Imaging light in motion and its application to tracking hidden objects

by
GENEVIEVE GARIEPY

Submitted for the degree of Doctor of Philosophy

Heriot-Watt University

School of Engineering and Physical Sciences

October 2016

The copyright in this thesis is owned by the author. Any quotation from the thesis or use of any of the information contained in it must acknowledge this thesis as the source of the quotation or information.

Abstract

It is well known that light, the fastest entity in the universe, moves at a staggering speed of 300 millions meters per second. The ability to stop its flight on a centimetre scale or lower requires a detector with temporal resolution of around a hundred picoseconds. Freezing light in motion at this scale is a feat worth achieving, as it leads to a variety of exciting applications, from observing dynamical light phenomena to measuring distances and depths with high precision, as in LIDAR technology. In the past decades, different technologies have been developed to image light in motion; in this thesis, we propose a new method that exploit a recently-developed single-photon detector technology to capture movies of light in motion at very low intensity levels. We use this method to develop novel imaging applications and detection techniques. In particular, this thesis reports on the observation and study of dynamical light phenomena such as laser propagation in air, laser-induced plasma, propagation in optical fibres and slow light. We also show how the ability to record light in motion can be used to capture light signals scattered from around an obstacle, leading to the ability to locate and track moving objects hidden from view.

"*Yet there is also something curiously exhilarating about the contemplation of deep time. True, you learn yourself to be a blip in the larger projects of the universe. But you are also rewarded with the realisation that you do exist - as unlikely as it may seem, you do exist. "*

Robert Macfarlane, *Mountains of the mind*

To all who will read this thesis. Thanks for taking this journey with me - I hope you enjoy it, and find it useful!

Acknowledgements

I would like to thank: my supervisor, Prof. Daniele Faccio, for helping me make this PhD not only a successful enterprise, but also a very instructive and a truly enjoyable one; Dr Jonathan Leach, who provided generously valuable insights and expertise to the work presented in this thesis; all members of the Extreme Light Group and the Quantum Optics Group who have helped in this work, or simply contributed to creating the great ambiance and environment we have in our groups. It has been a pleasure to work with and alongside all of you. Another thank you goes to Dr Robert Henderson and the members of the CMOS Sensors and Systems group, who have been truly helpful collaborators of this work.

Thanks also to my incredible family and partner for providing the backbone of love and happiness to my life that keeps me believing I can accomplish anything I set my mind to, including writing this thesis, and expatriating myself to a snow-deprived part of the world to do so (which is otherwise an absolutely fantastic place!)

Table of contents

Abstract **i**

1 Introduction **1**

2 The quest for imaging light in motion: motivation and context **3**
2.1 Prior art 3
2.2 Thesis contributions in context 12

3 A new tool for imaging light in motion: single-photon avalanche diode arrays **14**
3.1 How to detect single photons with picosecond time resolution 14
3.2 SPAD arrays 18
3.3 Characterisation of a SPAD array detector 24

4 Imaging light in motion in action: light-in-flight experiments **37**
4.1 Laser propagation in air 37
4.2 Laser-induced plasma 51
4.3 Laser propagation in fibers 55
4.4 Slow light 56
4.5 Imaging light through an opaque object 60

5 Application of imaging light in motion: detecting and tracking hidden objects **67**
5.1 Prior art 68
5.2 Detection technique and position retrieval method 71
5.3 Proof-of-principle experiment 76
5.4 Single pixel implementation 96

6 Conclusion **108**

A Light-in-flight and tracking videos **111**

B Increasing the spatial resolution of the SPAD camera **112**

C Imaging geometry of SPAD camera-based non-line-of-sight tracking system **114**

D Optimising the collected signal in a non-line-of-sight tracking system **117**

List of Publications

Papers

1. G. Gariepy, N. Krstajic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Single-photon sensitive light-in-flight imaging*. Nature Communications **6** (2015)

2. G. Gariepy, F. Tonolini, R. Henderson, J. Leach and D. Faccio. *Detection and tracking of moving objects hidden from view*. Nature Photonics **10**, 23-26 (2016)

3. E. Bolduc, G. Gariepy, J. Leach. *Direct measurement of large-scale quantum states via expectation value of non-Hermitian matrices*. Nature Communications, **7** (2016)

4. R. Warburton, C. Aniculaesei, M. Clerici, Y. Altmann, G. Gariepy, R. McCracken, D. Reid, S. McLaughlin, M. Petrovich, J. Hayes, R. Henderson, D. Faccio and J. Leach. *Observation of laser pulse propagation in optical fibers with a SPAD camera*. Under review, Scientific reports (2016)

5. S. Chan, R. Warburton, G. Gariepy, J. Leach and D. Faccio. *Non-line-of-sight tracking of people at long range*. Under review, Optics Express (2016)

6. G. Gariepy, F. Tonolini, R. Henderson, J. Leach and D. Faccio. *Tracking objects hidden from view*. Accepted for OPN Optics in December 2016 special issue (October 2016)

7. K. Wilson, B. Little, G. Gariepy, R. Henderson, J. Howell and D. Faccio. *Slow light in flight imaging*. Submitted to PRA (2016)

First author and speaker conference contributions

1. G. Gariepy, F. Tonolini, R. Warburton, S. Chan, R. Henderson, J. Leach and D. Faccio. *Detection and tracking of moving objects hidden from view*. Oral presentation, OSA Imaging and applied optics meeting (July 2016)

2. G. Gariepy, R. Warburton, S. Chan, R. Henderson, J. Leach and D. Faccio. *Tracking hidden objects at human scale*. Oral presentation, Photonics North 2016 (May 2016)

3. G. Gariepy, F. Tonolini, R. Warburton, S. Chan, R. Henderson, J. Leach and D. Faccio. *Tracking moving objects hidden from view*. Oral presentation, IONS Quebec 2016 (May 2016)

4. G. Gariepy, F. Tonolini, R. Henderson, J. Leach and d. Faccio. *Tracking moving objects hidden from view*. Poster presentation, Photonics North 2015 (June 2015)

5. D. Faccio, J. Leach, G. Gariepy and R. Henderson. *Imaging at the speed of light*. Invited oral presentation, URSI Atlantic Radio Science Conference (AT-RASC) 2015 (May 2015)

6. G. Gariepy, N. Krastjic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Single-photon sensitive light-in-flight imaging*. Invited oral presentation, Symposium on structured light and orthogonality in imaging, Rank Prize Funds (November 2014)

7. G. Gariepy, N. Krastjic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Light-in-flight imaging in free-space using arrayed single-photon detector technologies*. Oral presentation, Frontiers in optics 2014 (October 2014)

8. G. Gariepy, N. Krastjic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Arrayed single-photon technologies for free-space light-in-flight imaging*. Oral presentation, Photon14 (September 2014)

9. G. Gariepy, N. Krastjic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, C. Barsi B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Arrayed single-photon technologies for free-space light-in-flight imaging*. Oral presentation, Photonics North 2104 (June 2014)

10. G. Gariepy, N. Krastjic, R. Henderson, C. Li, R.R. Thomson, G.S. Buller, C. Barsi B. Heshmat, R. Raskar, J. Leach and D. Faccio. *Light-in-flight imaging in free space with single-photon technology*. Oral presentation, IONS NA-8 (June 2014)

Chapter 1 – Introduction

High-speed imaging and photography seems to be getting ever faster. From commercial cameras that operate at tens of thousands frames per second and allow to capture movies of exploding melons [1] and running cheetahs [2], to advanced experimental techniques that can freeze the flight of a bullet [3] or witness a material being blasted by a laser [4,5], slow-motion movies and ultrafast imaging do not only capture the world's imagination, they also promote great technological and scientific leaps in diverse fields. However fast a camera can be, the ultimate slow-motion feat is to stop light in its path. Bursting through air at 300 million meters per second, light is well known to be the fastest entity in the universe – recording its motion on a human scale requires very well designed detection technology. Imaging light itself as it moves on a centimetre scale indeed requires a temporal resolution of about one hundred picoseconds. Achieving this technological prowess is well worth the effort, as the ability to freeze the motion of light can lead to a variety of novel applications. Some can be obvious, as the dynamical study of simple light phenomena such as propagation, reflection and focusing [6]. More complex phenomena can also be observed at those timescales, such as laser-induced plasma dynamics [7], or nonlinear processes in an optical fibre [8]. On the other hand, some applications of the capacity to freeze light in motion can be rather surprising, like the ability to detect or see objects hidden from view, by capturing light signals scattered from around an obstacle [9, 10].

Since this specific type of high-speed imaging is rich in applications, various methods have been developed since the 1960s that have the temporal resolution necessary to image light in motion. These methods can be used to record videos of light pulses as they move through a medium or across surfaces in a scene. In this thesis, we present a novel way to capture light in motion, that fills a space left blank in the array of existing technologies and their capabilities. We will show that we can image light propagating through air or other medium without any type of interaction with the phenomenon under study, and record directly movies of light moving on a centimetre scale with an apparatus that is very practical and portable. To do so, we exploit a recently-developed detector technology that consists of a single-photon avalanche diode (SPAD) array. Its inherent single-photon sensitivity allows us to capture ultrafast dynamics at very low intensity levels, meaning we can capture dynamical events that have not been previously observable, like laser pulses propagating through air, or capture low-light events at a much faster rate than what used to be possible, allowing us to develop detection technologies that operate in real-time.

Before detailing the technology we use to capture light in motion in this work, we will first take a moment to overview the existing methods that have this capability. Chapter 2 will present how various technologies have succeeded in stopping light in its flight from early experiments in the 1960s all the way up to the flourishing field high-speed imaging

is today. This will paint a picture of the imaging capabilities developed to this day, what they can accomplish and what still remains to be done. Chapter 2 will show that the work presented here places itself in that picture as a very practical and portable method that has the unprecedented ability to capture movies of light in motion in a direct fashion, thanks to its direct three-dimensional acquisition and its single-photon sensitivity.

In chapter 3, we will demystify the technology on which our method is based: single-photon avalanche diode (SPAD) arrays. SPAD detectors have been used for a few decades in multiple applications, in particular in LIDAR (light detection and ranging), quantum experiments and biomedical imaging. However, they were always used as single-pixel detectors. The interest for building arrays of them has grown in recent years, and arrays have been successfully developed for applications mainly in biomedical imaging and 3D imaging. Chapter 3 will explain how these detectors work and what gives them their two signature characteristics: their high temporal resolution ($\sim$100 ps) and single-photon sensitivity. In this chapter, we will also detail the different experimental calibration and characterisation experiments that were done on the SPAD array we use throughout this work, which is a 32x32-pixel detector developed by the CMOS Systems and Sensors group at the University of Edinburgh [11].

Chapter 4 and 5 will cover the main results of this thesis. First, we exploit our ability to record slow-motion movies of light to observe and study a variety of light phenomena in action. We will show that we can capture the propagation of laser pulses through air by relying only on the detection of photons scattered by air molecules in the laser's path [7]. We will also present the measurements of more complex light phenomena: laser-induced plasma in air [7], laser propagation in optical fibres [8], slow-light in rubidium [12] and light diffusion in a scattering medium [13].

In addition to the fundamental study of dynamical light phenomena, the ability to capture light in motion can also be used to measure distances and position. LIDAR is a classic example of how capturing the time-of-flight of light with high temporal resolution can be used to measure precisely depth and distances. However, LIDAR is a technique that is limited to measuring objects in its line-of-sight. In chapter 5, we will show how we can use our camera technology to capture signals from an object hidden behind an obstacle and retrieve its position. Our ability to capture light in motion at low intensity levels allows us to record signals scattered from around an obstacle in only a few seconds, and retrieve the position of a hidden object in real-time [10].

The work presented in this thesis lays the foundations for further work in both these directions: the study of fundamental light dynamics and applications in detecting hidden objects. In chapter 6, the conclusion, we will discuss various directions of interest for future work.

Chapter 2 – The quest for imaging light in motion: motivation and context

The quest to capture light in motion started very soon after the invention of ultrafast lasers: as they consist of small bursts of light of a few millimetres or centimetres long (femtosecond/picosecond duration), the question of observing their position and aspect at a given point in space and time, or to record their motion, immediately became relevant and interesting to researchers at that time. Early experiments in the 1960s allowed to visualise picosecond pulses frozen in time, and work followed in the 1970s and up to today to create movies of light pulses evolving in different scenarios, allowing to observe in a new way phenomena that we normally observe statically. For example, capturing light in motion allowed to observe reflection, interference and focusing of light pulses in a dynamic fashion. Interest in capturing light illuminating a scene with temporal resolution also grew, and became known as transient imaging. Studies of light propagation at picosecond timescales also allowed to deepen our understanding of light transport, and develop novel applications for imaging objects normally hidden from our view, either by an opaque scattering layer, or behind a solid obstacle. We will discuss these specific applications in later sections of this thesis - for now, this chapter focuses on the different methods that have been developed to image light in motion and to perform transient imaging.

The main challenge in capturing light in motion is the temporal resolution needed to do so. Light moves at a speed that is beyond anything we can intuitively grasp. A speed of 300 millions meters per second means that, in the time of an eye blink, light has time to travel around the world about three times. To stop its motion on a centimetre scale, a detector with a temporal resolution of the order of picoseconds or nanoseconds is needed. This in itself is quite a technical challenge, and we will see in section 2.1 how different approaches have achieved these temporal resolutions to record the dynamics of light in motion. In section 2.2, we will then point out the contributions of this thesis to the ongoing quest to capture light in motion.

2.1 Prior art

Throughout the years, various technological principles have been exploited to achieve temporal resolution in the picoseconds and nanoseconds timescales. Figure 2.1 shows a classification of the main imaging methods presented in this section. The first methods developed in the 1960s and 1970s were based on some form of temporal gating. This type of high-speed imaging presents similarities with the use of a mechanical shutter, but

an optical implementation of a shutter allows for much shorter exposure time that what is achievable with a mechanical shutter, as discussed in sections 2.1.1 and 2.1.2.

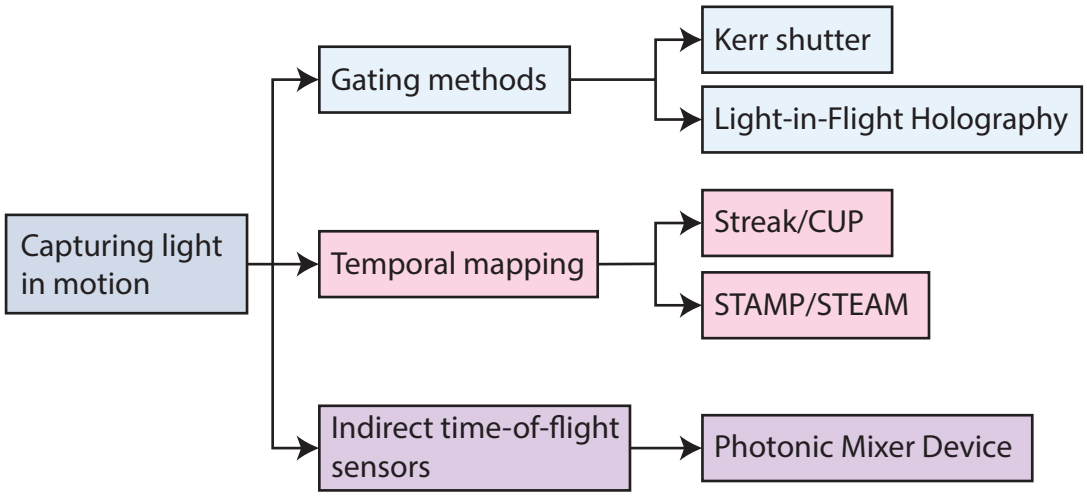

FIGURE 2.1: **Types of techniques for measuring time of arrival.** The methods developed to perform light-in-flight imaging are based on different techniques to measure the time of arrival of the detected light. These techniques can be classified in three categories: gating methods, in which a gate or shutter is used to collect the light during a specific exposure time; temporal mapping techniques, in which the temporal information is mapped onto another dimension, easier to detect (spatial dimension for streak cameras, wavelength and spatial dimension for STEAM/STAMP); and indirect time-of-flight sensors, that rely on modulation of the illumination and detection to infer the time-of-arrival of detected signals.

More recently, new methods have been proposed to achieve picoseconds temporal resolution by mapping the time information to record onto a different axis that is more straightforward to measure. In particular, methods based on streak cameras have allowed the capture of laser pulses travelling through a scattering medium with very high temporal and spatial resolution.

Finally, indirect time-of-flight sensors such as photonic mixers devices have also be used to record the propagation of light pulses across a scene. The following section covers these various techniques and their applications in more details.

2.1.1 Early methods for picosecond pulse display

Light-in-flight imaging and ultrafast laser pulses come hand in hand, as short bursts of light are needed to observe any type of dynamics of light - a continuous source of light, even photographed with picosecond exposure, would not display any spatial structure that a time-integrated measurement would not show. One of the first method to visualise light-in-flight was therefore a method to display picosecond pulses. In 1967, a group from Bell Labs managed to take a photograph of a light pulse by exploiting a nonlinear interaction of the pulse with a solution of organic molecules [14]. They sent two beams of

a picosecond laser into the solution from opposite directions, creating a standing wave in the middle of the solution. The two-photon excitation of the solution is then significantly stronger at the standing wave position, and very weak at any other point of propagation where there is no constructive interference between the two counter-propagating beams. The resulting fluorescence pattern in the solution was photographed and used as a way to visualise picosecond pulses for the first time.

Soon after, another group from Bell Labs developed a method that allowed to photograph picosecond pulses by relying only on a linear interaction between the pulses and a medium. They set a camera to look at picosecond laser pulses propagating through a milky solution, and installed a shutter in front of the camera. As a mechanical shutter could not give an exposure time short enough to photograph with accuracy the picosecond pulses in the milky solution, they used a shutter based on the Kerr effect: a birefringent CS_2 plate was positioned between two crossed polarisers, and optically activated with short laser pulses with a specific delay with respect to the pulses to observe: this would change the polarisation of the light passing through the shutter between the polarisers and create a short exposure time of about 10 ps during which light made its way through the crossed polarisers [16]. A simplified version of this imaging setup is shown in figure 2.2, as well as an example of a photograph taken with this technique. The publications detailing this technique ([15] and [17]) also discuss the different implications of visualising light in motion and the effects that could be observed from photographing objects moving at that speed, where the timescale of the object's propagation is of the same order than the time it takes for the light to travel from the scene to the camera. Section 2.1.6 will

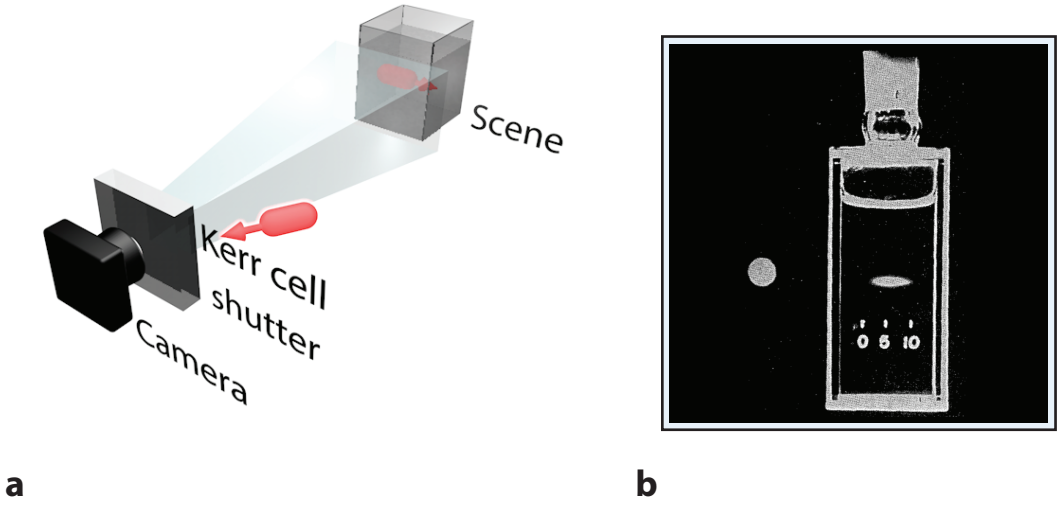

FIGURE 2.2: **Early method for picosecond pulse display.** a) Picosecond pulses pass through milky water in the field of view of a camera. The milky water assures that there is enough scattering from the pulses so that enough light is deflected towards the camera for detection. A Kerr cell shutter is placed in front of the camera, and is optically activated by a second, synchronous train of short pulses. b) During that short exposure, the propagation of the picosecond pulses was frozen in time, creating a photograph of a light pulse frozen in its path. Picture taken from [15].

discuss in more details these effects. This method was at that time a good tool to observe the temporal intensity profile of the photographed pulses. It was applied to gated picture ranging, where the temporal resolution is used to image scenes at different depths and even through some scattering material by selecting the light reaching the camera at the time of interest.

2.1.2 Light-in-flight holography

The term light-in-flight was coined by Abramson in 1978 [18]. He developed a holographic method for recording movies of light in motion. One of his first experiments is illustrated in figure 2.3. To create a hologram, one needs a reference beam and a signal beam coming from the scene of interest. When using a short pulse of light as a reference beam, you record in the hologram a picture carried in the signal beam with an exposure time corresponding to the reference pulse duration: you can hence acquire holograms with very high temporal resolution. Furthermore, Abramson illuminated his holographic plate with a reference beam at an angle, that had a short duration in time and a wide transverse profile, in such a way that the reference pulse reaches the left part of the plate before the right part. The hologram plate is therefore imprinted with a sequence of photographs corresponding to different times, hence recording a movie of light in motion directly on a single holographic plate. To observe a scene of light propagating, the same trick is used. In figure 2.3, we see that to record a reflection event dynamically, a beam is sent on a scattering surface at an angle: the beam therefore intersects parts on this surface at dif-

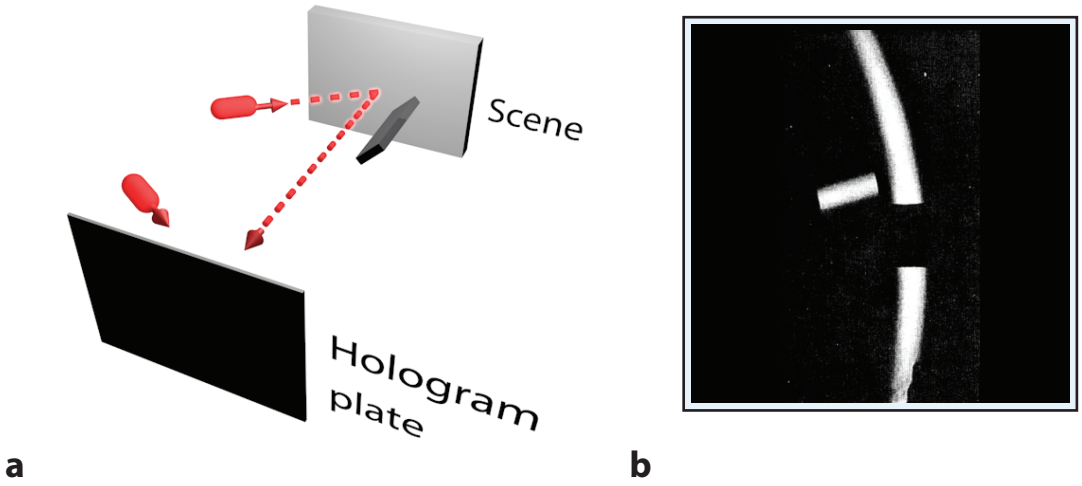

FIGURE 2.3: **Light-in-flight holography.** a) A beam that we want to observe in motion is sent on a scattering surface at an angle (scene), from which it is scattered onto the hologram plate, where it interferes with the reference beam. In this example, a mirror is placed in the middle of the scene to reflect part of the signal beam upwards. b) The movie recorded can be replayed by illuminating the hologram plate at different positions. This frame is an example taken from [18] where we see the curved wavefront of the signal beam as it intercepts the scattering plate, propagating from left to right, and the part of the beam that hit the mirror is now propagating upwards.

ferent times, and is scattered back to the holographic plate where it can interfere with the reference beam. By placing a mirror in the scene, one can observe how part of the signal beam is reflected upwards, then intersects with the scattering surface. Once the hologram is imprinted on the plate, one simply needs to move the observation point continuously to recreate a movie of the light propagating and reflecting. Figure 2.3b shows an example of a frame in such a movie, where the curved wavefront of the signal beam is recorded with very high temporal resolution (dependant on the geometry of the setup, of the order of 10ps). This holographic method was used to record a variety of light phenomena normally observed statically: reflection, focusing and interferometry [6]. The different effects of imaging at that speed were also discussed [19], and the ability to image objects at different depths by selecting the corresponding time of arrival was exploited to do 3D contouring of different objects, by sending pulses of light onto them and looking at these actively illuminated objects as the scene in the light-in-flight holography setup [6, 20].

This method was extensively used and improved throughout the years. Notably, Abramson presented work done in a single shot acquisition in 1989 [20], rather than relying on multiple exposures of repetitive events. More recently, this method has been extended to the recording of femtosecond pulses in three-dimensions, by sending the pulses to be observed in a scattering medium rather than relying on a 2D projection on a scattering surface as in the original work [21, 22]. It is now also possible to record digitally these holograms rather than having the measurement done on a physical holographic plate [23], hence giving access to digital light-in-flight videos.

2.1.3 Streak cameras

More recently, a team at MIT used a streak camera to develop a technique able to perform light-in-flight measurements and transient imaging [24, 25]. A streak camera is a type of detector that is used in different fields for its ability to measure temporal information with high resolution. It does so by mapping the temporal information onto a spatial axis: it consists of a 2D detector, which is imaging a single line in the spatial dimension (x axis) through a horizontal slit. The light captured through this slit is detected to generate photoelectrons that are accelerated towards a phosphorous screen while a sweeping transverse electric field is applied. As a result, light that is detected early at the entrance of the camera generates photoelectrons that are submitted to a low electric field and barely deflect in the vertical direction, while light that is detected at a later time generates photoelectrons that are submitted to a strong field which deflects them significantly in the y transverse direction as they travel to the phosphorous screen. The t dimension is therefore mapped in the y dimension of the detector [26]. Figure 2.4a shows a simplified illustration of the imaging setup for this method: a streak camera, which is quite a bulky apparatus, is imaging a horizontal line of a scene. The scene illustrated here is a femtosecond pulse propagating through a milky solution - in the reported experiment,

the pulses were propagating through a plastic bottle filled with a mix of water and milk. Since the streak camera can only acquire one line at a time, a system is in place to scan vertically the scene. The acquisition takes around one hour, and results in a movie of the scene recorded with a temporal resolution of 2 ps. Figure 2.4b shows an example frame of the movie obtained where the pulse is observed as it propagates through the bottle. A few different scenes are presented in [24] and [25], including some transient imaging measurements where a diffuse femtosecond pulse illuminates a scene containing multiple scattering and reflecting surfaces. This allows the clear visualisation, with high temporal and spatial resolution, of the multiple paths light takes as it travels through a scene. It also shows neatly the effects of imaging at the speed of light, that we will discuss in section 2.1.6.

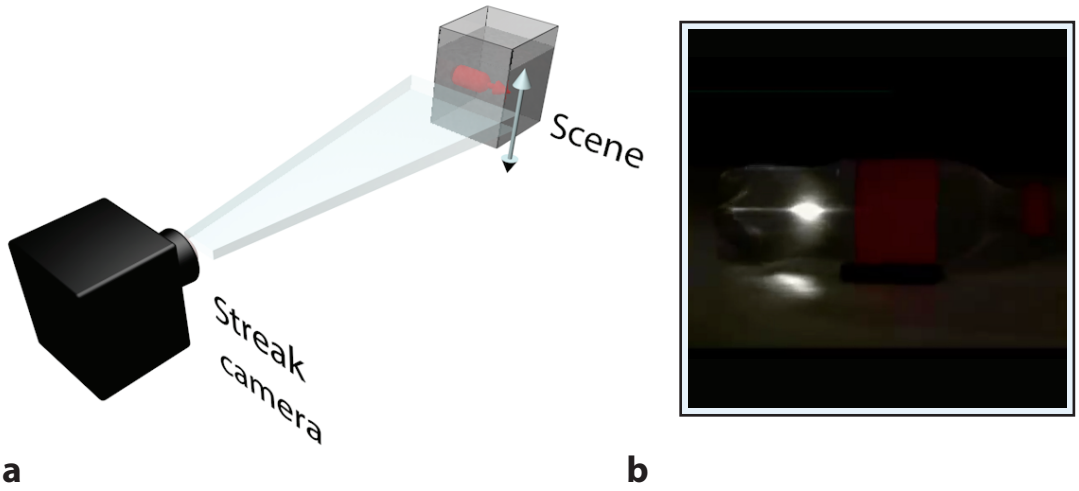

FIGURE 2.4: **Streak camera femtophotography.** a) Laser pulses propagate through milky water, to assure sufficient scattering of the pulses so that a signal can be detected by the camera. The streak camera images a horizontal line of the scene (x axis). To acquire a full (x, y, t) scene, the line of the field of view is scanned vertically across the scene. b) The work reported in [24] and [25] shows a recorded movie of a femtosecond light pulse propagating through a plastic bottle filled with milky water, with a resolution of 2 ps. Image taken from [25] supplementary video.

Since acquiring a scene with a streak camera requires scanning, an acquisition inherently needs to be an integration of the signal over multiple pulses. However, there is a lot of interest in the community to develop imaging methods that rely only on a single pulse acquisition. Gao $et\ al.$ have tackled this problem and demonstrated that it is possible to perform single-shot measurements with a streak camera [27] by incorporating techniques of compressed sensing into the measurement, in a method they named "Compressed Ultrafast Photography" (CUP). The principle is to get rid of the horizontal entrance slit in a streak camera. This will result in a y axis on the detector that contains blurred information about both time and the y spatial dimension of the scene. To reduce the amount of overlap between these two axes, they encode the spatial information coming from the scene with a sparse pattern using a digital micromirror device. This allows them to reconstruct the

full (x, y, t) information of the imaged scene by finding the solution to the inverse problem that takes into account the sparse pattern and the temporal mapping in y. Using this novel technique, they have recorded single-shot acquisitions of light propagation, reflection and refraction, as well as single-shot fluorescence measurements with a resolution of 10 ps. This opens the possibility to use streak cameras to perform measurements with high temporal resolution of events that are not repetitive.

2.1.4 Photonic mixer devices

The techniques presented above are all very powerful at visualising light in motion, but they are also quite intricate and expensive. A cheaper solution to perform transient imaging has been developed by Heide $et\ al.$ in [28]. This solution is based on the use of a photonic mixer device (PMD), a type of indirect lock-in time-of-flight camera, which performs indirect time-of-flight measurements by modulating in time a continuous light signal (e.g. from a laser diode) that is then used to illuminate a scene. The phase offset of the returning signal is then measured to infer the time-of-flight of said signal [29]. Several images are acquired with different modulation frequencies of the illumination source, and the (x, y, t) scene can be reconstructed by finding the solution to the corresponding inverse problem [28]. This device was used to perform transient imaging measurements of light sweeping across a scene, before it reaches an equilibrated distribution. This measurement was done for a cost of a few hundreds dollars, and in a relatively quick acquisition time ($\sim$90 s). The temporal resolution can vary as a function of the modulation frequency of the source, from a few nanoseconds [28] to 70 ps reported in [30].

2.1.5 Other ultrafast recording techniques

In the last few years, many novel detection techniques have successfully acquired impressive images and movies of different dynamic scenes with high temporal resolution. The ones reported below are not aimed at capturing light itself in motion, but are relevant to mention in the context of this thesis as they achieve a temporal resolution similar to the one needed to perform light-in-flight experiments. In particular, a method reported in 2009 in [4] called serial time-encoded amplified imaging (STEAM) uses a mapping technique to capture 2D frames at high repetition rates. It first encodes the spatial 2D information of a scene by sending an illumination of varying wavelength in x and y, by spatially dispersing a large bandwidth pulse. The signal reflected from a scene is then sent into an amplified dispersive Fourier transformer, a fibre component in which a pulse disperses so that its wavelength signature gets mapped into the temporal dimension. The apparatus can capture a frame with a single pulse of light, and its detected temporal profile can be directly reconstructed into a spatial 2D frame. This technique can therefore acquire frames at the repetition rate of the laser used, with an effective exposure time corresponding to the pulse length. A frame rate of 163 ns and exposure time of 440 ps have

been reported, which was used for the single-shot measurement of non-repetitive events such as microfluidic flow and laser ablation [4, 5].

Another interesting recent ultrafast imaging technology is "Sequentially timed all-optical mapping photography", or STAMP [31]. This method also relies on temporal mapping, and it can perform measurements with a single pulse. Here, a large bandwidth pulse is first dispersed and divided in a fast train of pulses of varying wavelength. These pulses are then incident on a scene and their transverse profile capture the spatial information of that scene at different times. The pulses are then separated in space by exploiting the fact that they have slightly different wavelength, and mapped onto a CCD detector using a grating system to send each pulse onto a different section of the CCD. This technique has been used to successfully observe the dynamics of plasma and lattice vibrational waves by capturing 6 frames of 450x450 pixels of these events. The frames' separation depends on the initial separation of the pulse trains, and can vary from tens of picoseconds to hundreds of femtoseconds. The effective exposure time is here too the duration of the pulses, around 100 fs. This powerful technique can be used in either microscopic or macroscopic configuration, and the reported limited number of six frames can be extended, reportedly to 100, by modifying the grating system that maps each pulse at different locations onto the CCD detector.

2.1.6 Particularities of imaging at the speed of light

When imaging an object that is moving at low speed (much smaller than the speed of light c), the time it takes for the light coming from the scene to reach the camera or detector is insignificant compared to the timescales of the motion observed in the scene. However, when the object being imaged is moving at relativistic speed, or is light itself moving at c, the timescale of the motion in the scene and the timescale of light from the scene to reach the camera are of the same order of magnitude, which gives rise to interesting effects. An example is illustrated in figure 2.5. Consider two pulses of light entering a scattering medium at different positions along the z axis (at different distances from the camera). They will take the same time t_0 to reach a given x position inside the medium. At that point, some of the light from the pulses will be scattered in the direction of the camera. However, since the scattering from the second pulse occurs farther away from the camera, the time it will take for the scattered light to reach the camera, t_2, will be slightly longer than the time it takes for the light scattered from the first pulse to reach the camera, t_1. If the objects imaged were moving much slower than c, the time t_0 would be significantly larger than t_1 and t_2, and the two objects would appear at the same place at the same time in the recorded movie. However, when imaging light moving at a speed of $\sim c$, there is a significant difference between the times $t_0 + t_1$ and $t_0 + t_2$. As illustrated in figure 2.5, for light to arrive at the same time at the camera, the pulse that is farther away needs to travel significantly less in x than the pulse closer to the

camera: in one recorded frame on the camera, the two pulses therefore appear to follow each other rather than appearing at the same position at the same time. This shearing of the observation was experimentally verified by Duguay and Mattick in 1971 [17] by photographing the motion of two pulses in different geometries. They also pointed out that a beam or relativistic object carrying a pattern in its transverse profile would show this pattern when photographed at the speed of light: you would for example be able to see the back of a coin moving transversely at $\sim c$. This effect was later observed in different light-in-flight studies [19, 21]. In particular, Velten *et al.* described this effect in the light-in-flight and transient imaging movies they recorded with a streak camera, and detailed how to correct for it and retrieve the scene with the real times corresponding to an event, in a process they named time unwarping [25]. Clerici *et al.* also recently exploited this effect to demonstrate that it is possible to record wavefronts that show an apparent motion at speeds greater than c, that appear to be moving backwards or that lead to the recording of creation and annihilation of pairs of images [32].

