

Kostengünstiges LIDAR

Machbarkeit eines einfachen und erschwinglichen 3D-Laserscanners

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Technische Informatik

eingereicht von

Michael Platzer, BSc

Matrikelnummer 1029376

an der Fakultät für Informatik
der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Peter Puschner

Wien, 25. Juni 2017

Michael Platzer

Peter Puschner

Low-cost LIDAR

Feasibility of a simple and affordable 3D Laser Scanner

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Computer Engineering

by

Michael Platzer, BSc

Registration Number 1029376

to the Faculty of Informatics

at the TU Wien

Advisor: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Peter Puschner

Vienna, 25th June, 2017

Michael Platzer

Peter Puschner

Erklärung zur Verfassung der Arbeit

Michael Platzer, BSc
Friedlgasse 45/13
1190 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 25. Juni 2017

Michael Platzer

Danksagung

Zuerst möchte ich meinem Betreuer, Prof. Dipl.-Ing. Dr. Peter Puschner, zutiefst danken, der die Idee für diese Arbeit von Beginn an unterstützt hat. Weiters bin ich Ing. Manfred Tomanek dankbar, der mir schon recht früh beim optischen Aufbau geholfen hat. Ich würde gerne auch Ing. Leo Mayerhofer danken, der mich in die Benützung der Heißluft-Lötstation des Instituts eingewiesen hat. Schließlich möchte ich Nicole Weinert meinen Dank aussprechen, für ihre steten Ratschläge und ihre Unterstützung beim Löten der elektronischen Komponenten.

Acknowledgements

First, I would like to sincerely thank my advisor Prof. Dipl.-Ing. Dr. Peter Puschner, who supported the idea of this thesis from the very beginning. Furthermore, I am grateful to Ing. Manfred Tomanek, who from an early stage on helped me with the optical setup. I would also like to thank Ing. Leo Mayerhofer, who introduced me into working with the institute's reflow soldering station. Finally, I would like to express my gratitude to Nicole Weinert for her constant advice and her assistance in soldering the electronic components.

Kurzfassung

Autonome Agenten, ob in der Luft oder am Boden, benötigen präzise Informationen über ihre Umgebung um sicher navigieren zu können. Laserscanner, auch als LIDAR (engl., LIght Detection And Ranging) bezeichnet, sind besonders geeignet um akkurate virtuelle Modelle der Umgebung zu erstellen und sind relativ unempfindlich gegenüber variierenden Umwelteinflüssen. Allerdings sind kommerziell erhältliche Laserscanner sehr teuer, was ihre Verbreitung einschränkt. Diese Arbeit schlägt ein Design für ein kostengünstiges LIDAR-System vor, bestehend aus einem optischen Aufbau, einem elektronischen System, einem Regler und Software zur Nachbearbeitung der Messungen. Das angestrebte Ziel ist die Konstruktion eines Laserscanner, der für das Erfassen der Umgebung und die Identifikation von Hindernissen auf unbemannten Luftfahrzeugen, häufig als Drohnen bezeichnet, geeignet ist. Während des gesamten Prozesses wird darauf geachtet, die Komplexität gering zu halten und ausschließlich leicht erhältliche Komponenten zu verwenden. Um die Gültigkeit des vorgeschlagenen Designs zu bewerten, wird ein Prototyp gebaut. Eine umfangreiche Evaluierung quantifiziert die Performance des Laserscanner in Bezug auf Reichweite, Genauigkeit, räumliche Auflösung und Scan-Rate. Der Prototyp ist in der Lage, grobe Scans der Umgebung zu produzieren, allerdings mit einer gewissen Verzerrung und einer langsameren Geschwindigkeit als kommerzielle Geräte. Auch wenn der Scanner derzeit noch nicht für den Einsatz zur Erkennung von Hindernissen auf unbemannte Luftfahrzeugen bereit ist, demonstrieren die Resultate die konzeptionelle Gültigkeit des Designs.

Abstract

Autonomous agents, whether in the air or on the ground, require precise information about their surroundings for safe navigation. Laser scanners, also referred to as LIDAR (Light Detection And Ranging), perform particularly well in building accurate virtual models of the surroundings and are relatively insensitive to varying environmental conditions. However, commercially available laser scanners are very expensive, which inhibits their widespread use. This work proposes a design for a low-cost LIDAR system, comprising an optical setup, an electronics system, a controller, and required software for processing the measurements. The ultimate goal is to construct a laser scanner which is suitable for environment scanning and obstacle identification on Unmanned Aerial Vehicles (UAV), commonly referred to as drones. Throughout the process care is taken to keep the complexity low and to rely only on readily available components. In order to assert the validity of the proposed design a prototype is built. An extensive evaluation quantifies the performance of the laser scanner in terms of range, precision, spatial resolution, and scan rate. The prototype is able to produce rough scans, albeit with some distortion and at a slower rate than commercial devices. Although the scanner is not yet ready to be used for obstacle detection on UAV, the results demonstrate the conceptual validity of the design.

Contents

Kurzfassung	xi
Abstract	xiii
1 Introduction	1
1.1 Applications and Current Status	1
1.2 Outline	3
2 System Design	5
2.1 Optical Design	6
2.2 Electronic Design	13
2.2.1 Photodetector Signal amplification	14
2.2.2 Noise Analysis	21
2.2.3 Square Wave Recovery	30
2.2.4 Time-of-flight Measurement	35
2.3 Controller Design	39
2.4 Post Processing	42
3 Evaluation	45
3.1 Validation of Calculations	46
3.1.1 Optical Setup Performance	46
3.1.2 Transimpedance Amplifier Evaluation	47
3.1.3 Precision of Delay Measuring Circuitry	50
3.2 Unexpected Effects	52
3.2.1 Amplitude Dependence	52
3.2.2 Ambient Light Effects	55
3.3 Test Scan	56
3.4 Discussion	61
4 Conclusion	63
List of Figures	65
List of Tables	69
Bibliography	71

Introduction

LIDAR (LIght Detection And Ranging) is a technology which uses time-of-flight measurement of light to infer the distance to objects. Performing multiple measurements while gradually varying the orientation of the light beam allows to build a virtual model of the surroundings [CW11, chap. 8].

The term LIDAR was initially a portmanteau of the words light and radar¹ [OED89], since a LIDAR system performs the same task as a radar system: it scans the surroundings and reports the location and shape of objects. By now, LIDAR has been in use for over 50 years. Its inception closely followed the first practical lasers², which is not by coincidence. Due to its monochromatic nature, laser light can be collimated to a narrow beam of light with only very slight divergence, hence laser light can be kept focused over a great distance. Therefore a laser is particularly useful in LIDAR, since it allows precise targeting over large distances. One of the first applications of LIDAR was to sense the shape of the moon's surface [Rin63].

Because laser light is very suitable for use in LIDAR and hence the majority of LIDAR systems use a laser beam to scan the environment, the terms LIDAR and 3D laser scanner or just laser scanner are used interchangeably throughout this work.

1.1 Applications and Current Status

An important application of LIDAR is the acquisition of topographic maps. Airborne laser scanners are used to measure the elevation of terrain from an aircraft, generating a 3D map of the landscape [CW11, sect. 8.3]. Terrestrial applications include the creation

¹ Radar was initially an acronym for RAdio Detection And Ranging, but has since entered english language and is no longer capitalized [MGH03].

² Laser likewise used to be an acronym, expanding to LIght Amplification by Stimulated Emission of Radition [Gou59].

of 3D models of construction sites or cultural heritage, replacing the previously used tachymeter [VM10].

When used in geodetics or construction surveying, laser scanners must perform accurate range measurements over large distances, requiring powerful lasers, sensitive detectors and precise circuitry. Such laser scanners are manufactured by only few companies and are quite expensive. The German firm SICK AG builds laser scanners for terrestrial applications [SICK]. A manufacturer of airborne LIDAR systems is the Austrian company RIEGL Laser Measurement Systems GmbH [RIEGL].

A recently emerging application of LIDAR is the detection of obstacles for autonomous vehicles. For instance, prototypes of the self-driving cars built by Waymo LLC (formerly known as the Google Self-Driving Car Project) have a LIDAR scanner mounted on top of the roof to generate a real time virtual 3D model of the surroundings, allowing the car to localize obstacles and steer adequately [Sim17][Ama17].

Laser scanners are particularly suitable for the sensing of the surroundings and detection of obstacles. They outperform other 3D sensing technologies in terms of range and precision. Radar has a poor resolution and can only roughly identify nearby obstacles [RH09]. Ultrasonic sensors equally suffer from poor resolution, work reliably only over short range and their performance can be severely compromised by nearby noise sources [BGO09]. Image processing based approaches, which build 3D models by matching key points across several images taken from different angles, greatly suffer from degrading visual conditions, such as low light or dampness [SK12]. LIDAR in turn has a great range and resolution and is quite insensitive to varying environmental conditions [RH09].

Despite its many advantages over other 3D scanning technologies, the widespread use of LIDAR in autonomous agents is inhibited by the high price. An early prototype of Waymo's self-driving car used a laser scanner built by US company Velodyne LiDAR, Inc. which apparently cost more than US\$ 75 000 [Kra16][Ama17].

Waymo identified this as a major hurdle in the inception of self-driving cars and since began to develop its own laser scanners [Sim17], cutting the costs by approximately 90% [Ama17][Muo17]. The transportation network company Uber Technologies Inc. is likewise developing LIDAR technology to be used on its autonomous vehicles. Recently Waymo filed a lawsuit against Uber and its subsidiary Ottomotto LLC, claiming that former employees had stolen proprietary designs of Waymo's LIDAR systems [Way17].

Velodyne itself is aiming at reducing the costs of its LIDAR units and released the PUCK VLP-16, which is advertised for its cost-efficiency with a list price of US\$ 7 999 [Vel16]. Numerous competitors are also working on small and cheap LIDAR scanners, such as US company Quanergy Systems Inc. or Israel-based Innoviz Technologies Ltd., both of which announced solid state LIDAR scanners for a few hundred US\$ [Kra16]. The Massachusetts Institute of Technology and US government agency DARPA are conducting research on even cheaper LIDAR sensors [PW18].

Affordable LIDAR technology would not only be beneficial to self-driving vehicles, but

could as well be used on Unmanned Aerial Vehicles (UAV), commonly referred to as drones [Li+16]. Research on UAV has been fueled by the availability of a large variety of flying platforms at low prices. Autonomous aircrafts are already used in a variety of applications, such as inspection and surveillance [SLC16], and many more are currently investigated, including search and rescue, assistance in disaster management and aerial delivery among numerous others [Qia+16].

In the light of recent research on autonomous UAV it is becoming increasingly apparent that reliable range scanning is essential for safe autonomous flight [Alm+16][SLC16]. LIDAR is predestined for such applications [SNH03]. The low travel speed and the fact that no people are aboard whose lives could be endangered greatly relieve the range and reliability requirements for a laser scanner to be used on UAV. If laser scanning technology were available at prices in the reach of hobbyists, it would certainly accelerate the deployment of airborne autonomous agents in various areas.

However, the majority of companies aiming to cut the costs of laser scanners target the automotive market, hence their devices will require sufficient range to be useful at a car's driving speed and will certainly have to meet automotive standards in terms of robustness and reliability. Therefore new LIDAR devices currently in development presumably will again exceed the range and precision requirements of UAV and probably retain a price which is out of reach for hobbyists.

1.2 Outline

Since it is unlikely that a laser scanner suitable for use on small multicopters will be available anytime soon, this work investigates whether it is possible to build such a laser scanner using simple means and a much lower budget. The goal is to design, build and evaluate a custom 3D LIDAR system and to assess relevant performance metrics such as range, precision and resolution.

Few details are known about the technology contained in commercially available laser scanners, therefore the present design is based upon the essential requirements of such a system, namely an optical setup which emits a laser beam and captures its reflection from remote surfaces, electronical circuitry which amplifies the received optical signal and measures its delay with respect to the emitted signal, a controller which orchestrates these measurements as well as the motion of the laser beam, and a post-processing software which aims to extract useful scan data from the measurements.

The ultimate aim of this work is to assess whether it is possible to construct a do-it-yourself laser scanner with simple means and limited time and budget and to determine the extent as to which such a self-constructed device has a performance comparable to its expensive commercial counterparts.

Following this introduction, chapter 2 gives an overview of the proposed laser scanner design. Chapter 3 presents the results of an extensive evaluation, first of individual system components and then of the range, precision, spatial resolution as well as the

1. INTRODUCTION

scanning rate of the whole device. The last chapter makes concluding statements on the feasibility of affordable LIDAR technology.

System Design

A LIDAR system scans the environment with a rotating laser beam, performing multiple time-of-flight measurements to determine the range of scanned surfaces, and based on these results builds a virtual map of the surroundings.

To enable this, a laser scanner has a number of tasks to fulfill, namely to generate a laser beam, to gradually rotate it so as to scan over the surrounding objects, to capture the reflected light and measure its delay with respect to the emitted light, to calculate the distance to the reflecting surface from this delay and to combine the results of multiple measurements into a virtual model of the environment.

Therefore a LIDAR system comprises an optical and an electronic part, as well as a control and processing system, each of which provides essential functionality for the operation of the device:

1. An optical setup rotates the laser beam in order to scan over the surroundings and directs the reflected light onto a photodetector.
2. Electronic circuitry drives the laser source, amplifies the signal received by the photodetector and measures the delay between emitted and received light.
3. A controller runs software which controls the electronics, digitizes the measurements and sends the generated data to a processing system.
4. A processing system filters the data, calculates the distance to the scanned surfaces and builds a virtual map of the surroundings.

This chapter explains the design of each of these parts as it was realized in the laser scanner built as part of this work.

2.1 Optical Design

As stated at the beginning of chapter 1, LIDAR senses the distance to a remote surface by time-of-flight measurement of light. A light source, usually a laser, emits light which is reflected by the surface and the reflected light is captured by a photodetector. The device measures the delay between the emission of a light pulse and the reception of its reflection. Multiplying this delay with the speed of light allows one to calculate the distance to the surface which reflected the light.

Measuring the distance to an object is called ranging and a device performing this task is termed a rangefinder. However, LIDAR is different from ranging in that it does not report the distance to a targeted object, but instead builds a virtual model of the surroundings by gradually varying the orientation of the light beam and performing multiple measurements.

Nonetheless, a LIDAR system is also a ranging system and in determining distances it operates no different from a laser rangefinder.

Any optical system that shall be used for ranging measurements essentially has two prerequisites:

- Light, usually from a laser, must be collimated in order to emit a non diverging beam which can be pointed at a surface whose distance shall be determined.
- The light reflected by the surface must be condensed so as to focus it onto a photodetector.

Both of these tasks are accomplished with lenses. A collimating lens is placed in front of the light source and a condenser lens captures the reflected light and concentrates it onto the photodetector.

Figure 2.1 shows the basic setup of such a system. The collimating lens can be kept small as its diameter does not need to be larger than that of the laser beam. On the other hand, a larger condenser lens captures more of the returning light, therefore it should be chosen as large as practicable in order to receive a strong signal.

Most surfaces are predominantly diffuse reflectors, which reflect incident light into any direction. Consequently only a small amount of the emitted light returns to the ranging device and this amount decreases as the distance to the reflecting surface increases.

The entire energy radiated by the laser source hits the surface, short of only some negligible losses in the optical components and due to atmospheric absorption. The surface absorbs some of this energy and reflects the remainder. The reflectance R of a surface is the fraction of incident radiant flux which the surface reflects, accordingly the total reflected energy E_r is given by $E_r = E_i R$, where E_i is the radiant energy incident to the surface.

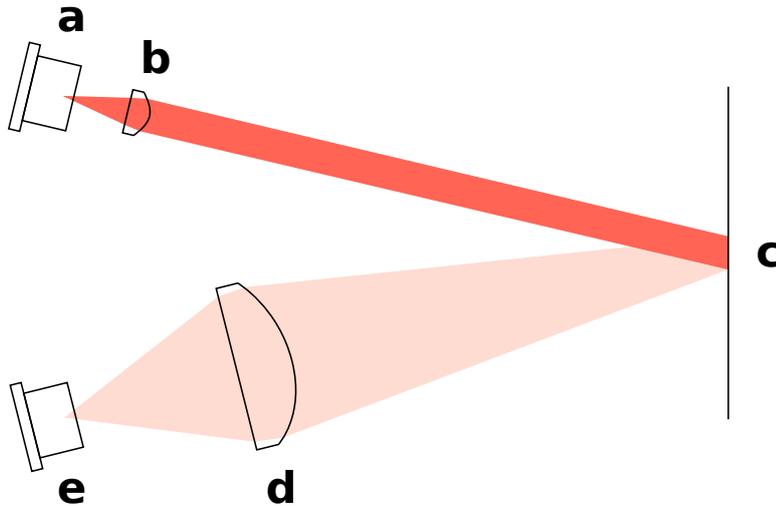


Figure 2.1 Basic setup for ranging measurements. The laser beam, illustrated in red, is emitted by **a**) a light source, **b**) collimated and directed onto **c**) the desired surface. The reflected light is partially caught by **d**) a condenser lens, which concentrates it onto **e**) a photodetector. Note that the bright red shape depicts only the portion of reflected light which is caught by the condenser lens, as a diffuse reflector will scatter incident light in any direction.

The reflected light is scattered within the hemisphere covering the top of the surface (reflected light does not penetrate the surface), hence the entire reflected energy is radiated within a solid angle of 2π steradians. Thus the average radiated energy per steradian, also called radiant intensity, is $E_r/2\pi = E_i R/2\pi$.

A condenser lens with an area A located at a distance d from the reflecting surface covers a solid angle of A/d^2 steradians. Disregarding losses from atmospheric absorption and imperfections of the optical components, the photodetector catches on average the following amount of radiant energy:

$$\frac{E_i \cdot R \cdot A}{2\pi \cdot d^2} \quad (2.1)$$

This is an average because diffuse reflectors do not scatter light evenly, but instead generally follow Lambert's cosine law. Nonetheless we want to retain from this, that the radiant energy received by the photodetector is proportional to the area of the condenser lens and inversely proportional to the square of the distance to the surface.

The optical design of a LIDAR system additionally introduces the following requirements:

- In order to sweep over the surrounding scenery, the laser beam must be continuously moved around. 3D scanning is usually achieved by a combination of deviating the laser beam with mirrors and rotating the entire scanner.

- The light reflected by potential obstacles must be caught from whichever direction the laser beam is currently pointing at and directed onto the photosensitive device.

For the present system, the optical design was split into two parts:

- A first part collimates the laser beam and condenses the returning reflected light. It is also tasked with aligning the optical paths of the emitted laser beam and the returning light.
- Once the paths followed by the issued laser beam and the received light are aligned, they can be deviated simultaneously with a mirror, which is done by a second part. The mirror equally reflects the emitted and the returning light, without affecting the alignment, and reflected light is caught from the same direction the laser beam is pointing at.

The setup chosen for the first part is illustrated in figure 2.2. The main difference to the basic ranging setup seen previously in figure 2.1 is a mirror with a small hole in the middle. The narrow laser beam passes through this hole, whereas most of the returning light is reflected by the mirror and caught by the condenser lens. This aligns the optical paths of the emitted laser beam and the incident light reflected by some surface.

Some of the returning light will not be reflected by the mirror, but pass through the hole. This fraction of light will not traverse the condenser lens and consequently not reach the photodetector. If the diameter of the laser beam and hence that of the hole is kept small in relation to the diameter of the condenser lens, this loss will be negligible.

A second mirror is placed behind the condenser lens, redirecting the light onto the photodetector located underneath. That second mirror is required because the laser source and photodetector will be placed on the same electronic board in the overall design.

As mentioned before, the alignment of the emitted laser beam and the returning light will allow both of these to be deviated correspondingly. The deviation is realized with an additional mirror, which reflects the laser beam and the incident light likewise, scanning the surroundings in vertical direction.

If just one rotating mirror is used, there will be lengthy interruptions in the scanning process while the mirror's reflective surface is oriented away from the laser beam. A wheel of consecutive mirrors reduces interruptions. When the laser beam sweeps over the edge of the reflective surface of one mirror, it is again deviated by the succeeding mirror. Short interruptions occur only while the laser beam is in the gap between two mirrors.

For the present design a mirror hexagon was chosen. The concept for the second part of the optical system and its interaction with the first part is depicted in figure 2.3. The effect of the mirror hexagon is shown for three distinct angles, demonstrating the

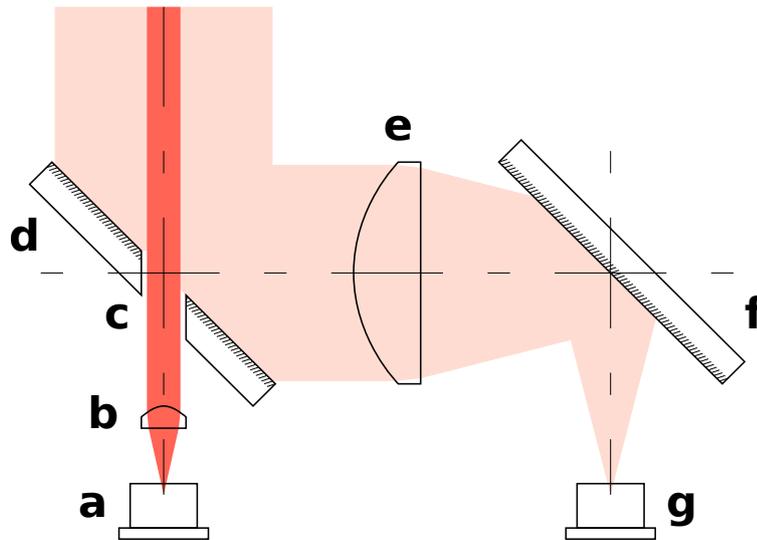


Figure 2.2 Setup to align the emitted beam and the returning light. This comprises **a**) a laser source, which emits a laser beam, shown in red, which is **b**) collimated by a small lens and **c**) escapes through a small hole drilled into a mirror. The laser beam eventually hits an obstacle, which casts back a portion of the emitted light. The returning light, represented in bright red, is **d**) reflected by the mirror with the hole, **e**) condensed by a large lens, **f**) reflected again by a second mirror and finally **g**) concentrated onto the sensitive surface of a photodetector.

deviation of the laser beam in various directions and how the returning light follows the same optical path as the emitted beam.

Since a hexagon is chosen, each mirror is exposed to the laser beam for a rotation angle of at most 60° . For specular reflection, which dominates for mirrors, the angle of incident light with respect to the surface normal equals the angle of reflected light. Hence the angular displacement of the redirected laser beam is twice that of the mirror. Therefore, the angle covered by the laser beam while being reflected by one mirror is at most 120° .

