WHITE PAPER

High-Resolution Time-of-Flight Cameras

Light travels at a speed of 299 792 458 meters per second. By measuring the time it takes for a light pulse to travel between two points, we can thus calculate the distance between them. This principle is at the heart of the measurement technique used in Time-of-Flight (ToF) cameras. Each pixel in a ToF camera is returned not just with a brightness value, but also with information about the distance between the light-reflecting surface and the camera, accurate to millimeters. This White Paper explores the concepts behind the Basler ToF Camera, particularly how it measures distance using light.

Content

1. Setup 1
2. Functional Principle for Time-of-Flight Cameras 3
3. Precision 3
4. Accuracy and Precision 3
    4.1 Multiple Reflections 3
    4.2 Scattered Light 4
    4.3 Working Area 4
    4.4 Solid Angle 4
    4.5 Ambient Light 4
    4.6 Degree of Reflection and Transparency 4
    4.7 Temperature 5
    4.8 Camera Configuration 5
5. Summary 5

The Basler ToF Camera features 640 × 480 pixels resolution and delivers 30 frames per second. That corresponds to 9.2 million distance measurements per second.

Before we can discuss the principles behind this measurement technique, the degree of precision it offers and relevant influencing factors, we must first talk about the setup.

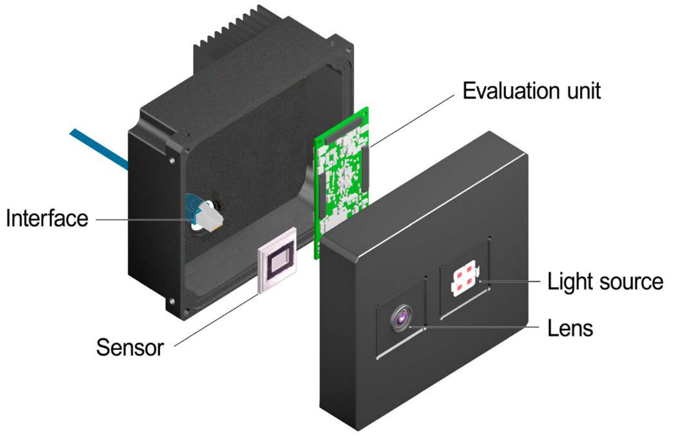

Illustration 1: Components of the Time-of-Flight Camera

1. Setup

The Time-of-Flight camera being presented here has the following components (see illustration 1):

- Light source
- Optics
- Sensor
- Evaluation Unit
- Interface

Let us take a closer look on each of these components.

Light source
The light source illuminates the surface to be measured and is harmonized with the lens's aperture opening. It is responsible for providing sufficient lighting, even for surfaces with a low degree of reflection at a significant distance, thus allowing for short exposure times and high frame rates. The light source is also an important component in the measurement process, and is controlled using a sophisticated technique by the ToF camera's electronics.

The light source switches on and off several thousand times during image capture. Each individual light pulse is just a few nanoseconds long. A camera parameter for the exposure time determines the number of pulses per image.

To measure precisely, the light pulses must be controlled precisely, with identical durations, rise and fall times. The brightness profile is calibrated for each camera individually, since even minor deviations of just a nanosecond can produce distance measurement errors of as much as 30 centimeters.

Because they are well equipped to provide these exact pulse characteristics, LEDs and laser diodes are particularly suitable as light sources for this difficult task.

Eye safety must be ensured while in the camera's range, even in the event of a fault. This is ascertained by validations in coordination with experts from testing organizations accredited for optical radiation protection.

Light in the near-infrared range is invisible to humans and does not bother us. However, basically any part of the entire light spectrum is in line for the measurement.

Optics

The lens reproduces the light pulse reflected from the surface onto the sensor.

Each camera is geometrically calibrated to correspond to spatial coordinates on the surface. Focal length, focus and aperture should thus not be modified once a calibration procedure has been undertaken.

An optical band-pass filter behind the lens allows only wavelengths used by the camera's own light source to pass through. This guards the sensor against overexposure by a disturbing extraneous light.

Sensor

All cameras use the latest generation of CMOS image sensors, developed by Sony especially for Time-of-Flight measurements.

The electronic shutter on the sensor runs in exact synchronization with the numerous light pulses until a sufficient charge has been generated in the sensor by the incoming light.

Once the charge has been accumulated, it is read out via a 12-bit analog/digital converter and transferred to the evaluation unit.

Lighting and readout speed for the sensor typically dictate the frame rates for this procedure. An additional idle time between the frames serves to cool off the illumination.

Evaluation Unit

Using the data from one single readout pass, the evaluation unit can calculate an accurate intensity profile and depth map as well as spatial coordinates for the recorded objects (point cloud).