In the work presented in this thesis, we will not correct the recorded light-in-flight videos for this effect. As we will see in chapter 4, our scenes generally lie in a single plane

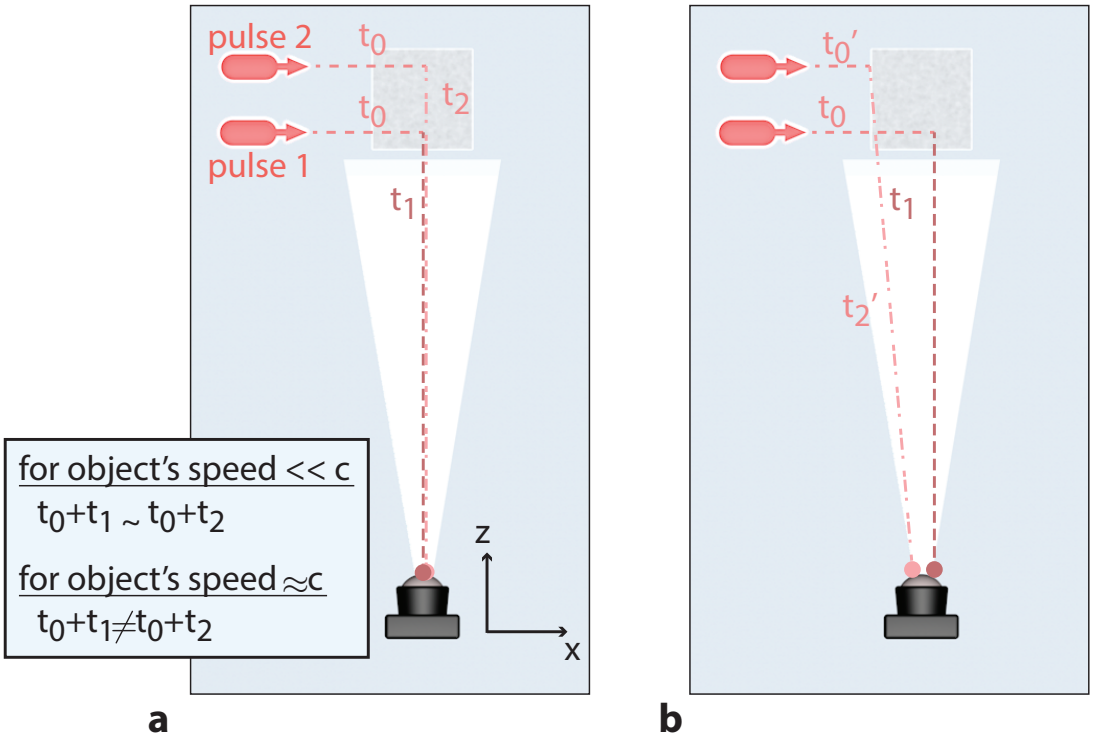

FIGURE 2.5: **Imaging at the speed of light.** a) The time it takes for two pulses to propagate to the same position in x is here t_o, the same for each pulse. The time it takes for their scattered signal to reach the camera however differs since they are propagating in different planes in z. If the objects move at the speed of light, the difference $t_2 - t_1$ becomes significant with respect to t_0. b) In a frame corresponding to a given time, the camera will therefore detect light that is coming from an object farther away after it has moved less in its propagation direction that an object located closer to the camera. The two beams will therefore appear to follow each other, even though in reality they propagate over the same distance in the same amount of time.

at a fixed distance to the camera, in a limited field of view; we therefore approximate that the time it takes for light to go from different (x, y) positions in the scene to the camera is essentially the same. An interesting study by Laurenzis *et al.* demonstrates how the effect explained above, when considered in a light-in-flight measurement where the pulses propagate at an angle with respect to the camera axis, can be used to infer more information about the recorded scene in three dimensions [33].

2.2 Thesis contributions in context

A summary of the characteristics of the main methods presented in section 2.1 is given below in table 2.1.

Method	Year	Applications	Limitations	Time resolution	Single shot?	Compact, portable
Light-in-flight holography	1978 [18]	• Imaging light in motion	• Requires reflections or scattering medium • Interferometric method	10-100 ps	Yes [20]	No
Streak camera	2011 [25]	• Imaging light in motion • Transient imaging	• Requires reflections or scattering medium • Need to scan to acquire a 2D scene	2 ps	No	No
CUP	2014 [27]	• Imaging light in motion	• Requires reflections or scattering medium • Requires post-processing and sparsity in the scene	10 ps	Yes	No
Photonic Mixer Device	2013 [28]	• Transient imaging	• Requires modulated illumination • Requires inverse problem solution for reconstruction	10 ns [28] up to 70 ps [30]	No	Yes

TABLE 2.1: **Summary of prior art.** Comparison of the characteristics and applications of the four main methods used to perform light-in-flight and transient imaging measurements, including the year they were first introduced.

We can see from this summary that all of these techniques require a scattering medium or surface to deflect enough light to the detector when imaging light in motion. So far, no technique has been developed that can directly record the motion of light in air, without any additional interaction than the one pulses undergo with air molecules in their path. Another conclusion that can be drawn from this summary is that most techniques are quite intricate to set up and use, the only compact and portable one being the PMD, which can

be used to perform transient imaging, but not to image light itself as it requires an active illumination of a scene. The aim of the work presented in this thesis was therefore to develop a novel light-in-flight imaging technique that could provide a direct measurement of light in motion, without interaction with the observed phenomenon, in a compact, portable and easy-to-use manner.

The main tool we use in this work to achieve this goal is an array of single-photon avalanche diodes (SPAD), that allows for the direct measurement of the spatial and temporal information of a scene, by detecting single photons with a temporal resolution of about 100 ps. This type of detector belongs to a different class of technology than those presented in figure 2.1: it relies on the direct measurement of the time of arrival of the detected photons. The next chapter will cover in details this type of technology, and explain how a SPAD detector works and why an array of SPADs is a great tool for imaging light in motion. In chapters 4 and 5, we will also further explore the use of a few other types of detectors for capturing light in motion. In particular, in chapter 4, we study the use of an optical gated imaging technique that has not been used yet to record light-in-flight: gated intensified CCD cameras. We will discuss how the results obtained with this type of camera compare to the results obtained with a SPAD array. In chapter 5, we will also study the use of single pixel SPAD detectors for the specific application of locating and tracking objects hidden from view, and discuss their advantage and limitations with respect to using an array of SPADs.

Chapter 3 – A new tool for imaging light in motion: single-photon avalanche diode arrays

As mentioned in chapter 2, the main tool we will use in this work to capture light in motion is a single-photon avalanche diode (SPAD) array. SPAD detectors operated in time-correlated single-photon counting (TCSPC) mode to measure the detected photons' arrival time are a well-known technology used in different fields that provide very high temporal resolution (picoseconds) and single-photon sensitivity. They have generally been used as single-pixel detectors, and the recent development of arrays of these detectors, which include the time of arrival measurement on-chip, gives the opportunity to exploit this type of technology for applications in light-in-flight imaging.

In this chapter, we first give an overview of the technology behind a SPAD detector, and explain how it achieves its single-photon sensitivity and high temporal resolution. We then discuss the different applications of SPAD detectors and the reasons for the growing interest in developing arrays of them. In section 3.2, we then focus on the specific detector we are using throughout this work: the MF32 detector developed by Dr Robert Henderson and the CMOS Sensors and Systems group at University of Edinburgh, a 32x32 pixel array of SPAD operated in TCSPC mode. Finally, section 3.3 details the calibration and characterisation of the specific detector that we use, in particular the characterisation of its temporal resolution.

3.1 How to detect single photons with picosecond time resolution

Single-photon avalanche diodes combined with time-correlated single-photon counting (TCSPC) measurements have been used as single-pixel detectors for a few decades in various fields, such as in LIDAR and 3D imaging [34], fluorescence lifetime imaging (FLIM) and quantum measurements [35]. In this section, we describe the basic principles of detection of single photons with a SPAD detector and time of arrival measurements with a TCSPC module.

The basic component of the SPAD detector is a p-n junction, a juxtaposition of two oppositely doped semiconductors, one with excess holes, or absence of electron (p), the other with excess electrons (n). At the junction, these free charges recombine so that the junction is depleted of charges in the conduction band, in a region known as the depletion

region. Depending on the voltage applied to a p-n junction, one can obtain different types of devices. Here, we are interested in the case where the p-n junction is operated in reverse bias voltage: a positive voltage is applied on the n-doped semiconductor (cathode with reverse bias) and a negative voltage is applied on the p-doped semiconductor. The effect of this voltage is that the free charges in the n part of the junction are attracted to the cathode, and the free charges in the p part are attracted to the anode, therefore extending the depletion region at the junction. Figure 3.1a illustrates a cross-section of a typical SPAD design.

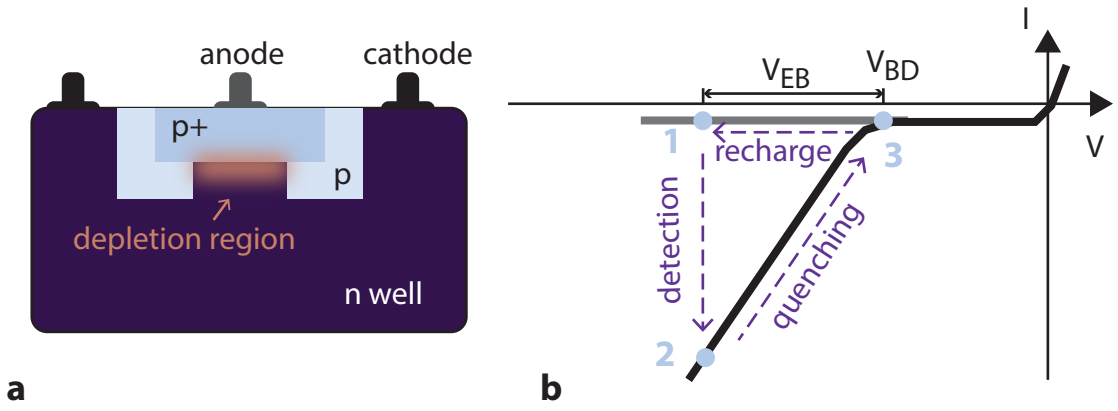

FIGURE 3.1: **Basic principle of a SPAD.** a) A SPAD consists of a p-doped semiconductor (p+ region) juxtaposed to a n-doped semiconductor (n-well), where a strong reverse bias voltage is applied through the anode and cathode. The two p-region represented act as a guard ring to prevent premature avalanche breakdown at the edges near the anode and the cathode. b) A current-voltage curve illustrates conceptually the behaviour of a SPAD detector. A voltage ($-[|V_{BD}| + |V_{EB}|]$) is first applied, creating a strong electric field across the junction (blue point 1). Upon detection of a photon, the generated avalanche creates a strong detectable current through the SPAD (blue point 2). In order to be able to perform a subsequent detection, the SPAD must be quenched: the voltage is brought back to the breakdown point (blue point 3) until no charges remain in the depletion region, at which point the SPAD is recharged back to point 1.

In figure 3.1b, the behaviour of the p-n junction used in a SPAD is illustrated in a graph showing the current in the junction as a function of voltage. When a strong reverse bias voltage is applied and there are no free charges in the depletion region, although there is a strong electric field across the junction, there is no current flowing in the device. The SPAD behaviour therefore lies on the grey line illustrated in figure 3.1b (point 1). When a photon is absorbed, it generates a free electron and hole in the depletion region. These free charges are then accelerated by the strong electric field across the junction. If the voltage is high enough, above a characteristic breakdown voltage V_{BD}, the free electron and hole will cause an avalanche to occur, by creating additional free charges by impact ionisation. A current will then pass through the junction and indicate the detection of a single photon by the SPAD (point 2 in figure 3.1b). To operate in a regime where a free charge creates a self-sustained avalanche, commonly known as Geiger mode, an

excess bias voltage V_{EB} is applied to the SPAD to bring the total applied voltage above the breakdown limit ($-[|V_{\text{BD}}| + |V_{\text{EB}}|]$).

Once a photon has been detected, the current in the SPAD needs to be stopped, so that the SPAD can detect a subsequent photon. To do so, the SPAD undergoes the process of quenching: the voltage is brought back to V_{BD} (point 3) for an amount of time (dead time) sufficient for all the charges to leave the depletion region. The simplest method to do so is through a resistor to lower the voltage, and is known as passive quenching. More complex active quenching techniques can also be implemented to decrease the dead time between subsequent detections. In the SPAD array used in this work, passive quenching is used and leads to a dead time of $\sim$ 100 ns. Once there is no charge in the depletion region, the voltage can be increased by V_{EB} once more, to recharge the SPAD back to point 1, where it is ready to detect another incoming photon. This process is what provides the very high sensitivity of a SPAD detector. The SPAD goes through this cycle for every photon detection [36–38].

The short signal in time generated by the detection of a photon in the SPAD junction can not only be used to indicate the detection of a single photon, it can also drive a counter to measure the time of arrival of the photon with precision. This type of measurement is known as time-correlated single-photon counting (TCSPC) [39]. There exist different types of counter, including analog ones, but the most common are time-to-digital converters (TDC). They are the basic building block of a TCSPC module. For the counter to perform a time-of-arrival measurement, the detection time of the photon needs to be clocked with respect to a trigger synchronised with the observed scene. Measurements with a SPAD are mostly done using an ultrafast laser, which emits a train of short pulses and can provide an electronic trigger synchronised to these pulses. A TDC then relies on the transit time of electronic signals through a series of logic gates to measure the time difference between the detected signal and the trigger [39]. In a linear configuration, such as the one presented in [40], the signal goes through a series of gates, which induce an increasing delay on the signal, and is compared to the trigger signal at the output of each gate. This allows to determine with precision the time delay between the signal and the trigger, and provide a time measurement with a few tens of picosecond precision. In a more advanced version, the linear series of gates can be looped into a ring oscillator, so that once the electronic pulse has travelled through the linear series of logic gates, in is sent back to the beginning of the linear series, and so on into a loop. In figure 3.2, we show a typical example of the main components of a TDC circuit, which is the specific architecture implemented in the MF32 array used throughout this work [11]. The SPAD feeds a sharp pulse to the TDC once a photon is detected, and an electronic trigger (TTL logic) synchronised with the laser source is also fed to the TDC. After going through some logic components, the signal from the SPAD circulates in the ring oscillator. Two components are then used to determine its exact delay with respect to the trigger. First,

a ripple counter counts the amount of complete rotations the signal completes in the ring oscillator with 7 bits precision. Second, a finer counter measures the delay between the signal and trigger within one loop of the ring oscillator (coder and fine memory). This provides a fine measurement stored with 3 bits precision. The output of the TDC is therefore a 10-bit time tag which corresponds to the number of gates, or time bins, by which the signal was delayed by before it coincided with the trigger. This time tag corresponds to the time-of-arrival measurement of the detected photon, and is typically performed with a resolution on the order of a few tens of picoseconds.

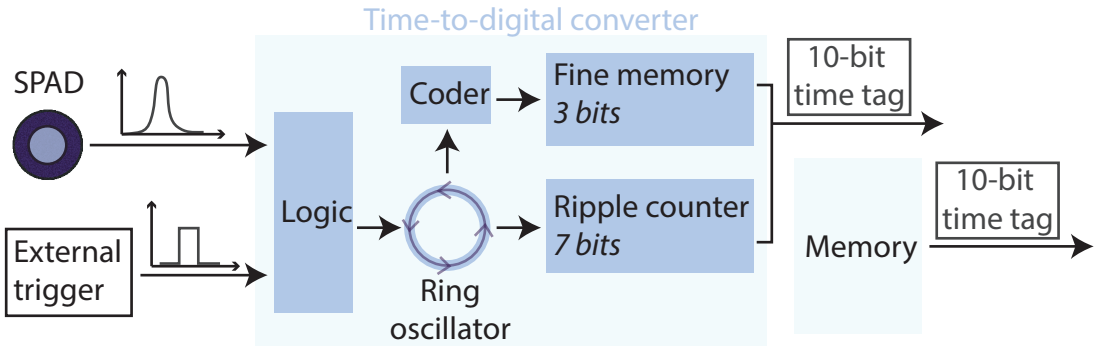

FIGURE 3.2: **Time-to-digital converter** The detection of a photon by the SPAD generates a sharp signal in time, which is sent to a time-to-digital converter that measures the delay between the signal and an external trigger. The output of the TDC is a 10-bit time-tag that corresponds to the time of arrival of the detected photon.

3.1.1 Applications and interest for arrays of SPADs

SPAD detectors used in combination with TCSPC modules provide the combined advantages of single-photon sensitivity and temporal resolution on the order of tens to hundreds of picoseconds. As mentioned above, single-pixel SPADs are commonly used in well-established imaging systems such LIDAR and 3D imaging, and fluorescence lifetime imaging.

In LIDAR measurements, a short pulse of light is sent to a target of which we want to acquire distance or depth information. If a target is at a distance d, a laser pulse will travel a distance d to reach it, will be reflected back and travel another distance d to reach the detector. A SPAD detector operated in TCSPC mode allows to measure the time of flight of this path t, and the distance can be recovered as $d = ct/2$, where c is the speed of light. The extreme sensitivity of SPADs allows one to reconstruct depth profile with centimetre resolution from targets located at kilometre distances, and the high temporal resolution of the TCSPC, to attain about 50 μm precision in close-range depth measurements (few meters stand-off distance) [41]. Since this detection system is based on a single-pixel detector and a single laser beam, one needs to scan the laser beam across the scene to recover a full 3D image of a scene.

In biomedical imaging, fluorescence lifetime imaging microscopy allows to not only image an object of interest, but also to distinguish different components by measuring their fluorescence lifetime [42]. To do so, one illuminates a single point in the region of interest at a time with a train of ultrafast laser pulses, and uses a SPAD detector and TCSPC module to measure the fluorescence emitted at that point [43]. Photons will be emitted and detected at different times across the fluorescence lifetime, and the temporal exponential decay can be measured and used as a contrast technique for imaging. Again, in this type of setup, one needs to scan the laser beam across the region of interest to recover a 2D or 3D image.

As can be seen by the short description of these two of the most common applications of SPAD and TCSPC, being limited to a single-pixel detector leads to the necessity of scanning a scene, as a 2D spatial scene and not only a single-point is usually of interest. There has therefore been a growing interest and multiplied efforts to miniaturise the detection system and implement a compact array of these detectors that could allow for the simultaneous measurements of a 2D scene. Acquiring multiple pixels at the same time would not only allow to simplify measurement setups by removing the need for scanning, it can also dramatically decrease the acquisition times to recover a 3D image or a fluorescence lifetime image of a scene.

3.2 SPAD arrays

One avenue that was taken to tackle the challenge of creating SPAD arrays was to design a detector that would combine the SPAD detector and the TCSPC module on a single chip in CMOS (complementary metal oxide semiconductor processes) architecture, to take advantage of the well-know design and fabrication techniques that are in place for CMOS components [37, 44]. Since the early 2000s, a few groups around the world have developed different types of detector arrays [37, 38]. While some detectors focused on the single photon detection for photon counting or indirect time-of-flight measurements [45–47], some have integrated time-to-digital converter (TDC) to their architecture to perform TCSPC measurements. Among them, in 2005, a group from Lausanne presented a 32x32 array of SPAD fabricated in CMOS, with a separate CMOS time-to-digital converter, and demonstrated 3D imaging of a human face at a range of 3 m [48, 49]. In 2012, a group from Milan have presented a building-block pixel that contains the TDC in-pixel [50], and later demonstrated a 32x32 pixel implementation [51]. Another successful implementation of single-photon avalanche diode and time-to-digital converter in a single pixel was realised by the CMOS Sensors and Systems group at University of Edinburgh, who created the MegaFrame detector in collaboration with STMicroelectronics, a 32x32 pixel array of SPAD detectors that have individual TCSPC modules directly on the pixel [11, 40, 52, 53]. The MF32 is the detector that is used throughout the work

presented here, as it is an all-in-one compact detector that allows single-photon detection with a resolution of about 100 ps. While different implementations of this detector have been presented, the structure of the one used in this work can be found in [11] and [53]. One of the goals of the MF32 was to simplify and accelerate FLIM measurements, and its application in this field has been demonstrated in [54]. Its usefulness is however not limited to FLIM applications and 3D imaging and, in this thesis, we detail how we exploit this powerful new detection tool to develop applications in different areas, namely in light-in-flight imaging and non-line-of-sight object detection and tracking.

In this section, we focus on describing the specific characteristics of the detection hardware and data acquisition of the MF32. A detailed representation of the architecture of the MF32 can be found in [11] and [53]. The work presented in this thesis was done in collaboration with the CMOS Sensors and Systems group, who provided the detectors we use in this work. In the last year, the MF32 has became available commercially through the newly-founded startup Photon Force [55]. Similar arrays operating at 1.05 μm and 1.55 μm are also commercialised by Princeton Lightwave (32x32 array with 250 ps resolution) [56], and silicon arrays operating in different modes (photon counting and indirect time of flight) are commercialised by Micro Photon Devices [57].

3.2.1 Hardware

The basic building block of the MF32 array is a smart pixel that contains a silicon SPAD detector and the TCSPC electronic circuit. The design and fabrication of the smart pixel in silicon allows for easy mass-production at reasonable costs, since silicon SPADs are a very common component in silicon CMOS devices. As explained in section 3.1, the detection of a single-photon occurs when a photon is absorbed in the active region of the SPAD. The efficiency of the SPAD, given by its photon detection probability (PDP), depends on the photon wavelength (silicon absorbs in the visible spectrum) and on the excess bias voltage, which affects the probability that a carrier will trigger an avalanche. The PDP of the MF32 is given in figure 3.3 as a function of wavelength and excess bias voltage. In the work presented throughout this thesis, the SPAD is operated at two main wavelengths: 532 nm, near the optimum detection efficiency, and close to 780 nm. The SPAD has an efficiency of around 23% at 532 nm and around 6.5% at 780 nm.

In addition to the SPAD detector itself, all of the TDC components illustrated in figure 3.2 are implemented on each of the 32x32 pixels of the MF32. The pixel architecture, shown in figure 3.4, therefore includes a detection area as well as the TCSPC electronics, so that each of the 1024 pixels of the MF32 can independently detect and time the arrival of single photons. The total size of the pixel is 50x50 μm^2. As can be seen from figure 3.4, the TCSPC components do take a significant amount of space on the pixel, and the active area for detection is limited to a diameter of 8 μm: the fill-factor of the pixel is

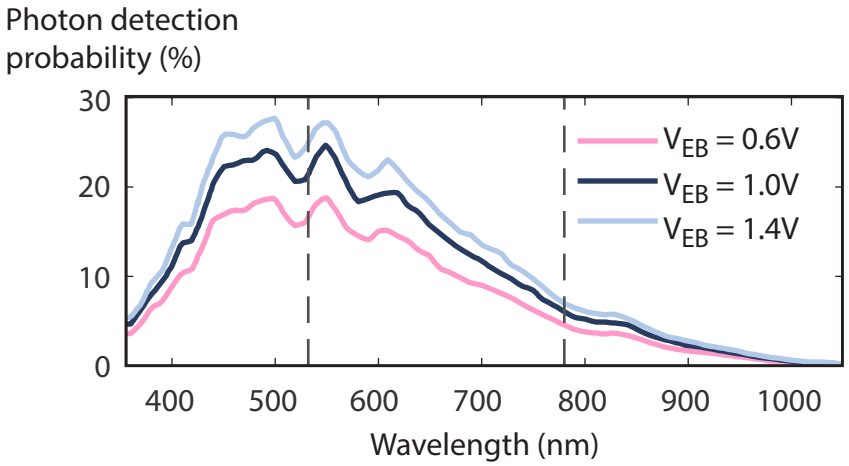

FIGURE 3.3: **Photon detection probability**. The detection efficiency of the MF32 silicon SPAD reaches around 25% in the green region of the visible spectrum, and tails down to a few percents in the near infrared. The detection probability also increases with excess bias, as a higher voltage leads to a higher chance that a free carrier in the depletion region will trigger an avalanche. In this work, the excess bias voltage used is 1.2 V. Data taken from [58].

therefore only 2%, which reduces the quantum efficiency of the detector by almost two orders of magnitude with respect to its photon detection probability. We will see that the MF32 nonetheless provides the necessary sensitivity to capture very low level intensity events, such as the propagation of laser pulses in air and signals scattered from around obstacles. However, there are ways to operate near the full photon detection probability of the SPAD in the MF32. A method that can currently be used is to incorporate an array of microlenses that focus the light reaching each pixel in the detector to its active area.

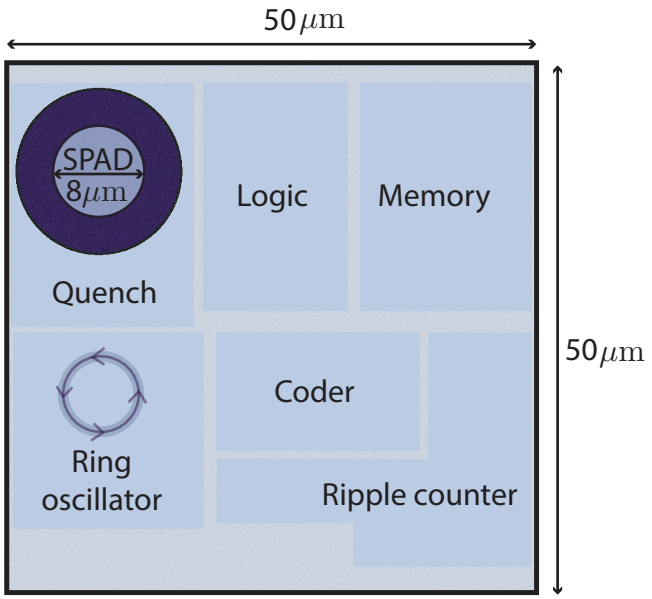

FIGURE 3.4: **The pixel architecture.** A SPAD detector and all detection and time-of-arrival measurement electronics are designed into each pixel of the MF32 array. The various components fit on a 50x50 μm^2 area. Figure adapted from [11].

It has been demonstrated that this technique can recover a factor 16 in efficiency, when applied to a different detector that had 3.14% fill-factor [59]. Another direction that is currently being investigated is to design and fabricate pixels that have a vertical stacking architecture, where the counting electronics is located on a layer under the SPAD rather than side-by-side.

Although every pixel's detection is independent of the others, all 1024 pixels are driven at the same acquisition and read-out rate. The camera is therefore operated at a given frame rate, and in one frame each pixel will either detect a single photon and provide the time-tag corresponding to its time of arrival, or not detect anything. It is important to underline that the frame rate is in no way dictated by the repetition rate of the laser used, which is the repetition rate of the external trigger. In figure 3.5, we show an example where the camera is operated at a frame rate that is lower that the laser/trigger repetition rate. Technically, the MF32 could be operated at 1 million frame per second (hence the name megaframe). In practice, the read-out of the data by the computer via a USB2.0 channel limits the frame rate to 3 kHZ. The laser used can have a wide range of repetition rate - in this work, we use lasers from 100 Hz to 80 MHz.

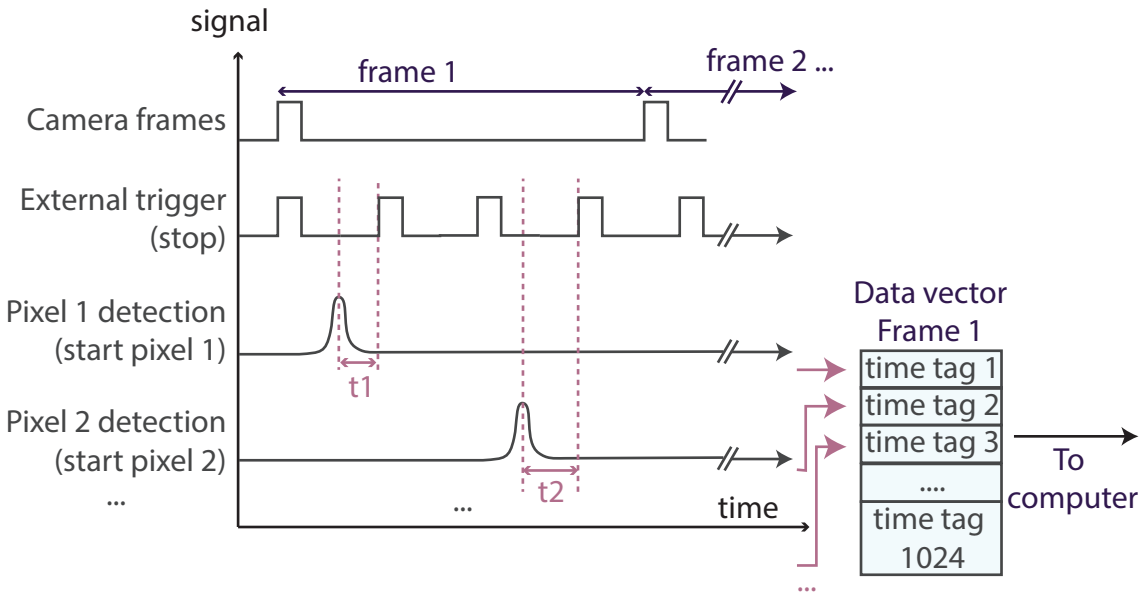

FIGURE 3.5: **Frame acquisition.** Each pixel is operated independently when it comes to detection and time-of-arrival measurement, but they are simultaneously activated for detection so that in every camera frame, only one photon is detected per pixel. Each pixel outputs a time-tag per frame, which is sent along with all other time-tags from that frame to a computer.

The SPADs on every pixel are then charged at the excess bias voltage and ready for detection at the beginning of a camera frame. If the repetition rate of the laser is higher than the camera frame rate, there is more than one pulse or laser cycle that gives the occasion for a photon detection. A photon might be detected in the first pulse sent during the frame, or it might take a few more cycles before a photon reaches the detector and is

absorbed. In figure 3.5, we show an example where pixel 1 detects a photon in the first pulse entering a time frame, at time t_1 from the trigger, and pixel 2 detects a photon at the third cycle within the frame, at a time t_2 from the corresponding trigger signal. As can be observed from the figure, the MF32 is operated in reverse start-stop mode, where the signal from the SPAD starts the clock, and the trigger is used to stop the clock. This allows to reduce false detection, and whenever no photon is detected during a camera frame at a given pixel, the time tag for this pixel will show zero. After one frame, all 1024 time tags are sent to a computer in the form of a 1024-long vector. Each time tag is a number between 0 and 1023 that corresponds to the measured delay between the photon arrival time and the trigger, in number of time bins. A time bin in the MF32 is around 50 ps, but varies from one detector to another and needs to be precisely calibrated as we will see in section 3.3.

3.2.2 Software

As detailed above, the output of the MF32 for a single frame is a 1024-long vector that contains the time tags for the photon detected at each of the 1024 pixels of the 32x32 pixel detector. As data is acquired by accumulating the signal over many frames, one can build up a 3D data set containing (x, y, t) information by defining a time histogram for each pixel, and populating it with counts as the number of frames acquired increases. As illustrated in figure 3.6, when the vector for frame 1 is read out by the computer, a single count is added to each of the histograms: if pixel 1 detected a photon at time t_1, the count in time bin t_1 of pixel 1's histogram is set to 1. Similarly, if pixel 2 detected a photon at time t_2, the time bin t_2 of pixel 2's histogram is set to 1, and so on, until all 1024 time tags have been accounted for and that each pixel have 1 count in their histogram. When the data from frame 2 is read, a count is added in the same fashion in each histogram, so that all pixels have an histogram with a total amount of counts equal to 2, and so on. A total acquisition is normally an accumulation of counts in the histograms over tens or hundreds of thousands of frames, and the total amount of counts in a pixel's histogram will always be equal to the amount of frames, including all the counts that will be set at time bin 0 when no photons are detected.

This build-up of the data leads to a final data set in the form of a 32x32x1024 three-dimensional matrix containing (x, y, t) information about an observed scene, with a temporal distribution for each position (x, y) of the detector. The temporal information is distributed over 10 bits or 1024 time bins. Events can therefore be observed over a window of time that depends on the size of the time bin of the detector: for a time bin of 50 ps, the total time window that can be observed is 50 ps/time bin x 1024 time bins $\approx$ 50 ns. If the laser used for the measurement has a pulse separation (1/repetition rate) that is shorter that 50 ns, the observation window will be set by the pulse separation. For example, when using a 67 MHz laser, as we do in the experiments reported in section 3.3,

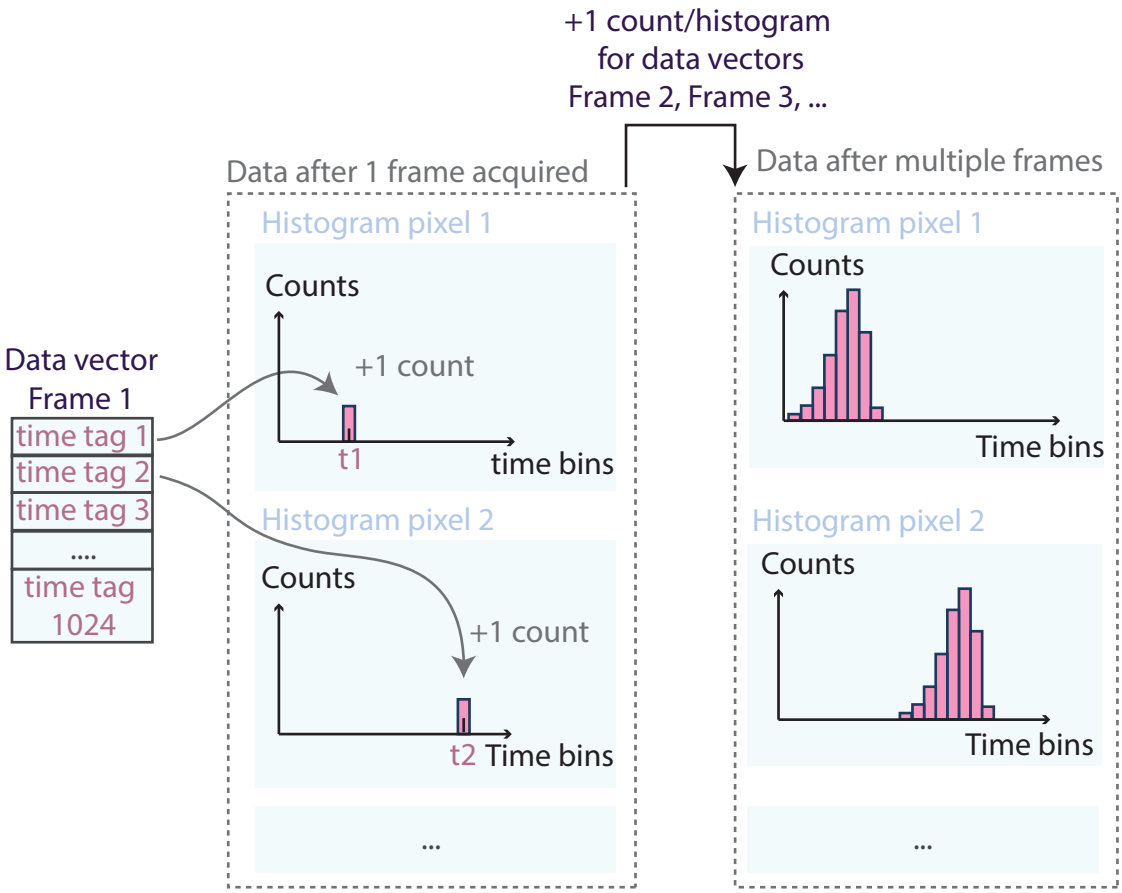

FIGURE 3.6: **Data acquisition.** At every frame, a 1024-long vector containing the time tags from each of the 32x32 pixels is read by the computer. All time-tags are placed as an additional count in the corresponding time bin in the pixels' histograms. After many frames, one obtain a distribution of counts for each pixel, hence obtaining a 3D dataset containing (x, y, t) information about the observed scene.

the pulse separation is 15 ns. The trigger pulses sent to the camera are therefore separated by 15 ns, and the maximum delay between an optical pulse (clock start) and a trigger pulse (clock stop) is 15 ns. The detected signal on the histograms will therefore span only the first 300 time bins rather than the full 1024 bins available. All remaining time bins will never be filled with any counts as no time tag will surpass 300 bins.

When a SPAD is operated in TCSPC mode, one wants to avoid the measurement of a photon at every single time a pulse is emitted: if the signal coming to the detector is too high and the chances of detecting it is near 100%, then the first photons arriving to the detector will always be the ones detected, leaving no chance for later-arriving photons to be detected. One therefore wants to operate the SPAD in a regime where less than 10% of the pulses (or frames) lead to a detection, to make sure the temporal statistics of the scene can properly be measured. This is called the photon-starved regime. To assure that a measurement is done in this regime, the power of the ultrafast laser used in the experiment has to be be adjusted. If the scene observed has very low signal intensity, the exposure time of a frame (corresponding to the camera frame rate in figure 3.5) can be set

to up to 1 ms to increase the probability of a photon detection in one frame, at the expanse of a slower data acquisition rate. The exposure of one frame on the MF32 can be set from 2 μs to 1 ms, so that if the signal coming from the scene is high, a low exposure time can allow one to acquire data faster with reduced noise in the histograms (see sections 3.3.4 and 3.3.5).

3.3 Characterisation of a SPAD array detector

The MF32 SPAD array is integrated on a small printed circuit board (see figure 3.7a). To use this detector as the sensing element of a camera, we integrated the circuit board into a custom-design metallic box that has an aperture at the position of the SPAD array itself, and a Nikon F-mount on which a standard camera objective can be attached. Figure 3.7 shows pictures of both the SPAD board and the camera implementation we use throughout the work presented here. This camera is compact, portable and easy to mount in any type of experiment. It requires only three connections: a 5 V power supply, the input of the external trigger used in the experiment (SMA cable), and the USB connection between the camera and the computer used to operate it. In addition to the TCSPC mode that has been described in the previous section, the camera can also be operated in photon counting mode, which gives simple 2D intensity frames as an output. This mode of operation is extremely useful for alignment of the camera in a setup, as it shows what is in the field of view of the camera in real-time.