Due to the short interruptions between subsequent mirrors, the angle effectively covered by the deviated laser beam is slightly reduced. An angle of 90° is more practicable, with the sub-figures 2.3a and 2.3c showing the lower and upper limit of the vertical range, with deviation angles of -30° and 60° respectively.

Given the hexagonal layout, the laser beam will scan the vertical range six times during one revolution of the mirror wheel. A larger number of mirrors would increase this ratio, at the expense of vertical range.

Without additional horizontal range, the scanner is only able to compute a 2-dimensional cut of the surroundings along its scanline. In order to generate a 3-dimensional model, horizontal coverage is added by rotating the entire scanner. In summary, the chosen optical design allows for a vertical range of -30° to 60° and a full 360° horizontal coverage.

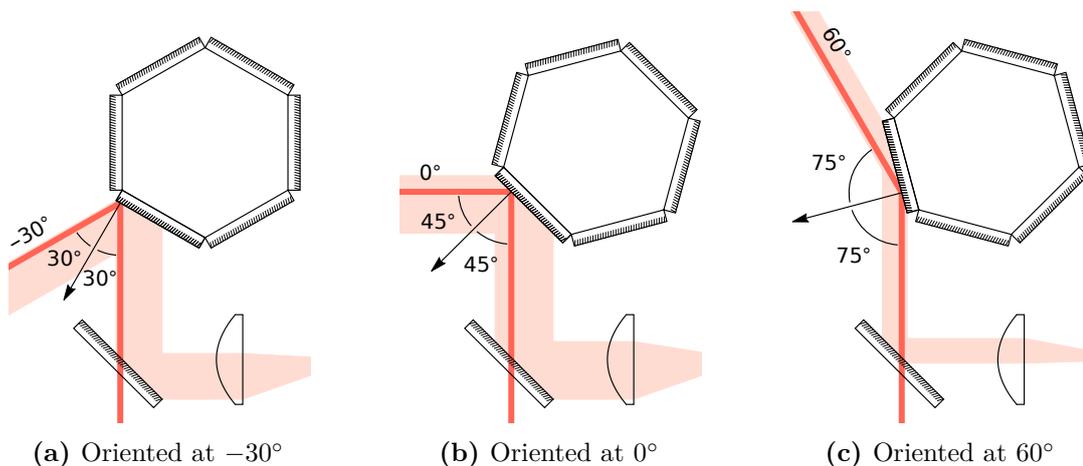


Figure 2.3 Mirror hexagon to redirect the emitted and returning light, illustrated for varying beam directions. This depicts how the emitted laser beam, shown in red, and the returning light, shown in bright red, are reflected by mirrors on a rotating hexagon, scanning the surroundings in vertical direction. Some details of the alignment setup are omitted.

Figure 2.4 shows a photograph of the assembled optical system. The only optical component which cannot be spotted is the lens located right above the laser source, which collimates the laser beam. That lens is hidden underneath the mirror with the hole, through which the beam escapes.

Finally, we want to estimate the amount of radiant energy which the optical layout will capture from light emitted by the laser source and reflected by a remote surface.

It should be noted that varying deviation angles will alter the amount of returning light received by the photodetector, since a steeper angle of the exposed mirror in the hexagon reduces the reflective surface seen by the condenser lens. This situation can be observed in sub-figure 2.3c, where only a narrow band of returning light can be reflected by the mirror hexagon and hence reach the photodetector.

The least light is captured when the laser beam is deviated 60° upwards, hence when the normal of the active mirror in the hexagon is at an angle of 75° with respect to the incident laser beam. Therefore the band of returning light reflected by the mirror in this position has a width of $l \cos(75^\circ) = l \cos(5\pi/12)$, where l is the length of the reflective surface of the mirror.

The scanner built as part of this work uses square mirrors with a side length of 1 inch (25.4 mm), thus the returning light reflected by the mirror hexagon has at least a width of $25.4 \cos(5\pi/12) = 6.574$ mm and a cross-sectional area of 167.0 mm². Not all of this light is captured by the condenser lens, some of it passes through the hole in the mirror below the hexagon. The hole has a diameter of 3 mm, thus an area of 7.07 mm² which needs to be deducted. The condenser lens is chosen large enough in order not to lose additional light. Consequently, the optical path which captures returning light and directs it onto

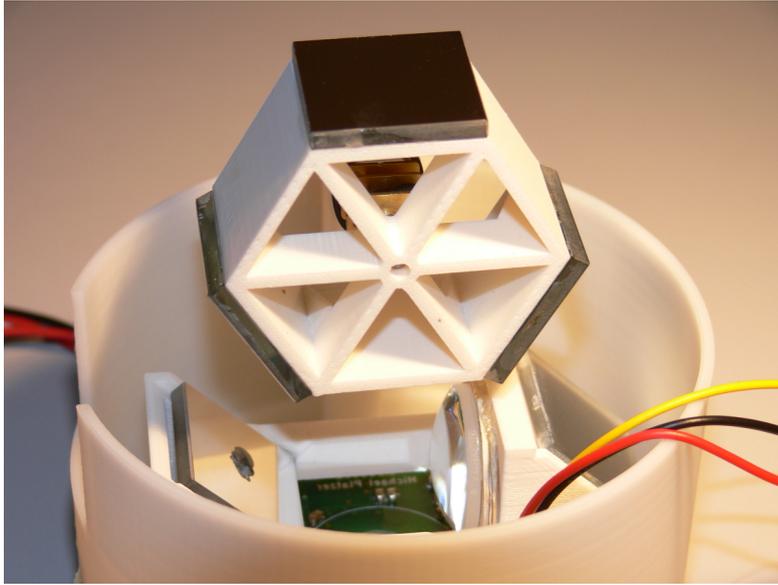


Figure 2.4 Photograph of the assembled optical system. The mirror hexagon is only equipped with 3 mirrors, instead of all 6. The small lens collimating the laser beam is hidden underneath the mirror with the hole. All optical components have been purchased at Thorlabs, Inc., the mounting parts have been 3D printed on a Stratasys, Ltd. Dimension Elite 3D printer.

the photodetector has a cross-sectional area of at least 159.9 mm^2 .

This area can be inserted into equation 2.1 to calculate the amount of radiant energy the photodetector will receive. To finally evaluate the equation, values for the distance d of the reflecting surface, its reflectance R and the incident radiant energy E_i must be selected. The reflectance of common surfaces such as concrete, soil or grass is generally between 0.25 and 0.45 for red and near-infrared light, hence it is assumed that all surfaces have a minimum reflectance of $R = 0.25$, which is used as a worst case value. Similarly a distance of 10 meters is set as the maximum distance up to which objects shall be detected. By inserting these values into equation 2.1, the radiant energy captured by the photodetector from light reflected by this surface evaluates to $6.37 \cdot 10^{-8} E_i$.

As stated earlier, the radiant energy incident to the surface E_i equals the energy emitted by the laser source, as minor losses in the optical system and the atmosphere are neglected. By using a powerful laser source one can make this value and hence the energy captured by the photodetector arbitrarily large. However, laser radiation is hazardous as it is capable of causing permanent damage to the human retina among other injuries. Therefore the energy emitted by the laser source is limited to a radiation power which is known to be safe.

The International Electrotechnical Commission (IEC) has published the standard IEC-60825-1 which sets limits on the radiation energy doses which are safe for the human eye, referred to as Maximum Permissible Exposure (MPE), and classifies laser sources

according to their hazard potential and the required precautions for safe use. The classification of a laser source essentially comprises radiation power, radiation density and exposure duration. Class 1 laser products are safe under all conditions since the radiation of these devices cannot exceed the MPE [IEC60825]. The present scanner shall conform to the limits for this class.

Although the laser source of the scanner is continuously emitting, the duration for which a human's pupil can be exposed to laser light is limited since the beam is rotated by the mirror wheel. According to IEC-60825-1 the human pupil has a diameter of at most 7 mm when fully opened, hence even if a human's eye is as close as 10 cm to the mirror wheel, its pupil only covers an angle of 4° . If the mirror wheel rotates at a speed of 10 revolutions per second, thus with an angular speed of $2\pi \cdot 10$ rad/s, the exposure time for each pass of the laser beam is 1.1 ms. For this exposure duration, an MPE of 10 mW/cm^2 is considered safe.

However, the exposure is repetitive, as each succeeding mirror on the wheel again diverts the laser beam along the same line and thus again across the pupil. For recurring exposures, the average radiation density additionally must be less than the allowed permanent exposure density [Sch15]. For permanent exposure, the MPE must not exceed 1 mW/cm^2 . Since the mirror wheel counts six mirrors, the angle between two consecutive mirrors is 60° and the laser beam rotates twice as fast as the wheel, hence from one wheel to the next it covers 120° . The 4° during which the laser beam hits a human pupil located 10 cm from the reflecting mirror occupy 3.3% of its total travel, thus the average energy density which hits the pupil is a mere 3.3% of the effectively radiated density, which therefore must be less than $1/0.033 = 30 \text{ mW/cm}^2$. This is a larger value than the MPE for one exposure and therefore needs not to be further considered.

To finally limit the power that a laser source is allowed to radiate, the standard introduces the concept of limiting aperture. Either the diameter of the laser source or the diameter of the human pupil limits the aperture by which light penetrates the eye. The diameter of the laser beam is 3 mm, which is less than the 7 mm of a fully opened human pupil, thus the limiting aperture is set to 7 mm, which corresponds to an area of 38.5 mm^2 . The MPE is multiplied by this limiting aperture to give the permissible radiation power for the laser source, which evaluates to 3.85 mW.

Setting the radiation power of the laser source with a 10% safety margin to 3.5 mW and inserting it into equation 2.1 finally gives the radiant energy captured by the photodetector from laser light reflected by a surface located 10 meters away and with a reflectance of 0.25, which is roughly 220 pW ($220 \cdot 10^{-12} \text{ W}$).

We want to retain the order of magnitude of the minimum captured radiant energy, as it is an important design parameter for the electronic subsystem, which is tasked with extracting the desired time-of-flight from this weak signal.

Finally, it should be noted that the reflection from the scanner's laser beam is not the only light captured by the photodetector. Ambient light is equally reflected by the scanned surface. Assuming that the optical path of light caught by the photodetector has

a divergence of 1° , all light reflected within a radius of 87 mm of the laser beam projected onto a surface which is 10 meters away adds to the laser's reflection, which is equivalent to an area 240 cm^2 . The spectral irradiance of sunlight is roughly $1.3 \text{ W/m}^2/\text{nm}$ throughout the visible part of the light spectrum, hence the spectral radiant energy hitting this area on a surface directly exposed to sunlight is 31 mW/nm . Therefore, additional radiation with a spectral power of 490 pW/nm is caught by the optical setup of the scanner when sunlight illuminates the previously defined surface with a reflectance of 0.25. By adding a narrow optical filter with a 10 nm bandpass region to the setup the resulting radiant power can be limited to about 5 nW ($5 \cdot 10^{-9} \text{ W}$).

2.2 Electronic Design

The electronic design comprises the circuitry required to drive the laser source, to amplify and filter the light signal captured by the photodetector and to compare the delay between the emitted and the captured signal. This delay is the time it takes the light to travel from the laser source to an obstacle and back to the photodetector, plus some additional delay arising in the circuits which will be accounted for during calibration. Knowing the time it takes light to travel from the scanner to a surface and its reflection back to the scanner will allow to calculate the distance to this surface.

In order to measure the delay between the emission and capturing of the light, the emitted signal needs to change over time, allowing these changes to be matched in the emitted and the captured signal. A simple approach is to use a square wave signal, merely switching on and off the laser source at a fast rate, to locate these switching transitions in the caught return signal and finally measure their delay with respect to the emitted signal.

The present work sticks with the square wave, as it allows for simple circuitry since this signal can be generated by switching action alone, whereas a signal with smooth transitions would require a more sophisticated circuit.

In addition to the signal shape, a signal period has to be chosen. The signal period is the time after which the signal waveform repeats, hence only delays shorter than this period are measurable. Therefore the signal period for a ranging system should be chosen large enough so that it is unlikely to capture light that has traveled for a longer time. However, the signal period should not be chosen too large either, since a large sample count and hence a high switching frequency are desirable for good coverage.

The present system is intended to measure distances up to 10 m. Light needs to travel this distance twice, once as it is emitted and once as it returns, hence the period must at least be chosen larger than the time it takes light to travel 20 m. However, a bright surface, i.e. a surface with high reflectance, might reflect enough light for it to be detected even if located at a greater distance which would lead to wrong measurements, as the delay would exceed the duration of one signal period and hence would be registered in

the next evaluation period. Therefore, the surface would be reported as being closer to the scanner than it really is.

Recalling equation 2.1, the energy of the received light is inversely proportional to the square of the distance. Hence a 10 times larger distance reduces the amount of captured radiant flux by a factor 100, which should be sufficient to assume that the amount of light reflected by surfaces located even further is too weak to be detected.

Accounting for these considerations the switching frequency of the laser source is set to 1 MHz for the present system, hence a signal period of 1 μ s, during which light covers approximately 300 m. Therefore, even for surfaces located up to 150 m from the scanner the delay of the returning light will not exceed the signal period. Capturing enough light to detect a surface situated so far away is unlikely.

As elaborated in section 2.1, the emitted radiant light power of the scanner must not exceed 3.85 mW to be compliant with the restrictions on class 1 laser products, hence the emitted light power will be set to 3.5 mW in order to retain a 10 % safety margin. This is the average power radiated by the laser source and since it is driven by a square wave signal and hence only lit half the time, the momentary light output power can be doubled to 7 mW. Note this is peak-to-peak amplitude, the effective amplitude of the square wave signal remains 3.5 mW. Now that all requirements for the emitted laser signal are set the actual circuit for the emission of the laser beam is considered.

The chosen laser source is an Arima Lasers ADL-65075TL AlGaInP laser diode with a wavelength of 650 nm, which corresponds to visible red light, and a light output power of up to 7 mW [ADL65]. Laser diodes are commonly used as sources for laser light and their behavior is similar to LEDs, likewise the light intensity is controlled by varying the current through the diode [YY97, sect. 6.7.3]. Choosing visible laser light allows for easier handling and debugging during prototyping and also increased safety.

Since a square wave signal has been chosen, the circuit driving the laser diode has only the task of turning it on and off. The microcontroller hosting the control subsystem generates the square wave signal which switches the laser diode. The current through the laser diode is controlled by a fixed series resistor and switching is done with a high-speed MOSFET transistor, as shown in figure 2.5.

Whereas the circuit driving the laser diode is quite simple, the circuit for processing the signal caught by the photodetector is more complex. Its task is to amplify this very weak signal with the goal to restore the square wave emitted by the laser diode, which will then allow to compare the phase of these signals.

2.2.1 Photodetector Signal amplification

The calculations in section 2.1 have shown that the amount of light captured by the photodetector for the present optical setup might be as low as 220 pW. It will be the task of the signal amplification circuitry to amplify this little amount of energy as cleanly as possible in order to restore the square wave signal for the subsequent measurements.

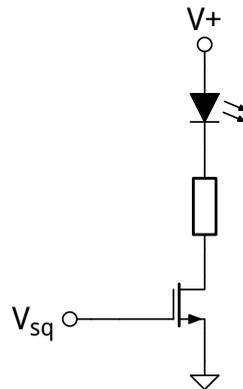


Figure 2.5 Circuit for driving the laser diode. V_{sq} is the square wave signal driving the diode by switching the MOSFET transistor. The current through the laser diode and hence its light output power is controlled by a fixed series resistor.



(a) Photodiode circuit symbol with an arrow indicating the direction of the photocurrent

(b) The photodiode can be modeled as an ideal diode in parallel with a current source

Figure 2.6 Operating principle of a photodiode. Incident light increases the leakage current, a current which flows in the reverse or blocking direction of the diode. This photocurrent is proportional to the intensity of light absorbed by the photodiode.

First a suitable photodetector has to be chosen. Semiconductor photodiodes, which are built on silicon substrate, are widely used for this purpose. A photodiode is similar to other types of diodes and as such has a $p-n$ -junction, which ideally allows current to flow only in one direction. Real diodes however are leaking, with a small amount of current flowing through the junction when a reverse voltage is applied. The absorption of photons from incident light in the $p-n$ -junction of a diode increases this leakage current. For a photodiode this effect is deliberately used and it therefore acts as a current source, with the current being proportional to the intensity of absorbed light [YY97, sect. 4.2].

This behavior is illustrated in figure 2.6a, which shows the circuit symbol of a photodiode, with an arrow indicating the direction of the current induced by captured light. As can be seen this current flows in the reverse direction of the diode. A photodiode can be modeled as an ideal diode in parallel with a current source, as shown in figure 2.6b.

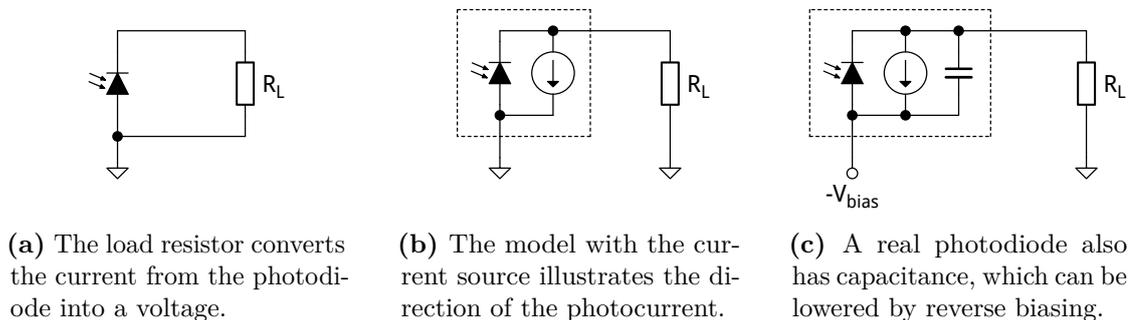


Figure 2.7 A load resistor as a simple photodiode front end.

So called PIN photodiodes have an intrinsic layer inserted in their p - n -junction, which greatly reduces their response time, therefore this type of photodiode is particularly suitable for high-speed applications. The name PIN photodiode refers to the sequence of layers: p -type silicon, *i*ntrinsic layer and n -type silicon [YY97, sect. 4.3].

The photodiode used for the present setup is a Hamamatsu S5972 PIN photodiode, suited for high-speed photometry, which has a stated photosensitivity of approximately 0.45 A/W at a wavelength of 650 nm [S5972], thus this photodiode will convert 220 pW of received radiant energy into a current of about 100 pA (0.1 nA, 10^{-10} A). The 5 nW of radiant power from sunlight reflection will be turned into a current of approximately 2.3 nA. According to the datasheet the photodiode might additionally experience up to 0.5 nA of dark current, which is not induced by light absorption. Hence the desired 0.1 nA signal rides on top of a base current of up to 2.8 nA.

The circuit between the photodetector and the signal processing system is referred to as front end. The front end converts the photocurrent into a voltage and performs amplification and filtering on the electronic signal [Hob09, ch. 18].

The front end for a photodiode might be as simple as a single resistor, connected as a load to the photodiode, as depicted in figure 2.7a, with figure 2.7b again illustrating the photodiode as an ideal diode in parallel with a current source. The load resistor converts the diode's current into a voltage, according to $V_o = R_L \cdot I_d$, where I_d is the photocurrent produced by the photodiode, R_L is the load resistor value and V_o is the output voltage, measured across the resistor. For instance a resistor of 1 G Ω would produce a voltage of 0.1 V when a current of 0.1 nA flows through it, which sounds like a reasonable amplification of the weak photocurrent and indeed this might be a suitable layout, if it were not for the poor frequency behaviour.

The photodiode has some capacitance, which must be added to the photodiode model as shown in figure 2.7c. This capacitance limits the available bandwidth for a given load resistance. The capacitance of the S5972 photodiode can be reduced to below 3 pF if a reverse bias voltage is applied to it [S5972]. Hence, disregarding parasitic capacitances

elsewhere than in the diode, the corner frequency for the simple load resistor front end is

$$f_{RC} = \frac{1}{2\pi \cdot R_L \cdot C_d} \quad (2.2)$$

This evaluates to 53 Hz for the exemplified $1 \text{ G}\Omega$ load resistor, which is far from the desired 1 MHz bandwidth that is required to amplify the square wave signal emitted by the laser diode.

As the capacitance of the photodiode cannot be further reduced, the only means of increasing bandwidth is to reduce the load resistance. For instance, with a load resistance of $50 \text{ k}\Omega$, equation 2.2 evaluates to 1.06 MHz, which is enough bandwidth for the present purpose. However, such a resistor would only give $5 \mu\text{V}$ of output voltage for 0.1 nA of photocurrent, which is a much lower value than would be suitable for further amplification, hence the load resistor front end is not satisfying and instead other front end types will be considered.

The most prominent front end for photodiodes is a transimpedance amplifier. A transimpedance amplifier, also called transresistance amplifier, converts current into voltage by forcing the current through a resistor with the help of an operational amplifier [JB10, sect. 10.9.2].

The basic circuit for a photodiode transimpedance amplifier is shown in figure 2.8. The photodiode is connected to the inverting input of an operational amplifier and a feedback resistor R_f is connected between this inverting input and the amplifier's output. The photodiode now draws its photocurrent from the node connected to the inverting input of the operational amplifier, which is usually referred to as the summing junction in an inverting amplifier circuit. The inputs of an operational amplifier do neither source nor sink current, hence all the photocurrent has to flow through the feedback resistor. At the same time, the operational amplifier adapts its output voltage to keep the inverting input at ground potential, as its non inverting input is connected to ground. Since the inverting input is kept at ground and all photocurrent flows through R_f , the output voltage is $V_o = -I_d R_f$, the same as for the load resistor front end, except with opposite sign.

As in the simple load resistor front end, the photocurrent still travels through a resistor, across which a voltage proportional to this current appears, but in the present configuration the work of the photodiode is facilitated by the operational amplifier, which sources the current while working to keep its inverting input grounded. This is called a virtual ground, since the node is not wired to ground, but actively kept at ground potential. Therefore the voltage across the photodiode is now nearly constant (the operational amplifier needs some time to adapt to changes in photocurrent), whereas it varied proportionally to the diode's current in the previous load resistor front end. The consequence is that the photodiode senses a much lower load resistance than the actual value of the feedback resistor. It would actually be 0 for an ideal operational amplifier [Hob01].