The intensity profile is representing the brightness as a 16-bit integer for each pixel. The wavelength of the camera's built-in light source is crucial here. Because near-infrared light is invisible to humans, the reflective properties of infrared light upon striking objects are not something we can experience directly. Because they vary greatly from visible light, the intensity profile can deviate significantly from our perception.

The depth map charts the distance between the light-reflecting surface and the camera, again as one 16-bit integer per pixel. Strictly speaking, these integers represent the time of travel for the light pulse, from transmission to return to the sensor, influenced by all the detours and shortcuts the light took along that path. In other words: it's not impossible for errors to slip into the results.

The depth map can also be output as a color image, with red, green and blue as 8-bit digits. This is primarily helpful as a visualization aid. Red surfaces are close, blue surfaces further away.

The confidence chart contains a scale for errors on the depth map: one 16-bit integer per pixel. The lower the number, the lower the precision for the measurement result for that pixel. This is the case, for example, with very dark surfaces, or for pixels where overexposure has occurred. The confidence chart provides users the chance to adapt their application to account for the reliability of the measurement values in the depth map.

The point cloud is comprised of three 32-bit floating point numbers per pixel. These are spatial coordinates for the reflecting surface in a Cartesian coordinate system. The units are indicated in millimeters. For precise calculation of these values, the cameras are geometrically calibrated individually at the factory.

Pixels with unreliable measurements have a low value in the confidence chart and are output as a distance of Null or NaN (Not a Number). This confidence threshold can be modified.

Interface

The camera is equipped with a Gigabit Ethernet interface, meaning it can be connected directly via an affordable standard cable to any host PC, even over extended distances.

The GigE Vision and GenICam communications protocols are proven standards for industrial cameras, independent of the manufacturer.

2. Functional Principle of Time-of-Flight Cameras

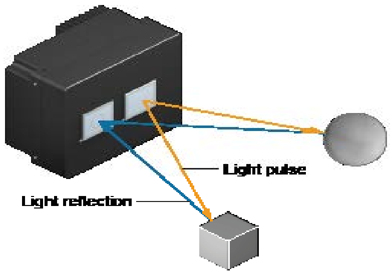

Illustration 2: Functional Principle of Time-of-Flight Cameras

For a ToF camera, there are a great number of potential variants for light pulses and electronic shutter timings, each with specific advantages and disadvantages. The following will detail a highly simplified but fundamentally correct method.

The control unit on the camera switches the light source on and off again, thus forming a light pulse. At the exact same moment, the control unit opens and closes the electronic shutter of the sensor. The charge generated in this way by the light pulse — let's call it S_0 — is stored in the sensor.

The control unit then switches the light source on and off again a second time. The shutter is opened later this time, namely at the point in time when the light source is switched off. The charge generated now (S_1) is also stored on the sensor (see illustration 3).

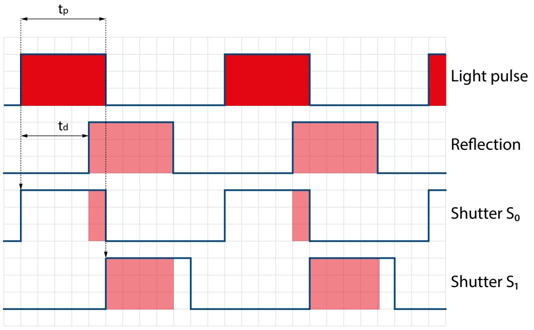

Illustration 3: Exposure Time

Because a single light pulse is very short, this procedure is repeated several thousand times until the configured exposure time has passed. The values in the sensor are then read out.

The exposure sequence has now produced two images. In the S_0 image, the closer surfaces appear to be lighter. The further away the surface is located, the less light it reflects back on the sensor during the period where the shutter is open. During the S_1 measurement, this is precisely the reverse: close surfaces are dark, as the shutter doesn't open until the light has already been moving for a while.

The actual distance can be determined from the ratio of these intensities. If the speed of light is defined as c, while t_p represents the duration of the light pulse, and beyond this S_0 stands for the collected charge of the earlier shutter and S_1 for the charge with the delayed shutter, then for distance d the ratio is:

$$d = \frac{c}{2} \times t_p \times \frac{S_1}{S_0 + S_1}$$

The smallest measurable distance is then measured once all charges during the earlier shutter period in S_0 and no charge during the delayed shutter period in S_1 are collected, i.e. $S_1 = 0$. The formula then produces $d = 0$.

The largest measurable distance is measured if all charges in S_1 and no charges at all in S_0 are determined. The formula then yields $d = \frac{c}{2} \times t_p$.

This is also an indicator that the light pulse width for this method determines the maximum measurable distance. For example, if t_p should total at least 47 nanoseconds, then up to 7 meters can be measured.