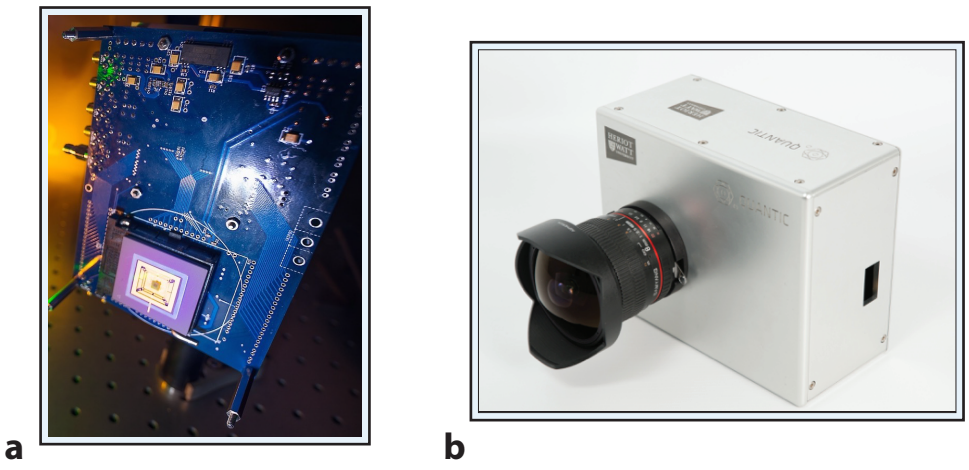

FIGURE 3.7: **Detector and camera.** a) The MF32 detector is a 1.5x1.5 mm detector placed onto a printed circuit board of 14x10 cm dimension. b) We fabricated a custom-designed box to hold the circuit board and easily mount a camera objective. This SPAD camera is easy to use in any experimental setup.

Before using this SPAD camera in any experiment, some of its parameters must be characterised - an essential one, as discussed above, is the temporal width of the time bins (temporal resolution of the ring oscillator). In the work presented in this thesis, we use

three different SPAD arrays that have the same design, but for which the specifications can vary. In this section, we present the complete characterisation of the main array we use (in all of chapter 5 experiments and in some of chapter 4 experiments), but any detector can be characterised in the same way.

To measure the bin width and other parameters of the SPAD array, we use the experimental setup shown in figure 3.8. The goal of this setup is to measure at each pixel of the SPAD array a single sharp peak in time. To do so, we send a laser beam directly onto the detector - in this implementation, no camera objective is placed in front of the sensor, and the beam is directly incident on the SPAD array. The laser we use here is a femtosecond oscillator (800 nm) sending 10 fs pulses of 10 nJ energy at a 67 MHz repetition rate. In figure 3.8, we see that a small part of the laser (8% reflection) is reflected to an optical constant fraction discriminator (Becker & Hickl OCF-401), which transforms an optical signal into a TTL electronic trigger signal, that is sent to the SPAD array for synchronisation of the acquisition. The laser then passes through a series of ND filters to decrease its intensity to reach the photon-starved regime of detection. Filters are placed to amount to a neutral density filter of optical density $OD = 9$; after passing through, each pulse carries only around 100 photons. The pulses are then sent to two mirrors that reflect the light directly on the detector – these two mirrors can be shifted to change the optical path of the beam, for reasons we will discuss in section 3.3.1.

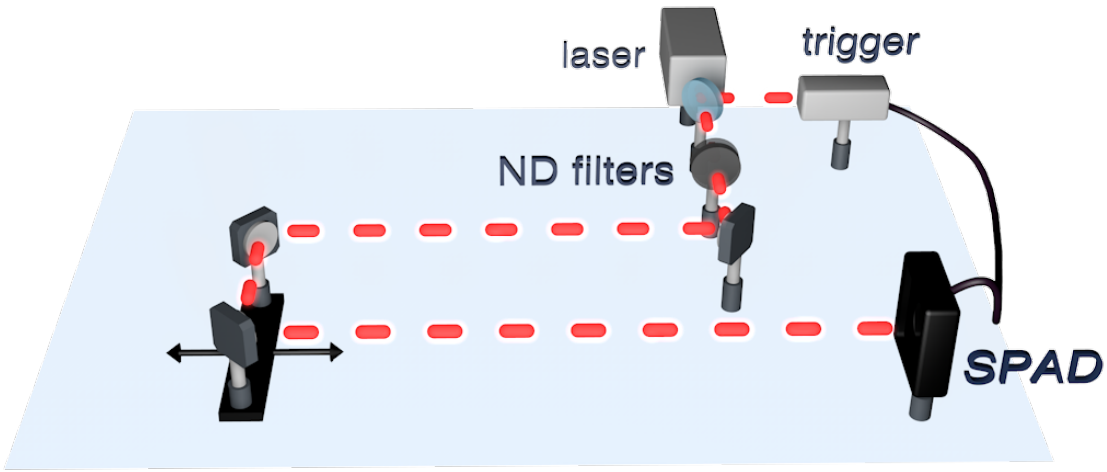

FIGURE 3.8: **Characterisation setup.** To measure parameters of the SPAD array such as the histograms time-bin width and the impulse response of the detector, we send ultrashort pulses directly onto the detector to measure a single sharp temporal peak on each of the array's pixels. Since the beam size is larger than the size of the detector array (1.5x1.5 mm) by 5-10 times, all pixels are illuminated.

We expect from this setup to measure a single peak in the temporal histogram of each pixel of the 32x32 SPAD array. This will be used to characterise the bin width of the

histograms, the impulse response of the detector and the uniformity of the array, which we detail in the following subsections.

3.3.1 Time bins

The width of every time bin in the histograms corresponds to the temporal resolution of the ring oscillator in the time-to-digital converter (see section 3.2.1). It depends on the design of the TDC and is specified to be around 50 ps for the MF32, but can vary slightly from one fabricated detector to the next. It is therefore important to measure precisely the bin width of any SPAD array used in experiments.

To do so, we use the setup illustrated in figure 3.8, where short pulses are sent directly at the detector, which measures their arrival time. This time of arrival is a measure of the delay between the arrival of the optical pulse and the arrival of the electronic trigger: it therefore depends on both the path taken by the trigger (length of the cable) and the optical path taken by the beam before reaching the detector. By changing the optical path by a known amount, we can change the measured arrival time by a controlled amount: to do so, we shift the two mirrors set on a rail from a position far from the detector (which we set to be the 0 cm position) to a position closer to the detector by 62.5 cm. The shift in the mirrors position Δd is related to the change in arrival time Δt:

$$\Delta t = 2\frac{\Delta d}{c}, \tag{3.1}$$

where c is the speed of light (299 792 458 m/s) and the factor 2 takes into account the return trip of the laser beam. In figure 3.9a, we show an example of four measurements taken at different mirror positions, taken from a single pixel of the SPAD array. At every position, a sharp pulse corresponding to the arrival of the 10 fs pulses is detected. We can observe from figure 3.9a that the pulse arriving from position 0 cm, which is the farthest from the detector, is measured at an earlier time bin than the pulse arriving from position 62.5 cm, which is closer to the camera. This is due to the fact that the time bin measurement is a clocked measurement that starts with the detection of a pulse and is stopped by the external trigger, as we have illustrated in figure 3.5: this mode of operation is called reverse start-stop, and it means that the temporal axis is reversed with respect to the temporal order of actual events. In experiments, this axis needs to be reversed to represent a forward direction in time, as we will see in chapter 4.

To measure the bin width from the measurements shown in figure 3.9a, we find the position of the peak signal in each measured histograms. A total of 13 mirror positions corresponding to 13 different optical path delays Δd are used, and the position of the peak signal is plotted against the expected time delay Δt in figure 3.9b. A linear fit of this data gives a measurement of the time bin for the given pixel. In the example shown in figure 3.9b, we find a bin width of 45.8 ps.

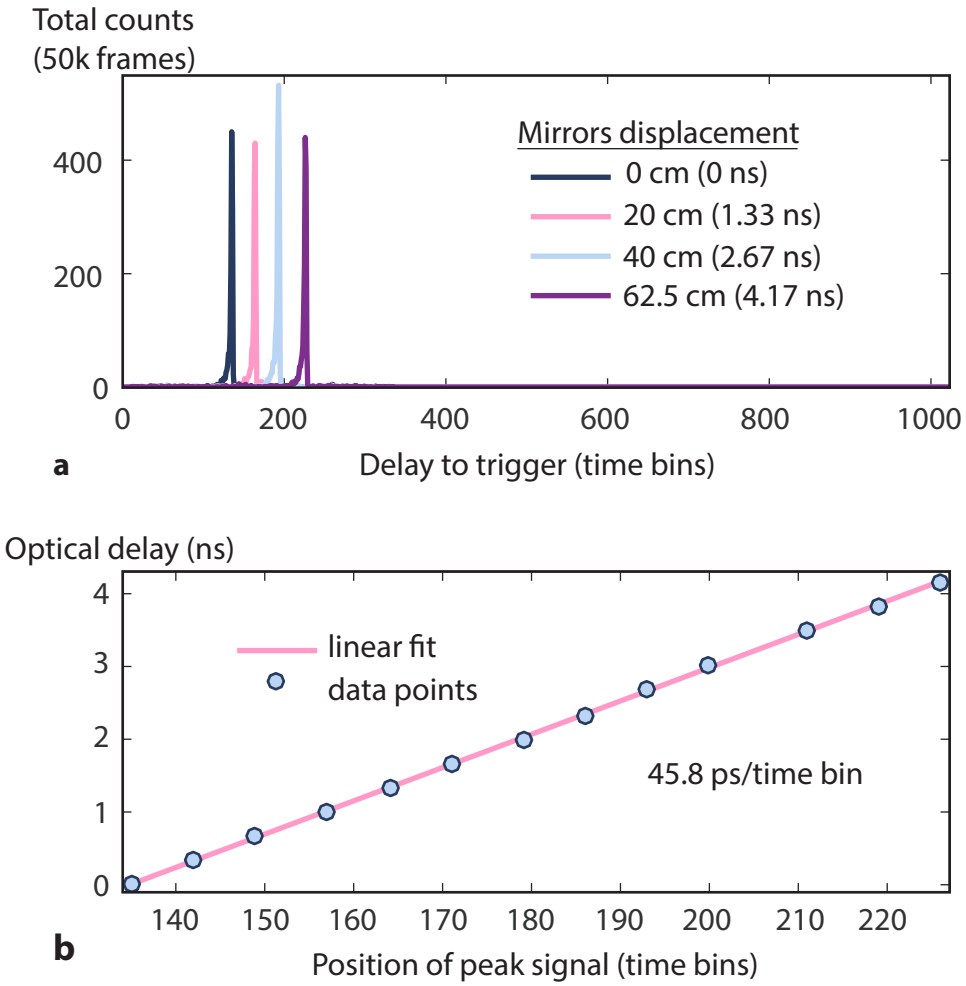

FIGURE 3.9: **Time bin measurement.** a) Example for a single pixel of histograms acquired for various optical path (varying mirror positions). The histograms are integrated over 50 000 frames. The detected peak is shifted as the optical path of the beam is changed in the experimental setup. b) Finding the peak position for every optical path allows us to determine the bin width of the histogram for a given pixel. The errors on the optical delay and peak measurement are smaller than the data points.

This measurement can then be done for every pixel of the 32x32 array, and we find that the temporal histograms' bin widths vary slightly from one pixel to the next. In experiments where a very high level of spatial and temporal accuracy is needed, a different bin width could therefore be used for every pixel. In fact, for even higher accuracy, one can calibrate how the bin width varies throughout the histogram, as some slight changes are expected. However, in the work presented here, there is no need for that level of precision, and using a single bin width for all the SPAD array gives very satisfactory results. To determine the global bin width, we therefore studied the obtained bin width for each pixel as a bar histogram, as illustrated in figure 3.10a. We see that the distribution is centred at 45.5 ps, which is also the mode of the bin width distribution over all pixels. In this thesis, we therefore use 45.5 ps as the bin width of this SPAD array.

In figure 3.10b, we also show the distribution of the bin widths on the detector. We can observe that the duration of the time bins varies from one pixel to another by about

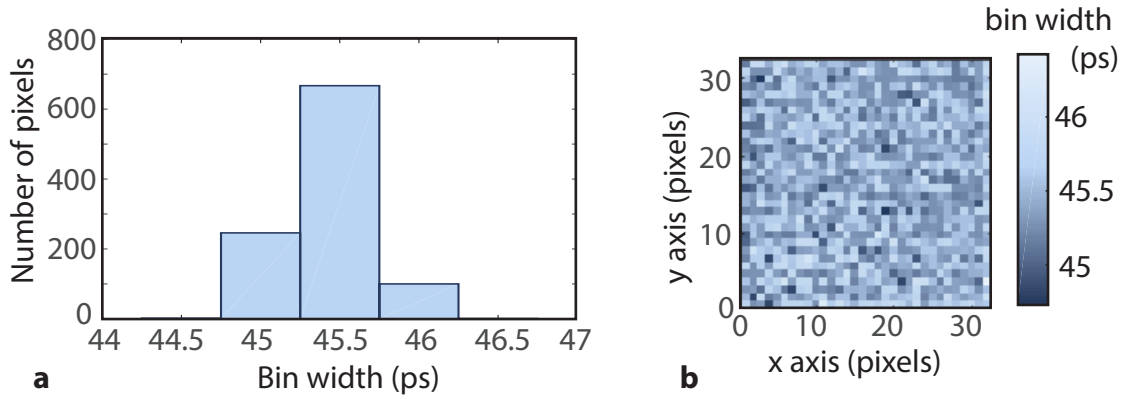

FIGURE 3.10: **Time bin distribution.** a) The time bin width measured on every pixels varies by about 1 ps, and the distribution is centered around 45.5 ps with a standard deviation of 0.3 ps. b) The spatial distribution of the bin widths of the 32x32 pixel array shows no structured pattern.

1 ps, and it does so in a random way (there is no structured pattern in the distribution of bin widths). We note here that this experimental characterisation allowed us to do a linear regression to find the bin width of 1016 of the 1024 pixels with a coefficient of determination R^2>0.9997. There were 8 pixels for which the noise level was very high (commonly referred to as hot pixels) and the peak measurements sometimes found a noise peak rather than the signal from the laser pulses – the result from these pixels are not included in the bar graph shown in figure 3.10a and a proper measurement of their bin width could be done by a more careful analysis of the data, such as windowing the histograms to find the proper peak of the signal or subtracting a background acquired without the signal. We use the latter technique in light-in-flight experiments as the first step of data processing to suppress the noise in the acquired data, as we will see in chapter 4.

3.3.2 Impulse response function

The characterisation measurements are performed, as mentioned previously, with a femtosecond oscillator sending pulses of 10 fs duration. This duration is much shorter than the width of the detector's time bins, so that the signal to be measured can be considered as a Dirac function. However, as can be observed in the histograms shown in figure 3.9a, the signal measured has a definite width, and is not simply a peak located at a single time bin. The shape of this signal is known as the impulse response function (IRF) of the detector: the function that will be obtained when measuring a temporal Dirac function. The resolution of the SPAD array, or its ability to distinguish adjacent peaks, is given by the width of this IRF, and not by the width of the time bins - the IRF is therefore a very important characteristic of the SPAD array. To determine to IRF of our detector, we can simply use one of the data sets acquired with the setup of figure 3.8. We show in figure 3.11 a zoomed-in section of one histogram where the shape of the IRF is clearly visible. From such an histogram, the full width at half-maximum (FWHM) of the response is

determined, which gives us a measure of the temporal resolution of the detector. In the example shown in figure 3.11, we measure a FWHM of 2.4 bins, or 109 ps.

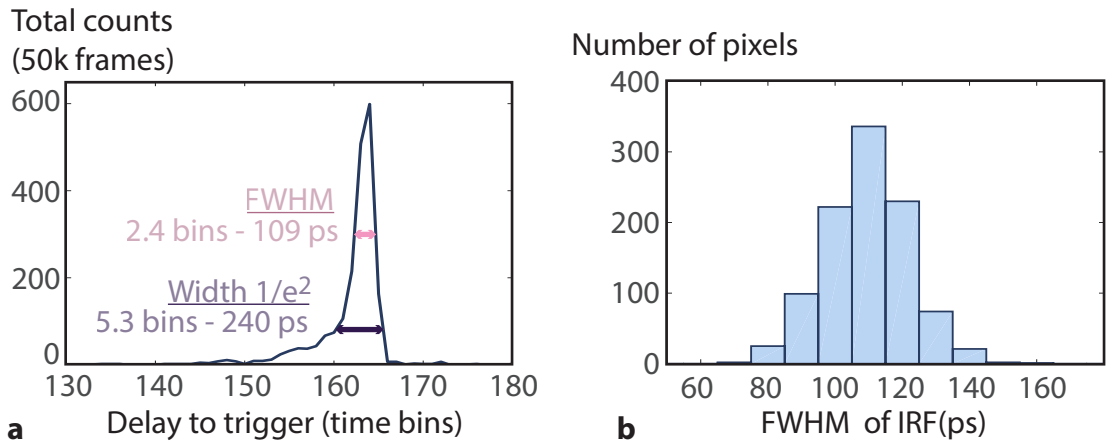

FIGURE 3.11: **Impulse response function.** a) We record the impulse response function of each pixel of the 32x32 array and measure its full width at half maximum (FWHM). b) The FWHM varies from one pixel to the next, as shown in this distribution of measured FWHM over all 1024 pixels of the array. The IRF distribution over the SPAD array is centred at 110 ps and has a standard deviation of 12 ps.

We can also observe that the IRF is not symmetric: it is very sharp on the right side (earlier times) and tails down slowly on the left side (later times). The width of the response is due to different phenomena, including jitter in the different electronic components of the detector. The main phenomenon that contributes to this widening of the response is the detection process in the SPAD junction. As we have explained in section 3.2.1, a photon is detected by a SPAD when it is absorbed in the p-n junction. Due to the high reverse bias voltage, the free carrier created by this absorption will trigger an avalanche, which in turn creates the sharp signal that is used in the TCSPC detection. This absorption of the photon can arrive at any position in the junction, and the free carrier can either immediately trigger an avalanche, or it can diffuse inside the junction before it does so. This diffusion process takes time, and contributes to widening the uncertainty of the photon's time of arrival. The impulse response is therefore characterised by a superimposition of a Gaussian response corresponding to the immediate generation of an avalanche in the p-n junction, and an exponential tail arising from diffusion of carriers in the junction [37].

As it was the case for the time bin measurement, the FWHM of the IRF differ slightly from pixel to pixel. In figure 3.11b, we show the distribution of the IRF's width over all 1024 pixels of the SPAD detector. The distribution is centred around 110 ps, value that we useto describe the detector's resolution as a whole in the work presented in this thesis. In specific experiments that require additional precision, the individual response function of each pixel can be used.

3.3.3 Uniformity of the array

We have seen in the previous two subsections that the time bin width and the impulse response function are two parameters that can vary slightly from pixel to pixel. There are two other important parameters for which the uniformity should be considered. One of them is the sensitivity - when illuminating the detector with a uniform intensity, the measured intensity can vary from pixel to pixel. We did not perform this calibration in the context of this thesis, as the work presented in the next chapters do not require a careful calibration of the pixel-to-pixel variation in the sensitivity. However, one very important parameter for us to consider in the context of our work is the uniformity in the time-of-arrival measurement. In the setup that we use, a beam is sent directly to the detector in a trajectory perfectly perpendicular to the plane of the sensor - we therefore expect the signal from the beam to be detected at the same time of arrival for every pixel of the array. However, we observe some significant shifts in the time of arrival of the laser pulses from one pixel to the next, as illustrated in figure 3.12a. In this figure, we show the histograms measured at 3 different pixels - since the illumination is uniform in time, we expect the laser beam to arrive and be detected at the same time at each pixel, but there is a difference in the peak position of the detected signal of several time bins. The difference in the peak position of the signal over the 32x32 pixels, as measured with our characterisation setup, is not random over the array as it was the case for the slight variations in time bins (figure 3.10b): we show its structure in figure 3.12b, where we assign the average peak position to be the reference position (0 time bins), and we show the shift in number of time bins for the histograms acquired by each pixel, with respect to this reference. Even though we expect all the pixels to detect a peak at the same position, there is a clear shift from top to bottom across the array of about 10 time bins (455 ps). This shift has been measured in

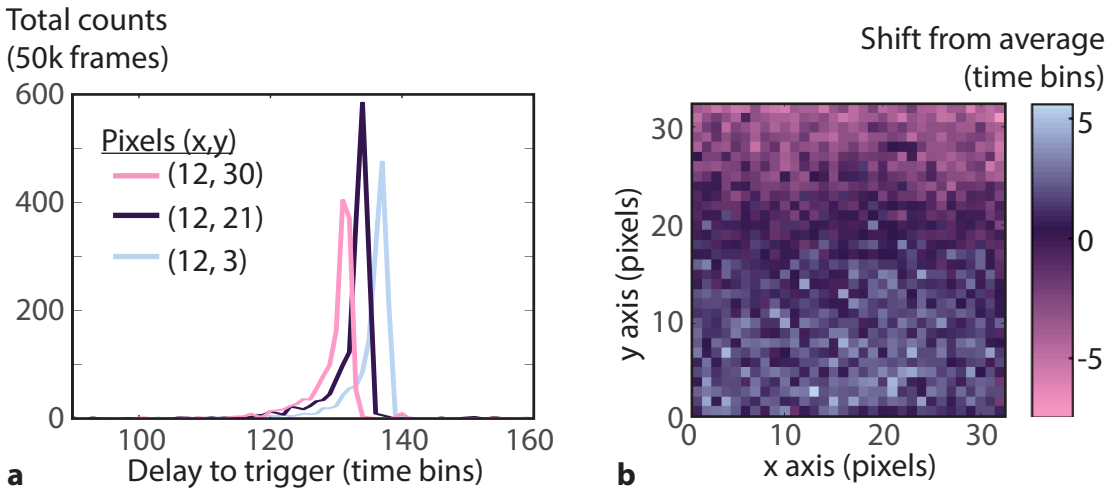

FIGURE 3.12: **Time of arrival uniformity.** a) An example of the response of 3 different pixels shows a significant difference between the time of arrivals of the detected signal at different pixels. b) Deviation of peak position of each pixel from the average peak position when the array is illuminated with a flat wavefront. Rather than obtaining a flat response in time, we observe a shift in the time of arrivals detected from the pixel to pixel in a clear pattern.

different configurations (orientation of the array, laser power, exposure time, etc.) and is independent of all these parameters. It is therefore straightforward to correct for this non-uniformity of the time-of-arrival detection: the difference over all pixels (figure 3.12b) is registered and in subsequent acquisitions, one can simply shift the histogram of each pixel by the correct amount to retrieve a temporal information which is consistent across all pixels. Alternatively, if this map is not known for a detector, a simple linear regression of an observed phenomenon evolving at a known fixed speed can provide a satisfactory correction to this shift in the histograms.

3.3.4 Signal-to-noise ratio and dynamic range

In section 3.1, we have seen how the absorption of a photon by the SPAD leads to an avalanche in the p-n junction, that results into a detection. However, an avalanche can also occur without being caused by an absorbed photon. Indeed, free carriers can be randomly generated, cause an avalanche, and being registered as a false detection, or noise. This is what causes the dark noise of the detector, which, as the photon detection probability (PDP), increases with the applied voltage on the p-n junction. As this dark noise is a thermal effect, it can also be decreased by cooling the detector. In the case of the MF32, the dark noise at excess bias voltage $V_{\text{EB}} = 1.2$ V is around 50 counts per second at ambient temperature, which is low enough for most applications so that the detector does not need to be cooled. In addition to the dark noise of the detector, ambient light can also lead to noise in the measured histograms. In this section, we therefore discuss the signal-to-noise ratio achievable in our measurements as a function of the noise level on the SPAD and other acquisition parameters.

As we discussed in section 3.2, a measurement performed with the SPAD array is an integration of the signal over multiple camera frames, typically between thousands and millions of frames. Logically, as the number of frames of the data acquisition increases, the integrated detected signal also increases. The measured histogram will also display some background level of noise, due to both the dark noise of the detector (free carriers thermally generated in the p-n junction) and to the experimental conditions (ambient or diffuse light). In figure 3.13a, we show an example of an histogram acquired over 1 million frames, where the inset zooms in on the background noise. To study the behaviour of both the acquired signal and noise as a function of number of integration frames, we repeat this measurement for varying number of frames between 10^4 and 10^6. In figure 3.13b, we show that the maximum count of the signal increases linearly as a function of the number of frames. However, we also observe that the background level also increases linearly as a function of number of frames, as shown in figure 3.13c. The ratio between the maximum recorded count in the signal and the background noise level is therefore constant as a function of number of acquisition frames (figure 3.13d). This result seems to indicate that there would therefore be no gain in integrating a signal over a larger num-

ber of frames, as it does not lead to an increased ratio between the signal peak and the noise level.

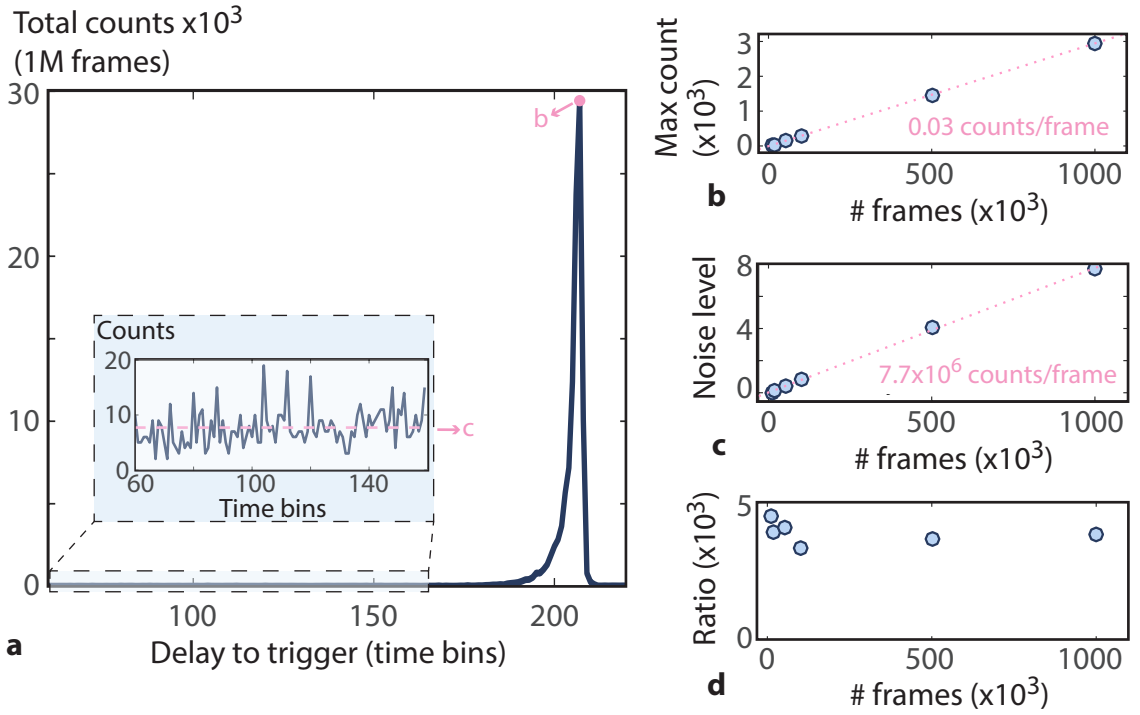

FIGURE 3.13: **Signal and noise level.** a) From a measurement of a single peak from a femtosecond laser, we can measure the peak signal (b) as well as the noise level (c). The inset shows a zoom over the noise level in a section of the histogram containing no signal from the laser. b) As we integrate over a larger number of frames, the peak of the detected laser signal increases linearly. c) The noise level increases in the same manner as a function of number of frames, so that the ratio between the peak signal and the noise level (d) does not vary with number of frames.

However, there is more to consider about the signal-to-noise ratio of a signal measured with a SPAD detector. We show in figure 3.14 that increasing the number of acquisition frames leads to a better defined signal. Indeed, both the signal and noise in an histogram acquired over one million frames show less variations in their normalised intensity than a signal acquired over ten thousand frames. This is due to the fact that the signal-to-noise ratio of a SPAD measurement is in fact characterised by the Poissonian statistics of the acquired signal and of the noise sources. Indeed, the detection of single photons emitted by a coherent laser source is characterised by its shot noise, or variation in discrete number of photons. The variance of a Poisson distribution being given by its mean value, a number of photons N_p can be measured with a SNR of $N_\mathrm{p}/\sqrt{N_\mathrm{p}} = \sqrt{N_\mathrm{p}}$, the SNR increasing as we increase the number of detected photons. In addition, a detector will also have dark noise that follows a Poisson statistics - for a number of counts N_n arising from the noise, the noise will have a standard deviation of $\sqrt{N_\mathrm{n}}$. The SNR of a measurement therefore depends on the fluctuations in both the signal and the noise [38,60]. In the case of multiple sources of noise that are independent and follow a Poisson or normal distribution, the SNR is expressed as

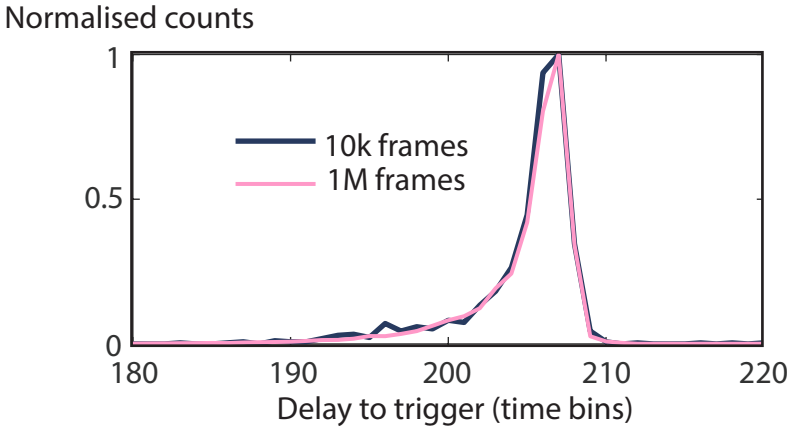

FIGURE 3.14: **Increasing the number of acquisition frames.** When increasing the number of acquisition frames, the Poisson noise on the counts at each bin of each pixel becomes less significant, leading to better defined measurement and increased signal-to-noise ratio.

$$\text{SNR} = \frac{N_\text{p}}{\sqrt{N_\text{p} + N_\text{n}}}. \tag{3.2}$$

In the case of the SPAD array, we can consider the SNR in an individual pixel's histogram. The number of photons in the signal N_p will be equal to the number of acquisition frames N_f multiplied by the rate of photon detection per frame in the signal n_p, $N_\text{p} = n_\text{p} N_\text{f}$. As discussed above, given the physics of detection of a SPAD, the maximum detection rate is n_p = 1 photon/frame, and it is preferable to keep this rate at about 10% ($n_\text{p} = 0.1$) to operate in the photon-starved regime.

The amount of noise counts registered can also be expressed in a similar fashion, with $N_\text{n} = n_\text{n} N_\text{f}$. Here, n_n represents the rate of noisy counts registered per frame. Our MF32 detector is specified to have a median dark count rate of 50 counts/second. The contribution of the dark counts to the noise rate n_n depends on two main factors: the exposure time of a camera frame and the number of time bins over which the measured histogram is distributed. For example, the exposure time in the measurement shown in figure 3.13 is 10 μs. There will therefore be $50 \frac{\text{counts}}{\text{second}} \times 10 \frac{\mu s}{\text{frame}} = 5 \times 10^{-4} \frac{\text{counts}}{\text{frame}}$ arising from the noise. This number of counts is also randomly distributed over the number of time bins contained in the histograms. Dividing the number of counts/frame by the number of bins will therefore give a value of the expected background noise level. Depending on the repetition rate of the laser used, the number of bins will either be 1024 or lower (see section 3.2.2). In the present case, the signal is distributed over 300 bins, so that the expected noise rate per frame per time bin is 1.7×10^{-6} counts per frame per bin. In figure 3.13c, we measure a rate of 7.7×10^{-6} counts per frame per time bin for the noise level. This higher value can be due both to the fact that the specified 50 counts/second may vary from pixel to pixel and from detector to detector, and to the fact that there is additional noise coming from ambient light.

From the expression of N_p and N_n detailed above, we can re-express equation 3.2 as

$$\text{SNR} = \frac{n_p N_f}{\sqrt{n_p N_f + n_n N_f}}, \tag{3.3}$$

which confirms that there is a benefit in acquiring signals over higher number of frames as it leads to an increased SNR following $\sqrt{N_f}$.

Depending on the noise rate n_n and the number of acquired frame N_f, there will be a minimum detectable signal rate $n_{p,min}$ corresponding to a SNR of 1. This minimum signal can be found by solving equation 3.3 for SNR = 1. When the contributions of the shot noise are smaller than the contribution of the dark and ambient noise, which is often the case, we can simplify the solution to:

$$n_{p,min} = \sqrt{\frac{n_n}{N_f}}. \tag{3.4}$$

Since the minimum detectable signal decreases with $\sqrt{N_f}$ and the maximum recordable signal $n_{p,max}$ is limited to 1, the dynamic range of the detector therefore increases with $\sqrt{N_f}$:

$$\text{Dynamic range} = \frac{n_{p,max}}{n_{p,min}} = n_{p,max}\sqrt{\frac{N_f}{n_n}} \tag{3.5}$$

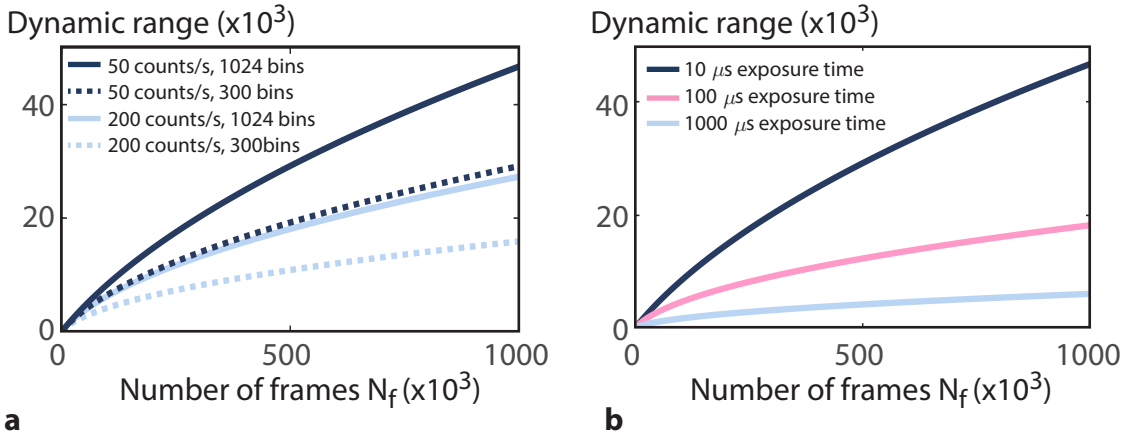

FIGURE 3.15: **Dynamic range.** The dynamic range of detectable signals is calculated with maximum recordable signal $n_{p,max} = 0.1$ and a minimum detectable signal $n_{p,min}$ given by equation 3.3 for a SNR of 1. a) The dynamic range is higher for lower noise rate (here shown for 50 counts/second and 200 counts/second), and when the histogram is distributed over the full 10 bits rather than on a limited number of time bins (here shown for 1024 and 300 bins). The exposure time for this calculation was set to 10 μs, as we used in our characterisation experiments. b) The dynamic range decreases as we increase the exposure time because of the increased noise rate (in counts/frame). Here we calculate the dynamic range for a dark noise rate of 50 counts per second distributed over all 1024 time bins, for an exposure of 10 μs, 100 μs and 1000 μs.

Since the number of counts contained in the signal arriving at a rate of n_{p} are distributed over about 5 bins (width at 1/e^2 of the IRF - see figure 3.11), the noise rate n_{n} considered for the calculation of the dynamic range should also be the noise rate per frame per time bin multiplied by 5 time bins. In figure 3.15, we show the expected dynamic range of our detector as a function of number of acquired frames, calculated with equations 3.3 and 3.5 for different parameters of acquisition. In figure 3.15a, we show the expected dynamic range for a 10 μs exposure time on each frame, for noise rates of 200 counts/second and 50 counts/second, distributed over either 300 or 1024 bins. The dynamic range shows an increase following $\sqrt{N_{\mathrm{f}}}$ and is of the order of 10^4 after a few hundred thousand frames acquisition. In figure 3.15b, we show how the dynamic range decreases as we increase the exposure time of a frame, due to the increased noise rate coming from a longer exposure time. For a 1000 μs exposure time, which we will use in chapter 4 for light-in-flight measurements, and a dark count rate of 50 counts/second distributed over 1024 time bins, the dynamic range of recordable signals is around 5000 after several hundred thousands of frames, which means we can detect signal rates from 0.1 photons/frame to 2×10^{-5} photons/frame.