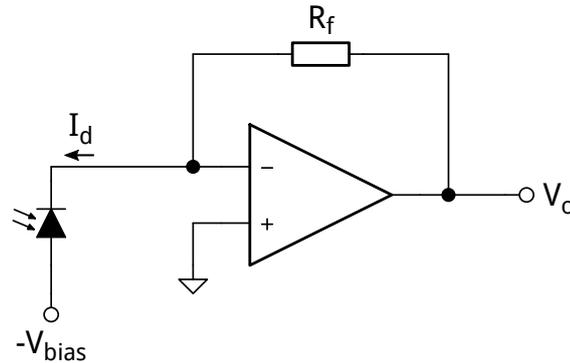


Figure 2.8 Circuit for a basic photodiode transimpedance amplifier. The photodiode current I_d cannot be sourced from the operational amplifier input, hence it is forced through the feedback resistor R_f . The operational amplifier works to keep the potential at its inverting input at ground, hence the output voltage is $V_o = -I_d R_f$. Simultaneously this holds the voltage across the photodiode nearly constant, which greatly reduces the resistance seen by the photodiode and therefore increases the bandwidth of the circuit.

As the photodiode sees less resistance than what its current is actually forced through, a larger resistor can be chosen as feedback resistor while retaining the desired bandwidth of 1 MHz. The frequency response of this circuit is somewhat more complex, as the bandwidth now largely depends on the operational amplifier [Hob09, sect. 18.4].

For an ideal operational amplifier the bandwidth would be infinite. However, real operational amplifiers do not only have limited bandwidth, but are affected by a number of other weaknesses. Before further evaluating the available bandwidth of the transimpedance amplifier circuit, it is advisable to consider these deficiencies, as some will require corrective measures which will further alter the circuit's frequency behavior.

When considering imperfect operational amplifiers, it is appropriate to remember that other devices in the circuit likewise are not ideal. The photodiode has capacitance and this capacitance connected to the inverting input of the operational amplifier introduces a major problem.

Unavoidable capacitances in operational amplifier circuits can cause undesired behavior in a variety of manners and often lead to instability, which manifests as wild oscillations. Capacitance at the inputs of an operational amplifier is generally particularly obnoxious [Kar00].

Signal amplification is characterized by amplitude gain and delay, which both depend on the signal frequency. In the frequency domain, this delay translates to a phase shift of the signal by some phase angle. The instability arises due to changes in the phase angle of the amplifier's closed loop gain, which are caused by the input capacitance. As the phase angle exceeds 180° for frequencies where the amplifier's gain is larger than 1, the

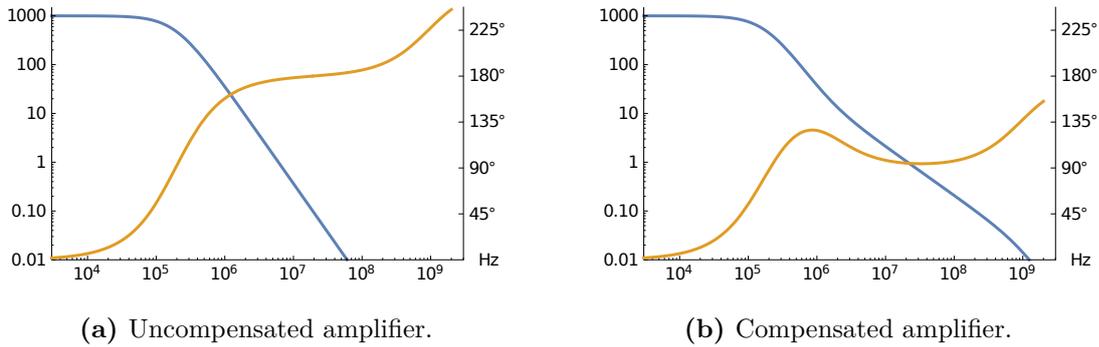


Figure 2.9 The effect of compensation on the loop transfer function of the transimpedance amplifier: The blue and yellow curves respectively show the gain and phase of the loop amplification, once for an uncompensated and once for a compensated amplifier. Instability occurs when the phase exceeds 180° at frequencies where the gain is larger than 1.

feedback of the output signal to the inverting input does no longer stabilize the amplifier, but instead leads to self-amplifying oscillations [Man04].

For the present transimpedance amplifier, input capacitance creates such a condition. The situation is visualized in figure 2.9a, which is a plot of the loop gain and phase of a transimpedance amplifier with a feedback resistance of $1\text{ M}\Omega$ and an input capacitance of 5 pF . As can be seen the phase shift nears 180° for a frequency of a few MHz, where the gain is still larger than 1.

In order to avoid instability due to input capacitance, the amplifier needs to be frequency compensated. For all frequencies which are amplified with a gain larger than 1, the phase angle needs to be smaller than 180° , preferably with some margin. The difference between the effective phase angle and 180° is called phase margin.

In Laplace space, the presence of input capacitance translates to the addition of a pole in the amplifier's gain transfer function. That pole is located at a frequency f_p given by

$$f_p = \frac{1}{2\pi \cdot R_f \cdot C_i} \quad (2.3)$$

Frequency compensation is realized by placing a compensating capacitor in the feedback network, in parallel with the feedback resistor, as shown in figure 2.10. The feedback capacitor introduces a zero in the gain transfer function, alleviating the effects of the pole. This zero should be located at a frequency f_z which is the geometric mean of the pole frequency f_p and the unity-gain frequency f_u of the operational amplifier in order to retain a 45° phase margin [MT059], which translates to following equation:

$$f_z = \sqrt{f_p \cdot f_u} \quad \Rightarrow \quad \frac{1}{2\pi \cdot R_f \cdot C_f} = \sqrt{\frac{f_u}{2\pi \cdot R_f \cdot C_i}} \quad (2.4)$$

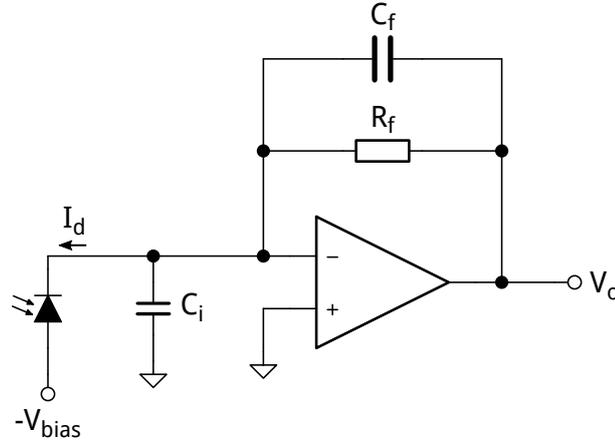


Figure 2.10 Frequency compensated transimpedance amplifier. The unavoidable capacitance in the operational amplifier's inverting input is modeled as the capacitor C_i in this circuit. It results essentially from the photodiode's and the amplifier's parasitic capacitance and requires frequency compensation, which is realized by placing a capacity C_f in parallel with the feedback resistor.

Hence, the value of the feedback capacitor C_f should be chosen as

$$C_f = \sqrt{\frac{C_i}{2\pi \cdot R_f \cdot f_u}} \quad (2.5)$$

The addition of this feedback capacitor changes the amplifier's loop transfer function, which is plotted in figure 2.9b. The phase is now kept below 135° up to frequencies of several 100 MHz, where the gain has long dropped below 1, thus matching the anticipated phase margin of 45° . This stabilizes the transimpedance amplifier and avoids oscillations.

Input capacitance as well as the addition of a feedback capacitor for compensation purposes alter the frequency response of the transimpedance amplifier. The transimpedance gain A_{TIA} , that is the amplification of input current to output voltage, of the compensated amplifier is given by

$$A_{TIA} = \frac{A_{VOL} \cdot Z_f}{1 + A_{VOL} + j \cdot \omega \cdot C_i \cdot Z_f} \quad (2.6)$$

where A_{VOL} is the open loop voltage gain of the operational amplifier, which depends solely upon it, and Z_f is the impedance of the feedback resistor and capacitor in parallel, hence $1/Z_f = 1/R_f + j\omega C_f$. Note that this gain has a complex value, as it contains amplitude and phase information, and depends on $\omega = 2\pi \cdot f$.

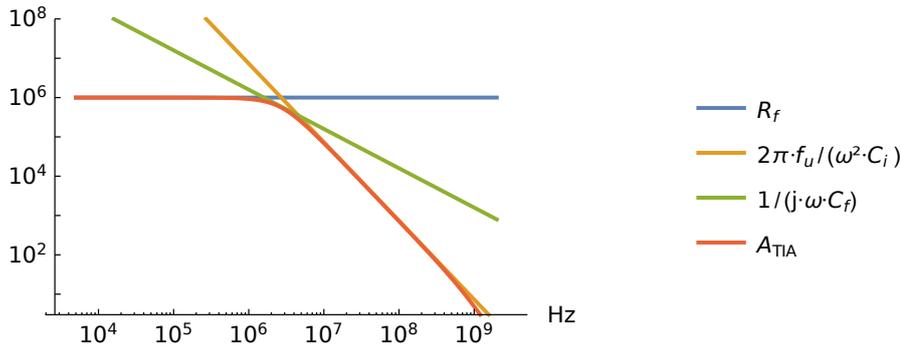


Figure 2.11 Transimpedance gain plot for a transimpedance amplifier with an OPA656 operational amplifier, 1 M Ω of feedback resistance, 5 pF of input capacitance and a feedback capacity of 0.1 pF. The blue, yellow and green curves show individual limits to the gain set by the feedback resistor, input capacitance and feedback capacitance, respectively. The red curve traces the resulting transimpedance gain. The 10^6 gain factor set by the feedback resistor is sustained up to a frequency of approximately 2 MHz, where feedback capacitance initiates the roll-off.

Figure 2.11 shows the magnitude plot of the transimpedance gain for a transimpedance amplifier built with a Texas Instruments OPA656 operational amplifier with a unity gain frequency f_u of 230 MHz [OPA656], a feedback resistor R_f of 1 M Ω and a total parasitic input capacitance assumed to be 5 pF (the capacity of the S5972 photodiode alone is 3 pF). In accordance with equation 2.5 a feedback capacitor C_f of 59 fF would be required, but since the feedback resistor has already a parasitic capacitance of 0.1 pF, no additional capacitor is used.

The plot allows to easily distinguish the frequency limitations caused by the capacitances in the circuit, which are shown by individual curves. The red curve is the resulting transimpedance gain, which is 10^6 up to a frequency of roughly 2 MHz, where the gain starts rolling off. At lower frequencies, the gain is set by the value of the feedback resistor as expected and for higher ones it is limited first by feedback capacitance and then also by input capacitance.

Despite all the capacitance and the large resistor value and thus the large gain, the transimpedance amplifier retains the required bandwidth to amplify the desired 1 MHz signal, giving an output amplitude of 0.1 mV for 0.1 nA of input current.

2.2.2 Noise Analysis

Up to now, it was assumed that an unobstructed photocurrent signal is cleanly amplified by the transimpedance amplifier. However, noise, which appears in all electrical circuits, mixes with and clouds the signal.

In electrical circuits the amount of noise is usually measured as noise density, which is the root mean square (RMS) amplitude of noise within a 1 Hz frequency band. Noise

density can be either expressed as voltage or as current density, which have the units $V/\text{Hz}^{1/2}$ and $A/\text{Hz}^{1/2}$ respectively. Multiplying these density values with the square root of the width of the frequency band gives the effective RMS amplitude. Note that filtering can reduce the bandwidth and hence the noise amplitude [HH15, ch. 8].

The relative amplitude between the desired signal and noise is called the signal to noise ratio (SNR), which is expressed in decibel (dB). For instance, if the signal amplitude is 100 times greater than the RMS amplitude of noise, the SNR is $20 \log_{10}(100) = 40$ dB.

The noise density of most noise sources varies over the frequency spectrum, gradually decreasing for larger frequencies. These types of noise, which are referred to as red noise and pink noise among others, are troublesome at lower frequencies, but for higher frequencies they become so weak that another type of noise dominates. Since the present circuit is dealing with a 1 MHz signal, the only type of noise which is relevant is white noise [HH15, sect. 8.1].

White noise has equal density over the entire frequency spectrum, hence it starts dominating all other noise types for all frequencies larger than some corner frequency, which is generally located within 100 Hz to 1 kHz. Two major noise sources are white: noise generated by thermal fluctuations in resistors, which is referred to as Johnson noise, and noise arising from statistical fluctuations of the discrete electrical charges, which is called shot noise.

In the transimpedance amplifier circuit there are three devices which are sources of white noise:

1. The photodiode itself generates shot noise as a result of the discrete nature of the electric charges which migrate across its p - n -junction. The amplitude of this noise depends on the current, hence it cannot be reduced for a given photocurrent by any means. Photodiode shot noise therefore sets an inescapable lower bound upon which the total noise obstructing the signal cannot be further reduced.

The current noise density of shot noise is given as

$$i_n = \sqrt{2q \cdot I} \quad A/\text{Hz}^{\frac{1}{2}} \quad (2.7)$$

where q is the electron charge¹ and I is the average current flowing [HH15, sect. 8.1.2]. For the 0.1 nA of signal current that shall be amplified by the present circuit this evaluates to $5.66 \text{ fA}/\text{Hz}^{1/2}$, however the shot noise level is set by the larger 2.8 nA current resulting from ambient light and dark current in the diode, which is $30.0 \text{ fA}/\text{Hz}^{1/2}$ ($30.0 \cdot 10^{-15} \text{ A}/\text{Hz}^{1/2}$). Assuming 1 MHz of bandwidth, the resulting RMS amplitude of shot noise is 30.0 pA, hence the SNR is a mere 10.5 dB.

2. The feedback resistor creates thermal noise, also called Johnson noise, of which the current noise density is given by

¹Electron charge: $q = 1.60 \cdot 10^{-19} \text{ C}$

$$i_n = \sqrt{4k \cdot T / R} \quad \text{A/Hz}^{\frac{1}{2}} \quad (2.8)$$

where k is the Boltzmann constant², T is the temperature and R the resistor value [HH15, sect. 8.1.1]. This formula shows that the current noise density produced by a resistor declines as the resistor value increases. Note that this is in opposition to the voltage noise density, which is given by

$$e_n = i_n \cdot R = \sqrt{4k \cdot T \cdot R} \quad \text{V/Hz}^{\frac{1}{2}} \quad (2.9)$$

and hence increases for larger resistor values.

3. Finally, the operational amplifier adds noise to the circuit, which is a mix of noise from various sources within the amplifier, consisting among others of shot noise from the many p - n -junctions in the device. Because of the complex nature of operational amplifiers, it is impossible to perform a complete analysis to calculate their noise behaviour. Instead one must rely on noise measurements conducted by the manufacturer.

The noise caused by an operational amplifier is modeled by two noise sources at the amplifier's inputs: one current source and one voltage source, which can be placed in either one of the inputs. The measured noise densities of both of these are usually listed in the amplifier's datasheet [AN358][AB76].

Figure 2.12 puts all these noise sources into perspective by adding them to the transimpedance amplifier circuit, modeled as one current and one voltage source. The current source is a combined model for the photodiode shot noise, the resistor's Johnson noise and the current noise of the operational amplifier. The total current noise $i_{n,tot}$ is

$$i_{n,tot} = \sqrt{i_{n,shot}^2 + i_{n,th}^2 + i_{n,amp}^2} = \sqrt{2q \cdot I_d + 4k \cdot T/R_f + i_{n,amp}^2} \quad \text{A/Hz}^{\frac{1}{2}} \quad (2.10)$$

where $i_{n,shot}$ is the photodiode's shot noise, $i_{n,th}$ is the feedback resistor's thermal noise and $i_{n,amp}$ is the operational amplifier's current noise. The voltage source models the operational amplifier's voltage noise and might be connected to either input for the purpose of noise analysis. It is arbitrarily placed into the non inverting input.

Nothing can be done about the shot noise in the photodiode, but the other noise sources should be quieted so as not to further worsen the situation. A photodiode amplifier is said to be shot noise limited when shot noise has the highest density level of all noise sources and hence limits the achievable SNR. This implies that the SNR cannot be further improved by circuit manipulation [Hob09, chap. 18].

²Boltzmann constant: $k = 1.38 \cdot 10^{-23}$ J/K

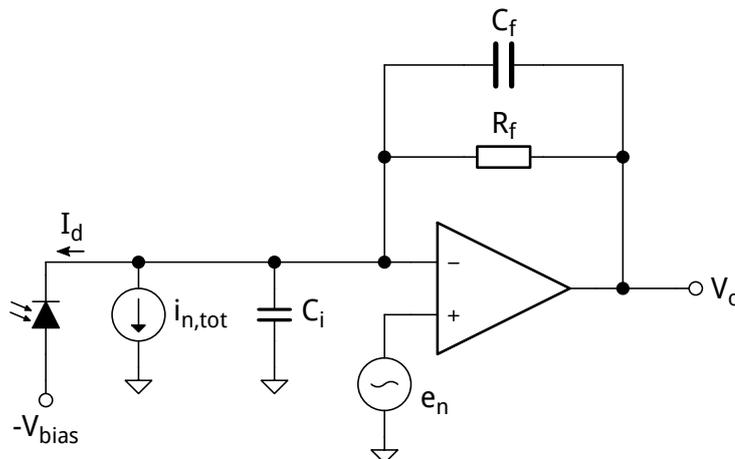


Figure 2.12 Noise sources in the transimpedance amplifier circuit. The current source $i_{n,\text{tot}}$ models the total current noise, which consists of photodiode shot noise, thermal noise from the feedback resistor and the operational amplifier's current noise. The voltage source e_n models voltage noise, which in this circuit arises solely from the operational amplifier.

For the present circuit to be shot noise limited all noise sources must have a noise density below the $30.0 \text{ fA/Hz}^{1/2}$ of photodiode shot noise. Concerning the operational amplifier, this means that a device with a FET or CMOS input stage must be selected, which generally have current noise densities of a few $\text{fA/Hz}^{1/2}$ only. Simultaneously, the feedback resistor R_f must be large enough to contribute less Johnson noise than the photodiode's shot noise, hence it must satisfy

$$\frac{4k \cdot T}{R_f} < 2q \cdot I_d \quad \Rightarrow \quad R_f > \frac{4k \cdot T}{2q \cdot I_d} \quad (2.11)$$

which evaluates to R_f being larger than $17.9 \text{ M}\Omega$ for the anticipated current of 2.8 nA .

While these requirements can be satisfied, although most likely with bandwidth issues, the impact of the operational amplifier's voltage noise poses a more significant problem. At first one might expect that the voltage noise, modeled as the voltage source e_n in figure 2.12, simply adds to the output signal, which is not of a big concern since the large feedback resistor turns the tiny photocurrent into a large output voltage which is no longer easily affected by noise. Although addition of the voltage noise to the output signal has only negligible effects, the voltage noise creates an additional noise contribution due to the presence of input capacitance, which is of much greater concern.

Input capacitance was already to blame for reduced bandwidth and instability problems and in combination with the operational amplifier's voltage noise it creates a noise problem on top of that. The operational amplifier forces the voltage at it's inverting

input to follow the voltage noise at its non inverting input, hence the voltage noise is equally present at the inverting input and the voltage fluctuations it produces contribute current noise, since a change in voltage induces a current in all attached impedances. The feedback resistor is large, thus the amount of current across it resulting from these voltage alterations is insignificant, however it is not the only impedance attached to the node. The parasitic capacitances contribute an additional impedance, which for high frequency signals is much lower than the resistance of R_f . This way, the amplifier's voltage noise contributes current noise to the photocurrent via

$$i_{n,e_n} = e_n \cdot C_i \cdot \omega = 2\pi \cdot e_n \cdot C_i \cdot f \quad (2.12)$$

where ω is the angular frequency and f the frequency of the signal [HH15, sect. 8.11.3]. The S5972 photodiode has a capacitance of 3 pF. Assuming that an operational amplifier with very low input capacitance is selected and hence the total capacitance at the inverting input can be kept below 4 pF, the operational amplifiers must have a voltage noise of less than $1.2 \text{ nV/Hz}^{1/2}$ in order to keep this noise contribution below the photodiode's shot noise at a frequency of 1 MHz.

While there are some bipolar operational amplifiers with such a low voltage noise, these have larger input current noise densities, in the order of several $\text{pA/Hz}^{1/2}$ and thus are not suitable. It turns out that there is not a single device which simultaneously meets both of these requirements [Hob09, sect. 18.4.4].

Since it is currently impossible to get an operational amplifier with both low enough current and voltage noise densities, something has to be done about the capacitance at the amplifier's inverting input. The parasitic capacitance of the individual devices, i.e. the photodiode and the operational amplifier, cannot be lowered, however the capacitance of the photodiode can be isolated from the amplifier by inserting a common base transistor in between them. This setup is illustrated in figure 2.13, where the Bipolar Junction Transistor (BJT) Q_1 separates the photodiode from the amplifier. The reason for adding resistor R_1 is discussed later, it can be ignored for the moment.

The transistor transmits any current drawn from its emitter to its collector, with only a small amount lost in base current. Therefore, the photocurrent flowing through the transimpedance amplifier is only marginally affected by this change. However, the capacitance of the photodiode is now isolated from the amplifier's input, which decreases the input capacitance. On the other hand, the parasitic capacitance between base and collector of the transistor has to be added, but by choosing a low capacitance transistor and a decent operational amplifier, such as a Linear Technology LTC6268 with an input capacitance of just 0.55 pF [LTC6268], the total input capacitance can be kept below 1 pF and the voltage noise under $5 \text{ nV/Hz}^{1/2}$, for a current noise contribution of approximately $30 \text{ fA/Hz}^{1/2}$ at a frequency of 1 MHz, hence right at the shot noise level.

While the common base transistor allows to overcome the issue of amplifier voltage noise which multiplied with input capacitance, it introduces its own complications, which are again bandwidth limitations and noise problems.