As mentioned above, this is a simplified explanation. In fact, the camera uses four instead of two samples, which leads to a more complex formula, but also to a more accurate measurement.

3. Accuracy and Precision

Accuracy is the mean difference between the measured distance and the actual distance. Precision is the standard deviation for that accuracy.

An accurate and precise measurement is possible under specific conditions, which will now be described.

4. Influencing Factors

There are several influencing factors that can impact the measurements of Time-of-Flight cameras and which can limit measurement precision. They include: multiple reflections, scattered light, ambient light and temperature.

4.1 Multiple Reflections

Distance measurement requires light that has been reflected just once. Light that is reflected multiple times

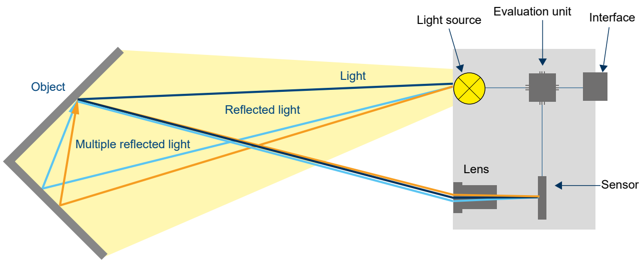

Illustration 4: Multiple Reflections

distorts the measurement. Things like corners of rooms and concave shapes (i.e. the inside of a coffee mug) will typically produce errors due to the multiple reflections of the light beam.

Mirrors and reflective surfaces (polished tabletops) deflect the beam of light. Always imagine how many paths the light pulse may take from the light source back to the camera. If the light is deflected at 100%, then no reflected intensity reaches the camera, preventing a proper measure of the distance of the reflective surface. If, conversely, the beam of light is mirrored into the sensor, then overexposure frequently occurs.

On the other hand, a bright, even and diffusely reflective wall in an otherwise empty black room is ideal.

4.2 Scattered Light

Scattered light (see illustration 5) occurs due to unwanted reflections within the lens or behind it. Even for the most carefully engineered setup, scattered light cannot be completely eliminated. Bright surfaces located very close to the light source quickly scatter too much light into the lens.

This surface doesn't even have to be within the sensor's field of view If, for example, a camera is placed directly in the center of a tabletop, then scattered light will strongly interfere with the distance measurement.

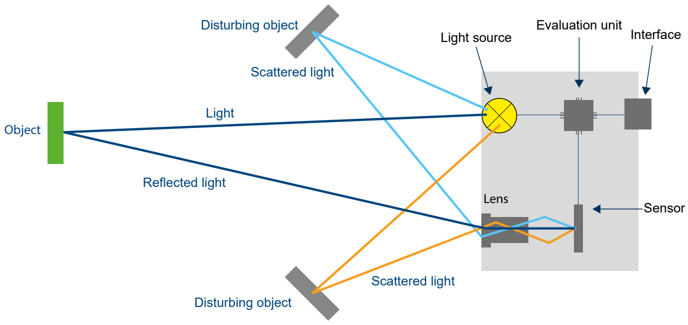

Illustration 5: Scattered Light

In the intensity image scattered light leads to a washed-out image with poor contrast. The problem has long been familiar in the photographic field. The solution is to ensure that the space directly in front of the camera is always free of strongly reflective objects.

4.3 Working Area

Four factors limit the working area for a ToF camera:

- The measurement method being used, such as light pulse width. Each temporal arrangement of light pulse and electronic shutter is designed for a specific distance range. Surfaces outside that range produce no measured values, or incorrect ones. The intensity profile is less affected, although here too, surfaces may be poorly illuminated or out of focus.
- The contrast range (dynamic range) for the camera: on the one hand, nearby surfaces appear very bright, while those at a larger distance are very dark. The intensity of the reflected light is reduced by the square of their distance. The goal is to measure at both ends of this brightness scale, which requires a good contrast range.
- Intensity of the light source: The light must be sufficient to illuminate surfaces even at larger distances. Otherwise the noise levels will be too high, producing imprecise measurements.
- Depth of focus of the lens: Indistinct borders between the foreground and background lead to invalid distances in the middle.

4.4 Solid Angle

The lens and light source intensities are reduced in the marginal regions of the image. For this reason, surfaces in those areas tend to be underexposed and the distance measurements become less precise.

Beyond this, shadowing in the picture margins tends to be larger, leading to individual light elements being recorded more strongly in the measurement than others.

4.5 Ambient Light

Although the camera measures the ambient light and subtracts it during evaluation, it nevertheless remains a physical problem. A pixel on the sensor can only hold a limited amount of charge. The more this capacity is filled up by the ambient light, the less capacity remains for the crucial desired light pulse to be recorded. In other words, the signal-to-noise ratio drops.