3.3.5 Data acquisition rate

The exposure time of the camera frames can be changed from 2 μs up to 1 ms, which will influence both the available dynamic range and the signal detection rate. However, as mentioned in section 3.2.2, the read-out rate of these frames cannot follow the detection rate up to a certain limit. In figure 3.16, we show a measurement where we look at the acquisition time of 2400 frames for varying exposure times. We expect the acquisition time to vary linearly with exposure (2400$\times$exposure time). We observe that the measurement follows the expected value for exposure times down to 300 μs. For shorter exposure times, the acquisition time is then limited by the read-out rate and reaches a plateau. This 3 kHz limit in the acquisition rate is set by the limited speed of the USB2.0 connection –

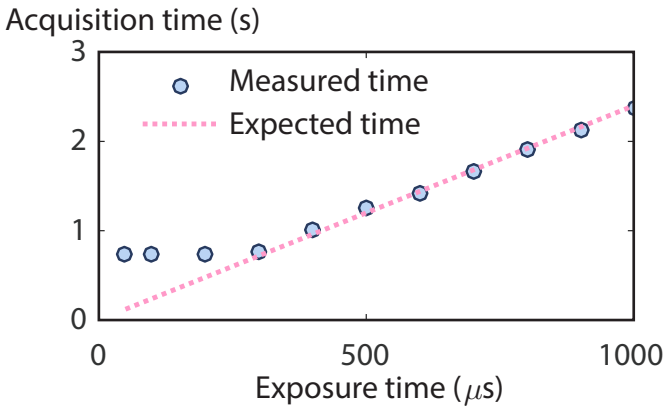

FIGURE 3.16: **Data acquisition rate.** We measure the acquisition time over 2400 frames for varying exposure times. The linear correlation stops at exposure time shorter than 300 μs, and the read-out rate is thereby limited to 3 kHz.

a future version of the MF32 will include a USB3.0 connection and will therefore allow faster acquisition rates [61].

This last section detailed the main characteristics of the MF32 SPAD array that we use as the main tool for the work presented in the next two chapters. Now that this novel detector has been explained and characterised, we have all the background information necessary to grasp the type of measurements that we will show in the following chapters as well as the intricacies of the important experimental parameters. The next chapter will demonstrate how this novel camera can be exploited to perform light-in-flight measurements at light level lower than ever before, keeping a temporal resolution of the order of 100 ps.

Chapter 4 – Imaging light in motion in action: light-in-flight experiments

We have discussed in the last chapters how a detector with picosecond timing resolution is needed to capture light in motion, and have proposed a novel solution to capture videos of light in motion at the single-photon level. The SPAD camera described in chapter 3 allows us to perform light-in-flight measurements without perturbing the dynamics we want to observe – previous techniques relied on a scattering medium or surface to enhance the signal deflected to their detector, but here, we rely uniquely on the scattering already present in the phenomenon we want to study. In this chapter, we will present a series of experiments that pushed the limits of the state-of-the-art in light-in-flight imaging. We will first demonstrate that we can record the motion of laser pulses directly in air, without any additional medium to enhance scattering, by detecting only the photons scattered by air molecules in the path of the laser. We will then present the study of a few more complex phenomena that can directly be observed and studied using the SPAD camera: laser-induced plasma, pulse propagation in optical fibres and slow light in a rubidium cell. We will finally discuss how our novel technology could be applied to observe the propagation of light in highly scattering medium and retrieve information on the medium.

4.1 Laser propagation in air

Although impressive visualisations of light propagation have been achieved in the past, we underlined in chapter 2 that they all relied on scattering from either a scattering surface, such as in light-in-flight holography [18], or a scattering medium such as a milky solution [25] (see figure 2.1). This surface or medium is typically needed to deflect enough light from the path observed to the detector. Here, we investigate the possibility of observing laser pulses propagating in air, without interaction with an additional scattering medium. As light propagates in air, it undergoes some elastic Rayleigh scattering events when it encounters air molecules (mainly N_2 molecules). The scattering cross-section of N_2 is 5.1×10^{-31} m^2 [62]; for a concentration of 2×10^{25} mol/m^3 in ambient conditions, this means that the fraction of light being scattering over 1 centimetre of propagation is about 10^{-7}. In a typical ultrafast laser, such as the one used in chapter 3 for calibration purposes, a single pulse is composed of 10^{10} photons. Therefore, only about a thousand photons are deviated from the path while a laser pulse propagates in 1 cm of air - these photons are scattered in every direction, so that very few will reach a detector imaging the pulse propagation. This explains why recording light propagation in air is extremely challenging, and why our proposed method relying on a SPAD camera is a great candidate to tackle this challenge.

To record the propagation of laser pulses in air, we use the experimental setup illustrated in figure 4.1. We use a 532 nm wavelength microchip laser (Teem Photonics STG 03E) emitting 500 ps long pulses of 3.5 μJ energy at 4 kHz (14 mW average power). Each pulse is therefore composed of 9×10^{12} photons. As it was the case in the setup used in chapter 3, a fraction of the beam (8% reflection) is deflected onto an optical constant fraction discriminator (Becker $\&$ Hickl OCF-401) to produce a trigger signal synchronised with the laser pulse emission. The pulses are then directed to propagate across the field of view of the SPAD camera, and are reflected by mirrors so that the pulses cross the field of view 3 times. As discussed above, light from the laser pulses is scattered from air molecules and is detected by the SPAD camera with picoseconds resolution thanks to the SPAD single-photon sensitivity and time-correlated single-photon counting (TCSPC) mode of operation where each photon's time of arrival is clocked and stored in a time histogram (see section 3.2).

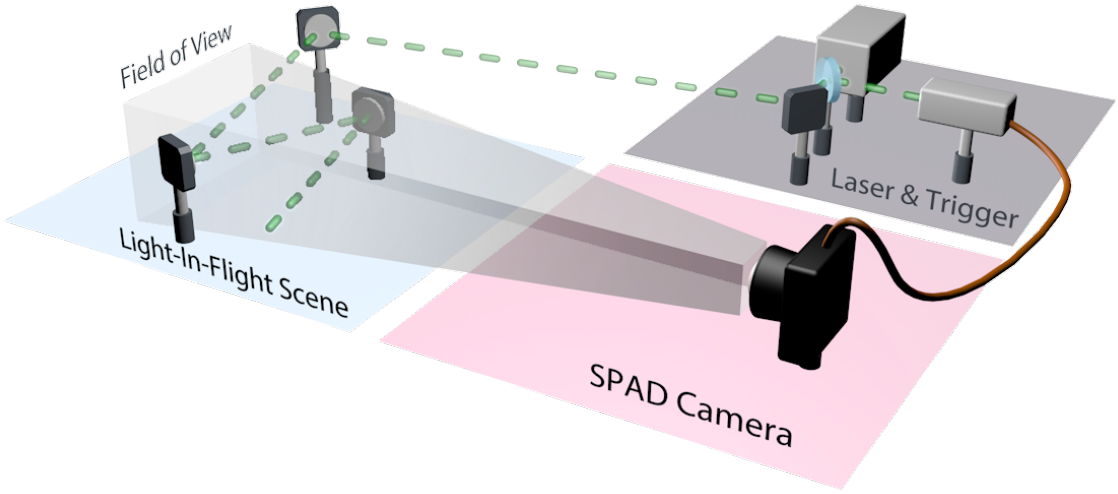

FIGURE 4.1: **Light-in-flight experimental setup.** To record the motion of laser pulses in air, the SPAD camera is imaging a 35 cm wide scene at a distance of 1.8 m. The laser pulses are directed across the field of view of the SPAD camera: they first propagate from right to left and are reflected by two mirrors, so that they propagate three times across the field of view.

The detector used in this specific experiment is another version than the one described in chapter 3 - although the design and fabrication process is the same, it has slightly different characteristics: it has 67 ps time bins and around 500 ps wide impulse response function. We use this detector in combination with a 8 mm fisheye lens with f/3.5 aperture to assemble the SPAD camera. To record the motion of the laser pulses across the field of view, the camera frames are set at an exposure time of 1000 μs (there are 4 pulses passing in every frame), and the data is acquired for 500 000 frames. The total acquisition is therefore an integration of the signal over 2 million identical laser pulses, and the acquisition time is about 10 minutes. As every pulse is exactly repeated and synchronised with the acquisition, the data will show the evolution of a single pulse, as integrated over 2 million repetitions.

In figure 4.2, we explain how we can directly interpret the acquired 3-dimensional data set as a recorded movie of the laser pulses in motion. For each pixel of our 32x32 array, we obtain a time-histogram, as explained in the previous chapter. An example is given in figure 4.2, where a 5 ns portion of the histogram recorded at pixel (22,21) is shown. The signal detected corresponds to the time at which the laser pulses crossed the field of view at that pixel position. If we consider all pixels at a particular time t, we obtain a (x, y) image representing the spatial evolution of the laser pulses at that time. By playing successive frames, we therefore obtain a movie of the light pulses as they propagate across the field of view. As discussed in section 3.3.1, the time axis of each histogram needs to be reversed in order to represent the time evolution in the forward direction. When reversed, an event unfolding later in time will appear later in time in the histograms, as it is the case for the results presented here. If we integrate over all frames, we obtain a time-integrated image showing the total path taken by the laser. A time-integrated image of the path can also be acquired with a very sensitive camera that does not possess any temporal resolution. We recorded an image of the path with an EMCCD camera at maximum gain for an exposure time of 7 s and obtained an image in agreement with the one obtained by integrating over all frames of the SPAD camera data, as shown in figure 4.2.

The frames and histogram shown in figure 4.2 are taken from post-processed data. As the raw data contains noise, we undergo a few processing steps to extract the signal. Having access to a 3-dimensional data set, we have the possibility to exploit both the temporal and spatial information to perform data processing and noise suppression. In figure 4.3, we show a time-integrated image of the data at different steps. First, the raw data contains not only the signal but also background noise from dark counts. In this experiment, the

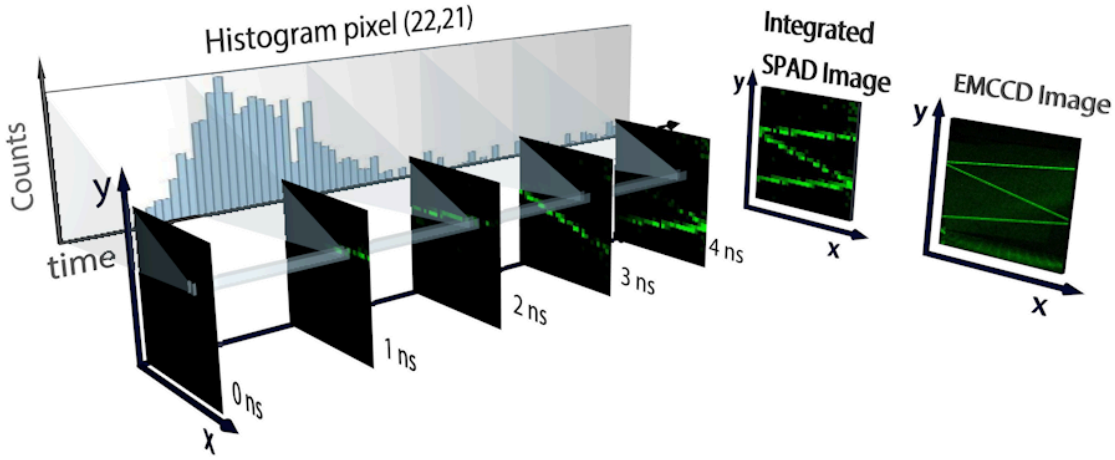

FIGURE 4.2: **3D data set and light-in-flight movie.** An acquisition with the SPAD camera gives a 3-dimensional (x, y, t) data set, where each pixel records a time-histogram of the signal. By looking at all (x, y) data for a particular time bin at time t (here for 0, 1, 2, 3 and 4 ns), we obtain a frame showing the evolution of the pulses in space at that time. Summing all the frames gives a time-integrated image of the path, consistent with an image of the path acquired with an EMCCD camera.

noise from ambient light is almost non-existent due to the fact that the experiment was performed in the dark and the camera was properly shielded from any stray light. In the time-integrated image of the raw data, the signal is barely visible – although the dark noise is generally low (83% of the pixels on this detector have a dark count rate below 50 counts/second), a few pixels have very high noise level (82 pixels have dark count rate higher than 900 counts/second). Luckily, a simple background subtraction is sufficient to suppress most of this noise: we therefore perform an acquisition with the same parameters, only without the pulses propagating across the field of view. This background is then subtracted from the raw data containing the signal - after this step, we can see that the laser path becomes visible in the integrated image (figure 4.3b). In a second step, we exploit the temporal profiles to further clean the data: each histogram is fitted with a Gaussian profile. This fitting gets rid of the remaining noise in the time bins containing no signal, as shown in figure 4.3c. It also conveniently puts to zero the pixels where no signal from the laser is detected, so that this temporal processing has a significant effect on the spatial profiles of the time frames. We then deconvolve the Gaussian fit with the IRF that we here approximate as a Gaussian, to recover 500 ps FWHM temporal signals, which is consistent with an independent measurement of the laser pulse duration. We obtain the de-noised image shown in figure 4.3e. We here perform a last processing step to enhance the quality of the visualisation and overcome the main drawback of our system: its limited spatial resolution at 32x32 pixels. To obtain a visualisation with more (x, y) pixels, we perform an interpolation step: this does not increase the information contained in our data, but it leads to a clearer and easier-to-interpret visualisation of the pulses motion. To increase the quality of our images, we therefore perform a linear interpolation in such a way that we connect the centre of mass of each line and column to reach a final (x, y) image of 310x310 pixels, as shown in figure 4.3f.

The final result of this measurement takes the form of a movie of a pulse propagating across the field of view (see appendix A for a link to the video). We show in figure 4.4 a time-integrated image from this video as well as a few selected frames separated by 0.7 ns. The data from the SPAD camera has here been overlaid with a picture of the setup taken with a standard DSLR camera. In the first frame, we can observe the pulse as it first enters the field of view. The pulse then propagates and is reflected by two mirrors, before exiting the field of view. This event unfolds over a few nanoseconds and we can observe the dynamics of the 500 ps pulses (15 cm long) as they travel 2 cm between each frame.

From the measurement, we can also estimate the signal detection rate on the SPAD detector. Summing the number of counts over the time bins of a pixel where a signal is detected gives us the number of counts corresponding to the detected signal. In this case, we obtain around 1000 counts detected per pixel. Since this measurement was acquired over 500 000 frames with exposure times long enough for 4 pulses to pass through, we therefore detect on average 0.002 photons per frame, corresponding to 0.0005 photons

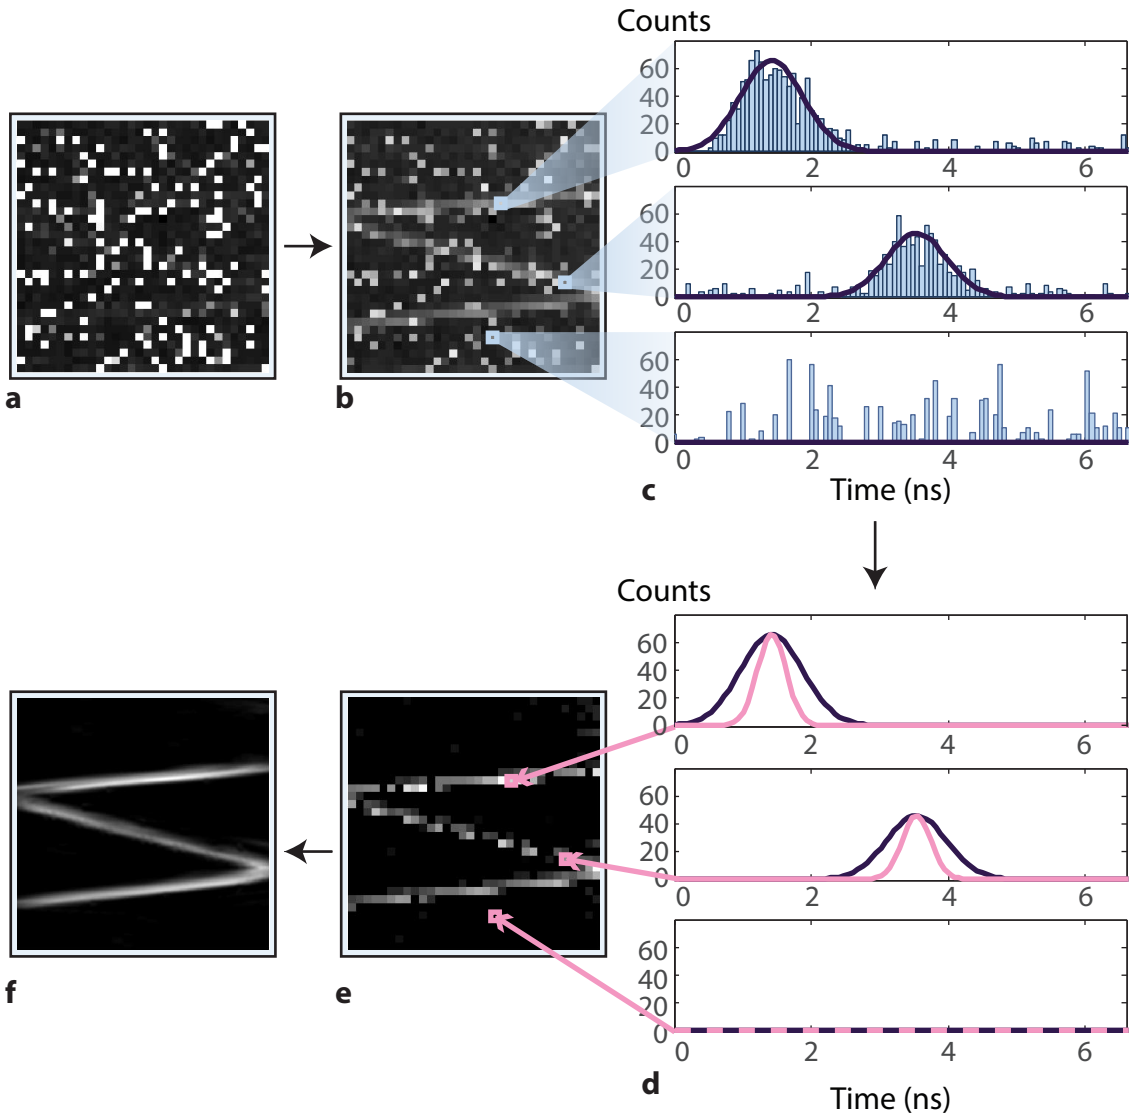

FIGURE 4.3: **Data processing.** We exploit the temporal and spatial information available to perform data processing and noise suppression. a) The raw data contains noise to a level that the signal is not visible in the integrated image. b) Performing a simple background subtraction suppresses most of the noise. c) A Gaussian is then fitted on each histogram. If the width at half-maximum of the fitted Gaussian is smaller than the pulse duration (500 ps) or higher than an upper threshold (here set at 3 ns), the fit is rejected and all values in the histogram of that pixel are set to 0, as in the third histogram shown. d) The Gaussian fits are then deconvolved back to the width of the laser pulses, 500 ps. e) The obtained frames and integrated image are cleaned temporally and spatially. f) To further enhance the quality of the visualisation, we perform a linear interpolation between the center of mass of each columns and frames to obtain frames that are 310x310 pixels.

per pulse, per pixel. At this detection rate, our measurement is definitely situated in the photon-starved regime. Taking into account the detector efficiency (23% at 532 nm) and fill-factor ($\sim 2\%$), this detection rate corresponds to a photon incidence rate of 0.4 photons per pixel per frame, which highlights the necessity for single-photon sensitivity to perform such an experiment.

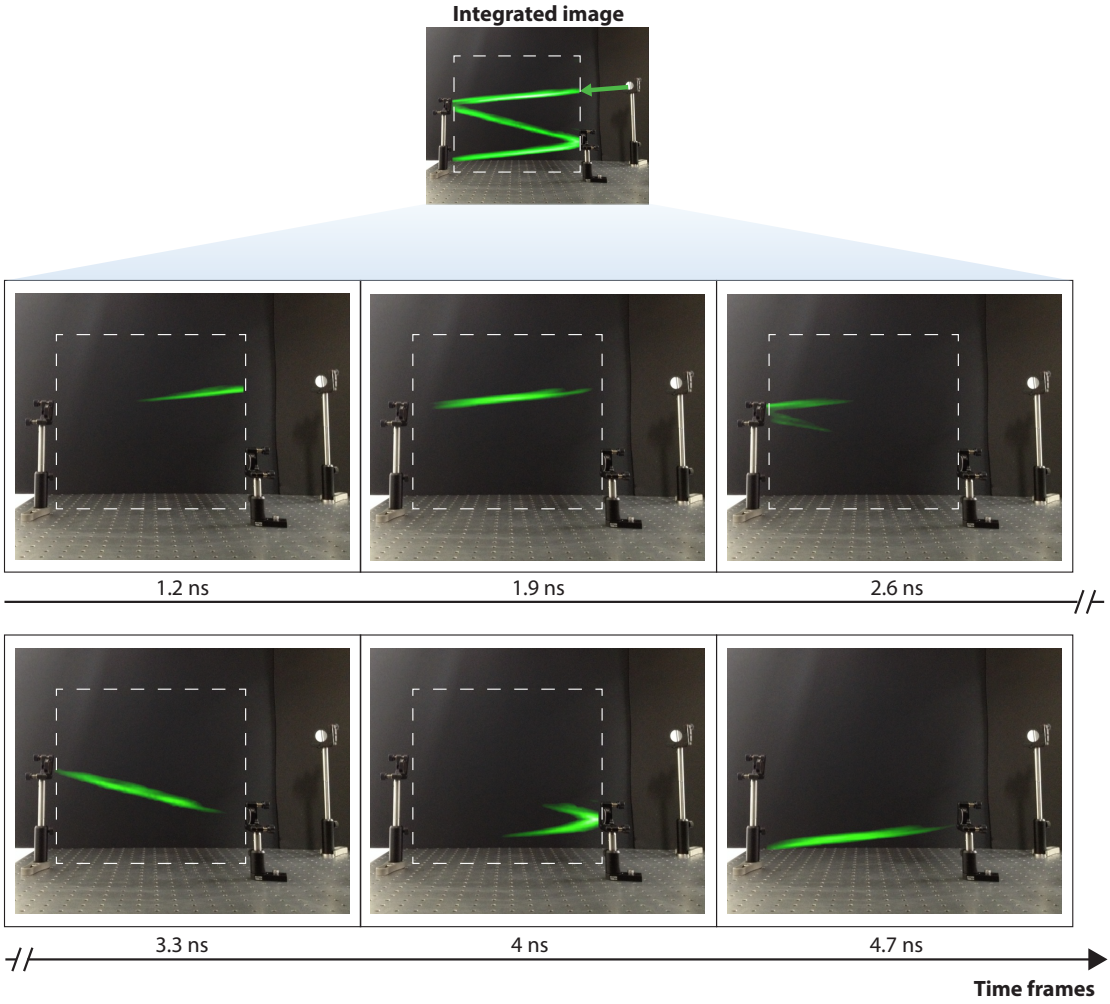

FIGURE 4.4: **Light-in-flight movie.** Our recorded data allows us to visualise a pulse of light as it travels in air across the field of view of the camera, indicated by the dashed square. We observe the pulse first entering the field of view and evolving over a few nanoseconds as it is being reflected by two mirrors before exiting the field of view.

This result demonstrates that it is now possible to record the dynamics of light propagation in air, and was published in [7]. This experiment shows a novel way to perform light-in-flight imaging at very low level of intensities, where the setup can be extremely simple as the detector is compact and easy-to-use and any light dynamics, including propagation in air, can be observed directly without the need for any additional components in the setup to enhance the signal. We will see in the next sections how this technology can be used to study more complex phenomena, but we will first discuss some details and subtleties of this first light-in-flight experiment.

4.1.1 Effect of polarisation

As the scattering the light pulses undergo as they propagate is Rayleigh scattering, the efficiency of scattering in the direction of the camera is dependant on the polarisation of the laser. Indeed, the angular distribution of Rayleigh scattering is isotropic in the plane perpendicular to the polarisation of the beam, and follows a $\cos(\theta)^2$ distribution in

the plane of the polarisation [63], as illustrated in figure 4.5a and 4.5b. We verified this dependence by capturing the propagation of laser pulses in air, as above, and changing the polarisation of the laser. To do so, we use a $\lambda/2$ waveplate and rotate its angle by steps of $10°$. We observe in figure 4.5c the expected decrease in signal as the polarisation goes from vertical (waveplate at $0°$) to horizontal (waveplate at $45°$) and increases again as we increase the angle of the waveplate to $90°$ and bring back the polarisation to vertical. This confirms the expected behaviour of the detected scattered light, and underlines the importance of setting the laser to the adequate polarisation to perform a light-in-flight experiment in air.

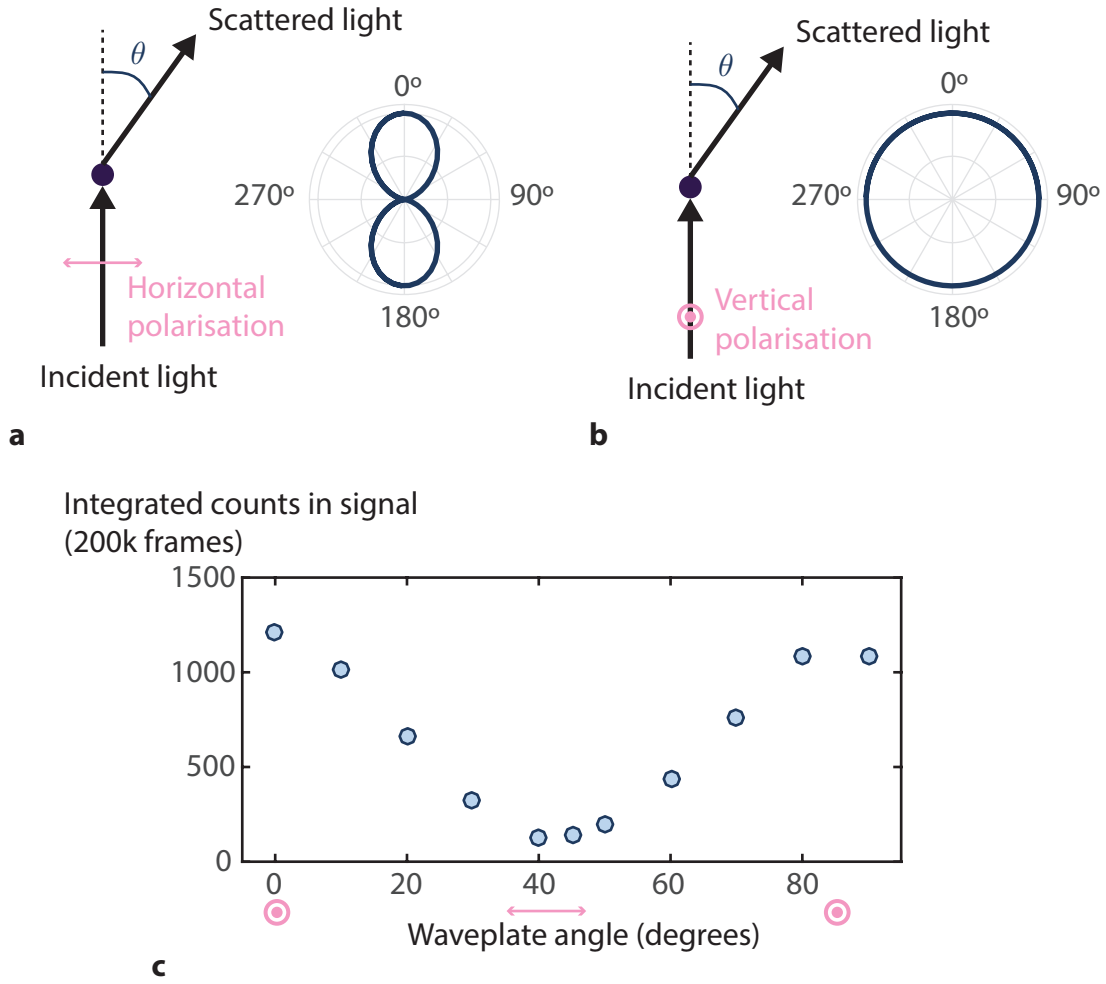

FIGURE 4.5: **Scattering efficiency versus polarisation.** The angular distribution in the horizontal plane of the scattering by air molecules (Rayleigh scattering) is dependant on the polarisation of the beam. a) If the beam is horizontally polarised, there will be no scattering at $90°$ from the direction of propagation and the scattering will be preferentially in the forward ($\theta = 0°$) and backward ($\theta = 180°$) direction. b) However, if the beam is vertically polarised, the angular distribution of the scattering will be uniform in a horizontal plane, and some of the scattered light will be deviated in the $90°$ direction towards the detector. c) The experimental measurement of the signal dependance on the laser polarisation shows the expected behaviour.

4.1.2 Measurement in pure N_2

Our experiment was performed in a lab which has a cleanroom 100 000 classification, meaning that there is at maximum 100 000 particles/ft^3 dust particles bigger than 0.5 μm in the air, and at maximum 3.6×10^6 particles/ft^3 bigger than 0.1 μm, orders of magnitude less than the number of N_2 molecules contained in the air. Our measurement is therefore likely to mainly rely on scattering off air molecules rather than scattering off dust particles. However, to confirm this distinction in the possible sources of signal, we repeated a light-in-flight measurement in pure N_2.

To do so, we filled a tube with pure N_2. The tube had one end open so that the pulses could be directed in the tube and not encounter any surface that would generate a strong signal detected by the SPAD camera. To compensate for the loss of gas in the tube by this open end, and make sure it did not get filled with air from the lab, we left the gas

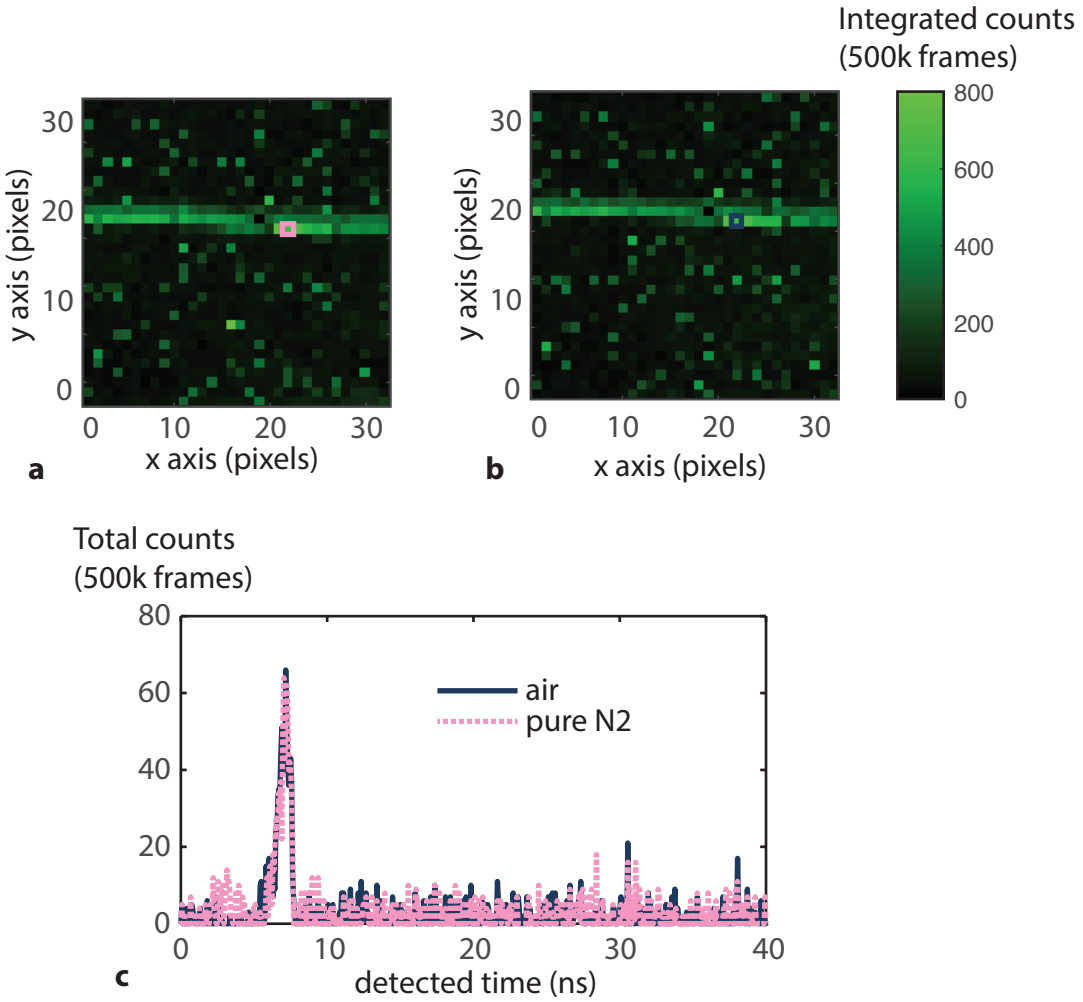

FIGURE 4.6: **Scattering from air molecules.** a) The integrated image over 20 time bins (1.3 ns time window) showing the path of the light in pure N_2 is similar to the same path measurement in ambient air (b). c) By looking at the histograms from the same pixel (shown by coloured square in (a) and (b)), we observe that we detect the same amount of light from a measurement in pure N_2 and in air, confirming that the main contribution to our signal when performing the measurement in air is the scattering from air molecules.

intake in the tube open with a pressure sufficient to constantly create a flow, such that the space in the tube was occupied by pure N_2 at all times. We then recorded the propagation of light inside this tube. In figure 4.6a, we show an image of this path integrated over 20 time bins, on which we can observe the linear path of the beam. We then repeat the same measurement in air, by removing the tube and the pure N_2 source. The image integrated over the same bins in the air measurement is shown in figure 4.6b. We also show a histogram taken from the same pixel in both measurements, which show that the two measurements are quantitatively equivalent. This comparison demonstrates that we obtain the same result when performing a light-in-flight measurement both in air and in pure N_2, which confirms that the light detected in the experiment done in air is indeed coming from scattering off air molecules. There might be some occasional scattering events from dust particles, but this is not the main contribution to our measured data. In fact, in section 4.1.4, we will discuss a light-in-flight measurement performed with a different type of camera, which is sensitive to intensity and does not operate in single-photon counting mode – in this measurement, we will see clearly the different contributions from light scattered from air molecules and from dust particles.

4.1.3 Measurement with a femtosecond laser

We consider here a light-in-flight measurement such as the one performed and detailed above, but using a femtosecond laser rather than the 532 nm microchip laser as above. The measurement discussed in this section was performed with the SPAD camera detailed in chapter 3, which has a 45.5 ps time bin width and a 110 ps impulse response function. The laser we use in this version of the light-in-flight experiment is a Ti:Sapphire oscillator emitting 10 fs, 10 nJ pulses at 67 MHz (670 mW average power). In this case, each pulse contains 4×10^{10} photons, which is two orders of magnitude less than in the green microchip laser, and the detector sensitivity drops from 23% to 6.5% around 780 nm. However, the repetition rate of the femtosecond laser is four orders of magnitude higher than the green laser used above, so that the expected detection rate with a femtosecond laser is 20 times higher than with the microchip laser.

The measurement presented in this section was acquired with a similar setup than the one presented above and published in [7], where the laser pulses propagate in air and cross the field of view of the SPAD camera 3 times as they are being reflected by two mirrors. The exposure on the camera is set to 1000 μs. In figure 4.7, we show an example histogram from the acquired data set, taken from the pixel indicated by the white square in figure 4.8, at frame 1.6 ns. This measurement shows clear differences to the histograms shown in figure 4.3, as it contains a lot more structure and the width of the peaks is much narrower. In this histogram, we find the signal coming from the laser pulse propagating across the field of view at this pixel position. This signal is indicated in figure 4.3. Its width is now limited by the impulse response function of the detector, as the width of

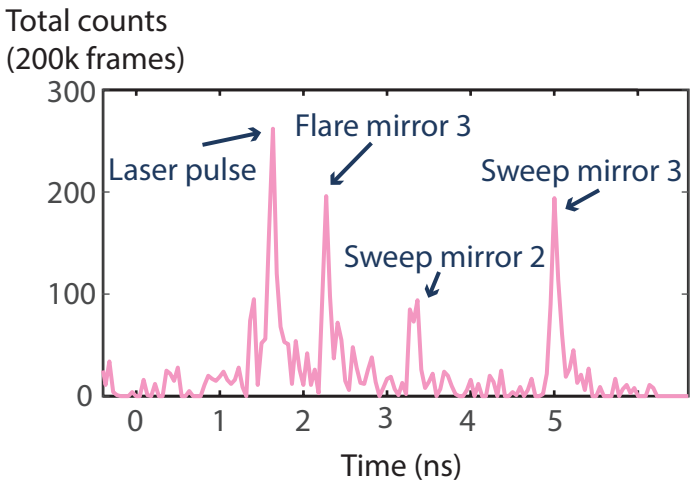

FIGURE 4.7: **Histogram from light-in-flight measurement of a femtosecond pulse.** An example histogram acquired over 200 000 frames shows a series of peaks corresponding to different paths taken by the laser pulses.

the actual pulse is only 10 fs, much shorter than the bin width of 45.5 ps. The other 3 peaks are caused by other sources of signal, as we will discuss below and show in figure 4.8. They are indicated here as "flare" and "sweep" signals originating from the different mirrors in the setup - these signal sources will become apparent in the frames of the movie shown in figure 4.8.