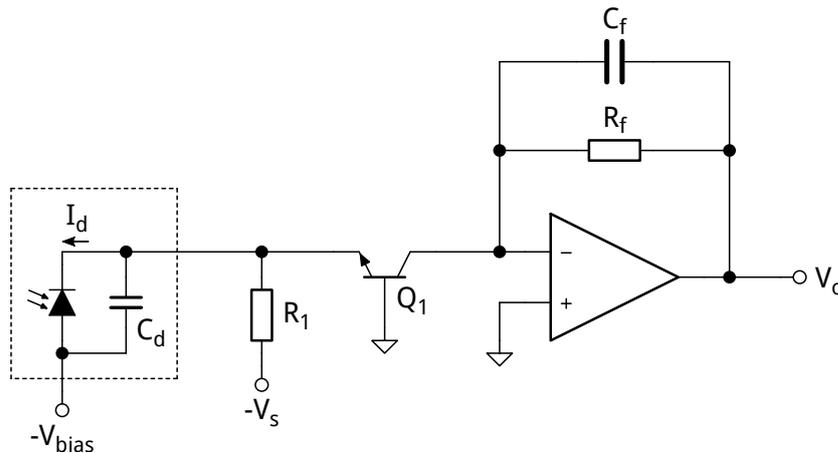


Figure 2.13 Cascoded transimpedance amplifier

First, the bandwidth is now limited by the input resistance of the transistor, which sets the corner frequency together with the photodiode's capacitance, according to equation 2.2. The small signal input resistance of a transistor, conforming to the Ebers-Moll model, is given by

$$r_e = \frac{k \cdot T}{q \cdot I_C} \quad (2.13)$$

thus it is inversely proportional to the collector current I_C . Therefore, a larger collector current means less resistance and hence a larger bandwidth [Hob09, sect. 18.4.4]. It turns out that for the photodiode's base current of 2.8 nA the transistor's resistance is 8.93 M Ω , which sets the corner frequency to approximately 6 kHz and hence is insufficient for the desired application.

One possibility of fixing this problem is to increase the collector current by adding resistor R_1 as shown in figure 2.13, thus creating a bias current across the transistor which lowers its resistance. For instance, a resistor of 10 M Ω connected to a -5 V voltage supply induces a current of 0.5 μ A through the transistor and thus reduces its resistance to 50 k Ω , increasing the bandwidth to 1.06 MHz. However, this resistor adds its thermal noise to the photocurrent, which has a density of 40 fA/Hz^{1/2}, hence more than the photodiode's shot noise. The resistor's noise contribution can be lowered by raising the supply voltage, which allows for a larger resistor value without reducing the current, but before addressing this, it is important to realize that the transistor also adds some noise of its own.

A transistor has p - n -junctions which produce shot noise, just as in the photodiode [Bax68]. First, base current shot noise is generated as a result of the base current $I_B = \beta I_C$ flowing through the base junction. Therefore it is important to select a transistor with a large

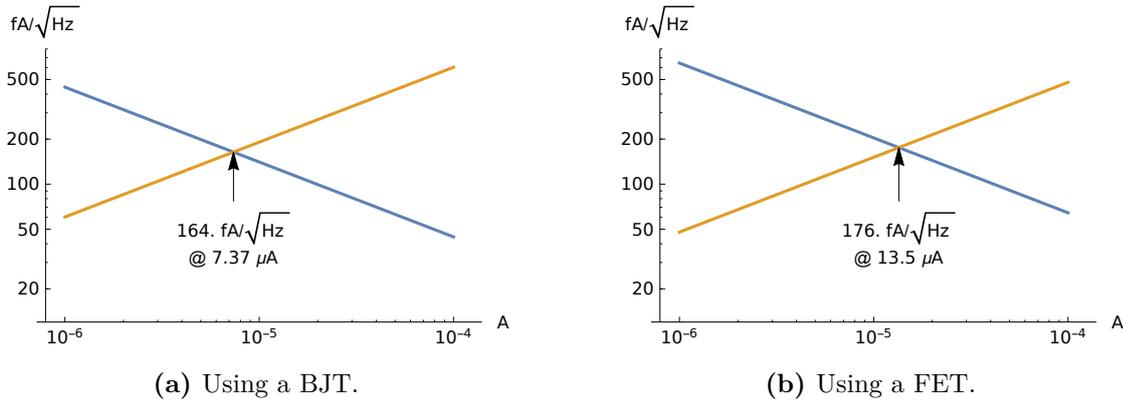


Figure 2.14 Noise in the cascode: The noise contribution of the transistor's voltage noise, represented by a blue curve, diminishes as the bias current is raised, however this requires reducing the value of the bias resistor, which increases its thermal noise, depicted by a yellow curve. These plots assume that the bias resistor is connected to a -5 V power supply.

β value, in order to keep the base current below the photodiode current of 2.8 nA and hence the resulting base current shot noise below the photodiode's shot noise. However, the β value of a BJT is proportional the collector current and for such a low current as the present one there is no BJT which retains a high enough β as would be required here. Using a Field Effect Transistor (FET) instead, which has a very low gate leakage current, overcomes this problem.

Secondly, a transistor has collector current shot noise (or drain current shot noise in the case of a FET), which shows up in parallel with the transistor's small signal resistance r_e and thus is turned into a voltage noise contribution given by

$$e_{nQ} = r_e \cdot n \cdot \sqrt{2q \cdot I_C} = n \cdot \sqrt{2} \cdot \frac{k \cdot T}{\sqrt{q \cdot I_C}} \quad (2.14)$$

where n is a factor which is 1 for a BJT and 1.05 for a FET, accounting for the different physical properties of these devices [HH15, sect. 3.3.1.C]. For a collector current (or drain current) of 0.5 μA , this evaluates to 20.0 nV/Hz^{1/2} for a BJT and 21.0 nV/Hz^{1/2} for a FET. As for the operational amplifier, the voltage noise multiplies with input capacitance to contribute to the total current noise density. For an input capacitance of 5 pF, they turn into current noise densities of 628 fA/Hz^{1/2} and 660 fA/Hz^{1/2} respectively at 1 MHz, which is more than 20 times stronger than the photodiode shot noise.

As can be seen in equation 2.14, the transistor's voltage noise can be reduced by increasing the bias current. However, this means a lower value for R_1 and hence more thermal noise from this resistor. This tradeoff is illustrated by the plots in figure 2.14, which show the noise contributions of the transistor and R_1 , once when using a BJT and once for a FET, with R_1 connected to a supply voltage of -5 V.

The plots visualize what has been shown arithmetically: as the bias current increases the transistor's noise contribution decreases, however this also raises the thermal noise of the bias resistor. The total noise contribution is the least where both noise densities are equal, hence at $7.37 \mu\text{A}$ for the BJT cascode and at $13.5 \mu\text{A}$ for the FET cascode, giving noise densities of $164 \text{ fA}/\text{Hz}^{1/2}$ and $176 \text{ fA}/\text{Hz}^{1/2}$ respectively, which is a lot more than the acceptable limit set by the photodiode's shot noise.

As mentioned earlier, raising the supply voltage allows to increase the transistor's bias current while retaining a larger resistance for R_1 , thus reducing the transistor's voltage noise and the resistor's thermal noise. However, for both noise sources to have a noise density below the photodiode's shot noise of $30.0 \text{ fA}/\text{Hz}^{1/2}$, a bias current $240 \mu\text{A}$ across the cascode transistor when using a FET and a resistance of R_1 of $17.9 \text{ M}\Omega$ is required. In order to force this current through the resistor, a supply voltage of 4300 V is needed, which is not realizable with ordinary components.

Another solution is to improve the cascode by adding another transistor which forms a feedback loop with the cascode transistor, as shown in figure 2.15. The second transistor Q_2 is used as a feedback amplifier which controls the gate voltage of the cascode transistor, such that the voltage across the photodiode changes less, thus lowering the input resistance of the cascode [SG90]. A similar circuit is used in fiber-optic receivers to achieve very high bandwidths in the order of a GHz [PY04].

The regulated cascode offers two important improvements:

- The input resistance of Q_1 is reduced by a factor equal to the voltage gain of Q_2 and hence is then given by

$$r_{in} = \frac{r_{e1}}{1 + \frac{R_2}{r_{e2}}} \quad (2.15)$$

This allows to reduce the bias current of Q_1 below $0.5 \mu\text{A}$ without limiting the bandwidth.

- The input noise voltage is no longer set by Q_1 , but instead by Q_2 , which can be operated at a higher drain current, thus reducing the voltage noise without increasing Q_1 's bias current.

For instance, when choosing a value of $20 \text{ k}\Omega$ for R_2 and connecting it to a positive supply voltage of 5 V while retaining the -5 V of negative supply voltage, Q_2 has a drain current of 0.5 mA , hence a resistance of $r_{e2} = 50 \Omega$ and a voltage noise of $664 \text{ pV}/\text{Hz}^{1/2}$. For R_1 a resistance of $20 \text{ M}\Omega$ is selected, which gives Q_1 a bias current of 250 nA and thus a resistance of $r_{e1} = 100 \text{ K}\Omega$. This sets the input impedance of the regulated cascode to 249Ω , giving a considerable bandwidth of 128 MHz . The current noise contribution of Q_2 evaluates to $20.9 \text{ fA}/\text{Hz}^{1/2}$ and R_1 adds $28.3 \text{ fA}/\text{Hz}^{1/2}$ of thermal noise.

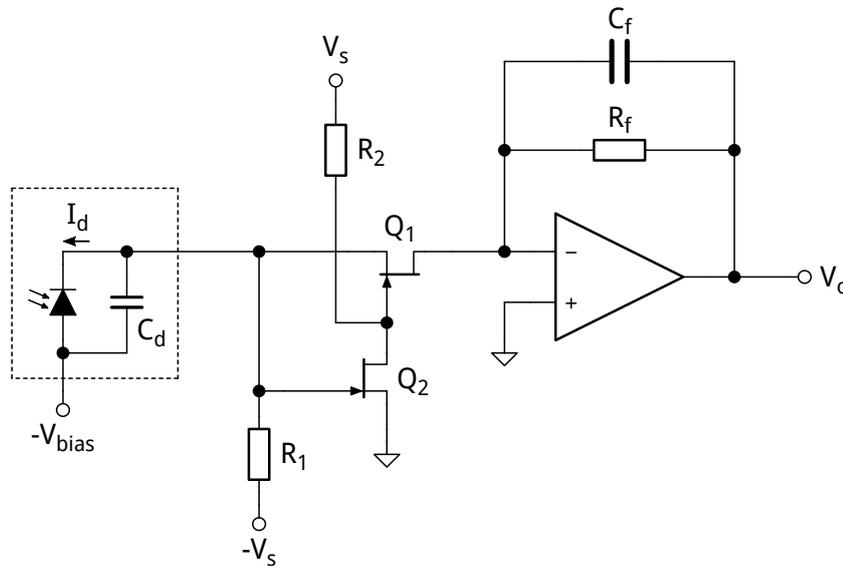


Figure 2.15 Transimpedance amplifier with a regulated cascode: The base of Q_1 is no longer connected to ground, thus its voltage noise is no longer of a concern. The noise contribution of Q_2 can be limited by applying a high collector current to it. The feedback loop formed by Q_1 and Q_2 also considerably lowers the input resistance of the cascode, hence the bias current can be reduced, thus allowing for a large and therefore quiet R_1 , without suffering bandwidth limitations.

All noise sources within this circuit add less noise density to the photocurrent than the photodiode's shot noise, hence it is shot noise limited. Therefore, the total current noise density is only insignificantly larger than the shot noise density and can be assumed to be equal. The cascode stage offers more than enough bandwidth for the anticipated 1 MHz signal. The last thing to reconsider is the frequency response of the transimpedance amplifier.

As revealed by equation 2.11, the feedback resistance must be larger than 17.9 M Ω to prevent the resistor's thermal noise from dominating over the photodiode's shot noise, however the gain curve plotted in figure 2.11 assumed a feedback resistance of just 1 M Ω .

Figure 2.16 again plots the amplifier's transimpedance gain, this time with a feedback resistance of 20 M Ω and the use of the LTC6268 operational amplifier, assuming an input capacitance of 1 pF and a feedback capacitance of 0.05 pF, which is achieved by using two resistors with a parasitic capacitance of 0.1 pF each in series. The minimum required feedback capacitance according to equation 2.5 is 4.0 fF, thus even this low value is largely sufficient. The plot reveals that the roll-off of the gain function now happens earlier. Although the feedback capacitance is reduced to a minimum, the gain already starts to decrease at 100 kHz.

However one should recall that the reason for choosing such a larger feedback resistance

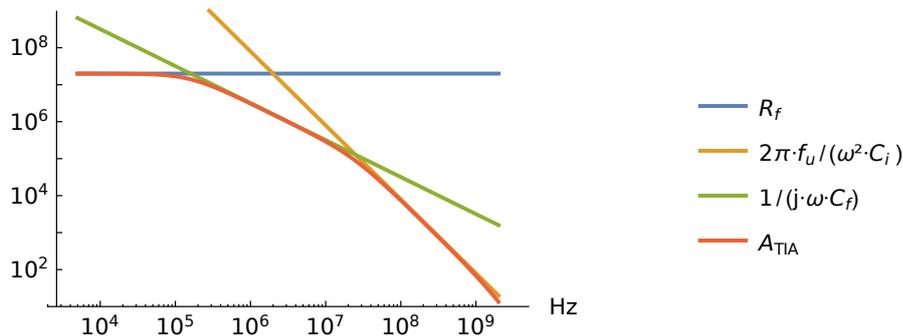


Figure 2.16 Transimpedance gain magnitude plot for a transimpedance amplifier with a regulated cascode input stage, using an LTC6268 operational amplifier for 1 pF of input capacitance, 20 M Ω of feedback resistance. Although feedback capacitance is lowered to 0.05 pF, it is to blame for the early roll-off which starts at about 100 kHz. Nevertheless, for a frequency of 1 MHz the gain is still larger than 10^6 .

was the reduction of its noise contribution, rather than the large gain it produces. For a frequency of 1 MHz, the transimpedance gain is still approximately $3.2 \cdot 10^6$, hence larger than the gain of 10^6 that could be achieved with the 1 M Ω feedback resistor of the initial amplifier layout. Therefore, the signal current of 0.1 nA is turned into a voltage of 0.32 mV and the 30 fA/Hz $^{1/2}$ of input noise current result in roughly 0.1 μ V/Hz $^{1/2}$ of output voltage noise density, however this is only true for 1 MHz noise. Because the transimpedance amplifier does unequally amplify different frequencies, the noise at the amplifier's output is no longer white as its spectral density is not constant anymore, but instead is given by

$$v_{n,TIA}(f) = A_{TIA}(f) \cdot i_{n,total} \quad (2.16)$$

and hence has a frequency behavior which corresponds to the amplifier's gain curve.

2.2.3 Square Wave Recovery

Before continuing with further processing of the photodiode's signal magnified by the transimpedance amplifier, it is appropriate to take a brief look at the alterations the amplifier causes to the waveform of the signal.

The laser diode emits a square wave by alternately being turned on and off, thus the light signal captured by the photodiode is also a square wave, assuming it originates from the emitted laser beam. As any periodic signal, a square wave is the sum of an infinite number of sinusoidal waves, each with a different amplitude and a frequency which is an integer multiple of the signal's frequency. These individual sine waves are called

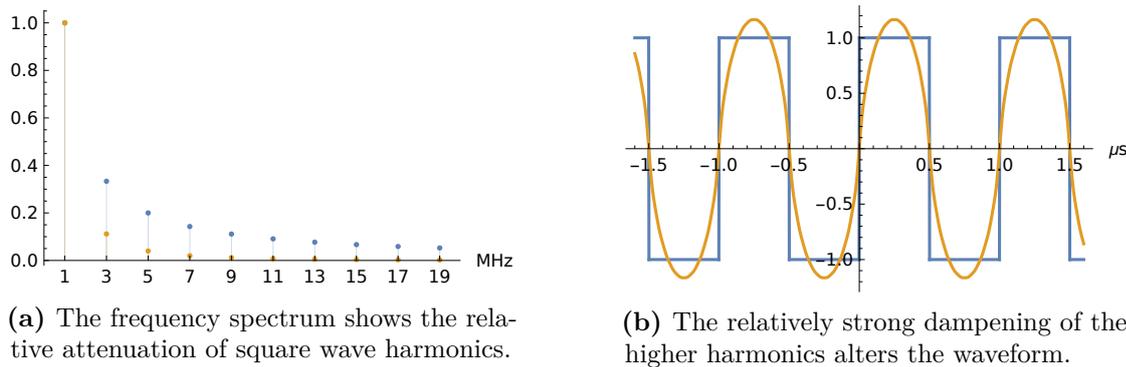


Figure 2.17 Alteration of the signal waveform due to the frequency response of the transimpedance amplifier: These plots show the harmonics and waveform of an unaltered square wave in blue, as well as the relative dampening of harmonics and the resulting waveform as the signal passes through the transimpedance amplifier in yellow.

harmonics. In the case of a square wave there are only odd harmonics, thus harmonics of which the frequency is an odd multiple of the square wave's frequency.

Given the very high bandwidth of the regulated cascode, the photocurrent signal from the photodiode remains a nearly perfect square wave until reaching the transimpedance amplifier. However, the frequency dependent gain of the latter unequally amplifies the harmonics of the square wave, thus changing its waveform, as can be seen in figure 2.17.

Note that in figure 2.17b the waveform of the dampened signal has a larger peak amplitude, which seems counter-intuitive. However, the signal power depends on the RMS amplitude, which is equal to the peak amplitude for a square wave signal, but for a sine wave the peak amplitude is larger than the RMS amplitude by a factor of $\sqrt{2}$. Therefore the peak amplitude of the resulting signal is larger although its RMS amplitude is slightly decreased.

The transimpedance amplifier has turned the current signal issued by the photodiode into a voltage. The original square wave signal emitted by the laser diode has been turned into a sinusoidal waveform, with an amplitude of at least 0.32 mV, depending on the amount of light which is reflected and caught by the photodiode, and a phase shift with respect to the initial signal as a result of the time-of-flight of the laser light and delays in the circuit.

The amplitude of the signal is independent from the time it takes the light to travel to some surface and back again, only the phase matters when it comes to determining the distance covered by the laser beam. Therefore it is desirable to recover the original square wave signal from the received signal, removing the unnecessary amplitude and only keeping the phase information, which facilitates the subsequent phase shift measurement.

This is done by means of a comparator circuit. As shown in figure 2.18a, a comparator is essentially an operational amplifier without feedback, hence upon the slightest voltage

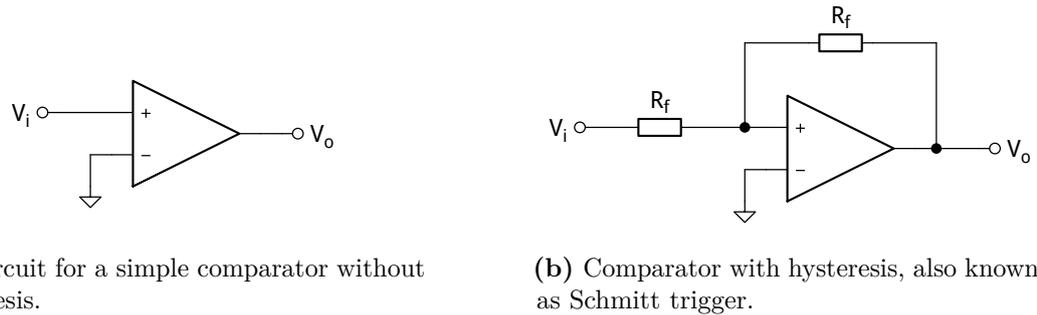


Figure 2.18 Basic comparator circuits.

difference between its inputs the amplifier will immediately drive its output to the highest or lowest voltage available.

For comparison figure 2.18b shows a comparator with positive feedback. The positive feedback creates a hysteresis, meaning that the comparator will not switch as the input signal crosses 0, but slightly later, as the input reaches a threshold voltage. There is one threshold voltage for positive switching and one for negative switching. While the input signal remains within these thresholds, the comparator does not switch, but remains in its current state. Only a sufficiently larger change of the input will cause the comparator to switch.

Figure 2.19 shows the benefits of using a hysteresis when a comparator is driven by a noisy input signal, by comparing the behavior of a comparator without hysteresis and one with hysteresis. The blue curve in the two plots is a sinusoidal input signal with noise added to it. The yellow curve shows the waveform of the comparator output. As the signal in figure 2.19a is close to 0, the noise causes the signal to cross 0 multiple times and therefore the comparator switches multiple times within a short period. In figure 2.19b the threshold voltages shown by two red lines cause the signal to switch only once per zero-crossing of the sinus. The noise is too weak to alter the comparator output on its own.

However, hysteresis also delays the switching, which does no longer happen at the zero-crossings, but slightly later, as the signal crosses the threshold values. Even worse, this delay in the output introduced by hysteresis depends on the signal amplitude. If the slope of the signal at the threshold voltage is steeper, the delay is shorter. If the slope is more gentle, the delay is longer.

Amplitude information is lost during the comparator stage, therefore a delay which varies with amplitude cannot be corrected and irreversibly corrupts the measurements. Hence the delay caused by comparator hysteresis must be kept to an acceptable minimum. The delay caused by a hysteresis voltage v_{th} is given by

$$\Delta t_{hys} = \frac{v_{th}}{SR} \quad (2.17)$$

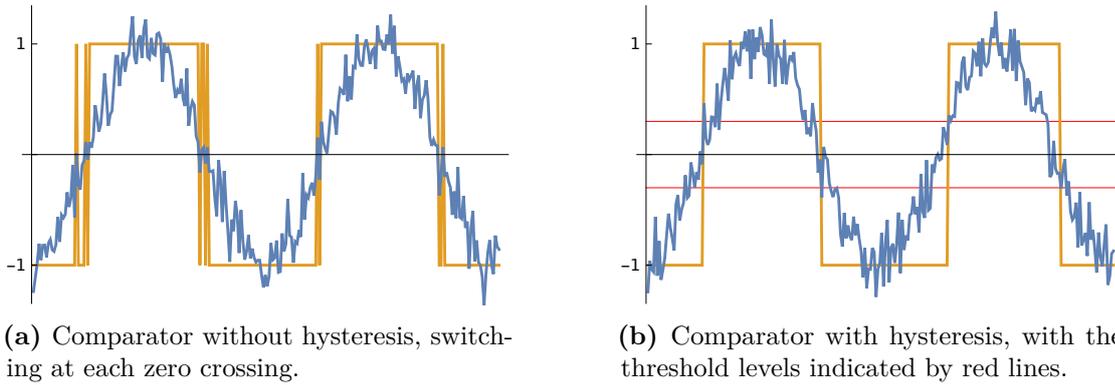


Figure 2.19 Effects of hysteresis on a noisy signal.

where SR is the slew rate, commonly referred to as slope, which is the rate by which the signal value changes with respect to time.

For $v_{th} \ll v_{sig}$ the switching happens when the signal is close to its zero-crossings. The slope of a sine wave at the zero-crossings is 2π times the signal period, multiplied with the signal's peak amplitude. For the given signal, the slope is slightly larger, since it is composed of multiple sine waves, but for a worst case estimation it is assumed to be equal. The peak amplitude of the signal is $\sqrt{2} v_{sig}$, hence the slew rate as the comparator switches is assumed to be $SR = 2\pi \cdot f \cdot \sqrt{2} v_{sig}$.