An optical band-pass filter only allows the spectrum present in the light source to pass and to reach the pixel. Artificial light is thus usually not a problem, since it contains only a small portion from that part of the spectrum.

Daylight, however, is active across almost the entire light spectrum, and in some cases, like on a sunny summer day, reaches a significant intensity. The camera thus requires additional protective mechanisms to ensure that the light source remains measurable.

A daylight robust camera works at 940 nm, with the sunlight being strongly absorbed by the atmosphere.

4.6 Degree of Reflection and Transparency

The degree of reflection for the surfaces to be measured, and their distances, determine the optimal exposure time. If surfaces that are a mix of highly reflective and less reflective are to be recorded at the same time, then the exposure time must be chosen very carefully. Otherwise over- and underexposure will occur.

It's recommended that the exposure time is configured in such a way that the most reflective surfaces that are closest to the camera are not overexposed when the ambient light is at its strongest. This can be checked effectively using the intensity profile and then afterwards with the depth map. The exposure time should be selected for as long as possible, without exceeding the saturation threshold. The position and alignment of the camera may need to be varied for this.Working with the same exposure time, you then check the low-reflecting surfaces at far distances for proper intensity and correct distance values.

If the effort to find the proper exposure time is unsuccessful, then two or more shots must be made with different exposure times and combined into one profile, potentially with the aid of the confidence values.

Mirrored and transparent surfaces are not conducive to accurate distance measurement, since the primary paths of the light are not corresponding to the shortest distance.

4.7 Temperature

The distance calculation is based on a model involving the timing of a light pulse and the shutter time. If for example the electronic shutter on the camera opens just 33 picoseconds later than intended, then the calculated distance will already be a centimeter too short.

Even high-quality electronic components struggle to meet the precision requirements for ToF measurement. High temperatures can lead to noise, and temperature fluctuations impact the timing behavior.

It is thus important that the camera is operated in thermally stable conditions. Extreme temperatures and/or fluctuations should be avoided. The camera should be cooled if possible, such as through an air flow or mounting to a massive metal bracket.

4.8 Camera Configuration

The camera has numerous configurable parameters to make it work with top efficiency. Beyond the exposure time and frame rate, there are also different filters to boost precision. Not every combination will work together to ensure the precision called for in the specification, however.

5. Summary

Distance measurement with a Time-of-Flight camera is handled quickly and efficiently. Unlike normal cameras, the control unit and light source are key components of a ToF camera, and proper usage has a direct impact on measurement precision for the camera.

A ToF camera must be calibrated as a measurement device, including the lens. It delivers optimal results only under specific ambient conditions and for a defined measurement range. The ToF camera delivers significantly more data than a standard camera, namely the intensity profile and the depth map as well as confidence values for each point on the depth map.

Even under optimal conditions, successful ToF measurement relies on numerous factors inside and outside the camera. For best results, the user of the ToF camera should:

- avoid multiple reflections
- avoid scattered light
- measure in the center of the working area
- measure in the center of the image
- avoid ambient light in the near-infrared
- keep the camera temperature constantly low
- avoid mirrored and transparent objects
- prioritize bright, diffusely reflective surfaces
- optimize the camera position and alignment
- use prior knowledge when processing images, such as „the shape is a cube,“ etc.
- use noise filters, both for space and time
- create an intensity profile
- operate the camera from a stable position and avoid changes to the parameters while operational

Users who can follow these rules can take advantage of 9 million distance measurements per second, accurate to millimeters! There's no quicker way to record a space.

Author

Martin Gramatke
Product Manager - 3D Business

Martin Gramatke is a Product Manager at Basler AG. He is responsible for the development of Basler's 3D products, their business strategy and success, the product specifications and launch of new features. Martin has been with Basler for more than 20 years. His previous work in R&D as Developer, Project Manager and Technical Architect contributes to the successful coordination of his current tasks.

About Basler

Basler is a leading manufacturer of high-quality cameras and camera accessories for industry, medicine, traffic and a variety of other markets. The company's product portfolio encompasses area scan and line scan cameras in compact housing dimensions, camera modules in board level variants for embedded vision solutions, and 3D cameras. The catalog is rounded off by our user-friendly pylon SDK plus a broad spectrum of accessories, including several developed specially for Basler and optimally harmonized for our cameras. Basler has three decades of experience in computer vision. The Basler Group is home to approximately 800 employees at its headquarters in Ahrensburg, Germany, and at other locations in Europe, Asia, and North America.

Contact

Martin Gramatke – Product Manager

Tel. +49 4102 463 211
Mobile +49 1511 630 7350
Email: martin.gramatke@baslerweb.com

Basler AG
An der Strusbek 60-62
22926 Ahrensburg
Germany