To create a clear visualisation of the recorded light-in-flight movie, we followed a few steps of processing that slightly differ from the processing steps discussed in figure 4.3. We first perform a background subtraction, as above, but rather than fitting a simple Gaussian, we perform a peak-finding step and position a Gaussian of width equal to the IRF width at every peak position. A more thorough fitting step could be developed for situations where we measure multiple peaks in a single histogram, and better noise removal obtained, but as this data is not the object of a deeper analysis and used only for visualisation and interpretation purposes, this step is sufficient to obtain satisfactory data without affecting the information contained in the histograms that we are interested in. We also skipped here the interpolation step, as the frames show more complex structures that cannot be properly interpolated with a simple linear interpolation as the one described above. The obtained movie is available online as indicated in appendix A, and selected frames separated by 0.4 ns are shown in figure 4.8.

The overall scene is shown at the top of the figure. The laser is directed in a path composed of three mirrors: mirror 1, out of the field of view, directs the laser to enter the field of view on the right moving towards the left, then mirror 2 reflects it back towards the right, and finally mirror 3 reflects it back again towards the left and outside the field of view. We observe in the recorded light-in-flight movie frames three different types of signal. First, we can see the laser pulse entering the field of view in frame 0.4 ns, and we

follow its evolution as it propagates in air and is reflected by the two mirrors. The pulse is clearly seen in frames 0.4 ns and 0.8 ns as it propagates from right to left, then in frame 1.6 ns it is seen travelling from left to right between the two mirrors, and finally traveling

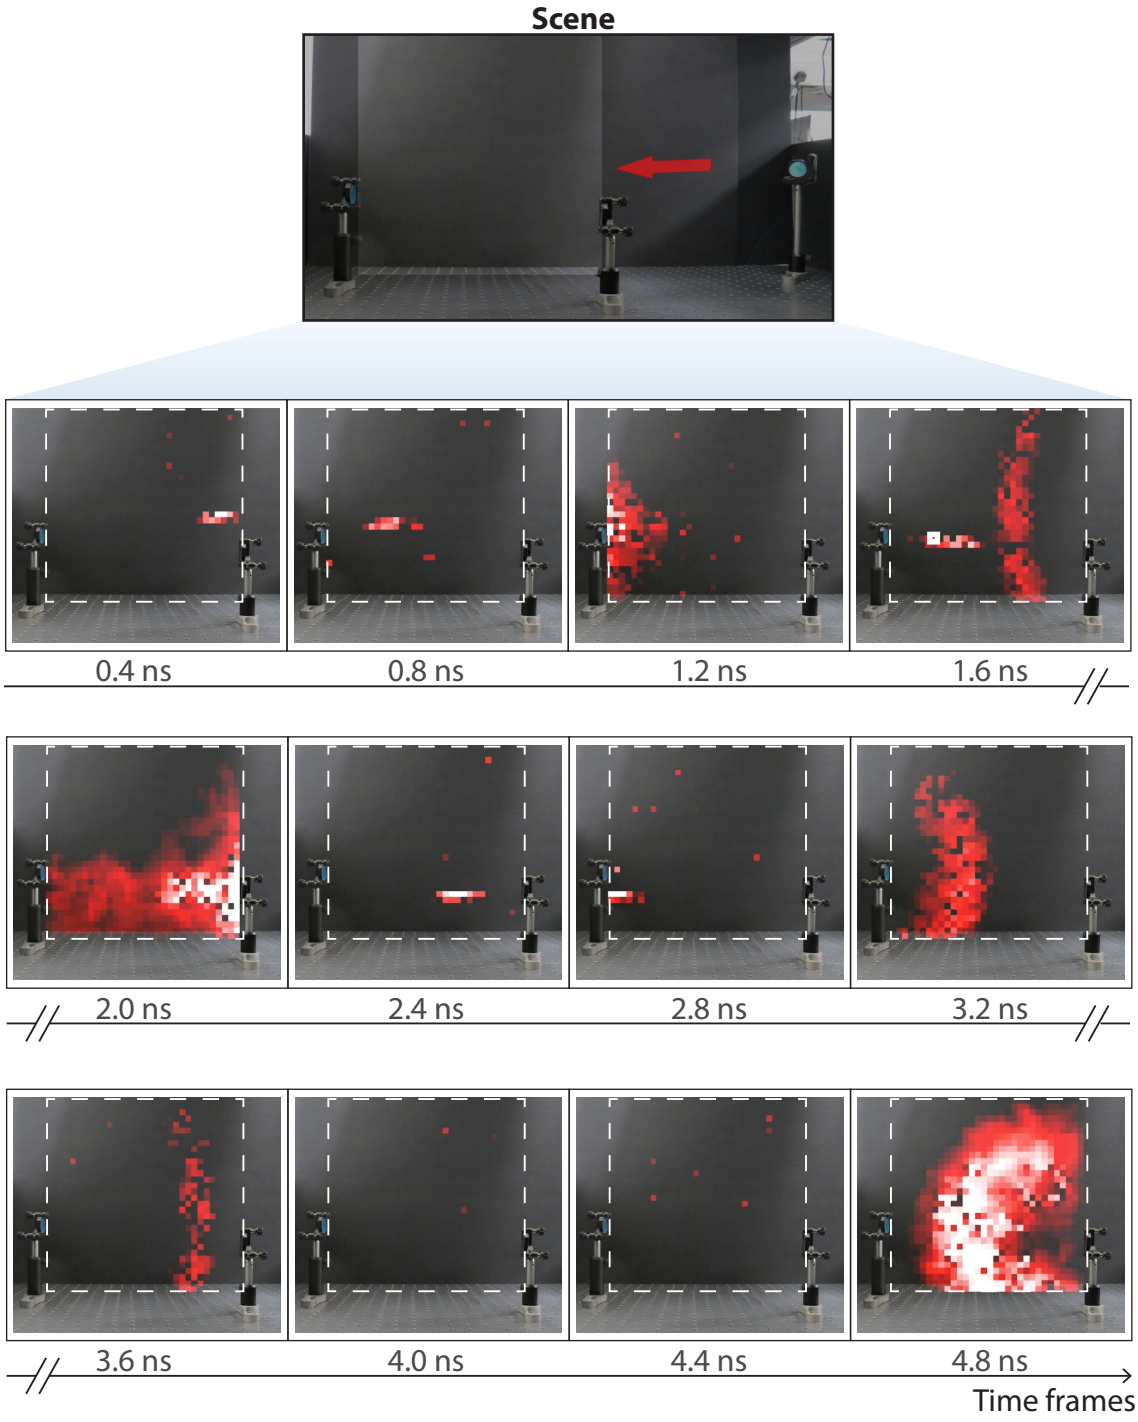

FIGURE 4.8: **Light-in-flight movie of a femtosecond pulse propagating in air.** The scene measured is a femtosecond laser pulse propagating across the field of view of the camera (entering on the right as indicated by the arrow on the top picture). The pulse is first seen as it propagates from right to left, and is subsequently observed after being reflected by two mirrors on the left and right of the field of view. We also observe signals coming from the fraction of the light that is scattered from the mirrors after a reflection: it first creates a flare effect, seen in frames 1.2 ns and 2.0 ns. The scattered wave then reaches the back end of the scene and intercepts the black cardboard: we can then observe slices of that wave as it sweeps across the cardboard.

to the left and out the field of view in frames 2.4 ns and 2.8 ns.

In frame 1.2 ns, we can already observe a second type of signal: when the laser hits a mirror, some part of it is scattered approximately as a spherical sphere, and some of the light scattered leaks in the camera, as you would observe with a normal camera when it records a flare from a nearby very bright source. In frame 1.2 ns, we can therefore observe a flare from the light scattered by mirror 2, and we can see the flare created from mirror 3 at frame 2.0 ns. The reason why we did not observe this type of signal in the results presented above (with the 532 nm laser) is that the mirrors were then placed a bit farther out of the field of view. Here, the mirrors are placed at the edge of the field of view and generate this flare effect when they reflect the laser. These flares coincide with the times at which the laser is seen reflecting on the mirrors.

The third type of signal can be observed in frames 1.6 ns, 3.2 ns, 3.6 ns and 4.8 ns, in which we observe a signal shaped as an arc of a sphere propagating across the field of view. These are also due to the scattering originating from the mirrors that occur when there is a reflection: the flare was associated with the portion of that scattered light going towards the camera, and this third signal is associated to the part of the scattered light that is directed to the back of the scene. Indeed, after a scattering event, light is dispersed in all directions and moves approximately as a spherical wave. Part of this wave will hit the black cardboard at the back of the scene, behind the mirrors. Since this happens within the field of view, the camera will record an arc of that spherical wave as it intersects the cardboard and sweeps across it. In frame 1.6 ns, the arc of circle is bent towards the position of mirror 1, where it originates - in the movie we can see it sweeping from right to left. In frames 3.2 ns and 3.6 ns, the arc originating from mirror 2 is seen propagating from left to right. In the final frame, we can see the scattered signal from mirror 3 as it intercepts the cardboard. This type of signal is of great interest as it can provide information about an object that is outside the field of view of the camera: the so-called sweep signals shown in figure 4.8 show a clear spatial profile in the form of a circle section, which centre indicates the origin of the signal. Combined with the temporal information contained in this signal, one can use it to pinpoint the position of an object hidden outside the field of view of the camera – this is the principle we use in chapter 5 when we describe one of the main application of the ability to capture light in motion: using the recorded information to detect the presence and position of hidden objects.

It is also interesting to note that two types of mirrors were used in this setup and that a difference in their flare and sweep signal can be observed: mirror 1 and 2 were dielectric mirrors, and mirror 3 was a silver mirror, which has slightly lower reflectivity around 800 nm [64]. By observing the frames in figure 4.8, we can see that mirror 3 creates a very strong flare and sweep signal compared to mirrors 1 and 2.

4.1.4 Measurement with an ICCD camera

As we have discussed above, the ability to capture light in motion as it travels through air relies on the extreme sensitivity of the detector so that it can detect the photons that are scattered by air molecules in the laser path. We have demonstrated that the SPAD camera is a great tool to perform such a measurement, and we have shown that a static measurement of the light's path is possible with an extremely sensitive camera such as an EMCCD. In this section, we consider an additional type of technology as a tool for capturing light in motion: an intensified CCD camera (ICCD). This is a type of gated measurement device that has not been exploited so far in the field of light-in-flight imaging. We therefore study here its potential for capturing light in motion in air. As illustrated in figure 4.9, such a camera is composed of two main components: a CCD camera, cooled to achieve low noise level, and an intensifier tube, that intensifies the incoming light signal and can provide a gain of about 4 to 6 orders of magnitude [65]. The first component of the intensifier is a photocathode, that transforms any incoming photon into a photoelectron. A voltage is applied between the photocathode at the front of the intensifier

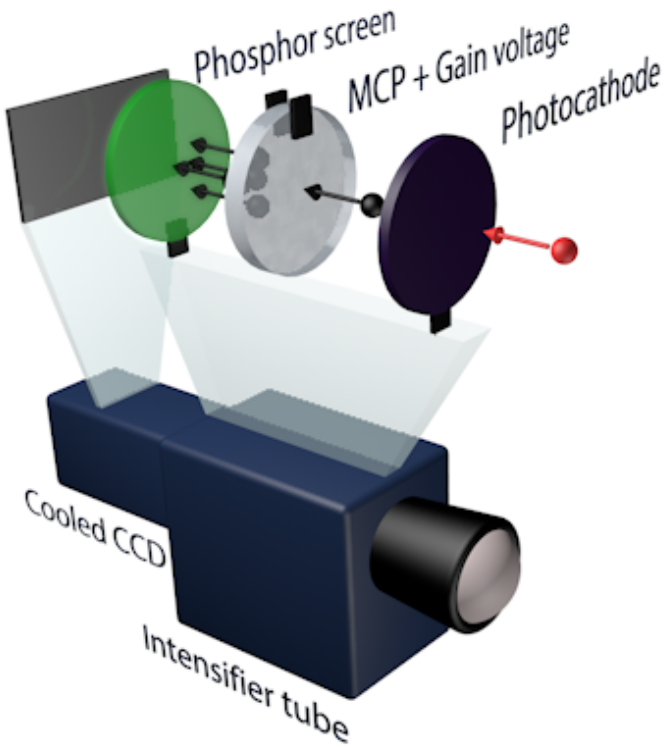

FIGURE 4.9: **Intensified CCD camera components.** An ICCD is composed of an intensifier tube, which transforms the optical input into photoelectrons to amplify the signal. The amplified electron signal is then converted back to an optical signal when the photoelectrons hit a phosphor screen. A cooled CCD detector then image the phosphor screen. Since the CCD detector is rectangular (520x688 pixels) and the intensifier tube is circular, the resulting field-of-view is the overlap between these two surfaces. Extremely high sensitivity can be achieved, and temporal resolution of around 200 ps can be obtained by applying a voltage on the micro-channel plate for a fixed period of time.

and the phosphor screen at the back so that electrons are accelerated through. They pass through a microchannel plate (MCP), composed of a multitude of independent electron multiplier channels, where, depending on the voltage applied, the signal will be amplified to the desired gain level. The electrons then reach a phosphor screen so that the signal is transfer back to an optical signal that can be detected by the CCD detector. This intensifier allows to reach sensitivities similar to that of an EMCCD or even our SPAD camera. Additionally, data can be acquired with some temporal resolution by applying a gated voltage on the MCP rather than a continuous one. The ICCD used in this subsection is the PicoStar by LaVision, and its intensifier can be gated with 200 ps gates: when operated in that mode, the CCD will detect only light in an exposure of 200 ps. The gate can be synchronised with the laser pulse emission. In one measurement, an ICCD therefore only provides one frame at a specific time - a movie can be produced by scanning the gate, and here we do so by steps of 20 ps.

We used the ICCD camera to record the motion of a femtosecond laser of 240 mW average power following a simple path: a straight line from left to right across the field of view. We set the exposure time for the acquisition at one gate position at 1 second, and the gain voltage of the MCP is set to 420 V (from an available range of 260 V to 850 V, where the gain increases exponentially with the voltage). The data acquired was compiled as a movie of the laser pulse's motion (see appendix A). In figure 4.10, we show three frames from that movie, separated by 0.3 ns. In the first frame, we can observe the pulse on the left as it enters the field of view. We can see in the three frames the evolution of a scattered wave sweeping across the back end cardboard and the floor, similarly to the waves observed in section 4.1.3.

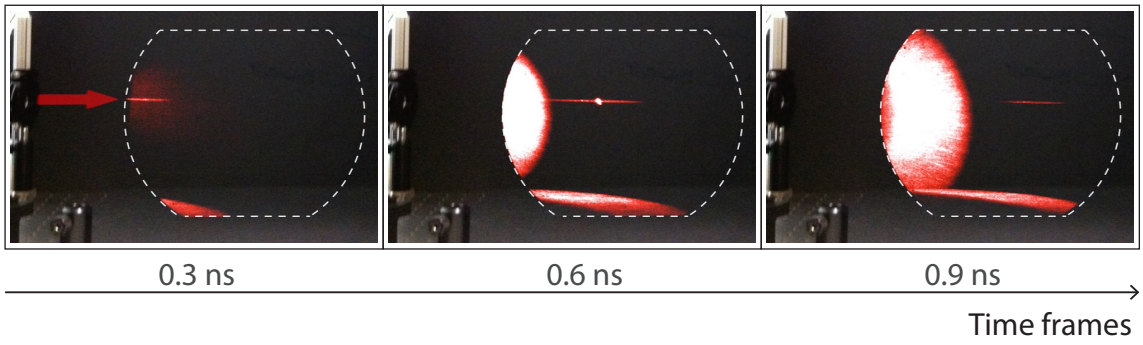

FIGURE 4.10: **Light-in-flight measurement with an iCCD camera.** The ICCD provides high enough sensitivity to capture the scattering from air molecule of laser pulses propagating through air, as we measured with a SPAD camera. In these frames, we can also observe the high intensity signal originating from scattering events from dust particles crossing the path of the laser pulses.

An interesting point to observe in the frames presented in figure 4.10 is that we can observe contributions from light scattered along the laser's path by both air molecules and dust particles. In frame 0.3 ns and 0.9 ns, the pulse is measured as uniform – the detected

light was therefore coming predominantly from scattering off air molecules. However, on frame 0.6 ns, we can observe a bright spot in the middle of the pulse: this corresponds to an event where the laser encountered a dust particle in its path, resulting in a strong scattering. If such an event happens when measuring with the SPAD camera, the SPAD can only count one photon scattered this way, as it operates on a photon-counting principle. Since the dust particles move around, the strong scattering at that specific point in space only occurs for one laser pulse, and is not repetitive as the scattering off air molecules is. Therefore, a scattering event from a dust particle will only contribute as one more count to the signal detected by the SPAD camera – with the ICCD, however, the intensity of that event is registered, and is much higher than the signal obtained from air molecules. We can therefore clearly distinguish the signals coming from these two types of scattering event – in the full movie, we can observe multiple scattering signals arising from encounters between the laser pulse and dust particles in the air.

This result demonstrates that an ICCD can also be a great tool for imaging light in motion at low intensity levels, although the signal for light-in-flight in air can be corrupted by strong signals generated by dust particles, and this type of camera is significantly more bulky and expensive than a SPAD camera. However, it provides great spatial resolution (520x688 pixels in the PicoStar) and similar temporal resolution (200 ps gate), so that it can be the ideal tool to use in some circumstances where high spatial resolution is paramount. The work by Clerici $et\ al.$ which studies the particularities in the apparent motion of light, as discussed in section 2.1.6, was performed using this ICCD camera [32], and further work is likely to be done in this field using this alternative tool.

4.2 Laser-induced plasma

One of the first more complex phenomenon we observed with our SPAD camera was the dynamics of a laser-induced plasma (see publication [7]). Such a phenomenon occurs when a laser pulse is intense enough to ionise air molecules, in particular N_2 molecules, forming a plasma that has a lifetime on the order of a nanosecond. When the free electrons recombine with the molecules, they emit a fluorescence signal in the blue region of the visible spectrum. To record this phenomenon, we used an amplified Ti:Sapphire laser (785 nm) emitting 12 mJ pulses of 90 fs duration at 100 Hz. Loosely focused with a 50 cm focal length lens, each pulse weakly ionises the air at focus and creates a plasma filament of about 5 cm long [66]. The fluorescence signal emitted from this plasma originates from the fluorescence of N_2 molecules and the emission is located in the 300 to 440 nm wavelength region [67]. To observe the dynamics of this plasma creation and fluorescence emission, we used our SPAD camera that has 67 ps time bin and 500 ps impulse response function. This first measurement does not provide very precise information about the temporal structure of the phenomenon, due to the relatively wide response of the specific

SPAD array used, but provides sufficient information to create a qualitative visualisation of the dynamics of this phenomenon. To observe both the femtosecond laser signal and the plasma fluorescence, we performed two successive measurements using different colour glass filters: one with a red filter placed between the camera objective and the detector array (Altechna RG15, highpass > 700 nm), to record the light-in-flight measurement of the laser pulses as they focus in the centre of the field of view; and one with blue filters to select uniquely the signal from the plasma fluorescence (Altechna BGG24 and BGG25, $\sim$300-700 nm bandwidth) [68]. Both measurements are acquired over 100 000 frames at 1000 μs exposure time – since the laser emits pulses at 100 Hz, only 1 in 10 frames is synchronised with a pulse propagating across the field of view, and a measurement therefore integrates the signal over about 10 000 pulses. Both measurements are then combined with a video-editing software where both signals are coloured in red and blue respectively, to create a visualisation of a laser-induced plasma evolving over a few nanoseconds (see appendix A). Selected frames separated by various times from 0.3 ns to 1 ns are shown in figure 4.11.

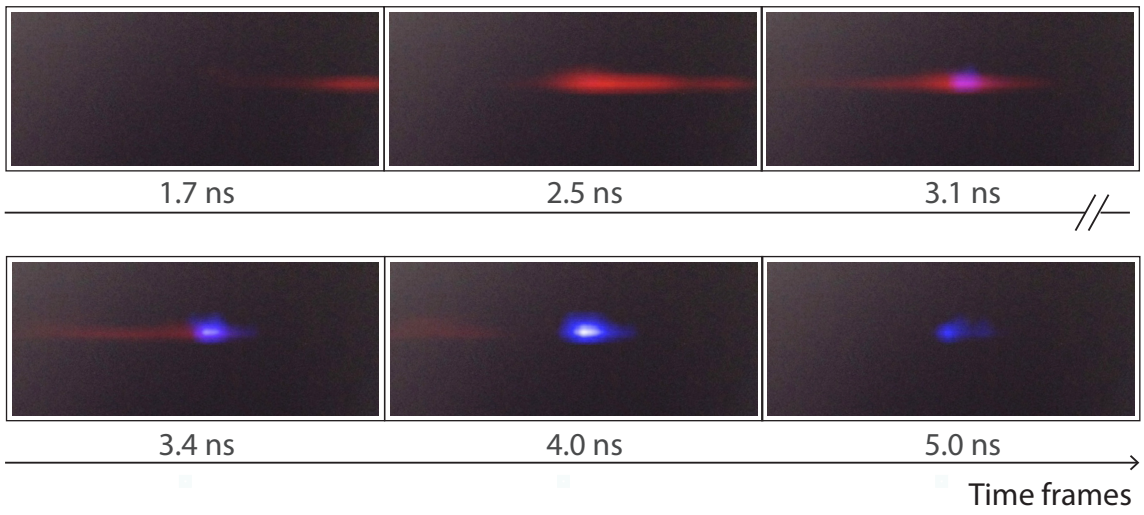

FIGURE 4.11: **Light-in-flight measurement of a laser-induced plasma.** We observe the propagation of a femtosecond pulse (in red) as it focuses in the centre of the field-of-view and creates a plasma, which emits a fluorescence signal in blue. The plasma fluorescence evolves over a few nanosecond after the propagation of the femtosecond pulse. There is no transverse geometrical effect observable during the focusing of the femtosecond laser beam, as the beam size is on the order of the pixel size.

A few interesting details can be observed in the frames shown in figure 4.11. First, we can observe the femtosecond pulse as it enters the field of view on the right, propagating towards the left. The first frame (1.7 ns) shows a similar behaviour than the one observed in the green laser light-in-flight measurement discussed in section 4.1. The behaviour then differs from the simple light propagation measurement in the subsequent frames: we can see a much stronger scattering signal at focus, as the pulse is scattered by the plasma created by the leading edge of the pulse. The beam then follows its path with a reduced

intensity due to this high scattering at the focus: the beam appears much dimmer in frames 3.4 ns and 4.0 ns as it moves out of the field-of-view. In blue, we observe the plasma being formed at focus and fluorescing well after the femtosecond pulse has passed.

To create this visualisation, the data was processed in a similar way than explained in section 4.1: we first subtract a background acquired without the laser propagation, then apply a Gaussian fitting to the histograms to clean the data both spatially and temporally. The 32x32 frames are then interpolated linearly in x and y. A Gaussian temporal profile is not the ideal fit to a fluorescence measurement, but since the impulse response function of the SPAD camera is 500 ps and the histograms were acquired over only 10 000 pulses, the temporal profile obtained with this detector does not contain very rich temporal information, and the Gaussian fit is a satisfactory approximation to create a qualitative visualisation of the plasma fluorescence. To obtain a more precise quantitative measurement, we later recorded the plasma fluorescence with the higher resolution detector that we have, the 45.5 ps time bin and 110 ps impulse response detector that we detailed in chapter 3. The temporal histogram recorded with that detector is much clearer and shows the exponential decay expected from a fluorescence measurement, as shown in figure 4.12a. The fluorescence decay can be observed over 2 ns after the peak signal.

A proper measurement of the fluorescence lifetime of the plasma needs to take into account the impulse response function of the detector. As measured for characterisation purposes in section 3.3, the IRF's FWHM of this detector is 110 ps. However, when operated at extremely low repetition rate (here at 100 Hz), the detector shows higher jitter in the measurements. A measurement similar to that reported in section 3.3 was therefore performed with the 100 Hz laser and an example of a recorded histogram is shown in figure 4.12b. The IRF width was here measured as 280 ps. We can also observe that the noise level is higher when operating the camera at 100 Hz, and shows a pattern where one in every eight bin systematically detects a higher signal than the others. This is a feature originating from the TDC design and the non-uniformity of the gates in the ring resonator, the last gate of the 3-bit resonator being slightly longer than the other bins. This feature could be calibrated, but in this experiment we simply reduce the noise level by performing a background subtraction. This noise pattern and the wider IRF are features arising when operating the SPAD camera at 100 Hz, as the detector was designed for higher repetition rate (kHz - MHz). The detector however still works at these low repetition rate, and here we approximate the impulse response function as a Gaussian of 280 ps FWHM to deconvolve the plasma fluorescence and obtain a fluorescence lifetime measurement. The fluorescence lifetime is found using the DecayFit software [69], widely used for fluorescence lifetime measurements and imaging, which was adapted to our specific measurement by Dr Nikola Krsatjic from the CMOS Sensors and Systems group at University of Edinburgh. This code finds with an iterative algorithm the exponential decay fluorescence function that best fits the acquired data after convolution with

our impulse response function. The fit performed on the histogram shown in figure 4.12a is shown in figure 4.12c - the lifetime found here is 600 ps, a value consistent with the lifetime reported in literature [66], on the order of 1 ns.

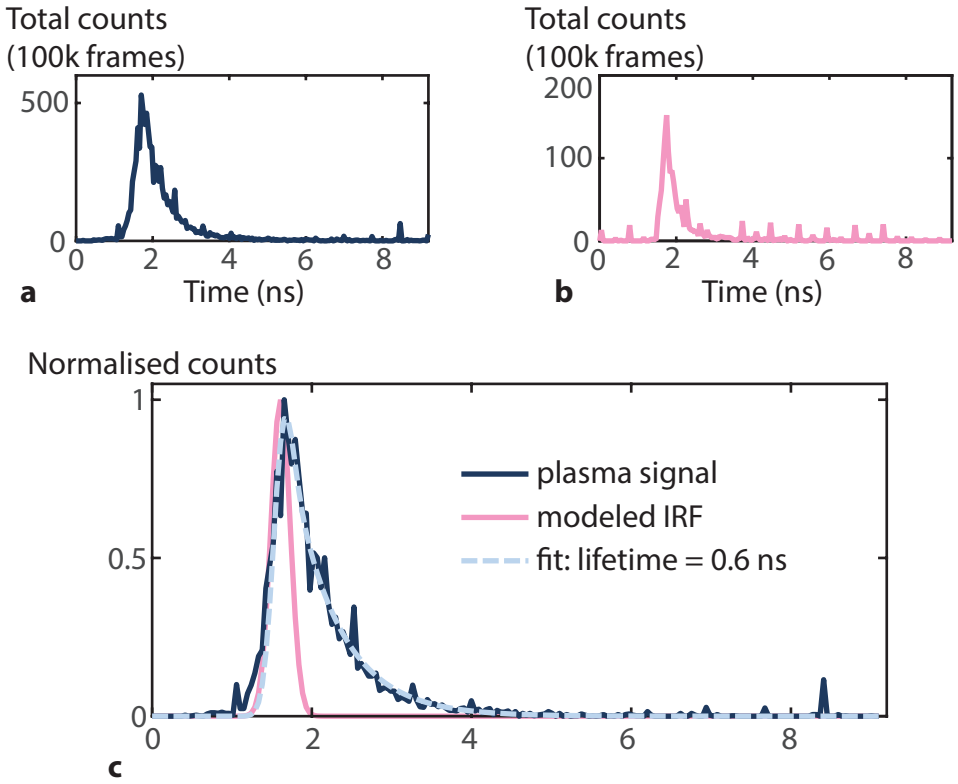

FIGURE 4.12: **Laser-induced plasma fluorescence measurement.** a) The measured fluorescence signal shows an exponential decay up to 2 ns after the main peak of the signal. b) The measurement of the exponential decay lifetime needs to take into account the impulse response of the detector, which has a temporal width significant with respect to the lifetime of the signal (280 ps). c) We find the best exponential decay that fits the data after convolution with the impulse response function. Here we show the best fit obtained over the few pixels where we measure the fluorescence of the plasma, and measure a lifetime (decay to 1/e of peak) of 600 ps.

The deconvolution fit was performed on several pixels of the plasma measurements, and it was observed that the obtained values varied from pixel to pixel and depending on the deconvolution starting parameters. When analysing signals acquired on 10 different pixels that had an acquisition rate of less than 0.25 detected photon per pulse, we found lifetime values varying between 400 and 600 ps (using an uncertainty of 2 times the standard deviation: 500$\pm$100 ps). The measurement presented here is therefore a good approximation of the laser-induced plasma fluorescence, but a more thorough measurement and analysis would fill the need for a more precise measurement and more reliable value of the lifetime. Having access to a powerful enough laser to produce a plasma at higher repetition rate would be of great advantage. An appropriate analysis of the lifetime should also include a better model function for the impulse response function. Here, we did not use the IRF measurement for the deconvolution as it is quite noisy, but a model function as a series of exponentials would be a better approximation than a Gaussian im-

pulse response. This type of function has been used with SPAD detectors for precise LIDAR measurements [41]. The measurement reported here is therefore a demonstration that the SPAD camera can be used to study the dynamics of laser-induced plasma, and we indicate the way forward to a more thorough analysis of the phenomenon.

4.3 Laser propagation in fibers

Given the SPAD camera allows us to observe light propagation in a very lowly scattering medium such as air, we evaluated the possibility of observing at a standoff distance the propagation of light in a different medium: an optical fibre. This is a type of optical waveguide formed of a core and a cladding of lower refractive index than the core. Light is guided along the fibre by total internal reflection inside the core. To couple a laser beam into the core of a fibre, one needs to focus the beam to a transverse size corresponding to the dimensions of the fibre core. Leaking out of the fibre arises during propagation, giving access to a signal to record the propagation, and is increased by a non-efficient coupling into the core and by strong bends in the fibre. The first measurements of a laser pulse propagating in an optical fibre were performed in collaboration with the Camera Culture group at MIT Media Lab, where we sent laser pulses from a femtosecond Ti:Sapphire laser oscillator directly into a standard multimode fibre, which resulted in an inefficient coupling into the fibre. The laser propagated in a few tens of centimetres of fibre before entering the portion of the fibre in the field of view of the SPAD camera. We then arranged the fibre into a series of three loops in one plane and imaged that plane with the SPAD camera (45.5 ps bin, 110 ps IRF). The light that leaks out of the fibre is detected by the SPAD camera and this signal is sufficient to track the propagation of the laser pulses across the fibre. Data for this scene was acquired over 111 480 frames with 100 μs exposure time.

In figure 4.13, we show a series of frames from the resulting movie (see appendix A for the full video). We can observe the total path taken by the pulses in the integrated image, and the evolution of the pulse in the fibre over a few nanoseconds in the individual frames. The data presented here was only processed by performing a background subtraction. We show in the top right inset an example of a pixel's histogram, where two peaks appear as the pulses pass one time through this pixel in the first fibre loop, then a second time about 3 ns later when the pulses go through the second loop.

This early measurement showed the possibility to use the SPAD camera to observe light phenomena happening inside an optical fibre at a standoff distance, without perturbing the phenomenon itself. A deeper and more detailed study of optical phenomena in fibres followed this work. This study, lead by Dr Ryan Warburton of the Quantum Optics Group at Heriot-Watt University, focused on two phenomena of interest: supercontinuum gener-

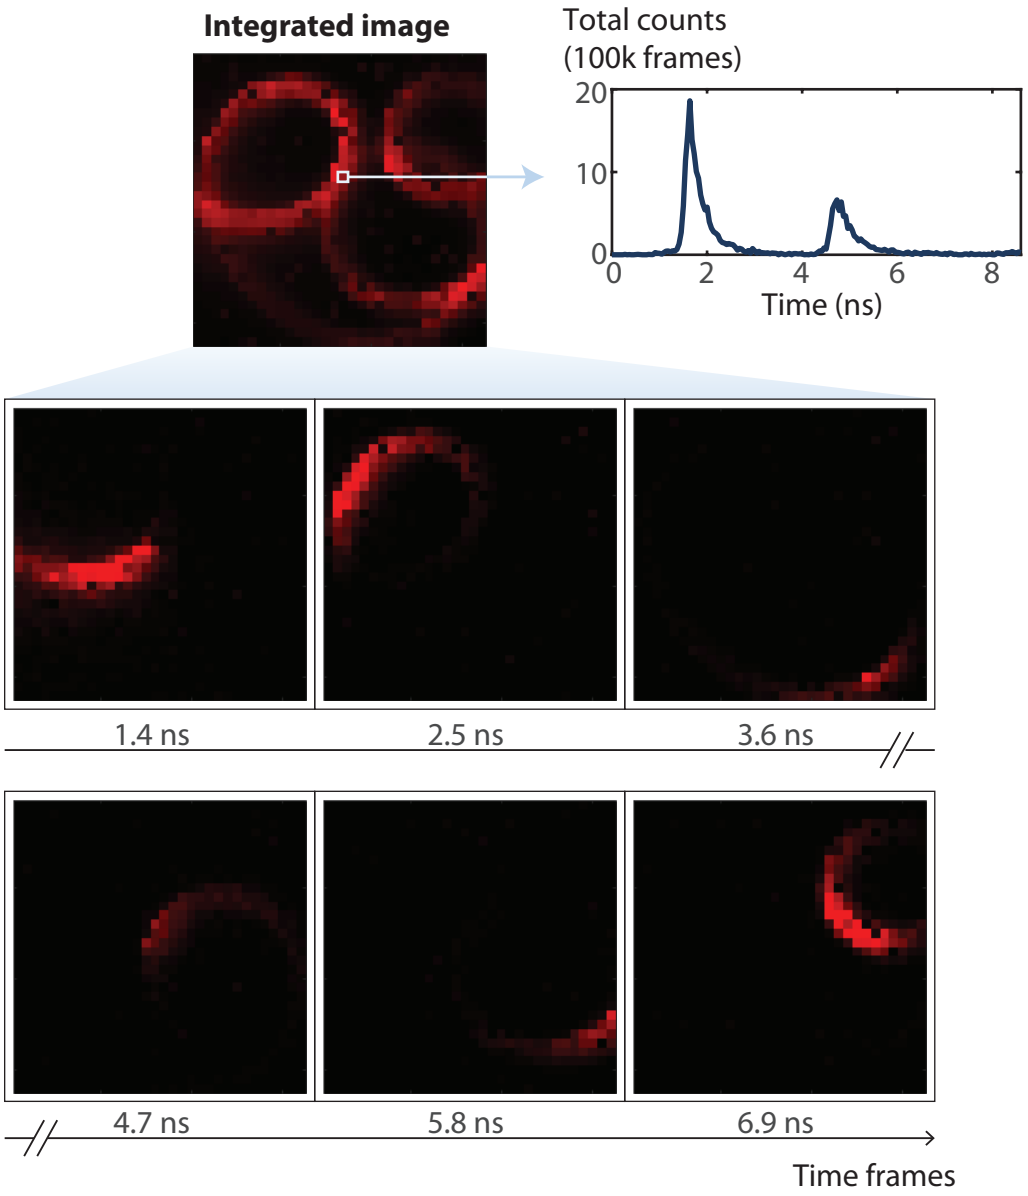

FIGURE 4.13: **Light-in-flight measurement of laser pulses propagating in an optical fibre.** Light can be observed as it propagates through an optical fibre by detecting the light that scatters out of the fibre during propagation. Here, we observe a laser pulse as it travels along the fibre and follows a series of loops.

ation in a photonic crystal fibre and the propagation of pulses at extremely high repetition rate (1 GHz), well above the 100 MHz limit of operation of the SPAD camera [8].

4.4 Slow light

A different medium of interest for light-in-flight measurements is a slow light medium, where a pulse of light can slow down to a velocity orders of magnitude slower than the speed of light in air c. We used our 45.5 ps bin SPAD camera to study the dynamics of a nanosecond-long 785 nm laser pulse as it propagates into such a medium and slows down to a speed of $c/150$. These experiments were performed by Dr Kali Wilson, from

the Extreme Light Group at Heriot-Watt University, together with the Howell Research Group at the University of Rochester, where this group have an experimental setup for slowing light in a heated rubidium cell.

