laser sheet is in focus, then the sensor's pixel response is the primary filter, which in turn is set by the design of the sensor. The frame filter is essentially a rectangle function in x with width proportional to the exposure time. In addition, sensors generally have a time interval when they stop capturing light while shifting out the recorded signals. The spectral bandwidth of the temporal filter is thus wider than the distance between samples. In both the spatial and temporal cases, some aliasing artifacts are bound to occur. This aliasing generally leads to reduction in image quality; however, it also means that multiple spacetime images may be combined to improve the accuracy of the range data. Techniques such as the one described in [Irani & Peleg 1991] could be used to combine these spacetime images. In order for such an approach to work, a sequence of spacetime images must be acquired after translating the object (or scanner) between scans. Rotating the object (or scanner) between scans generates an entirely new spacetime image that may not be combined with prior acquisitions. Increased resolution through combination of multiple spacetime images remains an area for future work.