As anticipated, the slew rate is proportional to the signal amplitude and hence the signal time shift caused by the comparator's hysteresis is inversely proportional to the signal's amplitude. Once the signal has passed through the comparator, the amplitude information is lost and thus there is no way to tell apart the delay caused by hysteresis from the delay resulting from the time-of-flight of light. Therefore, to avoid measurement errors, the hysteresis voltage must be limited.

In order to impose a limit on the maximum time shift due to hysteresis which can be tolerated, one needs to consider how this shift affects the distance calculations performed on the delay measurements. The difference in the calculated distance for a given time shift induced by hysteresis is $\Delta d_{hys} = \Delta t_{hys} \cdot c/2$, where c is the speed of light³. Hence, to keep this difference below a tolerable maximum error value Δd_{max} , the threshold voltage must be set as follows:

$$\Delta d_{hys} < \Delta d_{max} \quad \Rightarrow \quad \Delta t_{hys} < \frac{2 \cdot \Delta d_{max}}{c} \quad \Rightarrow \quad v_{th} < \frac{4\pi \cdot f \cdot \sqrt{2} v_{sig} \cdot \Delta d_{max}}{c} \quad (2.18)$$

For instance, if the error caused by the comparator's hysteresis shall be kept below 5 cm, this evaluates to $v_{th} < 2.96 \cdot 10^{-3} \cdot v_{sig}$, thus for an input signal with a RMS amplitude

³Speed of light in vacuum: $c = 299\,792\,458$ m/s

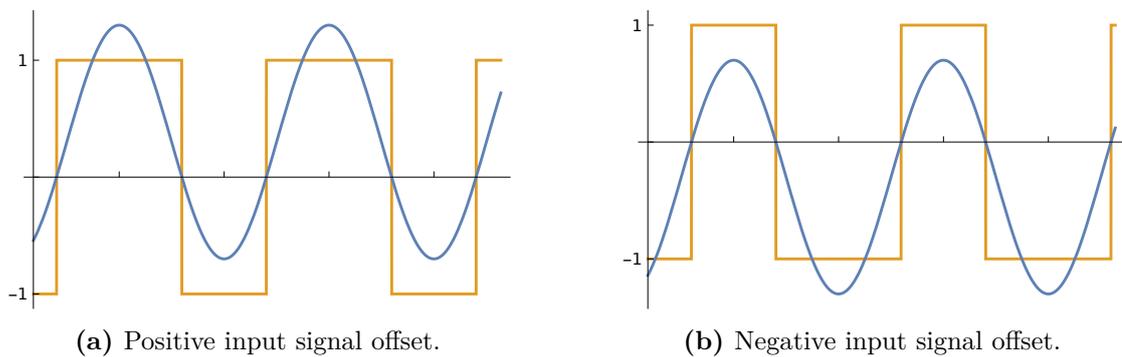


Figure 2.20 Effects of input signal offset on comparator operation. A positive offset of the input signal causes the output high state to last longer than the low state, a negative offset causes the low state to last longer than the high state. However, both states remain centered around the moments when the input reaches local extrema.

of 0.32 mV the hysteresis threshold voltage must be kept below 0.95 μV . This is a much lower voltage than the noise.

Since hysteresis causes an undesirable amplitude dependent delay of the produced square wave, it seems reasonable to simply use a comparator without hysteresis. The electronic design used for the present scanner indeed does not intentionally add a hysteresis to the comparator circuit. Nonetheless, the effects caused by hysteresis must be accounted for, since a non ideal comparator will always have a slight hysteresis which cannot be removed.

Another property of non ideal comparators is input offset. This means that the comparator does not switch as the input signal crosses 0 (or at threshold voltages equally far from 0), but instead it switches at a slightly different voltage. This effect can be modeled by an input signal which is offset by some amount.

Figure 2.20 shows the effects of input signal offset, once for an input signal which is raised and once for an input signal which is lowered. In both plots the blue curve is the comparator's input and the yellow curve is the output signal. Figure 2.20a shows the behavior for positive offset and figure 2.20b for negative offset. Due to input offset the produced signal is no longer a perfect square wave, as it would be if the input was a centered sinusoidal signal. Instead, the duration of the high state of the signal is unequal from the duration of the low state. Note though, that the output signal's high and low states are centered around the same temporal marks in both plots, which correspond to the moments when the sinusoidal input waveform reaches local extrema.

The comparator switches to its high state earlier if the signal offset is larger, whereas the switch to low state occurs later. The amount of time by which switching is shifted due to non zero input offset is similar to the delay caused by hysteresis, given in equation 2.17. As will be seen in the next section, one can take advantage of the fact that the high and low states remain centered around the same values for any input offset. This property

will allow to correct the error introduced by input offset in a later stage.

Whether the comparator has a hysteresis or not, noise on the input signal causes it to switch slightly earlier or later as it would without noise. This happens because the noise slightly lowers or raises the signal level. Therefore, the noise causes time jitter on the comparator output. This time jitter induced by noise is random, since the instantaneous voltage level of noise itself is random. The amplitude of time jitter depends on the signal slope, as for the delay caused by hysteresis, and the noise amplitude.

The RMS value of time jitter is given by

$$t_n = \frac{v_n}{SR} = \frac{v_n}{2\pi \cdot f \cdot \sqrt{2} v_{sig}} \quad (2.19)$$

where v_n is the RMS noise voltage.

Assuming the instantaneous noise voltage follows a Gaussian distribution, time jitter also is normally distributed, with its standard deviation given as $\sigma_t = t_n$.

Time jitter will manifest in the subsequent time-of-flight measurements and corrupt the distances calculated from these. The standard deviation inflicted upon the distance calculations is given by

$$\sigma_d = \frac{c \cdot \sigma_t}{2} = \frac{c \cdot v_n}{2 \cdot SR} = \frac{c}{4\pi \cdot \sqrt{2} \cdot f} \cdot \frac{v_n}{v_{sig}} \quad (2.20)$$

which is $\sigma_d = 16.9 \cdot v_n/v_{sig}$ for the selected frequency of 1 MHz. For a SNR of 20 dB the individual distance measurements will thus have a standard deviation of 1.69 m.

2.2.4 Time-of-flight Measurement

Once the waveform of the signal captured by the photodiode has been restored to its original square wave shape, the delay of the recovered signal with respect to the signal emitted through the laser diode can be measured. This delay is the sum of the time-of-flight of the laser light from the emission by the laser diode until the capture by the photodiode and additional delays due to the non zero response time of the various electrical components. The latter is assumed to be constant, hence the delay can be corrected by a constant term to get the time-of-flight of the laser beam.

The delay is measured by feeding both the emitted and the captured signal into an AND gate and integrating the result. Figure 2.21 presents an overview of the circuit used for this purpose. It comprises all the components encountered so far, although some details are omitted for simplicity. The voltage V_{sq} drives the laser diode, located in the upper left corner of the figure. Below, the photodiode captures the returning light. A single amplifier symbol is shown in place of the cascoded transimpedance amplifier and subsequent amplification and filter stages. The amplified photodiode signal is turned into

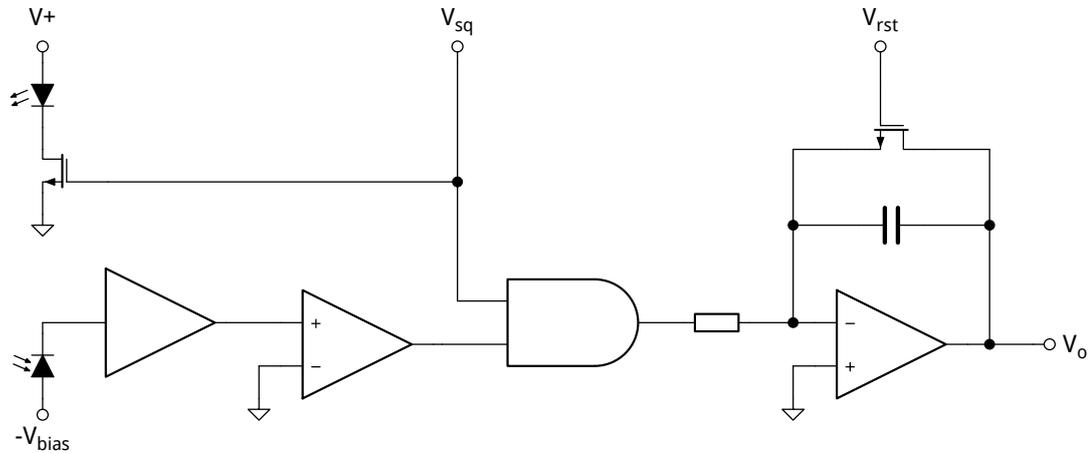


Figure 2.21 Overview of the time-of-flight measuring circuit. V_{sq} drives the laser diode in the upper left corner. Below, the photodiode captures the returning light. The photocurrent is amplified and passes through a comparator, recovering the initial square wave. The emitted and the received square wave signals are both fed into an AND gate and then integrated.

a square wave by the comparator and fed into the AND gate, together with V_{sq} . Finally, the output of the AND gate passes through an integrator, which is periodically reset by V_{rst} .

Figure 2.22 shows the various waveforms this circuit produces for two different signal delays. The red curve on top is the V_{sq} signal, the blue curve underneath is the recovered square form of the captured signal, measured at the output of the comparator. The orange curve shows the output $V_{\&}$ of the AND gate, which is the logical AND of the two signals above it. The green curve at the bottom is the output V_{int} of the integrator, which is reset at the beginning of each period of V_{sq} . As visualized in these plots, the voltage produced by the integrator is inversely proportional to the signal delay.

The output of the integrator is continuously raised while its input is high and kept nearly constant as the input is low. When the delay of the captured signal is small, the high states of the emitted and received signals overlap almost completely, thus the AND gate is kept high for almost as long and the integrator output has time to grow to a high value. As the delay increases, the overlap of the two high states shrinks, therefore the AND gate is high for a shorter time and the integrator output cannot grow to an as high value as before.

When the received signal is shifted by half a period with respect to the emitted signal, the high states do not overlap at all. In this situation the output of the AND gate is always 0, thus the integrator output equally remains 0. As the signal delay further increases, the high state of the captured signal starts overlapping with the high state of the emitted signal's next period. The AND gate output will again be high during the overlap and

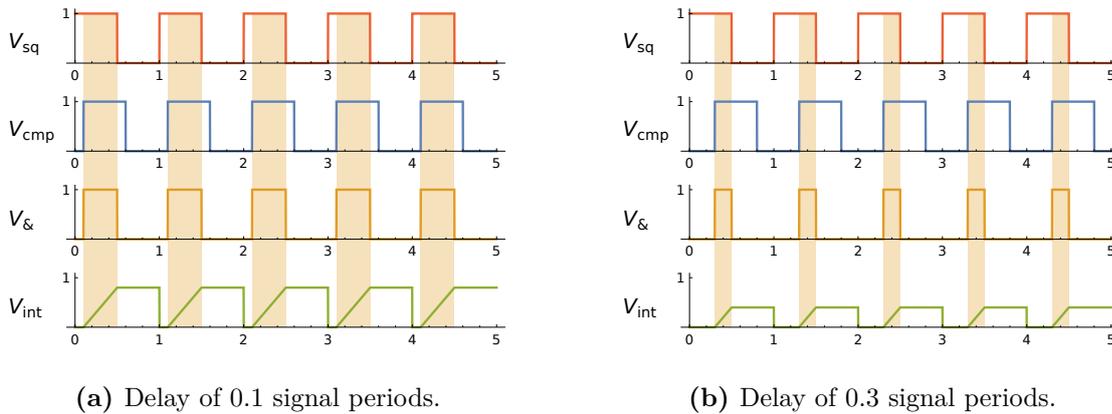


Figure 2.22 Signal waveforms in the delay measurement circuit: The red curve is the emitted square wave signal V_{sq} , the blue curve is the captured square wave signal at the comparator output V_{cmp} . Both are fed into an AND gate, the output $V_{\&}$ of which is the orange curve and also indicated by light orange bars. Finally, an integrator which is reset at the beginning of each period produces a voltage V_{int} inversely proportional to the delay.

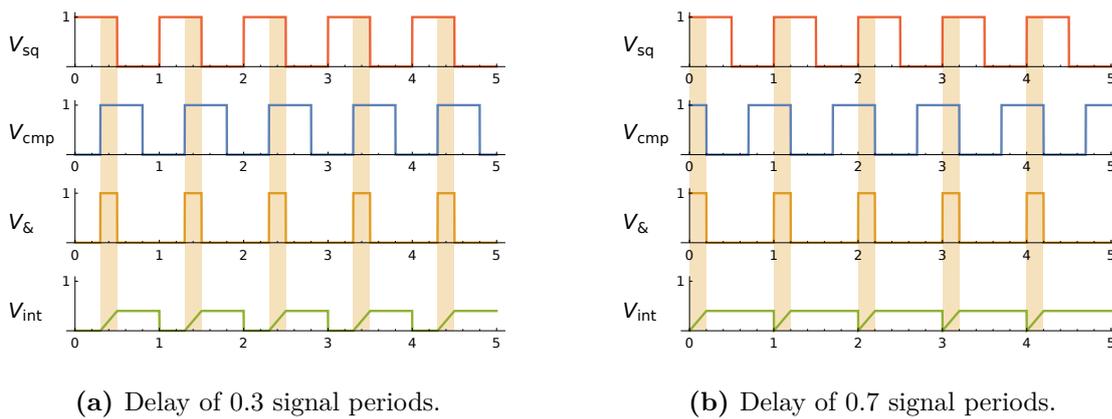
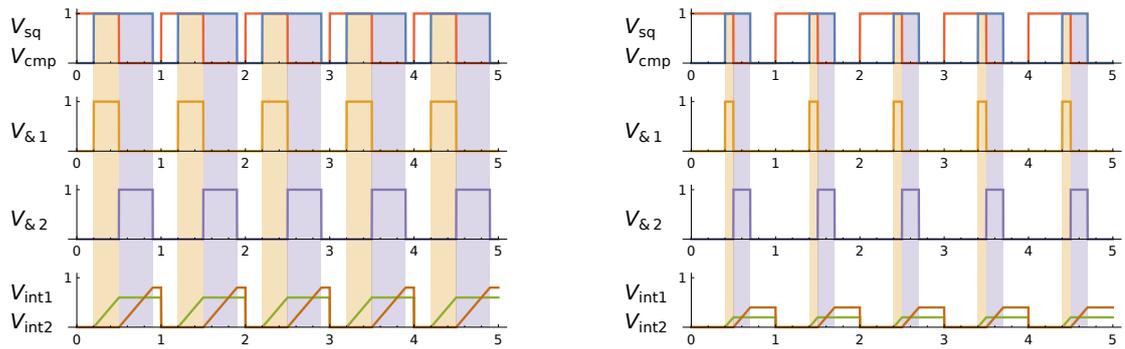


Figure 2.23 Limits of this delay measuring method: The output of the integrator reaches 0 V as the signal delay equals half a signal period. If the signal delay increases further, the voltage rises again, making it indistinguishable from smaller signal delays.

hence the integrator output grow again.

Figure 2.23 compares the circuit's response for a delay inferior to half a signal period and one delay larger than half a period, which both happen to produce the same voltage at the integrator output. Therefore, the maximum signal delay which can be measured is a shift by 180° .

As mentioned in section 2.2.3, it is likely that the comparator has a slight offset, which corrupts the conversion of the captured signal into a square waveform. This causes produced waveform to have unequally long high and low states, which would corrupt the



(a) Delay of 0.3 signal periods, with a high state lasting 0.7 signal periods.

(b) Delay of 0.3 signal periods, with a high state lasting 0.3 signal periods.

Figure 2.24 Correcting comparator offset during delay measurement. Comparator offset causes unequal duration of high and low states, however the signal remains centered around the same mean delay. Using two AND gates, with the second one taking the inversion of the output signal V_{sq} as input, allows to correct this offset. The output level of the two integrators may change, but the difference remains the same.

measurements produced by a single integrator as shown in figure 2.21.

However, the usage of a second integrator can correct this error. The second integrator is also fed by an AND gate. This second AND gate in turn is fed with the comparator output and the signal driving the laser diode, similar to the first AND gate, but with the difference that the signal driving the laser diode is inverted. Therefore, the output will be high during the overlap of the output signal's low state with the received signal's high state.

The waveforms produced by such a circuit are shown in figure 2.24. The red curve is again V_{sq} , the square wave driving the laser diode. It is shown in the same plot as the output of the comparator V_{cmp} . This signal is illustrated by a blue curve and the waveform shown on the left remains longer in high state than in low state, whereas the waveform on the right has longer low states than high states. However, the mean delay of both waveforms is the same. The orange curve is the output $V_{\&1}$ of the first AND gate, which is no different from the previous plots. The purple curve is the output $V_{\&2}$ of the second AND gate, which is high during the overlap of the output signal's low state with the captured signal's high state. Finally, the green and brown curves show the output waveforms of the two integrators respectively. The unequal duration of the captured signal's high state in the left and right plots causes the output of the integrators to be raised to different levels. Yet, the difference between the output of the two integrators is the same for both.

Measuring the difference between the output of the two integrators can be done in one step by taking advantage of a differential Analog-to-Digital Converter (ADC). This is an electrical device which has two differential inputs and which converts the voltage

difference between these two inputs into a digital measurement.

In contrast to analog values, which are continuous, digital values are discrete. Therefore the ADC will map the input voltage difference to a finite set of discrete measurement values, which limits their precision. If there is no delay between the emitted and the captured signal, the ADC will map the voltage to the lowest value in this set. If the delay is half a period, which is the maximum that can be measured with the present circuitry, the ADC will map the voltage to the highest possible value. Therefore, the signal phase shift required to raise the ADC output to the next higher value is $180^\circ/n$, where n is the amount of different digital values in this set.

The delay duration of one ADC step and the corresponding distance are given by

$$\sigma_{t,ADC} = \frac{1}{2 \cdot f_{sig} \cdot n} \quad (2.21)$$

$$\sigma_{d,ADC} = \sigma_{t,ADC} \cdot \frac{c}{2} = \frac{c}{4 \cdot f_{sig} \cdot n} \quad (2.22)$$

For the present circuit an ADC with 12-bit precision is used, thus mapping the analog input to $2^{12} = 4096$ discrete output values, which gives a precision of 1.8 cm for a 1 MHz signal frequency.

The ADC marks the conversion of the formerly analog signal to a digital signal. Here ends the signal conditioning performed by the electronic circuitry, any further processing of the measured values is done in software.

2.3 Controller Design

Before investigating any further processing of the time-of-flight measurements generated by the ADC, it is appropriate to take a look at the controller which orchestrates the various signals in the electronic circuitry and also directs the mechanics required to gradually vary the orientation of the beam in order to scan the surroundings.

The controller regulates the rotational speed of the mirror wheel and the entire scanner unit. The mirror wheel is a hexagon, as seen in section 2.1, thus for each revolution 6 vertical scans are carried out. The rotational speed of the mirror hexagon is higher than that of the entire scanner, thereby numerous vertical scans are performed for each revolution of the whole unit. As the scanner completes one revolution, the lined up vertical traces combine to form a 360° round scan.

The horizontal resolution of such a round scan is the number of vertical scans performed during one revolution of the scanner, thus it is given by $6 \omega_m / \omega_s$ where ω_m is the rotational speed of the mirror hexagon and ω_s is the rotational speed of the entire device.

The vertical resolution is the number of measurements generated during one vertical scan. As seen in section 2.1, a vertical scan covers an angle of 90° , with the mirror wheel

only moving by half that angle, thus 45° , which is $1/8$ of a full revolution. Hence the duration of one scan is $1/8\omega_m$ and the vertical resolution is $f_{meas}/8\omega_m$ where f_{meas} is the frequency by which measurements are generated.

Each distance measurement is associated with a corresponding horizontal and vertical angle. These three parameters are the spherical coordinates describing the location of a point in 3D space. The horizontal and vertical angle of the laser beam is set by the position of the mirror hexagon and the orientation of the scanner. Besides setting the rotational speeds the controller also reports about these positions, which allows one to associate the measurements with their respective angles.

In order to generate distance measurements, the controller must synchronously manage the electronic circuitry. This comprises the generation of the square wave driving the laser diode and being fed into the AND gates as well as the regular reset of the integrators and the triggering of the ADC, all of it in concert. The correct timing of these individual actions is critical to the functionality of the laser scanner. Therefore it is essential to choose a controller capable of executing the corresponding instructions at the very moment they are intended to.

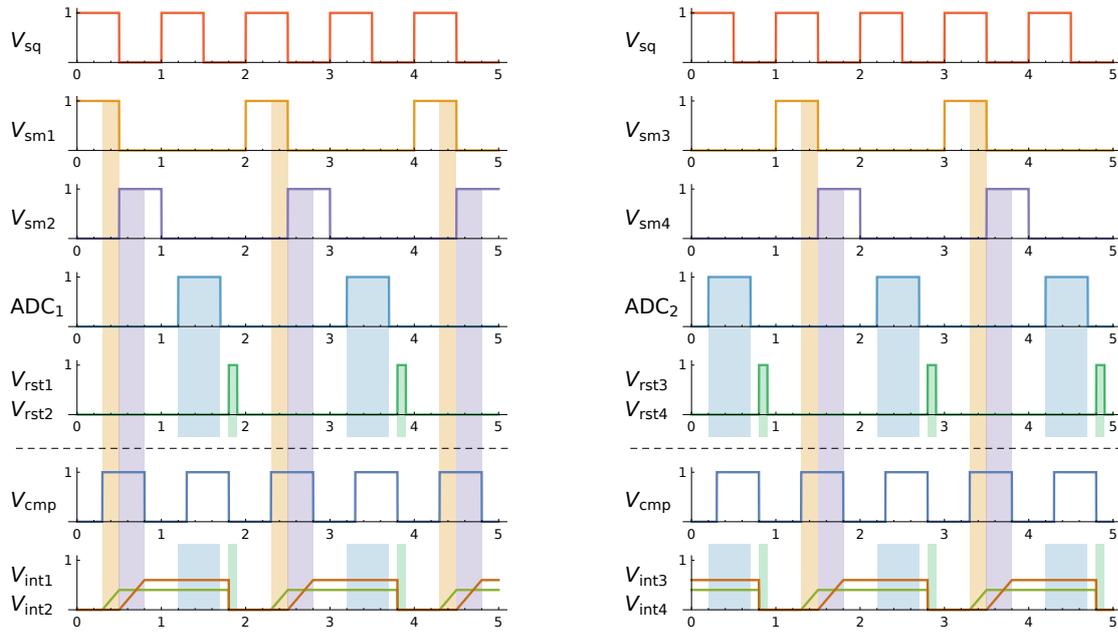
For the present laser scanner design an Atmel Corp. ATxmega32A4U⁴ MicroController Unit (MCU) serves as controller. This device has an integrated 2 MegaSamples Per Second (MSPS) differential 12-bit ADC and runs at clock rates up to 32 MHz. It also has sufficient General Purpose Input and Output (GPIO) pins to generate all the signals required by the circuit and features a 16 Mbit/s Universal Synchronous Receiver and Transmitter (USRT), which theoretically allows to transfer up to 1.3 million 12-bit measurements per second [XMEGA].