The phenomenon of slow light is not due to a dramatic decrease in the phase velocity of light in a medium, but rather to a decrease of its group velocity v_g. The group velocity of a pulse indicates the speed at which the pulse envelope moves through the material, and is inversely proportional to the derivative of the index of refraction n of the material with respect to the wavelength of the pulse λ [63]:

$$v_g \approx \frac{c}{\omega(dn/d\lambda)}. \tag{4.1}$$

To obtain very low values of v_g, one therefore needs a medium with a large change in refractive index over a short wavelength range. There are many ways to obtain a slow light medium [63, 70], and the method used here is to propagate a 785 nm pulse through a hot atomic rubidium vapour, where the pulse's wavelength is situated between two absorption resonances [71, 72]. The index of refraction near resonance varies rapidly, giving a high value of $dn/d\lambda$, hence a slow group velocity. The dispersion can be tuned by changing the temperature of the rubidium vapour – in this experiment, it was set at various temperatures up to 150°C. As the temperature increases, the group velocity at 785 nm decreases. The propagation of 1 ns pulses of 1 pJ energy, emitted at 25 MHz, was observed in rubidium cells of 7 cm length and 30 cm length using the SPAD camera, which imaged a 5x5 cm section of the cell (using a 18 mm camera objective with aperture f/1.8). To obtain a larger field of view, the camera was translated by steps of 4 cm along the axis of the pulse propagation so that the recorded movie covered the whole length of the rubidium cell used. At each position of the camera, we acquired data over 100 000 frames with 500 μs exposure time.

In figure 4.14, we show selected frames of the movies recorded in a 7 cm rubidium cell heated at 130°C and a 30 cm cell heated at 105°C. For comparison, a light-in-flight measurement was acquired in a $n = 1$ medium – since the pulses have four orders of magnitude less energy than the ones used to produce the light-in-flight movie in air discussed in section 4.1.3, and are spread over a longer duration, a scattering medium was here required to perform a measurement. In this case, smoke was directed in the path of the laser to enhance the scattered signal to the SPAD camera. From the measurement recorded in air, we observe a straightforward path from left to right - the pulse is seen exiting the 55 cm long field of view at frame 4.3 ns in figure 4.14. In the slow light measurements, the pulses are seen entering the rubidium cells around frame 1.9 ns and propagating at different speeds. A last frame from a much later time (21 ns) shows that the pulse in air is long gone, while the pulse in the rubidium at 105°C is still travelling in the 30 cm long cell and the pulse in the rubidium at 150°C is just exiting the 7 cm long cell. We also

observe in these frames the difference in spatial compression of the 1 ns pulse due to the different group index of the medium at different temperature. The pulse in air spans over 30 cm, and is compressed to less than a centimetre in the cell heated at 130°C.

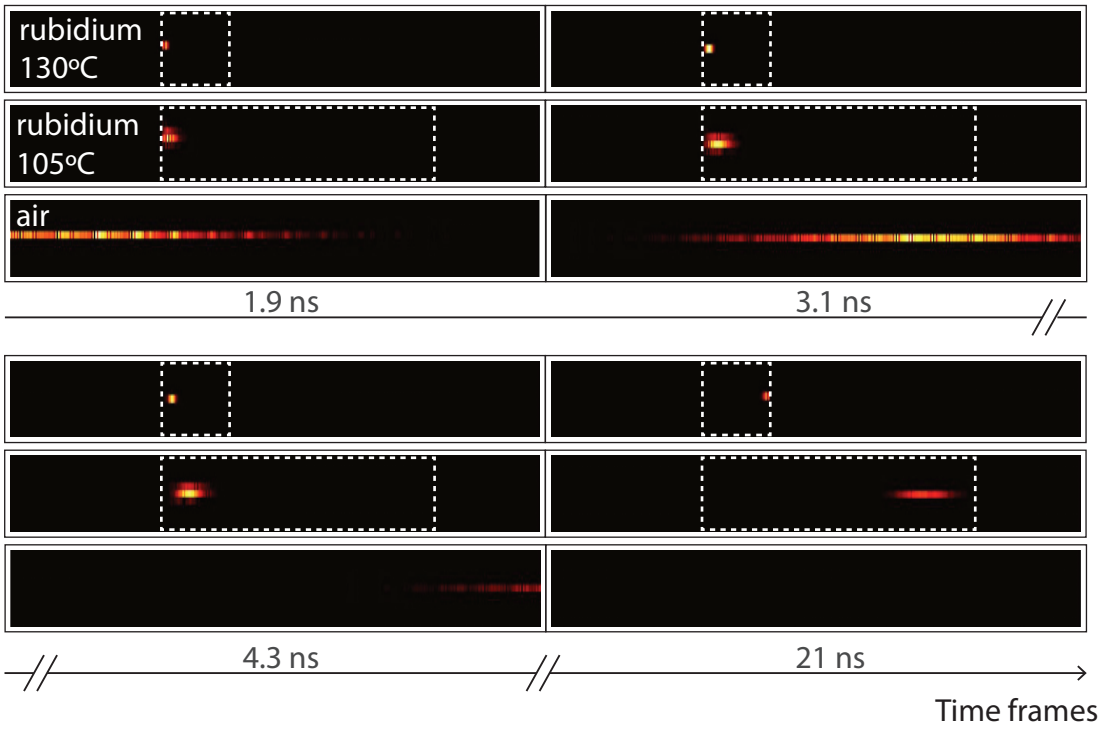

FIGURE 4.14: **Light-in-flight measurement of slow light in rubidium.** We record the propagation of a 1 ns long pulse through a rubidium cell heated at different temperatures and in air. The cells are indicated by the dashed white rectangles. The acquired data was processed by subtracting a background and performing a Gaussian fit on each recorded histogram. We can observe in the selected frames a significantly slower evolution of the laser pulse in the rubidium cell compared to the propagation in air, as well as a strong compression of the pulse due to the slower group velocity.

The group velocity of the light inside the rubidium cell at different temperatures T can also be measured from this data by extracting the position of the pulse at each time frames recorded (the time of arrival at a given position is given by the peak position of the Gaussian fit of that pixel's histogram). In figure 4.15, we show this position versus time as measured for $T \sim 150°$ and $T \sim 105°$. We observe as expected a linear relation which changes with temperature. The slope of the linear fit gives the group velocity of the pulse, which is $c/40$ at $T = 105°$ and slows down to $c/150$ at $T = 150°$. We can also observe that the linear fit does not cross the position 0 cm at a time of 0 ns, even though the time histograms were calibrated so that time 0 ns corresponds to the time the pulses hit the entrance face of the rubidium cell. In fact, the delay at which we observe the pulse first entering the rubidium cell increases with temperature. This is due to the fact that, as mentioned in chapter 2, the path taken by the detected photons between the scene and the camera is significant in light-in-flight measurements. Here, photons that are scattered from the pulses by the rubidium in the cell need to first travel out of

the cell before reaching the camera. The pulse propagating at the centre of a 2.35 cm diameter cell, the scattered light travels 1.175 cm in rubidium, at the same group velocity as the pulse, before exiting the cell and reaching the camera. This explains why there is a delay between the pulse entering the cell and the scattered light reaching the camera and being detected as positioned at the very beginning of the cell. This particularity in our measurement confirms that slow light is a phenomenon that not only applies to a laser pulse, but also to single photons [71]. In addition to the measurements reported above, covering my contribution to the experimental design and data analysis, a more complete and thorough analysis of all the particularities of this slow light measurement was performed by Dr Kali Wilson and is detailed in [12]. This experiment demonstrates

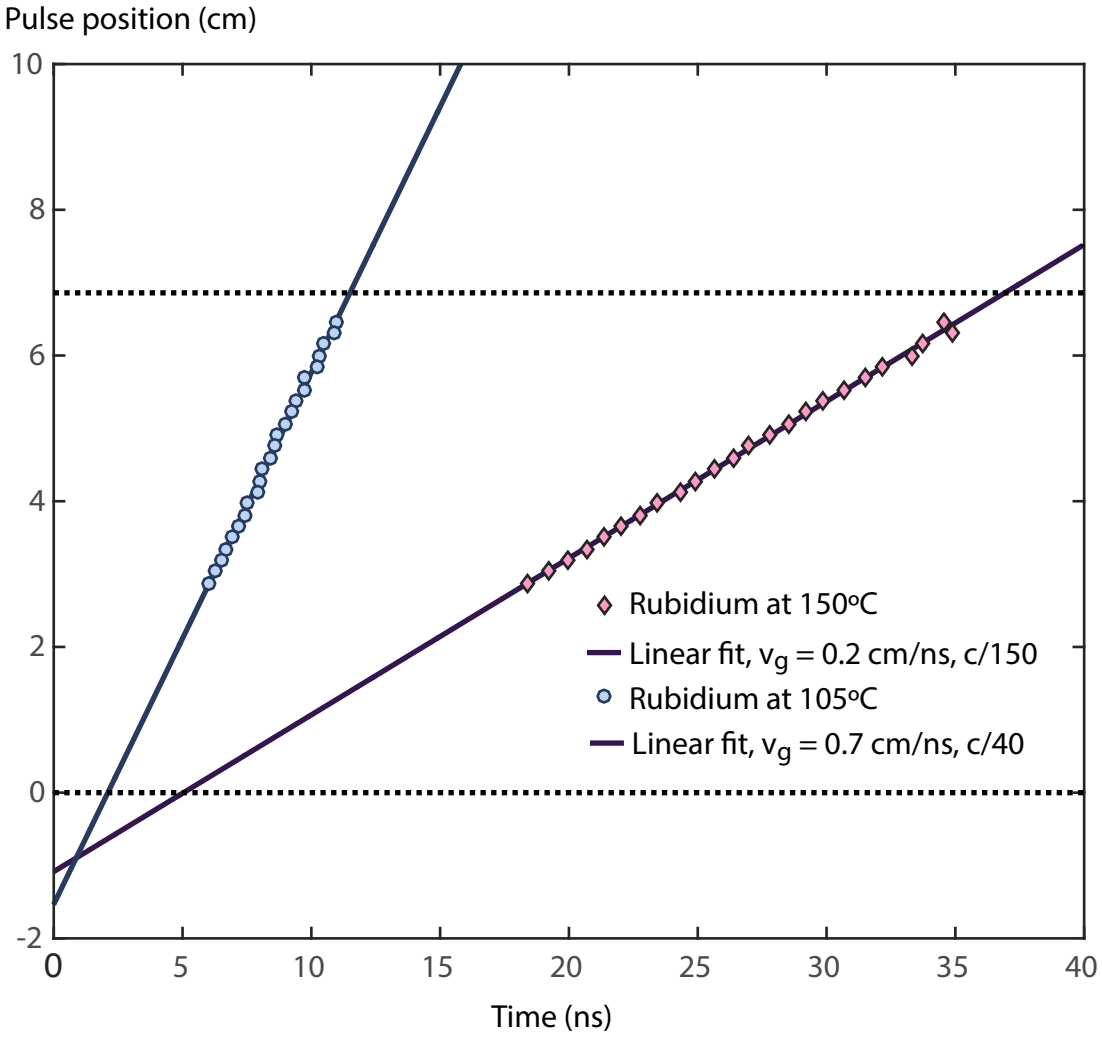

FIGURE 4.15: **Measurement of group index in slow light medium.** The recorded data allows us to extract the time of arrival of the pulse at each position along propagation in a 7 cm long rubidium cell, heated here at $\sim$ 150°C and 105°. The linear fit gives the group velocity of the pulse inside the rubidium, in centimetre per second. We also observe that the measurement at the entrance in the rubidium cell is delayed by a few nanoseconds, as we expect the pulse to enter the cell at time 0 ns. This is due to the time delay imparted upon the photons that scattered out of the rubidium cell to reach the camera and be detected: they also travel at a speed much slower than c, so they are delayed compared to a measurement where they would travel to the detector at the speed of light in air.

the possibility to study slow light as a dynamical phenomenon inside the medium, rather than only measuring its total effect at the output of the medium as slow-light experiments are generally done [73].

4.5 Imaging light through an opaque object

In this section, we discuss a type of experiment slightly different than the ones described above, where we were observing phenomena that were evolving on a plane perpendicular to the field of view of the camera, at a stand-off distance. Here, we are interested in assessing what type of signal we can detect when short pulses are sent towards the camera, but passing through an opaque object beforehand, as illustrated in figure 4.16. In such a configuration, the signal detected may carry information about the opaque object itself, or about a structure that is hidden behind or inside the object. The experiments reported here were performed during a visit in the Camera Culture Group at MIT Media Lab.

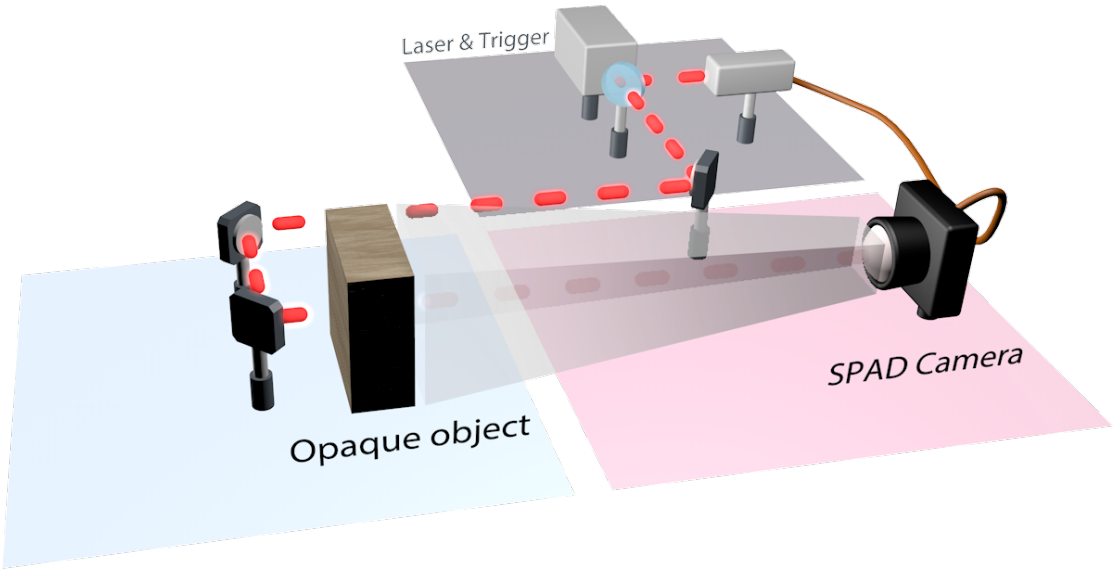

FIGURE 4.16: **Capturing a signal passing through an opaque object.** To gather information about an opaque object or what is hidden behind/inside it, we send short pulses directly through the object and towards the camera. Although the object is opaque, some photons will be transmitted through the object and be detected by the SPAD camera.

An object may be opaque due to two distinct types of interaction: absorption and scattering. In most objects, both interactions exist, but one of them dominates. We first assessed the possibility of imaging through a common opaque absorptive layer such as wood. In a rudimentary experiment, we sent femtosecond pulses at 80 MHz carrying an average power of 950 mW through a 9.1 mm thick pile of thin wooden pieces, as shown in figure 4.17a. We acquired data over 50 000 frames at 1000 μs exposure and, although the layer of wood is visibly opaque, we detected a signal from the transmitted laser beam

at a rate of 0.03 photons/frame (which corresponds to an estimated optical density of $OD = 12$). In figure 4.17b, we show a frame of the acquired data where the beam is clearly visible. Figure 4.17c shows the histogram of the pixel located in the centre of the beam. As the object is mainly absorptive, the beam detected has not significantly spread over time and the detected signal still shows a sharp peak at the time the beam passes through the wooden layer. This shows that it might be possible to use the camera to image through opaque absorptive objects – this idea has not been pursued further than this first demonstration during this PhD, but further experiments are underway to evaluate the possibility to image structure inside media such as wood and assess possible applications of this ability.

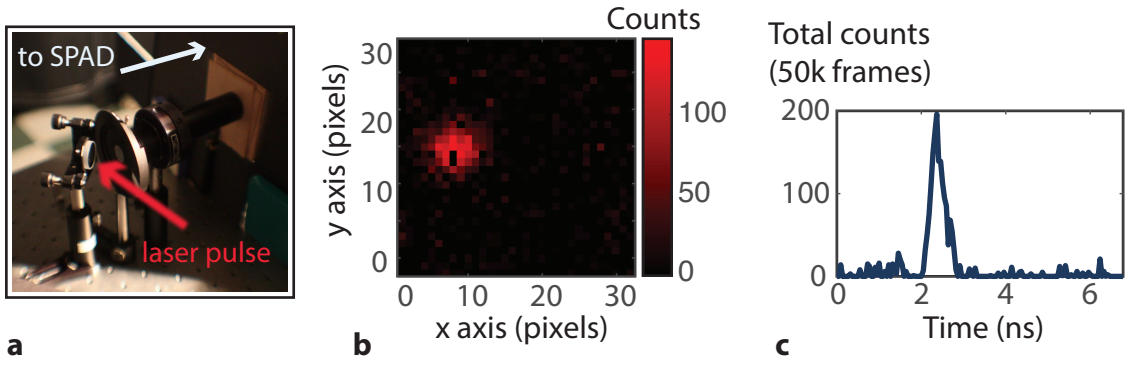

FIGURE 4.17: **Capturing a signal passing through a thick layer of wood.** a) When sending 950 mW of average power in the form of femtosecond pulses at 80 MHz through a 9.1 mm thick pile of wooden pieces, we still detect a signal transmitted using the SPAD camera. b) The beam is visible in the time frames acquired by the SPAD camera and c) the recorded histograms show a sharp peak in the signal. This ability to see through opaque material could be used to image structure through absorptive media.

The behaviour of light passing through a scattering object is very different than that observed in an absorptive object such as wood. We expect a strong spread of the detected signal over time as, when the light travels in a scattering opaque medium, it undergoes multiple scattering events in all directions and light travels over complex paths before exiting the medium. The resulting beam of light exiting the object is therefore composed of light having taken wildly different paths, leading to different times of exit out of the object. A first measurement in a scattering object was performed using a series of square plastic pieces, as shown in figure 4.18. To evaluate the possibility of detecting a hidden structure through the scattering layers, we placed a cross made out of aluminum foil in the middle of the plastic pile, so that the structure is hidden from view from both sides of the object. We then sent the laser pulses at the back of the object, using a diverging lens to illuminate a broad surface at the back of the plastic squares rather than a single point. We detected the signal transmitted by acquiring 100 000 frames of 500 μs exposure time, at a rate of 0.05 photons/frame (this object has therefore a similar optical density as the opaque layer of wood detailed above). When integrating over all temporal frames, we obtain the

profile shown in figure 4.18. We can see that the structure inside the plastic object is not totally hidden: we can detect even in the integrated image the shape of the cross, as well as some diagonal structure which comes from the design of the plastic pieces themselves.

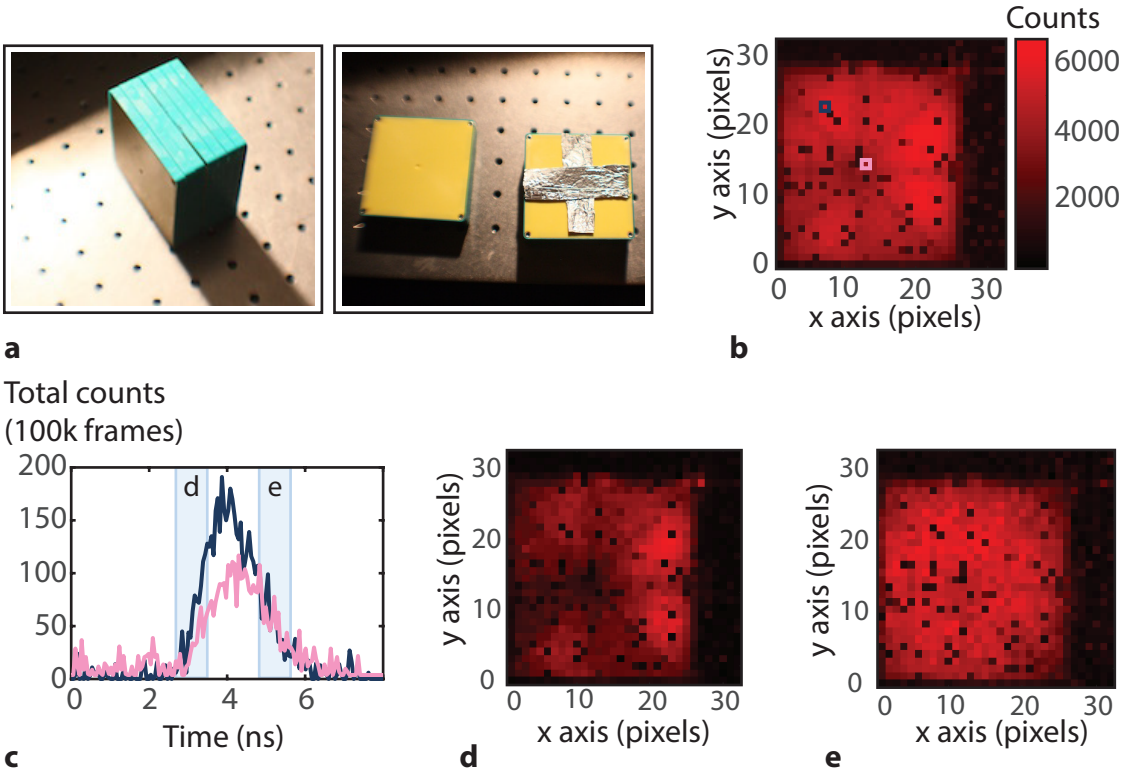

FIGURE 4.18: **Capturing a signal passing through a thick layer of plastic.** a) We send laser pulses through a thick layer of plastic pieces, hiding a aluminium cross in the middle of the layer. b) In the time-integrated image, the hidden structure is still visible, although with very low contrast. c) Histograms at two positions in the spatial profile are taken - the positions are indicated in (b) with a pink square at the centre of the cross and a blue square located in the top left quadrant. Both histograms show a signal spread over a few nanoseconds, due to the multiple scattering of the femtosecond pulse inside the plastic layer. The sections indicated by blue rectangles in (c) correspond to the temporal region over which we integrate to obtain the images in (d) and (e). In (d), we observe that the structure hidden in the plastic is clear when we integrate only over the photons that arrived at earlier times, hence followed a mostly straightforward trajectory in the plastic layer. e) The time-integrated image integrated over later arrival times shows no visible structure.

In figure 4.18c, we show two histograms taken at different places in the structure imaged: one at the centre of the cross and the other in the middle of the top left quadrant of the plastic square. Both signals are much wider than the femtosecond pulse that was sent through the plastic: they span over a little more than 3 ns, which corresponds to a flight path of about one meter, even though the plastic layer was only about 5 cm. This expansion is due to extended path taken by photons deviated by multiple scattering events inside the medium. We also observe clear distinctions between the two histograms. The one taken in the top left quadrant shows a strong signal earlier in time compared to the one taken in the centre: this is due to the fact that the paths that travel mostly in the forward direction are blocked in the centre by the aluminium cross. A very rich field of study, aimed

at gathering information through a scattering medium by observing selectively only the ballistic photons or the early photons (which have a mostly forward trajectory inside the medium) is the field of ballistic imaging [74]. In commonly used techniques of ballistic imaging, one can typically discriminate the early photons transmitted through a medium with a resolution of the order of 100 fs [74, 75]. Here, we have orders of magnitude less resolution in the selection of ballistic photons, but it is still interesting to assess the differences in the spatial information gathered in the early-arriving photons and the later ones. In figure 4.18d, we show the image integrated over the first 700 ps of the signal. We obtain a slightly clearer image of the structure than with the integrated image, as the wide cross becomes more visible. In contrast, when integrated the frames at later times, all the information about the structure is lost, as the photons arrived at those times have followed complex trajectory and blurred out any information about the internal structure of the observed layer. Even though the SPAD camera, in light of those results, is not an ideal tool for ballistic imaging at those scales, we can still detect rich information in the temporal profiles recorded. This rudimentary experiment therefore hints that the SPAD camera could be used to perform imaging through scattering medium, but would require a deeper analysis of the detected signal than simply considering spatial profiles at different time of arrivals.

The ability to image through scattering medium is of great interest to the biomedical imaging community, notably to image inside the human body, as we will discuss below. We report here on a first experiment performed with the SPAD camera to investigate the possibility to image through the body and recover information about its internal structure. We set up an imaging system to look through the hand and forearm by sending an extended laser beam to the back of the region of interest and imaging the transmitted light, similarly to the setup illustrated in figure 4.16 where the opaque object is now a region of the hand or arm. This configuration allows to directly image the structure of the transmitted light and gives the opportunity to visualise some structure in the arm without any complex retrieval or inverse problem algorithm.

In order to increase our spatial resolution for this measurement, limited by the 32x32 pixels of the detector, we installed the SPAD camera on a motorised stage (Zaber A LSQ075B; 75 mm travel range with 0.5 μm microstep size). Since the active area on each pixel is limited to a circle of 8 μm diameter on a 50x50 μm pixel, we can scan the position of the camera by steps smaller than the pixel pitch in a way that the active areas image sections of the field of view that do not overlap at every step of the scan. This way, we can acquire a higher spatial resolution dataset simply by acquiring data at different stage positions and combining all the datasets – there is no need for advanced processing as all the datasets contain distinct information that do not overlap with each other. An example of the increased resolution that can be obtained is showed in appendix B. Here, the SPAD camera was imaging a field of view of 5x5 cm so that each pixel imaged a

1.56x1.56 mm region - we therefore scanned the camera by steps of 312.5 μm to cover the area of a pixel with the active area in a 5x5 grid, leading to a total spatial resolution of 160x160 pixels rather than 32x32 pixels.

We performed four measurements on different regions of the arm: in the hand, the wrist and two adjacent regions of the forearm. For each acquisition we moved both the field-of-view of the camera and the illumination point on the arm, so that the illumination was always in straight line with the SPAD camera. The laser used for this experiment is a 80 MHz femtosecond laser, illuminating a 2.5 cm diameter region on the arm with an average power on the order of 10 mW/cm^2. At each position of the camera, we acquired 50 400 frames of 300 μs exposure time, so that data was collected for about 7 minutes on each of the four sections of the arm imaged.

In figure 4.19, we show selected frames of the resulting movie (see appendix A for the full video). We can clearly see the outline of the arm as the light propagates through it, and can even distinguish some structures. The observed contrast is likely to be originating from the different optical properties of the main components of the body such as bones, veins and arteries. At 800 nm, bones are the main scattering component in the body ($\mu_s = 1.8$ mm^{-1}, $\mu_a = 0.027$ mm^{-1}), while blood is strongly absorptive ($\mu_s = 1$ mm^{-1}, $\mu_a = 0.398$ mm^{-1}) [76].

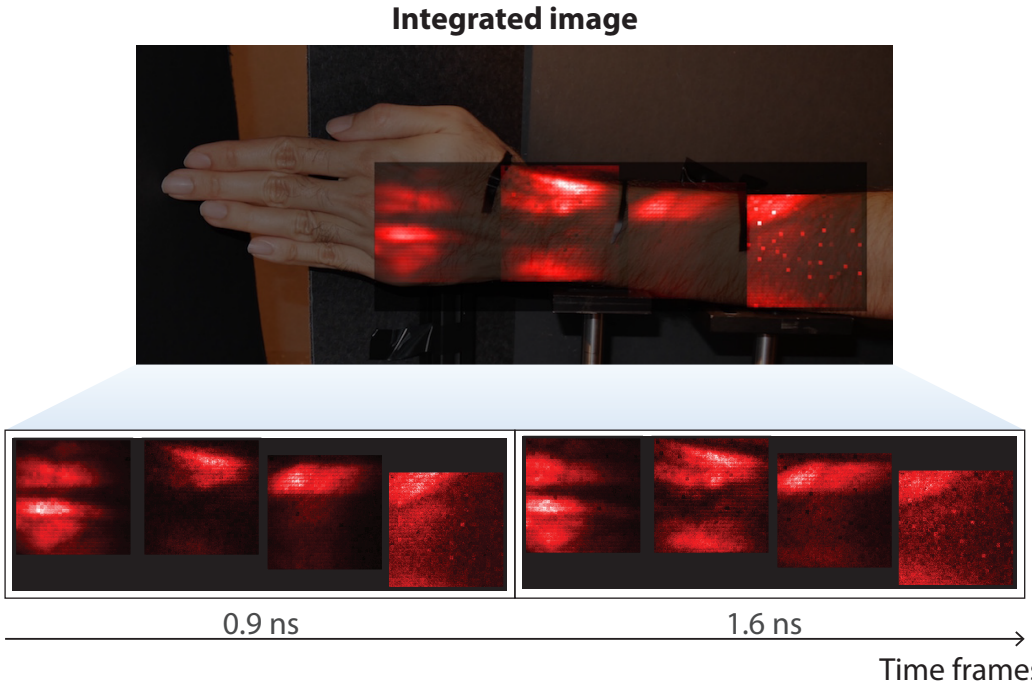

FIGURE 4.19: **Imaging through a hand and forearm.** We illuminate successively four sections of the hand and forearm from the back and image the transmitted light with the SPAD camera. The integrated image is overlaid with a picture of the arm as it was positioned in the imaging setup. We can observe some structure in the acquired images, showing darker and brighter regions in the transmitted light. This contrast is likely to be due to the presence of bones and veins in the imaged regions of the arm and hand.

As mentioned above, imaging through biological scattering media is a rich and very active field of study. We only brush this subject here as the full analysis of the data presented in figure 4.19 was not approached in the context of this thesis; it requires a deeper study to evaluate exactly what information we can gather about the internal structure of the arm from such a measurement. The development of an imaging technique in scattering media using the SPAD camera is a significant undertaking – following the work presented here, this project has been continued by other members of the Extreme Light Group and is currently underway to design an imaging and retrieval framework that fits the imaging capability of the SPAD camera. Imaging through scattering media such as biological tissue has received a lot of interest in the past two decades, and many techniques have been developed to face this challenge. An exhaustive review of the literature on this subject is out of the scope of this thesis, as many successful techniques with different aims have been developed in recent years to image through both thick and thin opaque scattering objects [77–82]. We give here a brief overview of the most widespread technique to image through a thick scattering medium, which shows the most resemblance with the experiments presented here and has been successfully used to image through the human body: diffuse optical tomography (DOT) [83, 84].

In a DOT system, one illuminates a scattering object from different positions and angles with a near-infrared laser (800 nm). The light then scatters in the medium, following multiple paths as it undergoes scattering and absorption interactions with the medium, and the light exiting the object is measured at different positions and angles. The distribution of output light versus input light is then analysed to reconstruct the scattering and absorption properties of the object under study. In a common configuration, only the outgoing intensity is detected, but an implementation richer in information is based on the measurement of the temporal profile of transmitted scattered light, similarly to the measurement we perform with the SPAD camera [85]. This technique does not offer a direct visualisation of the hidden structure: a complex forward model must be developed in order to simulate the data obtained with a DOT measurement. Usually, a uniform scattering medium is modelled with located intrusions having different scattering and absorption coefficients. From the measured data, one then needs to solve the ill-posed inverse problem to retrieve the structure of the imaged object. Even though this is a complex problem, DOT has proven a very successful tool to image with centimetre precision inside the human body, notably for imaging the breast and the brain [84]. Time-resolved DOT has also been used to image the human forearm, where circular cross sections were imaged and the retrieval shows the position of the main bone and artery components in the arm [86]. This imaging technique allows to image deeper in human tissue than other imaging techniques based on the detection of ballistic photons, reaching a depth of around 6 cm [83]. Recently, work has also been done to incorporate SPAD arrays in a DOT setup, exploiting the 2D imaging capability to acquire data from one detector position rather than placing multiple single pixel detectors around an object [87]. Many avenues can still be explored to best make

use of the SPAD camera to image through a thick scattering object, including exploring different imaging configurations and image retrieval algorithms.

In this chapter, we have covered various implementations of light-in-flight measurements that can be performed using a SPAD camera. This detector has the advantage of providing an easy and direct way to measure dynamics at very low light intensity with a resolution of about 100 ps. In addition to providing the first direct measurement of light propagating in air, we also studied more complex light phenomena such as plasma formation, propagation in fibre, slow light and propagation in opaque scattering media. One of the main drawback of the SPAD camera is its low spatial resolution (32x32 pixels), but in most cases a simple interpolation can provide a satisfactory visualisation. In the case where increased resolution is needed, it is straightforward to scan the SPAD camera and stack the measurements together directly. We have detailed in this chapter many advances in light-in-flight imaging measurements; in the following chapter, we will focus on one of the main application of imaging light in motion, that has practical applications in real-life scenarios: the detection and tracking of objects hidden from view.

Chapter 5 – Application of imaging light in motion: detecting and tracking hidden objects

In the previous chapter, we have discussed different implementations of light-in-flight measurements that allowed us to observe and study different fundamental phenomena evolving at the speed of light. A different application of observing light in motion is the ability to infer information about objects hidden outside the line of sight, which still scatter low intensity signals that can be detected and analysed, as we hinted in section 4.1.3. Hidden object detection is an exciting and promising avenue of research, as the ability to detect objects hidden from view could prove a great asset in a variety of real-life scenarios, such as search-and-rescue missions, surveillance and autonomous driving to detect incoming vehicles from around a blind corner. Given the extent of the challenge and applications linked to the ability to detect objects outside the line-of-sight, the detection and imaging of hidden objects has become in recent years a rich and active field of research on its own. In this chapter, we will therefore give an overview of the existing methods developed in this field before detailing our contributions to the state of the art. As we will see, some of the techniques aimed at detecting hidden objects are based on the same technologies as those aimed at imaging light-in-flight, so that some techniques reported in chapter 2 also find applications in hidden object detection. There is also another existing class of methods that are not based on time-of-flight measurements but purely on intensity detection. We will detail these different techniques in section 5.1.

As we will see, most existing methods are aimed at imaging objects hidden from view. Although impressive 3D reconstructions of a hidden object have been achieved with these techniques, an important challenge that remains to be faced is the ability to acquire information about a hidden scene in real-time, i.e. fast enough to track the motion of a hidden object. This is the challenge we set out to tackle in this thesis, and we developed a method based on the use of the SPAD camera to locate the position of a hidden object in only a few seconds acquisition time, without trying to retrieve the 3D structure of the hidden scene. Limiting ourselves to recovering the position allows us to greatly improve the speed at which data can be acquired and analysed and in many scenarios, the knowledge about wether or not one or multiple objects are present and moving in a hidden scene, as well as their position and speed, is a crucial asset and the exact structure of the hidden object might not necessarily be an essential information. The technique we developed was used in the first demonstration of a system that can track the motion of an object moving around a corner. We will detail the position retrieval method we designed for this task in section

5.2 and report the proof-of-principle experiment performed to demonstrate our method in section 5.3. We will then treat a slightly different implementation of our method, based on single pixel SPADs rather than the SPAD array, that allowed us to locate people hidden around a corner using a source and detector at a standoff distance of 50 meters. The work presented here paves the way to the development of an apparatus that could be used in real-life scenarios.

5.1 Prior art

A lot of interest has been shown in the last decade to develop optical methods to look around a corner. Methods to detect hidden objects based on radar technology already exist, but are limited to situations where one can look through a material that is transparent to radio waves [88] or using direct reflections off a surface in the line-of-sight to detect static or moving objects hidden by an obstacle [89, 90]. Methods based on the use of visible or infrared wavelength laser could therefore provide increased flexibility and resolution for use in various scenarios.

A simple configuration that allows to send a signal around an obstacle and detect a signal coming back is to use a time-of-flight imaging system to image a wall or surface in its line-of-sight, and send short laser pulses on that same surface. When hitting a scattering surface, the laser pulses are scattered in every direction so that some light reaches a scene that is hidden from view by an obstacle. Any object in the scene will then scatter light back onto the surface imaged by the system, and this signal can be used to infer information about the hidden object. This concept can be viewed as an extension of LIDAR technology, where short pulses are sent at an object in line-of-sight and a time-of-flight measurement gives the information about the distance of the measured object. For looking at objects outside the line-of-sight, time of flight information is still invaluable and additional information such as the spatial distribution of the scattered light is needed to pinpoint the hidden object's location. This concept is the basis of all time-of-flight based methods to detect hidden objects that have been developed so far, and we will use a version of this in the work presented below - the geometry described above is illustrated in figure 5.2.

Some early theoretical framework for the reconstruction of hidden objects based on transient imaging has been laid down from 2009 [91–93], and a method using a streak camera to record time-of-flight information has shown capable of reconstructing images that were hidden from view around a corner in a known plane [94]. In 2012, Velten *et al.* demonstrated a method that showed impressive 3D reconstructions of an object hidden around a corner. In their experiment, they used a streak camera to image a line on a wall in their line-of-sight and shone a femtosecond laser at different spots on the wall by scanning

the laser beam. The data acquired by the streak camera was then used in a back-projection algorithm [95,96]: for each set of laser position and pixel on the streak camera correspond a group of positions in the hidden scene that is located at the adequate distance to produce such a signal on that given pixel. All possible contributions to the detected signal were added and processed (high-pass filtering and thresholding) to obtain a visualisation of a 20 cm high mannequin where its main features were clearly visible [95]. This work showed that it was possible to retrieve 3D information on a hidden object located about a meter away from the detection system.

Further work was done following this demonstration to overcome two of the limitations of the streak camera technique: its cost and form factor, and its necessity for scanning that leads to long acquisition times. The acquisition for the reconstruction reported in [95] was on the order of half an hour. In particular, following the same measurement concept than that presented in [95], measurements based on a single-pixel SPAD and a scanning laser spot were presented in [97] in 2015. SPAD sensors are more practical to use than a streak camera and have the potential for mass-market applications thanks to their smaller size and cost. Three-dimensional information retrieval of a hidden scene has been demonstrated using a SPAD detector, however with lower resolution than what was obtained with the streak camera, as the SPAD temporal resolution is about an order of magnitude lower than that of a streak camera (250 ps versus 15 ps). On the other hand, the high sensitivity of the SPAD detector allowed the use of a lower power laser by about an order of magnitude (50 mW) and the data could be acquired in five minutes rather than half an hour. A different implementation also published in 2015 demonstrated the possibility to perform such a measurement in the short infrared wavelength range (1.545 μm) using an InGaAs SPAD array [98], where the position of a hidden object was retrieved with a precision of about 10 cm.

Another inexpensive method to measure time-of-flight, as discussed in chapter 2, is to use photonic mixer devices (PMD), a type of indirect time-of-flight camera, in combination with periodically modulated illumination. This detection system can be used to perform transient imaging measurements, so that scattered light containing information from an object outside the camera's field of view can also be detected and exploited. In [99], Heide *et al.* report on a measurement acquired in 5 minutes where the modulated light source was focused on a wall in the line of sight of the PMD and light scattered from a hidden scene, letters hidden by a wall about 3 meters away from the camera, was acquired. As in the transient imaging scheme, the retrieval of the hidden scene geometry is done through the solution of a linear inverse problem.

In figure 5.1, we show a classification of the methods to detect hidden objects and, as illustrated, the methods described above are all based on time-of-flight measurements. A few interesting methods have also been developed to image hidden objects without the

need for temporal measurements, in the aim to exploit only readily available cameras to detect information about a hidden scene. Such a technique was recently published in 2016 and showed that, given an illumination and detection geometry that is close enough to the hidden object, one can simply measure the intensity distribution of the scattered light on the wall rather than its full temporal evolution. Indeed, the scattered light from an object will have an intensity profile that is spherical with its centre coinciding with the location of the object. In [100], Klein *et al.* showed that it is possible to retrieve the position and angle of an object of known shape using an iterative algorithm that compares a simulated intensity measurement of a guessed scene with the actual measurement. If the object shape is not know, position retrieval that is still consistent with the object's actual location was obtained, and this technique also proved fast enough to allow tracking of the hidden object.

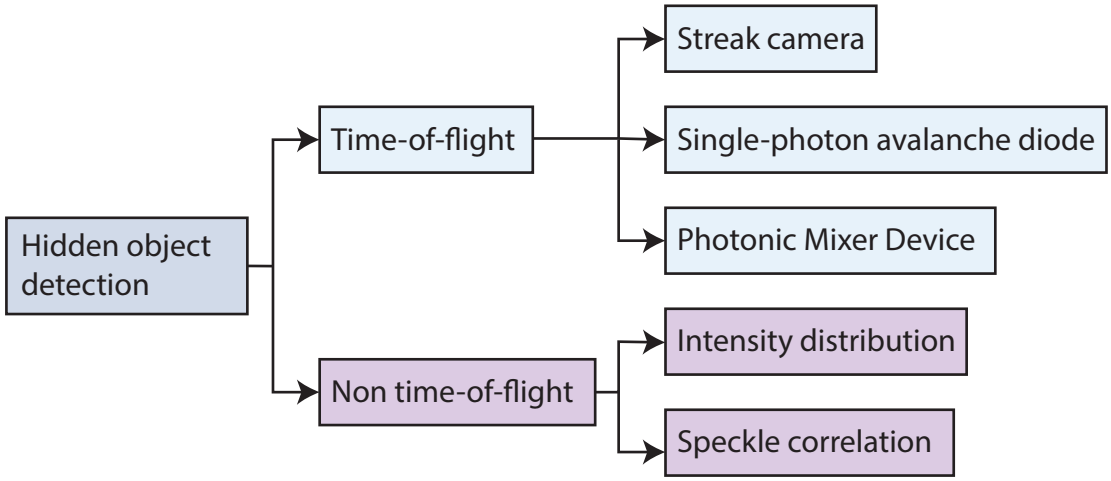

FIGURE 5.1: **Methods for detecting hidden objects.** The many techniques that have been developed to retrieve information about a hidden scene can be classified in two categories: those based on time-of-flight measurements and those which do not rely on acquiring temporal information but only intensity measurements. Time-of-flight techniques have been developed using different sensors such as streak cameras, single-photon avalanche diodes and photonic mixer devices. On the other hand, techniques using simple 2D cameras have shown capable of retrieving information about a hidden scene by analysing either the intensity distribution of a scattered signal from a hidden object, or its speckle pattern as the light from the object hits a diffuse wall.

Another class of detection techniques is based on the measurement of a speckle pattern. Scattering by either a thin scattering layer or a diffuse wall produces a speckle pattern that may seem random but still contains information about the object that is located behind the layer or around a corner. One way to exploit this information is to use adaptive optics to compensate for the effect of the diffusing layer or wall and retrieve the image of the hidden object. This wavefront shaping technique using a spatial light modulator has been shown to work for imaging objects through an opaque scattering layer [101], and later to image objects around a corner [79]. However, this technique requires a calibration measurement to find the correct wavefront correction given a particular scattering layer or wall. To

avoid the need for this extra step, which requires access to the hidden area that cannot be accessed in many scenarios, a technique was developed based on the property of speckle correlations. When light hits a scattering surface at a given point, it will produce a given speckle pattern. If the angle of incidence of the light is changed within a certain limit, this pattern will not change dramatically, but rather simply be translated - this is known as the memory effect. It has been exploited to image through an opaque scattering layer with no access to the hidden area [78], and Katz *et al.* demonstrated that the measurement of single speckle pattern provided enough information to look through a thin scattering layer, and around a corner using a diffuse wall [79]. A simple photograph of the speckle pattern can be taken, and its autocorrelation is in effect equal to the autocorrelation of the hidden object. In [79], the 2D image of an object hidden around a corner was retrieved using a simple camera phone. This method, simple and elegant, however has the disadvantage that the light source was placed on the same side of the corner as the object, and only the 2D image of the object was retrieved, and not its location in 3D space.

In light of the methods developed for looking around corners and their different advantages and range of applications, we set out to develop a technique that would be both compact and portable and provide information about the location of an object in a hidden area in real-time, so that the tracking of its motion would be possible. Our technique is based on time-of-flight measurements and is inspired by our work done in light-in-flight imaging, as we will detail in the following section.

5.2 Detection technique and position retrieval method

To explain our detection system and position retrieval algorithm, we consider the geometry of the first experiment we performed and detailed in section 5.3. Similar geometries can be treated in the same way. Here, we consider a situation where an object is hidden from the camera by an obstacle such as a wall, as illustrated in figure 5.2. In order to detect the object, we need to send a signal around the obstacle and detect a returning signal. To send light around the wall, we send short pulses of light on a scattering surface in the line-of-sight of the detection system - here we use the floor as a surface. When the pulses hit the floor, light is scattered in every direction so that it then propagates approximately as a spherical wave and reaches the hidden scene, as illustrated in figure 5.2c. If an object is hidden, it will scatter a signal back into the field of view of the camera, again as a spherical wave. In figure 5.2d, we show that a portion of this signal will be scattered back onto the floor by the object, crossing the surface area imaged by the SPAD camera.

The SPAD camera therefore captures the back-scattered signal from the hidden object. As this spherical wave intercepts the surface imaged by the camera, we record the corresponding section of the spherical wavefront as it sweeps across the field of view on

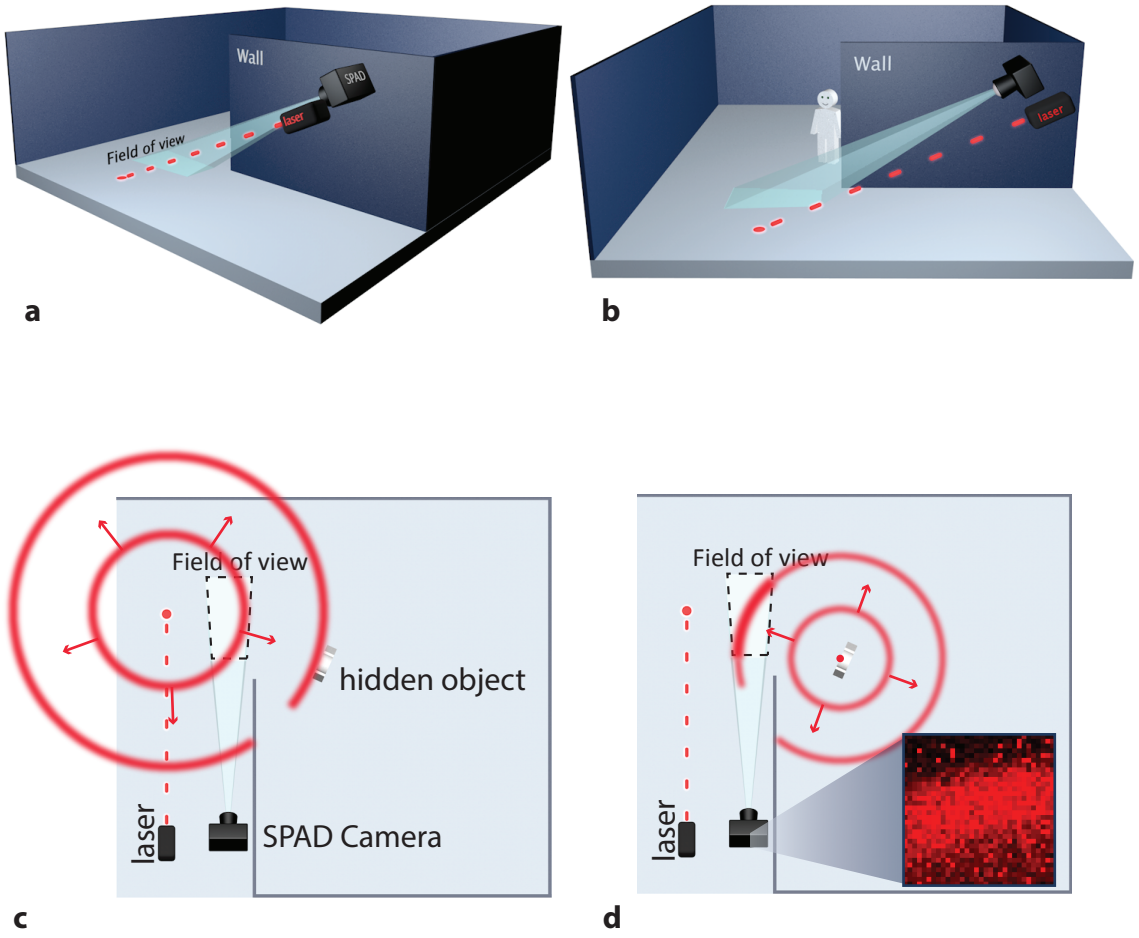

FIGURE 5.2: **Detection system to locate objects hidden from view.** a) The SPAD camera and an ultrafast laser are placed on one side of a wall. The camera is imaging an area of the floor beyond the edge of the wall. b) An object is hidden behind the wall, outside the line-of-sight of the camera. c) To send a signal in the hidden area, we direct short pulses of light at the floor beyond the edge of the wall. The pulses scatter on the floor and light reaches the hidden object. d) The object then scatters light back in every direction, so that a back-scattered signal reaches the surface of the floor imaged by the SPAD camera. The camera is used to capture the signal originating from the hidden object. Any signal coming from the object that has travelled a path with an additional scattering event will have a much lower intensity due to the spherical propagation, and is therefore insignificant with respect to the path illustrated here. The inset shows an example of a frame acquired where we can see a section of the spherical wave scattered back by the object as it intercepts the floor in the field of view of the SPAD camera.

the floor. We therefore record a 2D video of light coming back from the hidden object. This data contains two types of valuable information: temporal information in the form of a time-of-flight measurement at each pixel of the camera, and spatial information as we record the shape of the back-scattered wave as it intercepts the floor. The time-of-flight information, as in LIDAR measurement systems, gives information about how far the object is located. For a signal detected at time t, it means the light has propagated on a distance $d = c \times t$, where c is the speed of light. For a direct line-of-sight detection, this measurement gives directly the distance of the object. However, we consider here a case in which the light is following a trajectory where it scatters twice before reaching the imaged surface: once on the floor and once on the hidden object. There is therefore not

only one position in space where the object can be located to correspond to a measured time of flight t. In figure 5.3, we illustrate the case of a single pixel i of the SPAD camera, which images a position $\vec{r}_i$ in the field of view. The SPAD camera detects a signal coming back at time t_i at that position. That measured time of flight, when calibrated to take into account the flight path of the laser to the floor and of the signal to the camera, corresponds to the time it takes for a pulse to travel from the laser spot on the floor (at $\vec{r}_l$) to the hidden object at $\vec{r}_o$, and back onto pixel i. The total distance corresponding to the time-of-flight measurement is therefore

$$d = |\vec{r}_o - \vec{r}_l| + |\vec{r}_i - \vec{r}_o| = c \times t_i. \tag{5.1}$$

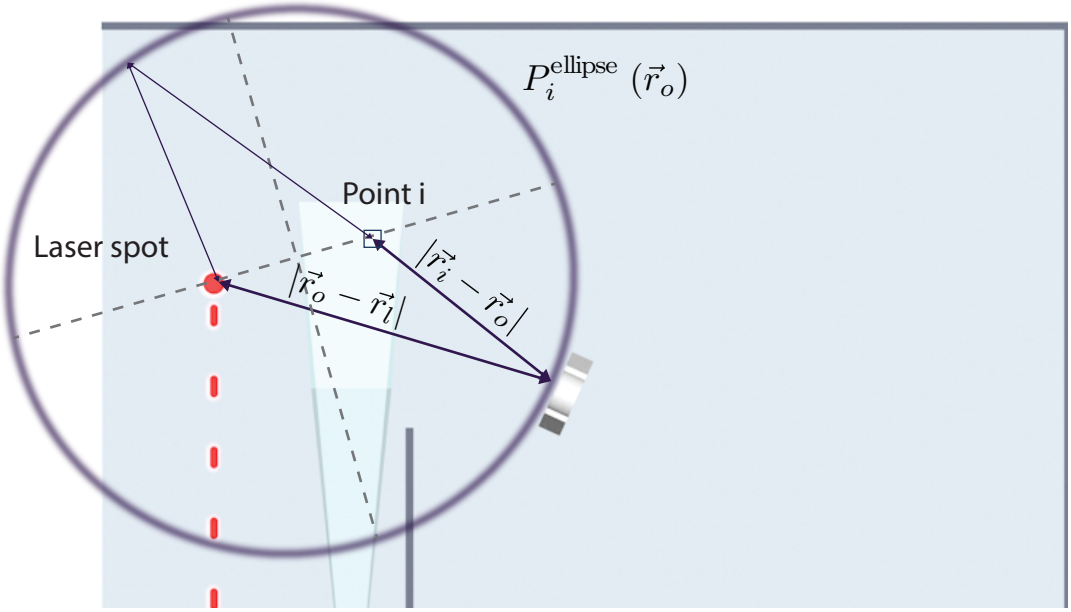

FIGURE 5.3: **Interpreting the time-of-flight measurements.** We measure a signal coming back at time t_i on a pixel imaging a position $\vec{r}_i$ on the floor. An object could be located anywhere along an ellipsoid P_i^{ellipse}, here illustrated as a two-dimensional ellipse in the plane of the floor, to produce the detected signal. The two foci of that ellipse are the laser spot and the pixel position. Because of the specific geometry of the experiment illustrated here, the major and minor axis (dashed line) have similar values so that the ellipse shape is near that of a circle. $P_i^{\text{ellipse}}(\vec{r}_o)$ is the probability density of finding the object at a position $\vec{r}_o$ in space given the measurement at pixel i.

This equation corresponds to an ellipsoid in three-dimensional space where the object can be positioned to give a signal coming back at time t_i on the pixel imaging a position $\vec{r}_i$. The probability density of finding an object at a position $\vec{r}_o$ in space given the detected signal of pixel i is then given by

$$P_i^{\text{ellipse}}(\vec{r}_o) \propto \begin{cases} 1 & \text{if} \quad |\vec{r}_o - \vec{r}_l| + |\vec{r}_i - \vec{r}_o| = c \times t_i \\ 0 & \text{otherwise.} \end{cases} \tag{5.2}$$

The projection of this ellipse in two-dimensional space is illustrated in figure 5.3. We do not write equation 5.2 as an equality here as we will consider the normalisation factor at the end of the development, in order to obtain a normalised probability density ($\int_S P_i(\vec{r}_o)d\vec{r}_o = 1$). The function above can be represented using the ellipsoidal coordinate ε and a Dirac function:

$$P_i^{\text{ellipse}}(\vec{r}_o) \propto \delta(\varepsilon - ct_i), \quad \text{where} \quad \varepsilon = |\vec{r}_o - \vec{r}_l| + |\vec{r}_i - \vec{r}_o|. \tag{5.3}$$

Our time-of-flight measurement does not give a single time of arrival t_i, but rather a histogram where the detected signal is approximately of Gaussian form, centred at $\langle t \rangle_i$ with a standard deviation σ_{t_i}. To take this uncertainty into account, we convolve the ellipsoid with the Gaussian form of the histogram $h_i(t)$:

$$P_i^{\text{ellipse}}(\vec{r}_o) \propto \int_{-\infty}^{\infty} \delta(\varepsilon - ct) h_i(t) dt = h_i\left(\frac{\varepsilon}{c}\right) = \exp\left[-\frac{(\varepsilon/c - \langle t \rangle_i)^2}{2\sigma_{t_i}^2}\right]. \tag{5.4}$$

This ellipsoid can be computed in three dimensions, but in our experiment we are interested uniquely in finding the position of the object in the plane of the floor, assuming the object to track does not move dramatically in the vertical direction. We therefore limit our treatment to a two-dimensional plane with $\vec{r}_o = (x_o, y_o)$, by taking an educated guess as to the height z_o of the object we aim at detecting and tracking. Restricting ourselves to a two-dimensional plane allows us to greatly reduce the retrieval time, so that it can be done in real-time (within a few seconds). An error Δz in estimating the height z_o will lead to an error in the determination of the distance in the (x, y) plane $r_{x,y}$ of approximately $(z_o/r_{x,y})\Delta z_o$. The error in the retrieved position is therefore at worst of the order of magnitude of Δz_o, and decreases significantly for objects farther away (larger $r_{x,y}$).

As mentioned above, we also obtain spatial information from our data acquisition as we image a portion of the spherical wave originating from the hidden object. This section of a circle that we record pinpoints to the position in space where the wave originated. To take advantage of this additional information, we choose to treat the time-of-flight information of all 1024 pixels independently and combine their corresponding probability density of finding the object at position $\vec{r}_o$. For each pixel i, we therefore compute the ellipse that corresponds to a measurement $\langle t \rangle_i$ at position $\vec{r}_i$ using equation 5.4. We obtain a different probability density $P_i^{\text{ellipse}}(\vec{r}_o)$ for each pixel, as illustrated in figure 5.4. We find that these ellipses only overlap at the actual location of the object. By multiplying the ellipses found for every pixel, we obtain a joint probability density that is maximum at the actual object position.

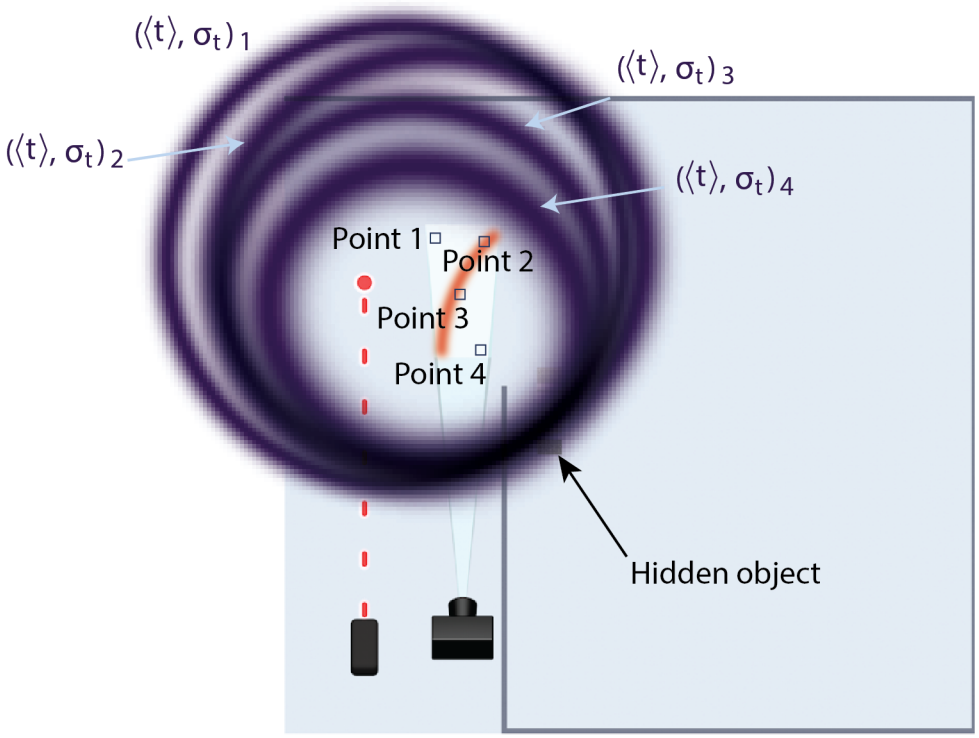

FIGURE 5.4: **Exploiting the spatial dimension of our measurement.** We measure a temporal histogram of the return signal from the object for all 32x32 pixels of the SPAD camera array. For each pixel, we compute the corresponding ellipse where the object is likely to be located to produce the signal recorded at that pixel. Ellipses computed from different pixels will differ in their spatial distribution, but will overlap at a specific point, where the object is truly located. Multiplying all computed ellipses will lead us to a probability density that is pointing to the location of the object, which allows us to retrieve the position of the hidden object.

Before computing the joint probability density, we make one last amendment to the expression of the probability density obtained with a single pixel. In the 1024 pixels of the array, there will be some pixels where the noise level is too high (hot pixels) or where the signal is simply too low to obtain a significant Gaussian fit to its histogram. The obtained values of $\langle t \rangle_i$ and σ_{t_i} obtained from those pixels would therefore not represent a signal coming from the hidden object. To avoid that these unreliable measurements contribute to the joint probability density, and affect the retrieval by multiplying relevant densities by zero, we define the probability density $P_i(\vec{r}_o)$ associated with the pixel i to be a linear combination of the probability density P_i^{ellipse} and a uniform probability density P^{uniform}. This last contribution will prevent any point in space to be multiplied by zero, even when the retrieved P_i^{ellipse} is unreliable due to a noisy measurement on a given pixel. We therefore use

$$P_i(\vec{r}_o) = \alpha_i P_i^{\text{ellipse}}(\vec{r}_o) + (1 - \alpha_i) P^{\text{uniform}} \tag{5.5}$$

as the probability density of finding the object at $\vec{r}_o$ given the measurement at pixel i. We use a coefficient α_i to regulate the relative contribution of P_i^{ellipse} and P^{uniform}. To

have a value of α_i that represents the reliability of the obtained P_i^{ellipse}, we set the value of α_i to be the overlap between the calculated ellipse and the search space in which we look for the hidden object and compute the probability densities P_i^{ellipse}. More precisely, we set $\alpha_i = A_i/A$ where A_i is the area of the ellipse above half-maximum contained in the search space, and A is the area of the search space. To obtain a normalised probability density $P_i(\vec{r}_o)$, we normalise both contributions P_i^{ellipse} and P^{uniform} so that their integral over the search space is equal to one.

Finally, we compute the joint probability density by multiplying the contributions of all 1024 pixels in our measurement:

$$P(\vec{r}_o) = N \prod_{i=1}^{1024} P_i(\vec{r}_o), \tag{5.6}$$

where N is a normalisation constant. In an experiment, as we will discuss in section 5.3, we compute this probability density by defining a search space in (x, y) where we expect the hidden object to be located and dividing that space into a grid of pixels. The obtained probability density $P(\vec{r}_o)$ is our measurement of the position of the hidden object.

5.3 Proof-of-principle experiment

We performed a first experiment at small scale to demonstrate that we can locate and track hidden objects using a femtosecond laser and a SPAD camera, together with the position retrieval procedure explained above [10]. We re-create a situation where a person is hidden behind a wall, at about a 1:5 scale. We use as a target a human form cut out in a piece of foam that is 30 cm high, 10 cm wide and 4 cm thick. We place the target behind a wall that hides it from the view of the SPAD camera, as illustrated in figure 5.2. All previous methods for detecting objects have relied on the presence of a wall in the field of view of the detector to use as a diffuse surface for sending and detecting signal around the obstacle. However, in some real-life scenarios, there might not always be a conveniently placed wall to use as a scattering surface - we therefore decide here to rely uniquely on the presence of a floor (using here a white cardboard on the optical table as a floor), and we image with our SPAD camera (45.5 ps time bin, 110 ps IRF) an area of the floor just beyond the edge of the wall that is about 16 cm wide and 30 cm high. The camera is located 67 cm away from the edge of the wall and is looking down at the floor from a height of 46 cm, which results in a trapezoidal field-of-view on the floor.

We then use a 800 nm femtosecond laser source to send a signal into the hidden area, as illustrated in figure 5.2c. The laser sends 10 nJ, 10 fs pulses at a repetition rate of 67 MHz (670 mW average power) at the floor, and the pulses scatter into the hidden scene

after hitting the floor. To avoid saturating the SPAD camera with a direct reflection from the laser, we direct the pulses out of the field of view, 15 cm to the left. The light that is scattered to the object and back is then recorded by the SPAD camera, which acquires data over 10 080 frames of 300 μs exposure time, for a total acquisition time of only 3 s. Each of the 32x32 pixels that image a point on the floor at position $\vec{r_i} = (x_i, y_i)$ therefore records a histogram of the photons arrival time. We discussed in section 5.2 how we can use the arrival time of the signal coming from the hidden object to infer its position. However, in the recorded signal, we collect signal not only from the hidden object, but also from every other surfaces in the environment of the experiment, including from the walls and the ceiling.

The first processing step we need to do to analyse the acquired data is therefore to distinguish the signal from the object of interest from all other signals detected. A straightforward way to do so would be to acquire a background signal without the presence of the hidden object in the scene, and subtract this background from any subsequently acquired

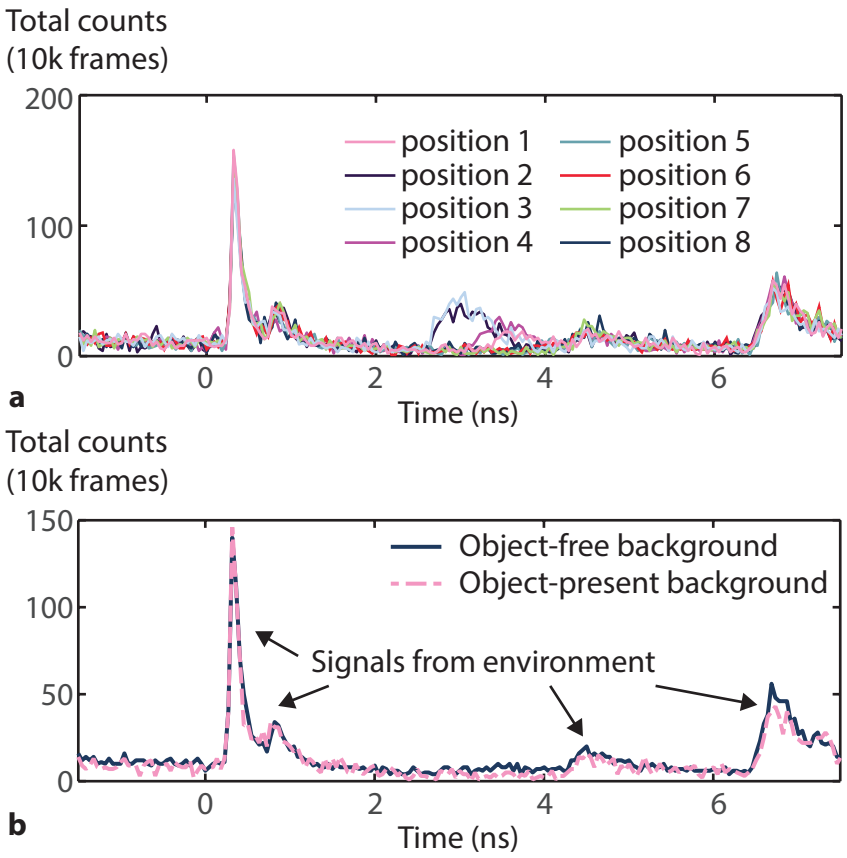

FIGURE 5.5: **Estimating the background with a non-cooperative moving target.** a) When acquiring data with a moving target, the signal from the target in the histogram will change from one acquisition to the next, while all signals from the fixed environment will remain the same. We therefore compute the median of the values acquired at each time bin of each pixel in different acquisitions to estimate the background signal coming from the environment. b) The obtained object-present background is in very good agreement with an object-free background acquired when the object is absent of the scene.

data. This could be an interesting solution for using this technique in a fixed and known environment, such as for surveillance applications. However, in most scenarios where the access to the scene is not possible and the target is non-cooperative, we must rely on a different method to isolate the signal of interest from all other signals. A promising avenue to achieve this is to rely on the motion of the hidden target: if we are interested in detecting targets in motion, we will collect with a given pixel a signal from the target that differs in its temporal distribution as the target is moving. On the other hand, all signals from the fixed environment will remain the same. To approximate the background signals coming from the environment in our experiment, we therefore use a set of eight acquisitions with the target at different positions in the scene - this could correspond to the first few acquisitions that are collected in a real-life implementation of this scheme, provided that an object is moving over a significant distance with respect to its size. These eight positions correspond to the ones shown later in this section in figure 5.7. In figure 5.5a, we show an example of the eight histograms coming from different 3 s long acquisitions, for a given pixel of the array. We observe that most peaks are fixed from one acquisition to the next, but that one part of the signal shifts by about 2 ns between different acquisitions. To estimate a background from this data, we consider each time bin of each pixel of the array and find the median value over the eight acquisitions [102–104]. This gives us a result that is in very good agreement with the object-free acquired background (see figure 5.5b). The median proves itself to be in this situation a better choice than the mean of the acquisitions, which tends to overestimate the background in some points where there is sometimes a signal from the moving target; and a better choice than the minimum of the acquisitions, which tends to underestimate the background as the counts fluctuates from one acquisition to the next, and taking the minimum leads us with background-subtracted data with a higher level of noise.

In the results presented here, we therefore use the median of eight acquisitions as a background and subtract this from each acquisition as a first step in data processing. To further clean the data, we apply a smoothing step in the temporal axis where the value at each bin of a histogram is averaged with its two nearest neighbour - this step is applied twice for additional smoothing of the histograms. We then use the spatial correlations in the frames to further enhance the signal detected: for each frame, the value of one pixel is taken to be the median of its value and those of its eight nearest neighbours. This median filtering in space greatly enhance the signal, and we finally fit a Gaussian to each histogram. An example of a raw histogram and the Gaussian fit to the processed data is shown in figure 5.6. This fit gives us the value of time of arrival $\langle t \rangle_i$, which we calibrate for a given imaging geometry to correspond to the flight path from the laser spot on the floor to the object and back to the field of view, as discussed in section 5.2. We also take into account the particular geometry of the imaging system to calibrate the time-of-arrival measurements from different pixels, which are situated at various distances from the camera since we image the floor at an angle - the geometry of the system and the

details of the correction to the measured $\langle t \rangle_i$ are described in appendix C. We also obtain the standard deviation of the Gaussian fit σ_{t_i} that we use as the standard deviation of the probability density P_i^{ellipse}.

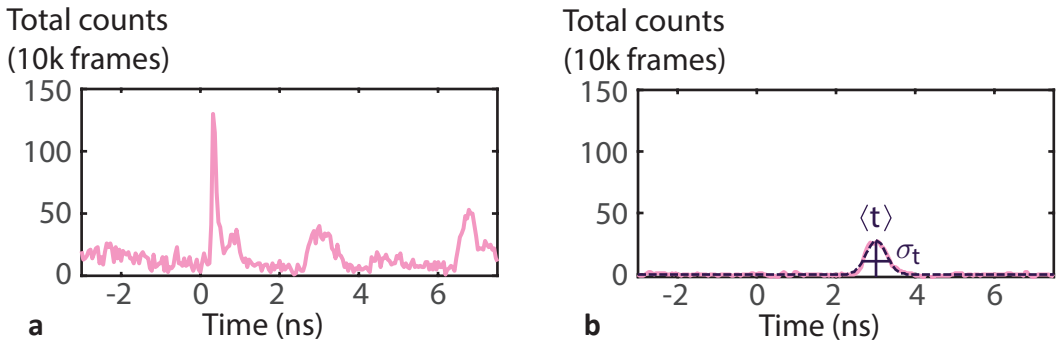

FIGURE 5.6: **Data processing.** a) An acquired histogram shows signals from multiple sources. We first subtract the background histogram taken with the object present in the scene, and apply both a smoothing step in the temporal dimension and a median filtering step in the spatial dimension to further clean the data and enhance the signal from the target. b) A Gaussian profile is fitted to each processed histogram to find the time-of-flight $\langle t \rangle_i$ and its uncertainty σ_{t_i}. We note that the width of the signals is much larger than the width of the impulse response function of the SPAD camera - this width is not limited by the detection system, but by the surface area of the object which, even though we treat it as a point in space, scatters light from every point of its surface.

Using the measured values of $\langle t \rangle_i$ and σ_{t_i} at every pixel, we compute for each pixel the probability density $P_i(\vec{r}_o)$ (equation 5.5 by using a guess height of $z_o = 17$ cm for the vertical position of our 30 cm high foam mannequin), and we divide the search space on the other side of the wall into a 90x110 pixel space with a pixel pitch of 5 mm in both x and y. We then compute the joint probability density $P(\vec{r}_o)$ of finding the object at a position $\vec{r}_o$ (equation 5.6).

In a first experiment, we placed the object at eight distinct positions in the hidden space and acquired data for 3 seconds while the object was fixed at each position. The retrieved probability densities are shown in figure 5.7 superimposed on the actual positions of the object and, to scale, the position of the wall, camera, field of view and laser spot. The target position is accurately retrieved with a precision of approximately ± 1 cm in x and ± 2 cm in y, corresponding to $\sim 20\%$ uncertainty with respect to the size of the hidden target. Those values are calculated as the square root of the weighted variance of the probability distributions in both x and y (see section 5.3.1). This precision is dependent on the position of the object for geometrical considerations that we will discuss in section 5.3.1. In the results shown in figure 5.7, we observe that our method provides a position retrieval not only when the target is located around the corner from the edge of the wall, but also when it is physically receded behind the wall, provided that the object is illuminated by the laser pulses scattered from the floor.