In section 2.2.4 it was assumed that sampling by the ADC and the reset of the integrators happened instantaneously at the beginning of each period of the emitted signal. However, the integrated ADC of the ATxmega32A4U requires at least $0.25\ \mu s$ to sample its inputs for a conversion [XMEGA, sect. 36.2.6]. Likewise the reset of the integrators needs some time to settle.

Therefore, the present design duplicates the AND gates and integrators, effectively increasing their number to 4, so that one pair can integrate the received signal while the other pair is sampled by the ADC and reset. For this purpose the replica of the emitted signal fed into the AND gates is replaced by dedicated sampling signals generated by the controller. Otherwise both integrator pairs would simultaneously sample the signal each period, without giving the ADC a chance to convert the measurements.

Figure 2.25 gives an overview of all the signals generated by the controller to drive the circuitry and the waveforms expected to appear at the outputs of the integrators. Subfigure 2.25a on the left shows the signals for the first pair of integrators and subfigure 2.25b on the right shows the signals for the second pair. The red curve on top is the square

⁴ Due to the recent purchase of Atmel by Microchip Technology Inc. this device might be rebranded soon.



(a) Signals driving the first integrator pair and resulting waveforms.

(b) Signals driving the second integrator pair and resulting waveforms.

Figure 2.25 Signals generated by the controller to control circuit behavior and the resulting waveforms for a selected input. The topmost red curve shows the square wave signal V_{sq} driving the laser diode, which is identical for both integrator pairs. The orange and purple curves are the sampling signals V_{sm} for the first and second integrator of each pair respectively. The light blue curve is the ADC sampling signal: the ADC samples its inputs while this signal is high. The turquoise curve below is the reset signal V_{rst} , driving the FET which simultaneously reset both integrators of a pair. The dashed line separates the signals generated by the controller from the two remaining plots, which exemplify waveforms produced by the circuit: The dark blue curve is a possible course of the comparator output V_{cmp} . The final plot at the bottom shows the corresponding waveforms at the output of the two integrators: Integrator output V_{int} rises as its input is high. The outputs are held as the ADC samples and reset immediately afterwards.

wave V_{sq} driving the laser diode, which is generated only once and identical for both subfigures. The orange and purple curves are the sampling signals V_{sm} respectively fed to the two AND gates of each pair. The second input, which is identical for all AND gates, is the comparator output V_{cmp} . A possible waveform of this signal is depicted by the dark blue curve just below the dashed line. The resulting high states at the outputs of the AND gates are indicated by light orange and purple rectangles in the background. The green and brown curves shown in one plot at the bottom are corresponding outputs of the two integrators. The light blue curve is a generic representation of the ADC sampling period. The turquoise curve is the reset signal V_{rst} for one integrator pair, which resets both integrators after the ADC has sampled their output.

The signals controlling both integrator pairs are identical but shifted relative to each

other, so that the pairs alternately sample the comparator output or get sampled by the ADC. By this way scanning is not interrupted, yet the ADC has sufficient time to sample the signal and the integrators can be reset.

As mentioned in section 2.2.4 the difference between the outputs of the two integrators in each pair is measured directly by taking advantage of the differential nature of the ADC. The ADC samples once per clock cycle of the emitted square wave and produces a measurement which is proportional to the delay of the captured signal relative to the emitted one. Hence for an output signal frequency of 1 MHz, 1 million measurements are generated per second.

These measurements, along with regular updates on the angular orientation of the laser beam, are transmitted to a processing platform via Universal Asynchronous Receiver/Transmitter (UART) interface. There final processing is carried out in order to extract the distance of surfaces which reflected the laser beam and combine individual results to form a 3D model of the surroundings.

2.4 Post Processing

One might be tempted to believe that post processing simply consists of turning the delay measurements obtained from the ADC into distance measurements one by one, applying a simple linear equation, and to associate each one with the corresponding angles to get the 3D coordinates of the respective point. However, the task is more delicate since for low signal amplitudes and hence low signal-to-noise ratio (SNR), the measurements suffer from a significant standard deviation, as revealed by equation 2.20. Table 2.1 shows the expected standard deviation according to this equation for a few signal to noise ratios.

SNR	Standard deviation		
	σ (68.3%)	2σ (95.5%)	3σ (99.7%)
30 dB	0.534 m	1.07 m	1.60 m
20 dB	1.69 m	3.38 m	5.06 m
10 dB	5.34 m	10.7 m	16.0 m
0 dB	16.9 m	33.8 m	50.6 m

Table 2.1 Standard deviation of distance measurements for various SNR

The table does not only list the standard deviation σ for each of the selected SNR values, but also 2σ and 3σ , the double and triple of σ respectively. For normally distributed samples, 68.3% of the samples are expected to lie within $\pm\sigma$ of the mean, 95.5% to fall within $\pm 2\sigma$ and 99.7% are expected to be within $\pm 3\sigma$ of the mean.

According to equation 2.7 the proposed shot noise limited amplifier circuit will yield a SNR of roughly 10 dB for the radiant energy expected to be captured from a surface 10 m

afar. Therefore, more than 30% of the measurements from such a surface are expected to differ by more than 5 m from the effective distance. Even worse, about 4.5% are expected to differ by more than 10 m, hence to be afflicted with an error larger than 100%.

Such a low precision is not satisfying and additional measures need to be taken to ameliorate it. The amplification circuit is already shot noise limited and hence the SNR cannot be improved. Therefore post processing will need to comprise techniques to improve the accuracy of the measurements.

In order to enhance the precision the standard deviation for one distance measurement must be reduced, which can be done by averaging multiple measurements. When taking the mean of n samples, the standard deviation improves by \sqrt{n} over the standard deviation of one sample. For instance, averaging 16 subsequent measurement reduces the standard deviation by a factor 4, hence when targeting a surface at a distance of 10 m 95.5% of the measurements will be accurate to ± 2.7 m.

Improving the standard deviation by averaging multiple samples is a valid method only if all samples are centered around the same mean, thus if the laser beam is targeting the same spot throughout all the measurements. It is therefore required that the orientation of the laser beam does not change substantially during this time, otherwise the measurement will be compromised. The angular deviation of the beam should be kept smaller than a threshold to be determined, which in turn restricts the available time for collecting samples or conversely requires the rotational speed to be reduced in order to allow the desired number of samples to be gathered.

Allowing for sufficient samples to accumulated within a certain angle is not the only factor which might require to slow down the mirror wheel. If multiple subsequent samples are averaged in order to get more precise values, then fewer measurements are generated over time. Therefore, the point cloud for one round scan will be less dense if the rotational speeds of the mirror hexagon and the scanner unit are left unchanged. The density of the point cloud can be sustained by reducing these rotational speeds, however this increases the time required for one scan and hence reduces the rate at which round scans are effectuated.

These considerations lead to a trade-off between the different performance metrics: The precision of the distance measurements, the angular resolution both in horizontal and vertical direction and the scan rate need to be balanced according to the needs and requirements of specific applications. The relationship between these metrics is given by

$$f_{scan} = \frac{f_{meas}}{n_{samp} \cdot n_{hori} \cdot n_{vert}} \quad (2.23)$$

where f_{meas} is the rate at which the circuit produces measurements, n_{samp} is the number of samples gathered for one distance measurement, which results in one point in the final point cloud, n_{hori} is the horizontal resolution of the point cloud, n_{vert} is the vertical

resolution and f_{scan} is the scan rate, i.e. the frequency by which the scanner generates 360° round scans.

For instance, if one chooses n_{samp} to be 16 in order to sufficiently enhance the precision of distance measurements and wishes to retain a resolution of 1° , both in horizontal and in vertical direction, thus setting $n_{hori} = 360$ and $n_{vert} = 90$, then the scanner is able to perform 1.93 scans per second for $f_{meas} = 1$ MHz. Reducing the horizontal and the vertical resolution to 2° , with $n_{hori} = 180$ and $n_{vert} = 45$, increases this number to nearly 8 scans per second.

However, if a higher precision is required, either the spatial resolution or the scan rate decrease. When 100 samples are required to achieve the desired precision, then the scanner can produce only 10 000 measurements per second, which reduces the scan rate for a spatial resolution of 1° in horizontal and vertical direction to only 0.31 scans per second. Thus, a single scan takes almost 3 seconds to complete. With 2° of spatial resolution the scan rate can be increased to 1.23 scans per second and with 5° the scan rate can even be raised to 7.7 scans per second.

Evaluation

An extensive evaluation is conducted to assess the performance of the proposed laser scanner design in terms of precision, range and spatial resolution. For this purpose, a prototype implementing the proposed design has been built.

Building the laser scanner comprises the layout of a Printed Circuit Board (PCB) with the entire circuitry and the controller aboard as well as the design and assembly of a mechanical structure to mount the optical components and the motors to rotate the mirror hexagon and the entire unit.

The design of the circuit and the layout of the PCB was carried out carefully, in order not to compromise the performance of the electronic design with a poor layout [Ard05][AN1258]. Separate supply sources have been used for the circuit driving the laser diode and the circuit amplifying the photodiode signal. Particular attention was given to proper decoupling of the critical components [AN202] as well as the layout of the signal return paths, which can severely degrade a circuit's precision [AN47].

The design of the circuit and the PCB layout has been realized with the open source KiCad Electronics Design Automation (EDA) suite. The PCB has been manufactured by ITEAD Intelligent Systems Co. Ltd. and all electronic components were purchased via Premier Farnell Ltd. and Digi-Key Electronics Corp.

Optical components, i.e., the mirrors and lenses as described in section 2.1, have been purchased at Thorlabs, Inc. All parts required for the mounting of optical components as well as the motors have been 3D printed on a Stratasys Ltd. Dimension Elite 3D printer. The motors required to rotate the mirror hexagon and the scanner were also bought via Digi-Key Electronics Corp. The overall cost of all parts and components was approximately US\$ 300.

3.1 Validation of Calculations

A number of assumptions have been made regarding some of the metrics of the laser scanner. Specific tests have been conducted to assess the validity of these assumptions and replace them with measured quantities where they failed.

3.1.1 Optical Setup Performance

In section 2.1 it was assumed that losses in light intensity from atmospheric absorption and imperfections of the optical components can be neglected. Furthermore the light incident to a surface was supposed to be evenly scattered within the hemisphere atop. In order to confirm these assumptions, a first evaluation shall quantify the amount of light effectively captured by the photodetector.

The fully assembled laser scanner unit, as shown in figure 2.4, is placed in front of a bright surface, once at a distance of 60 cm and once at 100 cm. The mirror wheel is oriented such that the laser beam is reflected horizontally and the entire unit is rotated such that the laser beam perpendicularly hits the surface. The exact reflectance of the surface is unknown, however it is assumed to be larger than 0.5. The amount of light captured by the photodiode is measured by attaching an oscilloscope probe to the output of the transimpedance amplifier, thus indirectly measuring the photocurrent produced by the photodiode.

For a surface with a reflectance of 0.5 the expected amount of light captured by the photodiode is approximately 150 nW when the scanner is separated from the surface by 60 cm and 50 nW for a distance of 100 cm, according to equation 2.1 and the aperture area and emitted energy levels calculated in section 2.1. These values translate to photocurrents of roughly 60 nA and 20 nA of photocurrent. The measured photocurrent is 64 nA and 17 nA respectively, thus very close to the expected values. Table 3.1 sums up all these calculations as well as the measured values.

Distance of the surface	60 cm	100 cm
Expected radiant energy captured	150 nW	50 nW
Expected photocurrent	60 nA	20 nA
Measured photocurrent	64 nA	17 nA

Table 3.1 Expected and measured photocurrent from light reflected by a bright surface

Whether the losses in the optical path are effectively very low or the reflectance of the surface is higher than 0.5 cannot be determined unless a surface with precisely known reflectance is used for testing. Nonetheless the experiment shows that the losses in the optical path are likely to be low and do not reduce the amount of light captured by the photodetector by a significant amount.

3.1.2 Transimpedance Amplifier Evaluation

The transimpedance amplifier is the part of circuitry which has received the most attention during the design of the electronic subsystem. It has the very delicate task of amplifying the weak photocurrent produced by the photodiode while ensuring that noise obstructing the received signal is kept to a minimum.

Unfortunately, due to various constraints the present electronic design sets on the transistors and the physical limitations of real devices, most notably the very high leakage current of most low capacitance transistors, the transimpedance amplifier could not be implemented as planned. The cascode stage had to be designed with only one cascode transistor, omitting the regulating transistor shown in figure 2.15. Consequently the transimpedance gain had to be reduced to $2\text{ M}\Omega$ instead of the planned $20\text{ M}\Omega$ in order to fulfill the bandwidth requirements. The noise penalties resulting from these changes are taken into account in the subsequent calculations.

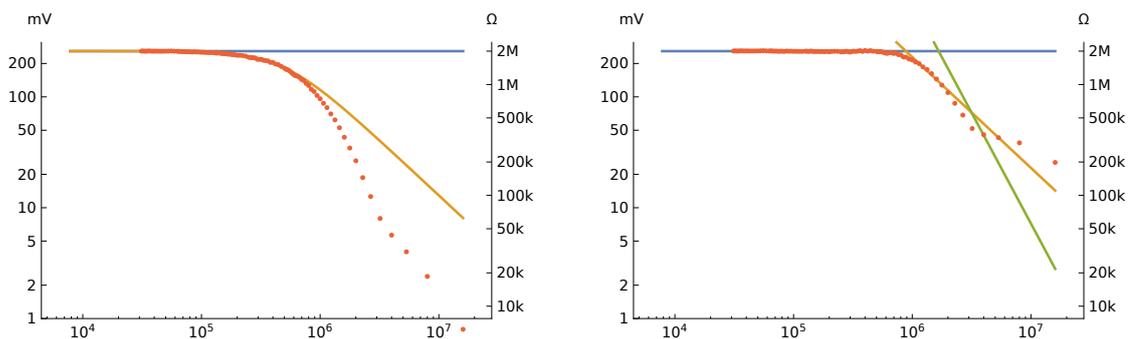
During the design of the transimpedance amplifier a number of assumptions have been made regarding some component properties which might have considerable effects on the bandwidth, stability or signal-to-noise ratio. In particular the performance is very sensitive to the parasitic capacitances at the input of the cascode transistor as well as at the inverting input of the operational amplifier.

It might seem reasonable to simply measure the amplitude of noise produced by the circuit and to verify whether it corresponds to the calculations. However, if higher noise levels than anticipated are measured, then this does not help in understanding their origin. Before measuring the noise, another test allows to estimate the effective parasitic capacitances of the realized circuit. This will further allow to adjust the expected bandwidth and noise behavior and help understanding the results of subsequent tests.

The various capacitances in the circuit can be estimated by measuring the frequency response of the cascoded transimpedance amplifier. Recalling figure 2.16, the input capacitance and feedback capacitance of the transimpedance amplifier can be interpreted graphically from the gain spectrum. Similarly, the bandwidth of the cascode stage is marked by a single roll-off frequency in its gain spectrum. The cascode and transimpedance amplifier stages are subsequent, hence their gains are multiplied. Localizing the cascode's roll-off frequency allows to restore the individual frequency responses of both.

To generate the circuit's frequency response, the laser diode was driven by a square wave with varying frequency and the RMS amplitude at the output of the transimpedance amplifier is measured with an oscilloscope. The test is run automated by a script for some hundred frequencies and the results are plotted in figure 3.1.

Figure 3.1a on the left shows the RMS voltages measured at the output of the transimpedance amplifier. At low frequencies the transimpedance gain corresponds to the value of the feedback resistor, hence the measured voltages can be matched to the corresponding gain values. A second scale is added to the right of the plot which indicates



(a) RMS voltages measured at the output of the transimpedance amplifier for various input signal frequencies. The blue line indicates the value of the feedback resistor, the yellow curve is the gain of the cascode stage.

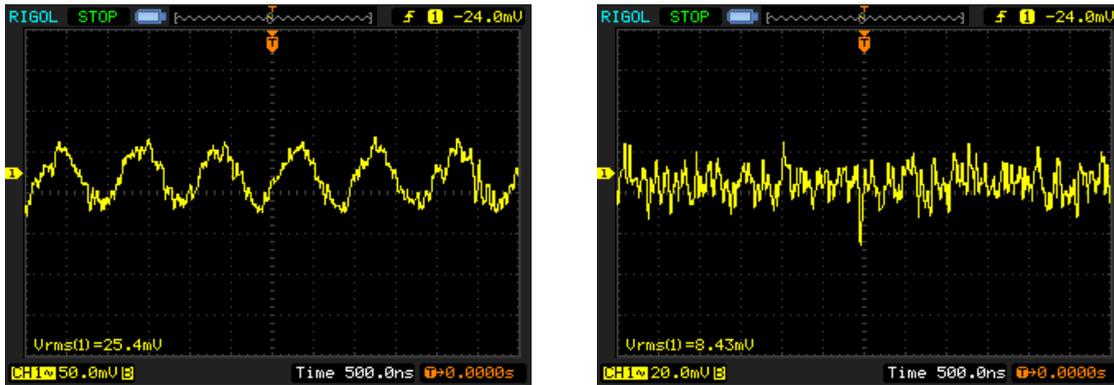
(b) Gain of the transimpedance stage only: the measurements have been corrected to remove the cascode gain. The yellow and green lines show the roll-offs due to feedback and input capacitance respectively.

Figure 3.1 Measured frequency response of the transimpedance amplifier: The gain for various frequencies is derived from RMS voltage measurements at the output of the transimpedance amplifier. Curves representing simulated roll-off and gain behavior are underlaid and adjusted to line up with measurements, thus allowing to infer the effective values of the circuit.

the transimpedance gain in Ω . The blue line shows the value of the feedback resistor. The first roll-off observed in this plot is not due to the transimpedance stage, but results from the cascode stage. The yellow curve shows the gain of the cascode stage. The input capacitance of the cascode has been adjusted so that this curve lines up with the measurements, thus allowing us to identify the effective capacitance. The second roll-off seen in this figure is from the transimpedance stage. Starting at this frequency the measured values no longer line up with the cascode gain.

Figure 3.1b is the gain of the transimpedance stage alone. These values are obtained by dividing the measurements by the relative gain of the cascode stage. Here the yellow and green curves indicate the roll-offs due to feedback capacitance and to input capacitance respectively, as previously seen in figure 2.11. The values of these capacitances have again been adjusted in order to line up with the measurements to identify the real values. Note that the last few of the measured values are very close to the noise floor of the oscilloscope and hence imprecise, however the roll-off due to input capacitance can be briefly observed at approximately 2.5 MHz. The feedback capacitance has been adjusted to 0.09 pF, which is very close to the rated value of 0.1 pF of the used capacitor. However the input capacitance had to be adjusted to 10 pF to line up with the measurements, which is much more than the anticipated 1 pF.

From these measurements it follows that the effective input capacitance of transimpedance amplifier is higher than predicted and therefore the noise calculations need to be updated. Since the input capacitance is effectively 10 pF, the noise contributed by the multiplication of the operational amplifier's voltage noise with its input capacitance becomes the dominant noise source and raises the noise level to approximately $300 \text{ fA/Hz}^{1/2}$ of



(a) Waveform captured when targeting a surface situated 1 m from the scanner. The sinusoidal waveform, although cluttered with noise, is still easily discernible.

(b) Waveform captured when targeting a surface which is more than 7 m distant. For such a large distance the signal is so weak that the capture essentially consists only of noise.

Figure 3.2 Captured output of the transimpedance amplifier. The waveforms show that the received signal is obstructed by noise with an RMS amplitude of approximately 8 mV. The signal has not been filtered yet and therefore still contains high frequency noise, which is attenuated by the subsequent band-pass filter stages. Note the different vertical scales for the two captures.

equivalent input noise current density at 1 MHz and even higher levels at higher frequencies. The equivalent input noise RMS amplitude will therefore be larger than 0.3 nA, resulting in over 0.6 mV of voltage noise amplitude at the output of the transimpedance amplifier.

Figure 3.2 shows waveforms of the transimpedance amplifier output captured with an oscilloscope, once for a bright surface located 1 m from the scanner and once for a surface which is more than 7 m afar. These waveforms show an even higher noise amplitude than anticipated. For the latter the photocurrent generated is less than 0.5 nA, which is converted to at most 1 mV of signal amplitude. Therefore, the recorded RMS amplitude of 8.4 mV is almost entirely due to noise.

A voltage noise amplitude of 8 mV at the transimpedance amplifier output corresponds to 4 nA of current noise density at its input. This is higher than the estimated 0.3 nA, however this value was calculated for a bandwidth of 1 MHz. Since the output has not been filtered yet the noise is likely to originate from a larger frequency spectrum. In fact it is apparent from the wobbly waveforms in figure 3.2 that the signal is obfuscated by high frequency oscillations of several MHz. Therefore the high RMS amplitude of noise is not surprising and subsequent filtering helps reduce the noise level and improve the SNR.

Nonetheless, the larger noise levels impact the SNR and drastically reduce the performance of the laser scanner. Assuming that filtering the signal limits the noise amplitude to 1 nA of current noise amplitude, the SNR drops to 26 dB for 20 nA of photocurrent, as has been found to be generated when a bright surface is 1 m distant from the scanner in section 3.1.1. According to equation 2.20, this results in a standard deviation of the measurements of roughly 0.84 m. For a surface located at a distance of 2 m this value

increases to 3.4 m and for a surface 3 m afar the standard deviation of the measurements is larger than 7 m. Table 3.2 sums up the anticipated standard deviations for surfaces located at various distances from the scanner.

Distance to surface	Standard deviation		
	σ (68.3%)	2σ (95.5%)	3σ (99.7%)
0.5 m	0.21 m	0.42 m	0.63 m
1 m	0.84 m	1.7 m	2.5 m
2 m	3.4 m	6.8 m	10 m
5 m	21 m	42 m	63 m

Table 3.2 Anticipated standard deviation of distance measurements for surfaces at various distances. The first, second, and third standard deviation are given for each distance and the percentage of samples which are expected to lie within that far from the actual distance is given in parenthesis. Note that the standard deviation increases as the distance to the surface increases and that it quickly becomes larger than that distance.

The larger noise levels heavily deteriorate the performance of the scanner. The standard deviation becomes already larger than the measured distance for surfaces located slightly farther than 1 m. Therefore a single measured value has absolutely no significance and does not even allow for a rough estimation of the effective distance. A large number of samples must be averaged in order to estimate the effective distance of a surface, which decreases the temporal resolution of the scanner.

Since the SNR and hence the noise level have a linear influence on the standard deviation, the 10 times larger noise amplitude compared to the initial calculations similarly increases the standard deviation by a factor of 10. Recalling section 2.4, the standard deviation of the mean of multiple measurements decreases only by the square root of the number of samples. Therefore 100 samples need to be averaged in order to achieve the same precision as initially anticipated.

It should be noted that the range of the scanner is theoretically infinite, however since the amplitude of the captured signal decreases by the square of the distance, as seen in equation 2.1, the standard deviation likewise increases by the square of the distance. Consequently, the number of samples required to retain a certain standard deviation grows by the power of 4 of the distance. Therefore, although the range has no strictly defined limit, the rapidly decreasing precision renders measurements over large distances unfeasible.

3.1.3 Precision of Delay Measuring Circuitry

As seen in section 2.2.3 the phase information of the amplified photodiode signal appearing at the output of the comparator is afflicted with time jitter, as a result of noise from the transimpedance amplifier. This increases the standard deviation of the individual delay

measurements and thus reduces the precision of the laser scanner. However, time jitter on the comparator output only limits the precision if the delay measuring circuit itself is sufficiently precise.

Therefore it is appropriate to first quantify the precision of the delay measuring circuit. This is done by feeding it with a square wave signal with known delay and with practically no time jitter. Deviations of the measurements from the effective delay are inherent to the delay measuring circuit and set a limit on the achievable precision of the laser scanner.

The execution of this test showed that the standard deviation introduced by the delay measuring circuit is close to the resolution of the ADC, which is 1.8 cm as determined from equation 2.22, hence it is much lower than the expected standard deviation caused by noise as presented in table 3.2. In fact the test was repeated multiple times and the standard deviation did never exceed twice the ADC resolution.

Given the high precision of the delay measuring circuitry its impact on the accuracy of the measurements can be ignored. Indeed, as long as the standard deviation introduced by the delay measurement circuitry does not get close to that caused by the noise in the amplification circuit it has very little influence on the overall precision and can be neglected.

For a final validation the completed scanner was aimed at bright surfaces at various distances and the standard deviation of the measurements was compared to the values listed in table 3.2. Since the reflectance of the surface impacts the amplitude of the signal and therefore the standard deviation of the measurements, an exact verification of the calculations is not possible. Nonetheless, the calculated standard deviations have been found to be quite accurate, thus confirming the validity of the corresponding calculations.

An important consequence of these high standard deviations is that a larger number of samples must be averaged in order to achieve some level of precision. As seen in section 2.4, the standard deviation of the mean decreases only by the square root of the number of values. For instance, in order to reduce the standard deviation to 0.1 m while measuring the distance to a surface 5 m afar, an improvement by a factor of roughly 20 is required over the standard deviation of single value, thus 400 samples must be averaged in order to reach this precision.

If a larger number of samples is required to produce a single distance measure, then the time to produce such a measure increases. Therefore the temporal resolution, i.e. the rate at which the scanner generates distance measures, decreases. Consequently either the spatial resolution or the scan rate decline, as seen in equation 2.23. When 400 values are combined to one distance measure, the temporal resolution drops to 2500 measures per second. If a spatial resolution of 1° shall be retained, where a full scan consists of $360 \cdot 90 = 32400$ individual measures, then generating a single scan takes 13 seconds. Reducing the spatial resolution to 5° increases the scan rate to roughly 2 scans per second.

3.2 Unexpected Effects

In the course of the evaluation of the built laser scanner unexpected effects were encountered, which had not been foreseen by the calculations carried out during the design of the laser scanner. The first effect to be found is an amplitude dependence of the produced measurements, causing the measurements to differ for equally far surfaces with different reflectance levels. The second effect is a variation in the measurements caused by changes in ambient light.

These effects have potentially devastating consequences on the performance of the scanner. Therefore, this section presents an attempt to analyze the origins of these effects and quantify their impact on the accuracy of the distance measurements.

3.2.1 Amplitude Dependence

An effect that was not taken into account in the calculations and which revealed during testing is a slight amplitude dependence of the time-of-flight measurements. This could not have been predicted since it was assumed that all the devices are linear. If all components in the circuit amplifying the photodiode current up to the comparator were perfectly linear, the phase would be left unaffected by changing signal amplitudes. However, since real components are burdened by some degree of non-linear behavior, the signal's phase is marginally altered by its amplitude [HH15, ch. 3].

The alterations, despite being quite small, influence the measurements. It is therefore required to analyze and quantify the phase variations resulting from differing signal amplitudes, in order to judge their consequences on the produced measurements.

The test setup used to assess the amplitude dependence compares the measurements generated by the circuit for surfaces which are equally far from the scanner but have different reflectances. Per equation 2.1 the amplitude of light received by the photodiode is proportional to the reflectance of the targeted surface. If there was no amplitude dependence of the signal's phase, the measurements would be equal despite the different reflectances.

The test setup is shown in figure 3.3. Sheets of paper with several shades of gray serve as surfaces with various reflectance levels. The individual regions on a single sheet are all equally far from the scanner, nonetheless the measurements differ significantly.

Figure 3.4 shows the results of this test. Numerous measurements are carried out for each region and the mean is plotted against the standard deviation of the samples. Measurements which result from the same sheet of paper and hence should all report similar values are shown in one color, with the legend indicating the effective distance between the sheet and the scanner. Note that the values are directly taken from the ADC, which produces 12-bit signed integers. Therefore, the resulting values vary between -2048 and $+2047$, with larger values indicating a greater distance. As seen in section 2.2.4, one unit is equivalent to approximately 1.8 cm.

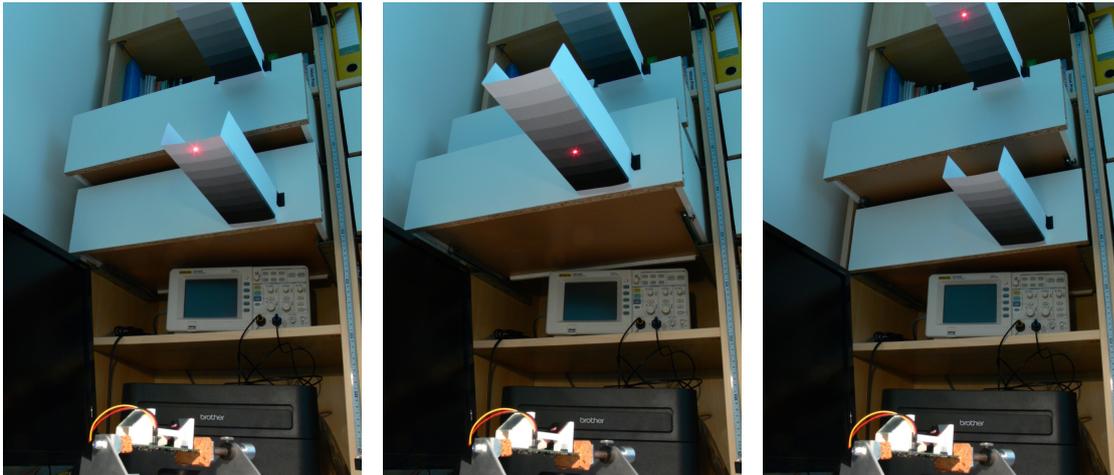


Figure 3.3 Evaluating the amplitude dependence: Various shades of gray printed on paper have differing reflectance levels. Ideally this has no influence on the measurements. However the measured values differ as a result of the circuit's amplitude dependence, which is analyzed and quantified using the depicted setup.

As seen in section 2.2.3, a larger standard deviation is caused by a lower signal to noise ratio. As predicted by calculations in section 2.2.2 and confirmed by measurements in section 3.1.2, the noise level is nearly constant and hence any variations of the standard deviation are due to changes in signal amplitude. Therefore, a higher standard deviation of the measurements indicates a lower signal amplitude.

Darker surfaces reflect less light, thus the signal amplitude is lower for such surfaces and

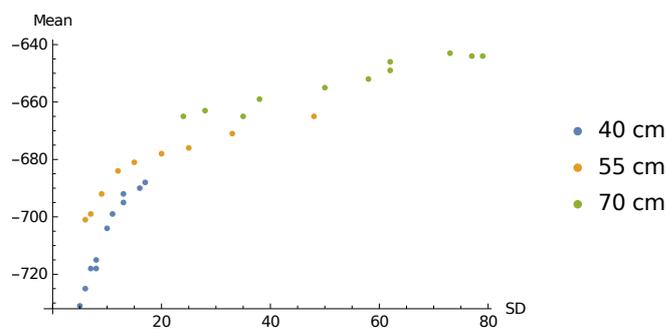


Figure 3.4 Amplitude dependence of the measurements. The standard deviation, shown on the horizontal axis, is an indirect measure of the signal's amplitude. A low standard deviation is evidence of a stronger signal, which is less vulnerable to noise. Conversely, a larger standard deviation indicates a weaker signal, which is more strongly affected by noise. Measures drawn in the same color emanate from surfaces equally far from the scanner, but with regions with varying reflectances. The experiment shows that surfaces which reflect less light are reported as being located further away than equally distant brighter surfaces.

consequently the standard deviation increases. The results depicted in figure 3.3 further reveal that the measured values increase for lower signal amplitudes, such that darker objects are reported as being further away than equally distant lighter objects.

This effect cannot be weakened by averaging multiple measurements, since each measurement is equally affected by it. In order to alleviate the effect an attempt was made at finding a function which describes the dependency between effective distance, mean and standard deviation of the measurements.

The plot in figure 3.3 shows that the measurements produced by equally distant surfaces are all lying on similar lines, which suggests that it might be possible to find a function which predicts these lines. If such a function is found it could be used to calculate a better approximation of the effective distance of a surface based on the mean and standard deviation. In fact, comparing the mean and standard deviation of measurements to the plotted values allows to graphically estimate the effective distance of a surface. However, the attempts at finding a functional dependence were unsuccessful.

Another option might be to gather a huge number of mean and standard deviation values for surfaces of different reflectances and located at various distances, to store these values in table, and to estimate the effective distance by comparing the mean and standard deviation of measurements to the stored values. An attempt at implementing such a table was not carried out as part of this work.

The amplitude dependence leads to a distortion of the scans produced by the prototype. Less light is captured from surfaces located further away, as seen in section 2.1, thus the amplitude of the signal caught from more distant surfaces is smaller which causes the measurements to report an even greater distance.

Additionally, the angle of incidence of the laser beam onto a surface influences the signal amplitude. During the calculations in section 2.1 it was assumed that a surface radiates incident energy equally in all directions. Yet, it was also mentioned that real surfaces rather follow Lambert's cosine law and radiate more light in normal than in off-normal direction. Consequently the amount of reflected light captured by the photodetector does not only depend on the distance and the reflectance of a surface, but also on the angle of the surface with respect to the viewing axis.

Figure 3.5 shows the results of a test conducted to demonstrate the distortion caused by the amplitude dependence. This is the horizontal planar scan of a straight white wall from a distance of 50 cm. The scan has been acquired by collecting more than 100 000 measurements for each direction and manually rotating the scanner by 4° between subsequent acquisitions. The large dark spot at the bottom of the figures is the location of the scanner and the horizontal gray line is the effective surface of the wall. The blue points mark the location of the wall surface according to the measurements and the yellow lines show the range of the standard deviation. Both figures are identical, except that in figure 3.5b the standard deviation has been divided by 10 for better readability, which is equivalent to the standard deviation of the mean of 100 individual measures.

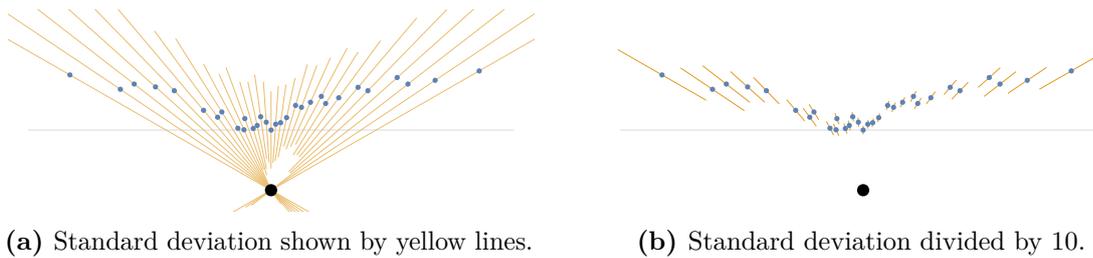


Figure 3.5 Distortion caused by the amplitude dependence: Horizontal scan of a straight white wall from a distance of 50 cm. The location of the scanner is shown by the black point at the bottom and the effective location of the wall is shown by a gray line. The blue points are the averages over all captured samples and the yellow lines show the range of the standard deviation.

The wall seems to be bent backwards on the left and right. This is in accordance with our expectations: The point of the wall which is closest to the scanner is the point straight in front it. As the laser beam is oriented sideways, it travels a larger distance before it hits the wall. Additionally, the angle of incidence of the laser beam with respect to the wall decreases. Both of these factors result in less light being reflected by these parts of the wall, thus the signal amplitude is lower and the measured distance is greater than the effective distance.

Another interesting consequence of the amplitude dependence, which can be observed in figure 3.5, is that slight variations in the reflectance of the wall result in small offsets of the measured distances which appear to be noise. However, since a large number of samples have been averaged to obtain these distance measures, they can be assumed to be fairly accurate and reproducible. Therefore, if the line across the locations of the measurements is not smooth, this must be caused by varying reflectance levels of the wall.

3.2.2 Ambient Light Effects

The optical path of the built prototype scanner is exposed, therefore ambient light easily reaches the photodetector. Nonetheless, this should have no effect on the measurements, since the pass-band filter stages subsequent to the transimpedance amplifier comprise high-pass filters, which remove any constant offset or low frequency disturbances from the signal.

Surprisingly, changes in ambient lightning cause variations of the measured values. Most remarkably, the standard deviation of the measurements apparently is not affected by it, but the mean varies. The prevalent lightning conditions cause a constant offset of the measured values, which changes by as much as 70 to 80 cm as the lightning conditions change.

The exact cause for these alterations has not been identified. However, in a linear system a constant offset has no influence on the phase of a signal. Therefore, the effect is most



Figure 3.6 Photograph of the living room which was scanned with the built prototype scanner.

likely caused by non-linearities in the amplification circuit, just as for the previously analyzed amplitude dependence.

No further analysis and quantification of the variations due to changes in ambient lightning have been carried out. However, care was taken to keep the lightning conditions constant during all other tests.

3.3 Test Scan

The previous sections in this chapter have thoroughly analyzed the design of the laser scanner and tested the validity of the calculations conducted during the design phase. The standard deviation of the produced measurements was identified as the key performance parameter in section 2.4 and the calculated values listed in table 3.2 have been shown to be accurate. Unexpected effects, which further degrade the performance have been analyzed and characterized. Finally, a last test attempts to visualize the performance of the device and render these numbers more tangible.

The prototype scanner is used to produce a 3D scan of an indoor space. The chosen space is a living room with white walls, wooden floor and basic furniture. A photograph of the room is shown in figure 3.6.

The scanner was placed in the center of the room, approximately 80 cm above the floor. This height was chosen because it allows for a good coverage of the floor as well as of the ceiling, since the scanner can scan 60° upwards but only 30° downwards as seen in section 2.1.

The prototype is equipped with a motor rotating the mirror hexagon and thus can perform vertical scans, however it lacks the second motor which would rotate the entire

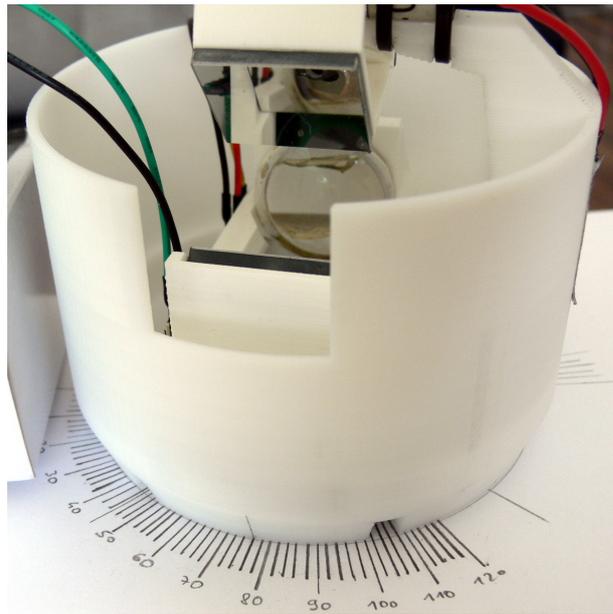


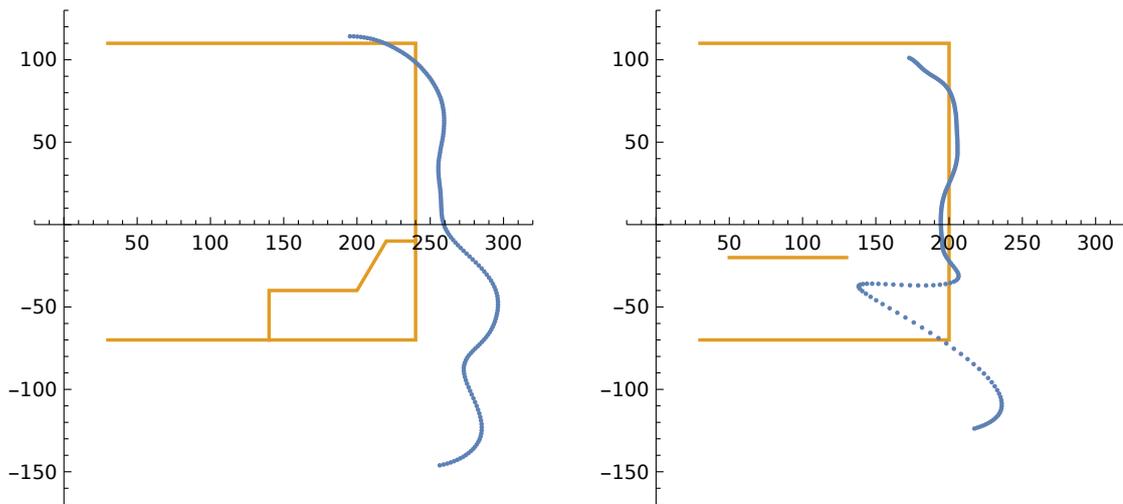
Figure 3.7 Photograph of the setup used to orient the laser scanner for 3D scanning. The mirror hexagon, visible on the top, is rotated by a motor, thus automating vertical motion. However, horizontal motion, which involves rotating the entire scanner, is not automated on this prototype. Therefore, the scanner is manually rotated to provide horizontal coverage. The lines on the support platform, spaced by 2° , are used to precisely orient the scanner.

device, adding horizontal motion to the laser beam. Therefore the scan was carried out by producing individual vertical slice scans and manually rotating the scanner by a small angle between scans. The individual slices were then combined to a 3 dimensional model.

Figure 3.7 shows the setup used to move the scanner. A marker on the bottom of the scanner indicates the direction of the laser beam, which is aligned with the desired line on the support platform. The lines are spaced by 2° , which is the horizontal resolution chosen for this test scan.

As the mirror wheel rotates the laser beam sweeps vertically over the surroundings and the recorded measurements show a repetitive pattern as the beam sweeps along the same line over and over again. A trigger condition is used to synchronize the beginning of each sweep. The border between two consecutive mirrors on the wheel produces a characteristic pattern in the sequence of acquired measurements, which has proven to be a robust trigger point to identify the beginning of the next sweep.

The trigger condition occurs for a known angle of the mirror wheel and thus is a reference point which is used to identify the angular orientation of the laser beam for the individual measurements. The rotational speed of the mirror hexagon is assumed to be constant and the angular direction corresponding to each measurement is linearly interpolated from the start of one sweep up to the next.



(a) Cross-section of the couch seen in the center of figure 3.6.

(b) Cross-section of the table seen on the left in figure 3.6.

Figure 3.8 Cross-sections of a living room space, obtained from vertical scanning, in blue and an approximation of the actual outline of each slice in yellow. The amplitude dependence causes darker surfaces to appear further away from the scanner and also distorts the dimension of the room and the shape of the walls, the ceiling and the floor.

The rotational speed of the mirror wheel is trimmed to the relatively slow value of 3 revolutions per second. Therefore, the laser beam sweeps over the 90° aperture of the scanner within 42 ms, since the angle of the laser beam varies twice as fast as the angle of the mirror wheel as seen in section 2.1. Within this period the scanner produces 42 000 measurements, far more than the 90 distance values required for a vertical resolution of 1° . However, the large number of values is necessary to achieve a decent result, since the standard deviation of the measurements is quite large for surfaces located 2 to 4 m afar, as it is the case for the chosen space. 200 subsequent measurements are combined to one distance value, yielding 210 values for one vertical slice.

Although 200 individual measurements are averaged to obtain one distance value in a vertical scan, these values are still not precise enough to produce a decent scan because of the large standard deviation at the involved distances, as listed in table 3.2. Therefore, multiple of such vertical scans are carried out for each orientation and combined to further improve the accuracy.

Figure 3.8 shows 2-dimensional representations of individual vertical slices captured by this setup. The slice in figure 3.8a shows the result of a vertical scan along the couch and the wall behind it. The slice in figure 3.8b is a cross-section of the table seen on the left in figure 3.6. The origin of these plots is the location of the scanner and the units correspond to the ADC resolution, which is approximately 1.8 cm as seen in section 2.2.4. An approximation of the effective outline of the room is underlaid for each scan.

The effects of the amplitude dependence analyzed in section 3.2.1 are clearly visible in these scans. In figure 3.8a the couch appears as if there was a hole in the wall. This is due to the lower reflectance of the gray couch, compared to the higher reflectance of the white wall, which makes the couch appear to be further away than the wall, despite the fact that it is placed in front of it. Similarly, the dark floor appears to be much further away than it effectively is. The scanner seems to be closer to the ceiling than to the floor, although it is actually much closer to the floor.

The amplitude dependence also causes a distortion of the dimensions of the room, as expected from the results seen in section 3.2.1. This causes the walls to appear as if they were bent.

Figure 3.8b also shows the importance of the angle of incidence of the laser beam. Although the table is close and therefore should reflect more light, it is detected as being much further away than it effectively is. The reason for this is that the angle of incidence between the laser beam and the top surface of the table is very low, thus most of the light is reflected upwards to the ceiling instead of back towards the scanner. Consequently, the amplitude of the returning signal is low and thus the table appears to be further away, although it is actually quite close to the scanner.

Despite all these distortions caused by the amplitude dependence analyzed in section 3.2.1, the shape of the room and the furniture can be roughly estimated in these slices.

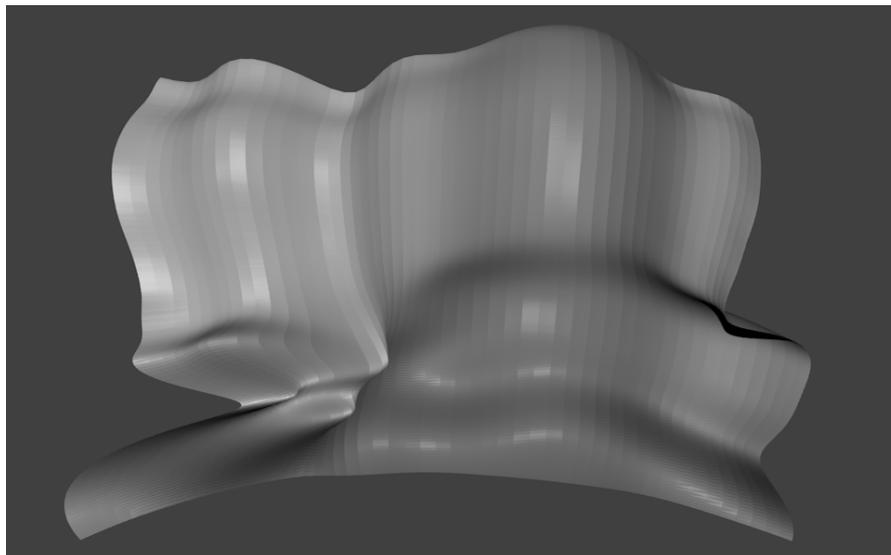
The last step required to build a 3-dimensional model of the scanned space is to combine the individual slice scans. The points of each slice are rotated around the vertical axis to match the direction the scanner was aiming at while the slice scan was generated. Figure 3.9 shows a rendering of the obtained model.

The same distortions already observed in the individual slices also apply to the entire 3D model. The gray color of the couch causes it to be set back into the wall, such that it is difficult to discern. Similarly, the floor is further from the scanner than it actually is. The walls are not straight, but seem to be bended.

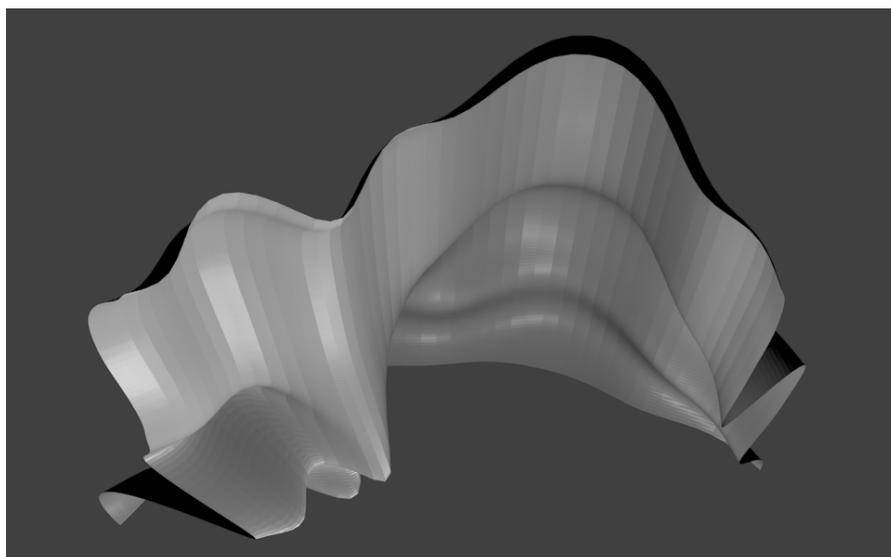
However, some details seen in the photograph in figure 3.6 are easier to distinguish. For instance, the table clearly stands out on the left and even the chair standing next to it can be identified. Further, it is possible to identify the wall corner of the room, which is seen in the middle of the photograph. Also, the front edge of the couch is visible in the 3D model.

Although the amplitude dependence severely corrupts the distance measurements, the produced 3-dimensional scan is not completely unintelligible. The shape of the room as well as some furniture can be identified.

This test scan demonstrates that despite the strong distortions, the 3-dimensional scans produced by the prototype scanner allow a rudimentary description of the shape of an indoor space and a rough localization and characterization of furniture. Although the evaluation revealed a number of weaknesses which need to be addressed before the



(a) View from the front.



(b) View from the top.

Figure 3.9 Renderings of a 3-dimensional model built from scanned data. Despite the strong distortion due to the amplitude dependence, some details from the photograph in figure 3.6 are distinguishable: The top surface of the table is clearly visible on the left of both images. A second flat surface underneath and slightly right to it has been produced by the seat. The front edge of the couch can likewise be identified. Also, the shape of the room can be roughly estimated from the virtual model. The wall corner can be seen particularly well in the lower render.

scanner can generate precise scans of its surroundings, the present implementation shows that the proposed design is conceptually feasible.

3.4 Discussion

The built prototype is not ready to be used on autonomous UAV yet, since the precision as well as the achievable scan rate or the spatial resolution are too low. However, a number of potential improvements might render this possible.

The major problem to be solved is the amplitude dependence of the measurements. As seen in section 3.2.1, the distance measurements produced by the scanner depend on the signal amplitude and therefore on the amount of light captured by the photodetector. The amount of light reflected by a surface depends on its reflectance, a parameter which should have no influence on the measured distance of a surface. This undesired dependence causes less reflective surfaces to be reported as being further away than equally distant more reflective surfaces.

As seen in section 3.2.1, the standard deviation of the measurements is an indicator of the signal's amplitude. Therefore, the standard deviation of some number of samples potentially allows to correct the error introduced by the amplitude dependence. Attempts to find a functional dependence between these parameters have failed, but if sufficient data is gathered for surfaces at various distances and reflectances, it might be possible to recover the effective distance of a surface from the mean and the standard deviation of multiple measurements.

Once the amplitude dependence has been eliminated or at least reduced, the next issue to be addressed is the high standard deviation of the measurements. As listed in table 3.2, the standard deviation quickly exceeds the measured distances, rendering single measurements totally insignificant. In order to estimate the actual distance of a surface with a decent level of certainty, multiple measurements need to be conducted and the mean of the individual values must be calculated. The standard deviation of the mean of multiple values is only reduced by the square root of the number of values, as seen in section 2.4. Therefore, a large number of samples is required to achieve sufficiently precise distance measures.

If a lot of measurements must be combined to one single reliable distance measure, then the temporal resolution of the scanner is reduced, which in turn decreases either the achievable scan rate or the spatial resolution, as shown by equation 2.23.

To improve the performance of the scanner, it is required to reduce the standard deviation of the individual measurements. As seen in section 2.2.4, the standard deviation depends upon the signal-to-noise ratio of the captured signal. Therefore, the noise level of the amplification circuitry must be reduced to improve the accuracy of the measurements. As seen in section 2.2.2, the transimpedance amplifier is the circuit element which sets the noise level. Although a lot of attention has been dedicated to this segment during

the design of the electronic system, additional improvements are required to ameliorate the signal-to-noise ratio and thereby the precision of the scanner.

Limitations of the transistors used in the cascode stage of the transimpedance amplifier forced a design change which raised the noise amplitude and therefore reduced the performance. If suitable devices, which overcome these shortcomings, can be found, then it might be possible to implement the electronic design with two cascode transistors, as originally planned, which should reduce the noise amplitude and consequently the standard deviation by a factor of 10. Consequently, 100 times fewer samples would need to be averaged in order to achieve the same precision as with the current device.

While the initial goal of building a scanner which can be used for obstacle detection on UAV was not reached, the built prototype shows that the design is conceptually feasible. Some remaining issues need to be addressed. Nonetheless, reiterations of the proposed design might be sufficiently precise and fast to meet the requirements of a surround scanning system for UAV.

Conclusion

The goal of this work is to investigate the feasibility of a low-cost LIDAR system suitable for use on UAV. A prototype has been built and tested to evaluate the performance of the proposed design by quantifying the range, precision, spatial resolution, and scan rate.

Issues caused by the imperfections of real electrical components prevent the built scanner from reaching the desired precision and distort the produced scans.

The standard deviation of the measurements increases by the square of the measured distance, which limits the range of the laser scanner. Calculating the mean of multiple measurements improves the accuracy, however it reduces the temporal resolution, necessitating either a reduction in spatial resolution or in scan rate. These performance parameters are all dependent on each other, creating a compromise.

Despite some efforts, no method could be found to alleviate the effects of an amplitude dependence of the produced measurements, which causes darker surfaces to be reported as if they were further away than equally distant brighter surfaces.

Although the insufficient precision as well as the slow scanning rate render the prototype scanner inappropriate for use on a UAV, this work shows that the design is conceptually feasible and further improvements might eventually lead towards a low-cost laser scanner suitable for obstacle detection on autonomous drones.

List of Figures

2.1 Basic optical setup for ranging measurements	7
Created in Inkscape.	
2.2 Optical setup to align emitted beam and returning light	9
Created in Inkscape.	
2.3 Mirror hexagon to redirect the emitted and returning light	10
Created in Wolfram Mathematica 10.	
2.4 Photograph of the assembled optical system	11
Captured with a Panasonic Lumix DMC-FZ30, edited in GIMP.	
2.5 Circuit driving the laser diode	15
Created in KiCad EDA and Inkscape.	
2.6 Operating principle of a photodiode	15
Created in KiCad EDA and Inkscape.	
2.7 A load resistor as simple photodiode front end	16
Created in KiCad EDA and Inkscape.	
2.8 Basic photodiode transimpedance amplifier circuit	18
Created in KiCad EDA and Inkscape.	
2.9 Effect of compensation on the loop transfer function	19
Created in Wolfram Mathematica 10.	
2.10 Frequency compensated transimpedance amplifier	20
Created in KiCad EDA and Inkscape.	
2.11 Gain plot of the transimpedance amplifier	21
Created in Wolfram Mathematica 10.	
2.12 Noise sources in the transimpedance amplifier	24
Created in KiCad EDA and Inkscape.	
2.13 Transimpedance amplifier with cascode	26
Created in KiCad EDA and Inkscape.	

2.14	Noise contribution of the cascode	27
	Created in Wolfram Mathematica 10.	
2.15	Transimpedance amplifier with regulated cascode	29
	Created in KiCad EDA and Inkscape.	
2.16	Gain plot of the transimpedance amplifier	30
	Created in Wolfram Mathematica 10.	
2.17	Frequency dependent signal waveform alteration	31
	Created in Wolfram Mathematica 10.	
2.18	Basic comparator circuits	32
	Created in KiCad EDA and Inkscape.	
2.19	Effects of hysteresis on a noisy signal	33
	Created in Wolfram Mathematica 10.	
2.20	Effects of input signal offset on comparator operation	34
	Created in Wolfram Mathematica 10.	
2.21	Overview of the time-of-flight measuring circuit	36
	Created in KiCad EDA and Inkscape.	
2.22	Waveforms of the delay measurement with AND gate and integrator	37
	Created in Wolfram Mathematica 10.	
2.23	Limits of the delay measuring approach	37
	Created in Wolfram Mathematica 10.	
2.24	Correcting comparator offset during delay measurement	38
	Created in Wolfram Mathematica 10.	
2.25	Signals generated by the controller to control circuit behavior	41
	Created in Wolfram Mathematica 10.	
3.1	Measured frequency response of the transimpedance amplifier circuit	48
	Created in Wolfram Mathematica 10.	
3.2	Captured output of the transimpedance amplifier	49
	Captured with a Rigol DS1052E Digital Storage Oscilloscope.	
3.3	Evaluation of the amplitude dependence	53
	Captured with a Panasonic Lumix DMC-FZ30, edited in GIMP.	
3.4	Amplitude dependence of the measurements	53
	Created in Wolfram Mathematica 10.	
3.5	Distortion caused by the amplitude dependence	55
	Created in Wolfram Mathematica 10.	

3.6	Photograph of the living room	56
	Captured with a Panasonic Lumix DMC-FZ30, edited in GIMP.	
3.7	Photograph of the setup used to orient the laser scanner	57
	Captured with a Panasonic Lumix DMC-FZ30, edited in GIMP.	
3.8	Cross-sections of a living room space, obtained from vertical scanning . .	58
	Created in Wolfram Mathematica 10.	
3.9	Rendering of a 3-dimensional model built from scanned data	60
	Created in Blender.	

List of Tables

2.1	Standard deviation of distance measurements for various SNR	42
3.1	Expected and measured photocurrent from light reflected by a bright surface	46
3.2	Anticipated standard deviation of distance measurements	50

Bibliography

- [AB76] Texas Instruments Inc. “High Speed Amplifier Techniques: A Designer’s Companion for Wideband Circuitry”. In: *Burr-Brown Application Bulletin* 76 (Feb. 1994).
- [ADL65] *ADL-65075TL AlGaInP Visible Laser Diode*. Manufacturer datasheet. Arima Lasers Corp.
- [Alm+16] Randa Almadhoun et al. “Aircraft Inspection Using Unmanned Aerial Vehicles”. In: *Proceedings of the International micro air vehicle competition and conference* (Oct. 2016), pp. 43–49.
- [Ama17] Ron Amadeo. *Google’s Waymo invests in LIDAR technology, cuts costs by 90 percent*. Ars Technica by WIRED Media Group. Jan. 10, 2017. URL: <https://arstechnica.com/cars/2017/01/googles-waymo-invests-in-lidar-technology-cuts-costs-by-90-percent/> (visited on 03/28/2017).
- [AN1258] Kumen Blake. *Op Amp Precision Design: PCB Layout Techniques*. Application Note AN1258. Microchip Technology Inc.
- [AN202] Paul Brokaw. *An IC Amplifier User’s Guide to Decoupling, Grounding, and Making Things Go Right for a Change*. Application Note AN-202. Analog Devices, Inc.
- [AN358] Lewis Smith and D. H. Sheingold. *Noise and Operational Amplifier Circuits*. Application Note AN-358. Analog Devices, Inc.
- [AN47] Jim Williams. *High Speed Amplifier Techniques: A Designer’s Companion for Wideband Circuitry*. Application Note 47. Linear Technology Corp. Aug. 1991.
- [Ard05] John Ardizzoni. “A Practical Guide to High-Speed Printed-Circuit-Board Layout”. In: *Analog Dialogue* 39.9 (Sept. 2005), pp. 1–6.
- [Bax68] P. J. Baxandall. “Noise in Transistor Circuits”. In: *Wireless World* 74.1397 (Nov. 1968), pp. 388–392.
- [BGO09] Antoni Burguera, Yolanda González, and Gabriel Oliver. “Sonar Sensor Models and Their Application to Mobile Robot Localization”. In: *Sensors* 9.12 (Dec. 2009), pp. 10217–10243.

- [CW11] James B. Campbell and Randolph H. Wynne. *Introduction to Remote Sensing*. 5th ed. Guilford Press, 2011.
- [Gou59] R. Gordon Gould. “The LASER, Light Amplification by Stimulated Emission of Radiation”. In: *Proceedings of the Ann Arbor Conference on Optical Pumping* (June 1959), p. 128.
- [HH15] Paul Horowitz and Winfield Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015.
- [Hob01] Philip C. D. Hobbs. “Photodiode Front Ends: The Real Story”. In: *Optics & Photonics News* 12.4 (Apr. 2001), pp. 42–45.
- [Hob09] Philip C. D. Hobbs. *Building Electro-Optical Systems: Making It All Work*. 2nd ed. Wiley, 2009.
- [IEC60825] *IEC 60825-1 Safety of laser products - Part 1: Equipment classification and requirements*. 3rd ed. International Electrotechnical Commission. May 2014.
- [JB10] Richard C. Jaeger and Travis N. Blalock. *Microelectronic Circuit Design*. 4th ed. McGraw-Hill, 2010.
- [Kar00] James Karki. *Effect of Parasitic Capacitance in Op Amp Circuits*. Application Report. Texas Instruments Inc. Sept. 2000.
- [Kra16] Donald Krambeck. *The Affordable LiDAR Movement - Researchers Aim to Drastically Reduce Sensor Costs*. All About Circuits by EETech Media, LLC. Sept. 29, 2016. URL: <https://www.allaboutcircuits.com/news/affordable-lidar-movementresearchers-aim-to-drastically-reduce-sensor-costs/> (visited on 03/28/2017).
- [Li+16] Jiaxin Li et al. “Real-time Simultaneous Localization and Mapping for UAV: A Survey”. In: *Proceedings of the International micro air vehicle competition and conference* (Oct. 2016), pp. 237–242.
- [LTC6268] *LTC6268/LTC6269 500MHz Ultra-Low Bias Current FET Input Op Amp*. Manufacturer datasheet. Linear Technology Corp.
- [Man04] Ron Mancini. “Op Amp stability and input capacitance”. In: *Analog Applications Journal* Q1 (2004), pp. 24–27.
- [MGH03] Sybil P. Parker, ed. *McGraw-Hill Dictionary of Scientific and Technical Terms*. 6th ed. McGraw-Hill Education, 2003.
- [MT059] *Compensating for the Effects of Input Capacitance on VFB and CFB Op Amps Used in Current-to-Voltage Converters*. Tutorial MT-059. Analog Devices, Inc.
- [Muo17] Danielle Muoio. *Google just made a big move to bring down the cost of self-driving cars*. Business Insider, Inc. Jan. 8, 2017. URL: <http://www.businessinsider.de/googles-waymo-reduces-lidar-cost-90-in-effort-to-scale-self-driving-cars-2017-1> (visited on 03/28/2017).

- [OED89] John Simpson and Edmund Weiner, eds. *The Oxford English Dictionary*. 2nd ed. Oxford University Press, 1989.
- [OPA656] *OPA656 Wideband, Unity-Gain Stable, FET-Input Operational Amplifier*. Manufacturer datasheet. Texas Instruments Inc.
- [PW18] Christopher V. Poulton and Michael R. Watts. *MIT and DARPA Pack Lidar Sensor Onto Single Chip*. IEEE Spectrum. Aug. 4, 2018. URL: <http://spectrum.ieee.org/tech-talk/semiconductors/optoelectronics/mit-lidar-on-a-chip> (visited on 03/28/2017).
- [PY04] Sung Min Park and Hoi-Jun Yoo. “1.25-Gb/s Regulated Cascode CMOS Transimpedance Amplifier for Gigabit Ethernet Applications”. In: *IEEE Journal of Solid-State Circuits* 39.1 (Jan. 2004), pp. 112–121.
- [Qia+16] Yuying Qian et al. “Design and Implementation of a Configurable Multi-rotor UAV”. In: *Proceedings of the International micro air vehicle competition and conference* (Oct. 2016), pp. 95–100.
- [RH09] Julian Ryde and Nick Hillier. “Performance of laser and radar ranging devices in adverse environmental conditions”. In: *Journal of Field Robotics* 26.9 (Sept. 2009), pp. 712–727.
- [RIEGL] RIEGL Laser Measurement Systems GmbH. *Airborne Scanning*. 2017. URL: <http://www.riegl.co.at/nc/products/airborne-scanning/> (visited on 03/28/2017).
- [Rin63] James Ring. “The Laser in Astronomy”. In: *New Scientist* 344 (June 1963), pp. 672–673.
- [S5972] *Si PIN photodiodes S5971, S5972 and S5973 series*. Manufacturer datasheet. Hamamatsu Photonics K. K.
- [Sch15] Karl Schulmeister. “Analysis of pulsed emission under Edition 3 of IEC 60825-1”. In: *Proceeding of the International Laser Safety Conference, March 23-26 2015, Albuquerque, New Mexico* (Mar. 2015), pp. 78–84.
- [SG90] Eduard Säckinger and Walter Guggenbühl. “A High-Swing, High-Impedance MOS Cascode Circuit”. In: *IEEE Journal of Solid-State Circuits* SC-25.1 (Feb. 1990), pp. 1–10.
- [SICK] SICK AG. *3D laser scanners*. 2017. URL: <https://www.sick.com/us/en/product-portfolio/detection-and-ranging-solutions/3d-laser-scanners/c/g282752> (visited on 03/28/2017).
- [Sim17] Tom Simonite. *Self-Driving Cars’ Spinning-Laser Problem*. MIT Technology Review. Mar. 20, 2017. URL: <https://www.technologyreview.com/s/603885/autonomous-cars-lidar-sensors/> (visited on 03/28/2017).
- [SK12] D. Skarlatos and S. Kiparissi. “Comparison of Laser Scanning, Photogrammetry and SFM-MVS Pipeline applied in Structures and Artificial Surfaces”. In: *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 1.3 (Sept. 2012), pp. 299–304.

- [SLC16] Mo Shan, Feng Lin, and Ben M. Chen. “Salient Object Detection Using UAVs”. In: *Proceedings of the International micro air vehicle competition and conference* (Oct. 2016), pp. 249–255.
- [SNH03] Hartmut Surmann, Andreas Nüchter, and Joachim Hertzberg. “An autonomous mobile robot with a 3D laser range finder for 3D exploration and digitalization of indoor environments”. In: *Robotics and Autonomous Systems* 45.3 (Dec. 2003), pp. 181–198.
- [Vel16] Velodyne LiDAR, Inc. *PUCK VLP-16*. 2016. URL: <http://velodynelidar.com/vlp-16.html> (visited on 03/28/2017).
- [VM10] George Vosselman and Hans-Gerd Maas. *Airborne and Terrestrial Laser Scanning*. Whittles Publishing, 2010.
- [Way17] Waymo LLC. *A note on our lawsuit against Otto and Uber*. A Medium Corp. Feb. 24, 2017. URL: <https://medium.com/waymo/a-note-on-our-lawsuit-against-otto-and-uber-86f4f98902a1> (visited on 03/28/2017).
- [XMEGA] *ATxmega32A4U Device Datasheet*. Manufacturer datasheet. Atmel Corp.
- [YY97] Francis T. S. Yu and Xiangyang Yang. *Introduction to Optical Engineering*. Cambridge University Press, 1997.