Since the acquisition of a signal from the hidden object is very fast (3 seconds), it is possible to acquire data in real-time as the object is moving in the hidden scene. We can track a motion at a speed where the target moves on a distance of the same order of magnitude than its size within the 3 seconds acquisition time - in our case, we performed an experiment to track our mannequin foam moving at a few centimetres per second. To do so, we placed the mannequin on a track parallel to the wall, located at different positions in x (20, 30 and 40 cm). We then moved the target back and forth in the y direction, following the track, at a speed approximately constant (the target was moved by hand). As the target was moved, we took 20 subsequent 3 seconds acquisitions, so that

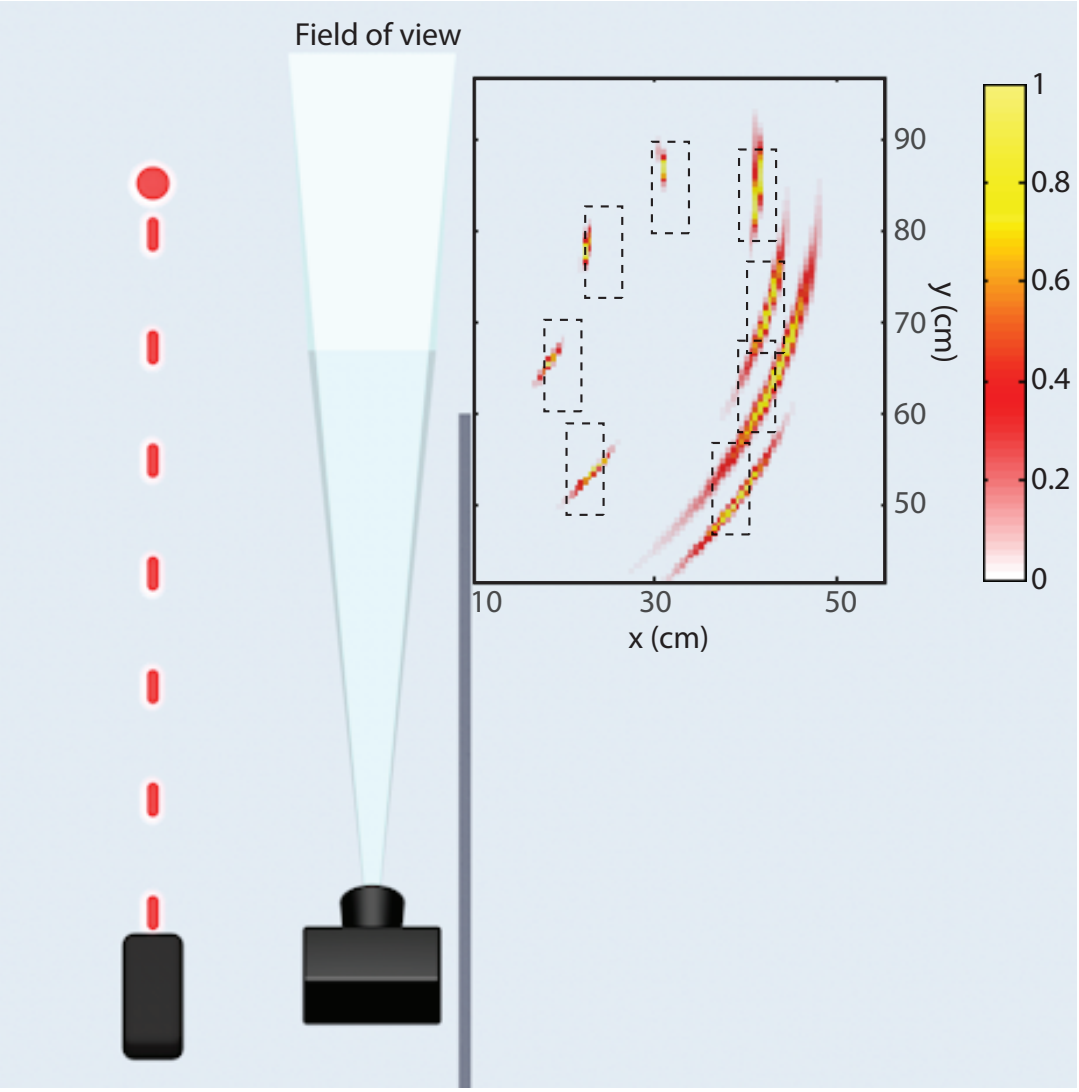

FIGURE 5.7: **Experimental implementation of our hidden-object detection technique using a SPAD camera.** The experimental layout is shown to scale and the retrieved probability densities, each normalised to a peak signal of 1 for clear visualisation, are overlapped with the actual position of the target during the eight acquisitions (dashed rectangles). The search space indicated by the black rectangle is divided into 140x140 pixels to compute the retrieved probability densities. The positions x and y are given with respect to the camera position, and we show that we can retrieve the position of a hidden object located about a meter away from the camera.

every 3 seconds we gather a signal to update the location of the moving hidden target. We processed the data the same way as above at a later time, but the processing of one data set takes only about 2 seconds on a laptop running Labview, so that it could easily be implemented to run in parallel to create a system that both acquires and analyses the data in real time.

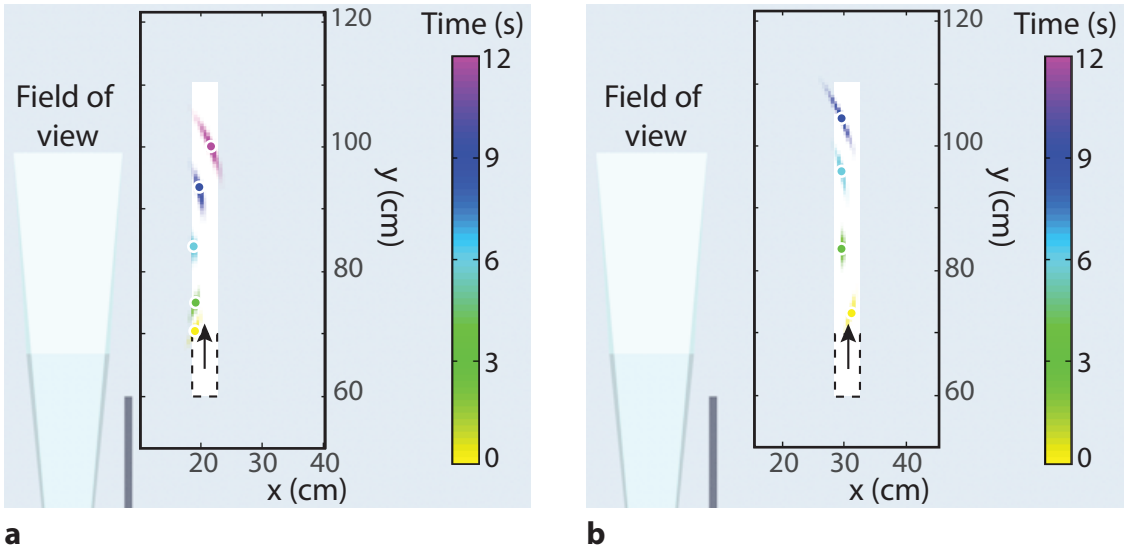

FIGURE 5.8: **Tracking the motion of a hidden object in real-time.** In the selected frames from the acquisition, the target is moving upwards in the y direction. Our acquisition scheme allows us to update the position of the target every 3 seconds, and the position retrieved corresponds to an average position of the target during that 3 second period. The coloured areas show the probability densities each normalised to a peak signal of 1, and the colour code indicates the time at which the signals were acquired, each separated by a 3 second interval. We show selected frames from the acquisition done with the target at a) $x = 20$ cm moving at 3.9 cm/s and b) $x = 30$ cm moving at 5.3 cm/s.

We show a subset of the results acquired with a moving target in figures 5.8 and 5.9. In figure 5.8, we show the retrieved positions as the target was traveling once from bottom to top in the y direction along the track positioned at $x = 20$ cm and $x = 30$ cm. In figure 5.8a, we show that we retrieve the positions with centimetre precision, similarly to what was obtained with the static target. We show five different positions acquired over 15 seconds, as the target was moving at a speed of 3.9 cm/s. In figure 5.8b, the target is moving at a higher speed (5.3 cm/s), and it covers the length of the track in only 12 seconds - we show here the four retrieved positions during that time. In figure 5.9a, we show 5 of the 20 positions acquired when the target was moving on the farthest track ($x = 40$ cm), back and forth in the y direction at a speed of 2.8 cm/s. A movie containing all the recovered positions is available online (see appendix A). In figure 5.9b and 5.9c, we show the full evolution of the recorded motion as it was recovered from subsequent 3 seconds acquisitions, giving the recovered maximum probability position as a function

of time over the 60 seconds total acquisition period. We see that the retrieved motion is in good agreement with the actual motion of the target along the track. These results, published in [10], demonstrate that our method can be used to reliably track the position of a hidden object in real-time.

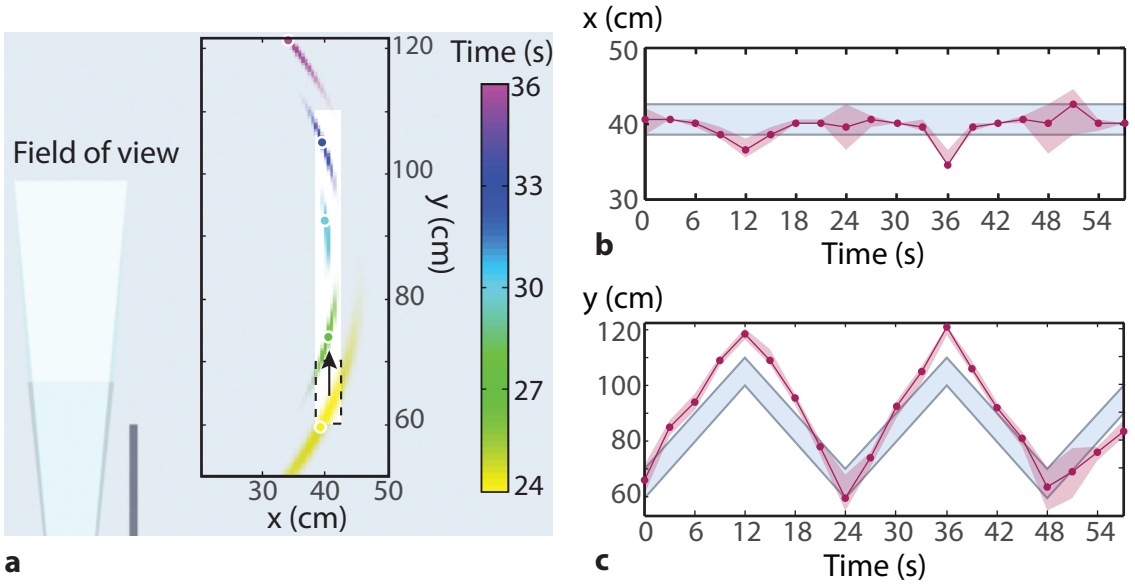

FIGURE 5.9: **Retrieving the path of a moving hidden object.** a) The motion of the target is acquired in successive 3 second long acquisitions for a target moving along the y direction at 2.8 cm/s at $x = 40$ cm. In (b) and (c), we show the retrieved maximum probability position as well as the width of the 50% probability value in both x and y. The blue area is the position of the track where the target is moved. All positions are given with respect to the camera position.

This first experiment provided a proof-of-principle for our detection system, as well as giving us the opportunity to gain a deeper understanding of some parameters of the technique, such as its uncertainty, its signal-to-noise ratio and the prospect of scaling up the experiment, as we will discuss in the following parts of this section. It showed a promising route to a non-line-of-sight tracking system that could become portable and implemented in different real-life scenarios. Further work along that direction will be discussed in section 5.4.

5.3.1 Position retrieval uncertainty

As can be observed in the results shown above, the precision of our retrieval seems to depend on the position of the hidden target in the (x, y) plane. We calculate the precision of the uncertainty, or precision, of our retrieved probability densities $P(\vec{r}_o) = P(x_n, y_m)$

by computing their mean square weighted deviation σ_x^2 and σ_y^2:

$$\sigma_x^2 = \frac{\sum_{n=1}^{N}\sum_{m=1}^{M} P(x_n, y_m)(x_n - \bar{x})^2}{\sum_{n=1}^{N}\sum_{m=1}^{M} P(x_n, y_m)} \tag{5.7}$$

$$\sigma_y^2 = \frac{\sum_{n=1}^{N}\sum_{m=1}^{M} P(x_n, y_m)(y_n - \bar{y})^2}{\sum_{n=1}^{N}\sum_{m=1}^{M} P(x_n, y_m)}, \tag{5.8}$$

where $\bar{x}$ and $\bar{y}$ are the mean positions in x and y of the probability densities, and the deviation is calculated over the $N \times M$ search space grid. The uncertainty is then taken as the weighted standard deviation σ_x and σ_y

To study the dependence of the uncertainty on the position of the target, we analyse the data of the eight positions shown in figure 5.7. In figure 5.10, we show the obtained uncertainties in x and y as a function of the distance between the actual object position and the centre of the field of view $|\vec{r_o} - \vec{r_i}|$. The uncertainty increases as the object moves farther away.

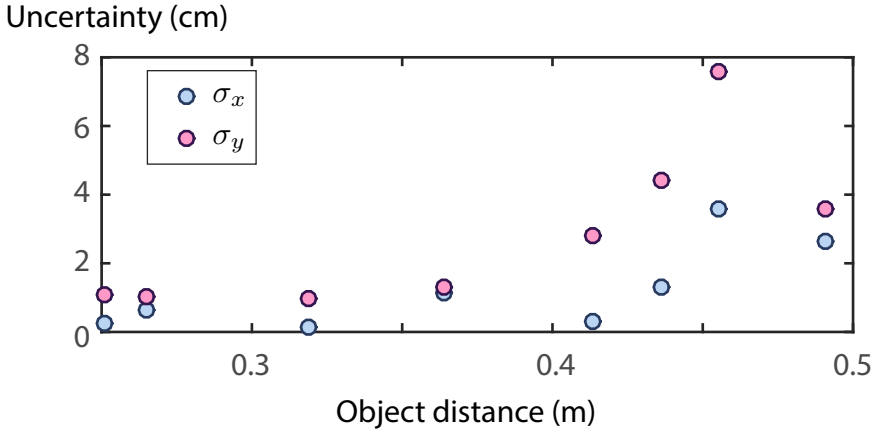

FIGURE 5.10: **Uncertainty of experimental data**. The uncertainties σ_x and σ_y of the retrieved probability densities increase as the target moves farther away from the field of view.

The uncertainties in our experimental results are also seen to vary as a function of the specific x and y position of the target. To study this dependence and confirm that this is inherent to our retrieval method rather than depending on experimental conditions, we perform a simulation of a single-point target moving on a 25x25 grid of (x, y) coordinates by steps of 2.5 cm. We simulate the signal coming from the target as a simple spherical wave with its centre corresponding to the position of the object, with a Gaussian temporal profile of $\sigma_t = 218$ ps standard deviation, which is the average deviation that we measure in our experimental data. From this simulated ideal signal, we apply the position retrieval

algorithm detailed in section 5.2, with a 300x200 search space with 5 mm resolution. We then calculate the weighted standard deviations σ_x and σ_y on the obtained probability densities $P(\vec{r}_o)$. We show in figure 5.11 the obtained uncertainties σ_x and σ_y for the different (x, y) target positions, where the uncertainty is colour-coded. We observe a complex dependence of the precision on the target position which is inherent to our retrieval scheme. The uncertainty in x varies as a function of the y position of the target: it is lowest along the line that corresponds to the middle of the field of view, and increases as the target moves up and down this line in the y dimension. When the target moves away from the field of view along the x axis, the uncertainty σ_x decreases rather than increasing. The uncertainty in y, on the other hand, shows a dependence along the x position of the target and increases as the target moves away.

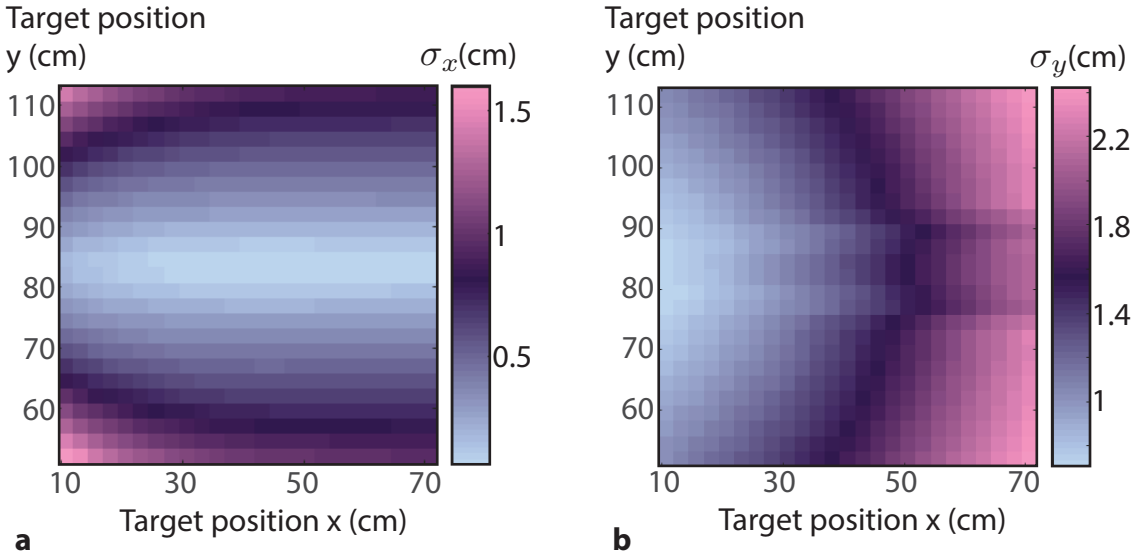

FIGURE 5.11: **Simulated uncertainty of the position retrieval algorithm**. We simulate an ideal signal coming from a single-point hidden target located at various (x, y) positions and apply our position retrieval algorithm. The obtained uncertainties σ_x (a) and σ_y (b) depend on the exact position in (x, y) of the hidden target. The simulation was done using the same geometry as the experiment described above, and the position of the target is given with respect to the camera position.

This dependence can be simply explained as a geometrical effect by considering the shape of the elliptical probability densities that we calculate for each pixel. In figure 5.12, we illustrate two situations, one where the target is located close to the field of view and the other where it is further away along the x axis. We trace three of the ellipses that correspond to the probability densities obtained with three of the pixels on the SPAD camera. We observe that, when the target is close, the three ellipses differ quite significantly from each other and we obtain an overlap that is small, giving a good precision in both x and y. As the target moves farther away, the probability ellipses get larger and similar to each other (the ellipses tend towards a perfect circle), which leads to an overlap between the ellipses that is larger in y and smaller in x, explaining the behaviour of σ_x and σ_y observed

in figure 5.11.

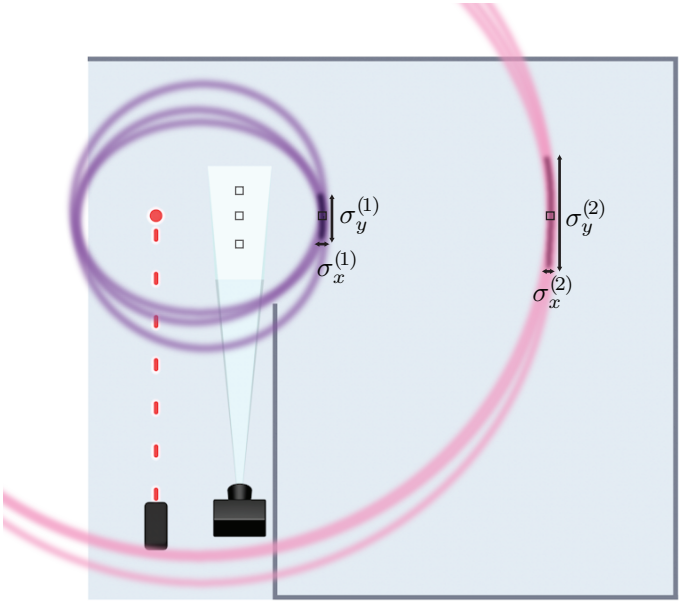

FIGURE 5.12: **Geometrical explanation of the uncertainty dependence on target position**. The retrieval algorithm relies on the computation of ellipses of probabilities corresponding to the detected signal at a given pixel. The size and shape of these ellipses, as well as their overlap, depends on the target position in the hidden scene. When the target is close, we obtain smaller ellipses with a stronger eccentricity (purple in figure) than when the target is farther (pink in figure). We therefore obtain a better precision in y when the target is closer, but a slightly better precision in x when the target is farther. A similar argument can be made to explain the behaviour of the uncertainties as the target moves along the y dimension, where the uncertainties will vary according to the curvature of the ellipses at the position of overlap.

Another parameter that may affect the retrieved probability density and its uncertainty is the number of pixels of the array that have a contribution in the retrieval. As discussed in section 5.2, some pixels have inherently a higher noise level, and as the target moves farther away some pixels will not have a strong enough signal-to-noise ratio for us to retrieve a reliable time-of-arrival measurement from their histogram. Again studying the measurements done at the eight positions shown in figure 5.7, we find that the number of pixels used in the retrieval algorithm is close to the 1024 available when the signal from the target is strong, but a decrease in signal-to-noise ratio leads to a diminution of the number of pixels used in the algorithm, down to only 150 of the 1024 available. We can also observe a dependence in the experimental uncertainties obtained for these positions as a function of number of pixels used, as shown in figure 5.13a. To confirm that this effect is actually due to the lower number of pixels used, and not just correlated to the position of the target as detailed above, we perform a simulation with a single-point hidden target at a fixed position. We simulate an ideal signal as above for the 32x32 pixels of the SPAD camera, and randomly choose a limited number of pixels for the retrieval algorithm, ranging from 10 pixels to 1024. In figure 5.13, we observe a clear dependence of the uncertainties σ_x and σ_y on the number of pixels used in our simulation. This

dependency shows how a loss of a portion of the 32x32 pixels of the SPAD camera due to lower signal will affect the precision of our retrieval. This dependency is also relevant to consider when designing a system based on a low number of single-pixel detectors rather than a SPAD array, as discussed in section 5.4. In such a case, the relative distance and positions of the pixel will also play an important role: spacing the pixels used in order to collect significantly different signals on all pixels is a good strategy to compensate for the loss of precision coming from using a lower number of pixels.

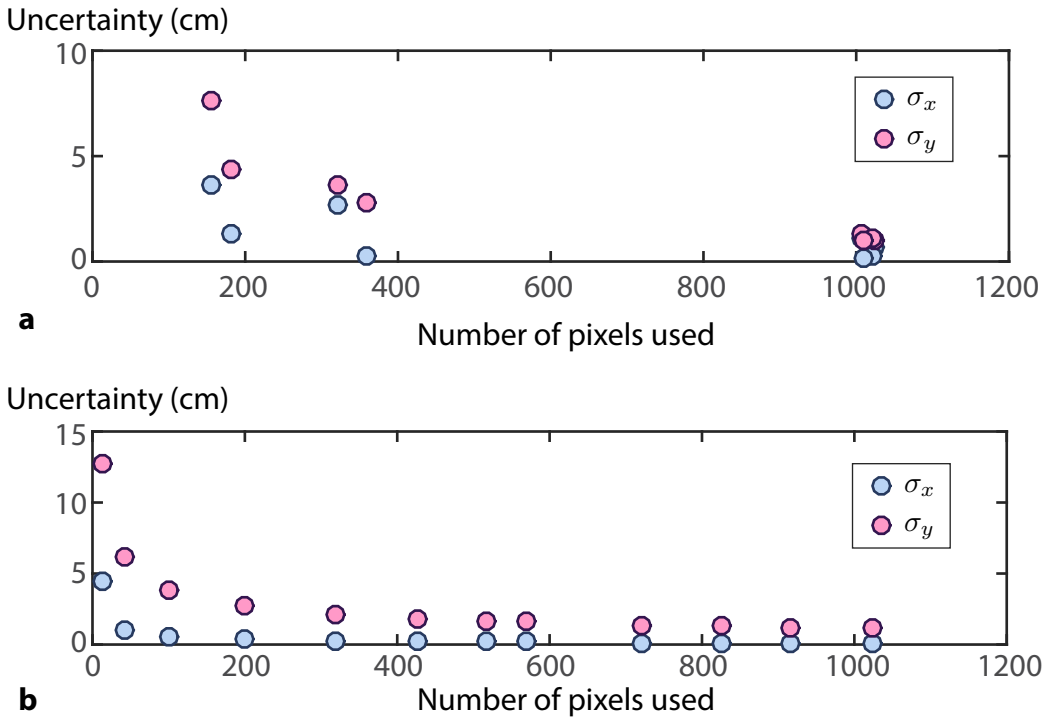

FIGURE 5.13: **Uncertainty of the position retrieval as a function of number of pixels**. a) In our experimental data, we find that the uncertainty of the retrieved position is higher for retrievals performed with a lower number of pixels. b) To confirm this apparent dependence, we simulate an ideal signal captured with a 32x32 camera, coming from a target at a fixed position. We randomly select a varying fraction of the 1024 pixels to perform the retrieval, and find the same dependence as observed in the experimental data: the uncertainties σ_x and σ_y increase as the number of pixels used in the retrieval decreases.

5.3.2 Detecting multiple targets

A capability of great interest for a hidden object detection system is the ability to detect more than one target in a hidden scene. In the sections above, we have developed a technique and a retrieval algorithm that is well adapted for a single target, and have demonstrated experimentally that this technique provides accurate position retrieval for a target hidden from view about one meter away from the camera. However, in many real-life scenarios, there will be more than one object moving in a hidden scene and the information about multiple targets could be of great value. We therefore investigate here the possibility of locating multiple targets with our detection system.

The main requirement to achieve a multi-target detection with our system is our ability to distinguish signals coming from different targets. If we can measure two distinct signals in the temporal and spatial dimensions of our recorded data, we can then isolate them and apply our retrieval algorithm as it is described in section 5.2 on both sets of signals separately. To produce distinct signals, two targets must be separated by a distance larger than their size and larger than the distance equivalent to our system temporal resolution, so that their respective returned back-scattered wave do not blur into a single visible wave in our measurement. In a proof-of-principle experiment, we placed two similar targets in the hidden scene, with the same imaging geometry as in the experiment reported above with the single target. The two targets are separated in the y dimension by about 45 cm, and are both located at $x \approx 25$ cm. We acquire data using the same parameters as in the previous experiments, i.e. we acquire a signal over 10 000 frames of 300 μs exposure time.

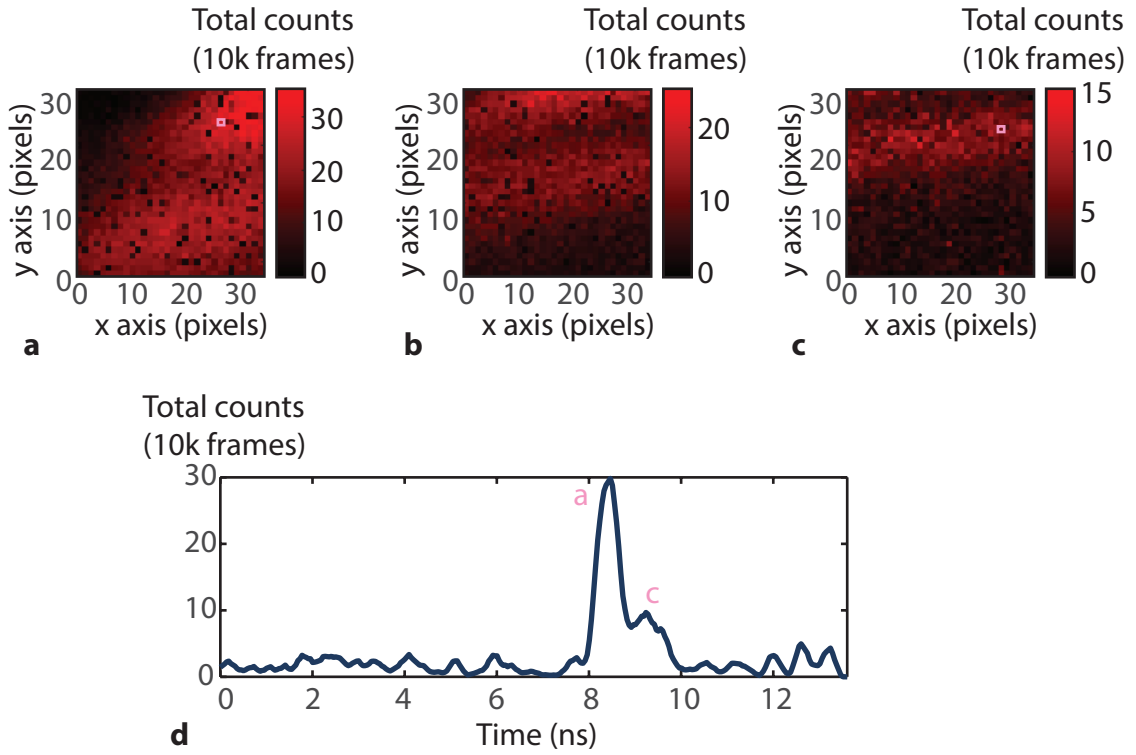

FIGURE 5.14: **Detection of two hidden targets.** We can distinguish in the acquired frames two back-scattered waves crossing the field of view, originating from two distinct targets in the hidden scene. The frames shown are at (a) 8.3 ns, (b) 8.8 ns and (c) 9.3 ns. a) In the first frame, we see the wave coming from the first target in the bottom of the field of view, and a second wave coming from a second target is visible in the top right quadrant of the frame. b) 500 ps later, both waves are present in the frame as they travel upwards. c) At 9.3 ns, we still see the wave coming from the first target, and the second wave has already exited the field of view. d) A histogram taken at the pixel indicated by a pink square in (a) and (c) shows the two peaks corresponding to the two waves – the signal from the first target is indicated as c in the histogram and the signal from the second target as a.

We subtract from the acquired data a background that was acquired without the objects present in the scene, although a similar strategy to that described above could be taken to estimate the background with moving non-cooperative targets. We then perform the same data processing as above, both in the temporal and spatial dimension. In figure 5.14, we show three frames of the recorded data after background subtraction and temporal smoothing. The frames are separated by 500 ps. We can clearly distinguish two waves travelling across the field of view of the camera: in figure 5.14a, one of the wave is visible across the bottom of the field of view, and the second is visible on the top right corner of the field of view. In figure 5.14b, they are both present and almost parallel, while in figure

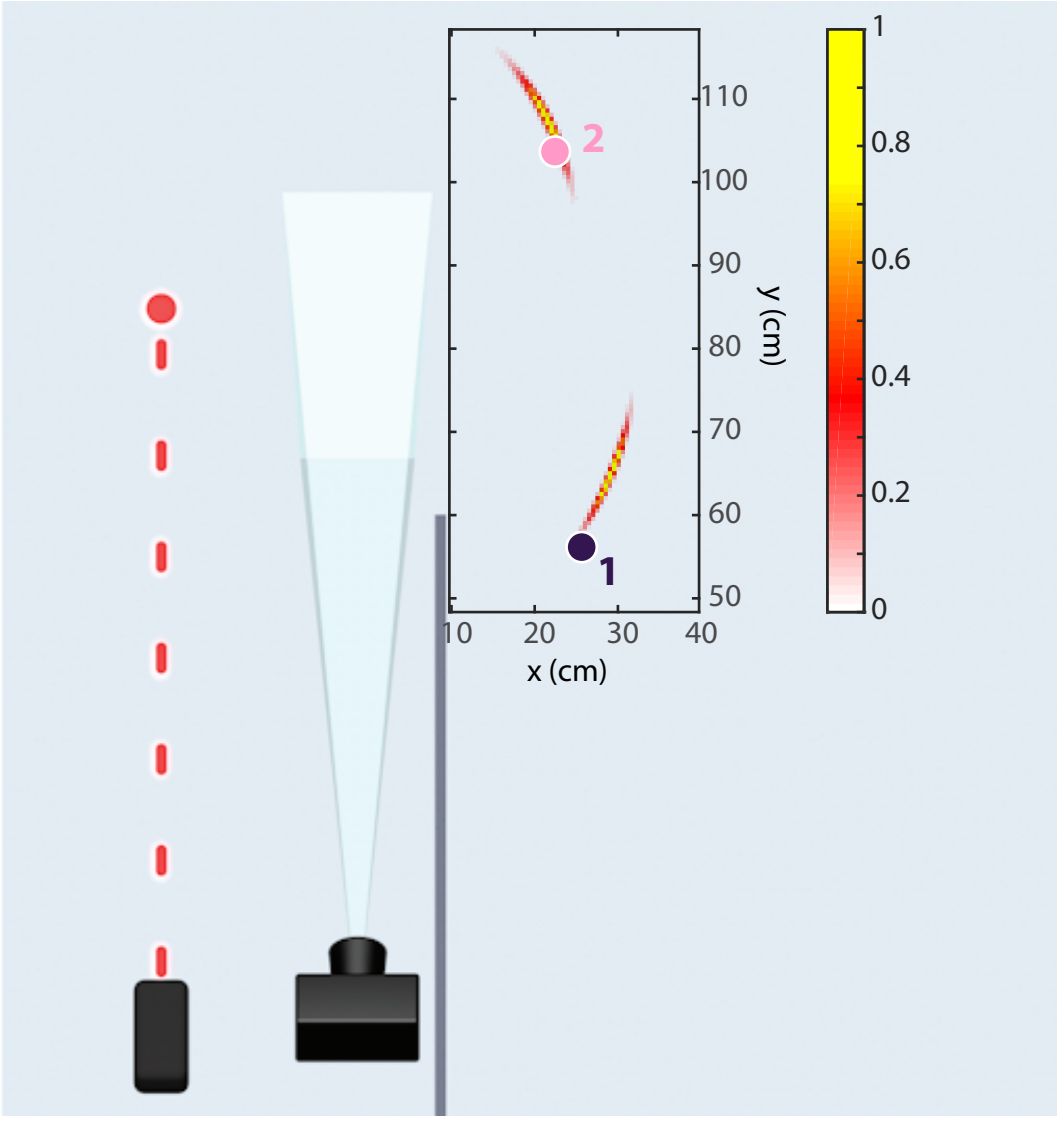

FIGURE 5.15: **Position retrieval of of two hidden targets.** By analysing separately the signals coming from the two hidden targets, we can retrieve a probability density for each of them. Both probability densities have been normalised to a peak signal of 1 for clearer visualisation. The retrieved positions when two targets are present are in good agreement with the positions found with single-target measurements (indicated by the circles). The signal from target 1 corresponds to the wave that can be observed in the lower section of the frames in figure 5.14a,b and is the one present in figure 5.14c and indicated as c in the histogram of figure 5.14d. The signal from target 2 is the second one that can be observed in the frames and is indicated as a in figure 5.14d.

5.14c, we can still see the first wave, and the second has already exited the field of view.

In order to retrieve the time-of-flight information from the two different targets, we need to find the peaks present in the histograms (see figure 5.14d) and assign them to the correct target. As the experiment presented here aimed simply at demonstrating that the detection of multiple targets is possible, we manually determined the limits in space-time between the two detected waves by analysing the limits between the two waves in the acquired frames. Designing an automatic algorithm to distinguish the signals is yet to be done and would be essential for a practical multi-target detection system. Here, we restricted our analysis to finding peaks in each histograms above a certain threshold, and classifying them into two data sets (for the two targets) depending on where the peaks lied in our space-time measurement with respect to the limits we determined. We kept only one peak per histogram per target and used this arrival time in our position retrieval. The retrieval algorithm is then used exactly as described in section 5.2 for each of the two signal sets separately. We therefore obtained two distinct probability densities $P(\vec{r}_o)$, as shown in figure 5.15. We also indicate by points in figure 5.15 the position of maximum probability found when the targets were placed in the same positions as in this measurement, but one at a time, so that only one signal per measurement was present. The multi-target retrieval process gives a result that is consistent with the single-target retrieval, confirming that is it possible to detect and locate multiple targets with our detection system provided that we can distinguish their signal in our measurement.

5.3.3 Signal-to-noise ratio

Another feature of great interest for a non-line-of-sight tracking system is its ability to detect objects at human-size scales rather than in the smaller laboratory scenes as presented above and in all previous experiments in the literature. Our ability to scale the distance at which we can detect objects depend on the signal-to-noise ratio of detected signals as a function of both the target distance from the field of view and the camera distance from the surface it images beyond the obstacle hiding the object from its view. The signal will scale differently as a function of these two distances.

We address in this subsection the signal-to-noise ratio obtained in our experimental measurements as a function of the target distance to the field of view. As the light travels approximately as a spherical wave from the laser spot to the object, and again from the object to the field of view, we expect the intensity of the signal to scale as

$$\frac{1}{|\vec{r}_o - \vec{r}_l|^2 \times |\vec{r}_o - \vec{r}_i|^2} \rightarrow \frac{1}{r^4}, \tag{5.9}$$

where we take the distance $|\vec{r}_o - \vec{r}_i|$ to be the distance from to object to the middle of the field of view, and we obtain a general scaling in $1/r^4$ with r being approximately the

distance between the object and the section of the floor used to scatter the laser pulses and image the signal coming back.

To verify this behaviour, we measure the signal-to-noise ratio (SNR) obtained in the eight measurements done with the fixed target positions shown in figure 5.7. To do so, we first measure the number of counts in every histograms contained in the target signal, by summing the processed histogram (background subtraction, temporal smoothing, spatial filtering) over the full signal width (20 bins). We average the signal at each position over the 100 pixels with the most signal, as to not underestimate the signal level by taking into account pixels that do not contribute to our retrieval, as discussed above. The level of noise is taken by summing over 60 bins in a region of the histograms where no signal is present, and we find an average of 50 counts over 20 bins as the noise level. The signal-to-noise ratio is therefore the ratio of the number of counts in the measured target signal over the 50 counts/20 bins noise level. We show in figure 5.16 the measured SNR as a function of distance $1/r^4$ as expressed by equation 5.9. We observe that the measured values of SNR follow a linear increase as a function of $1/r^4$, as expected.

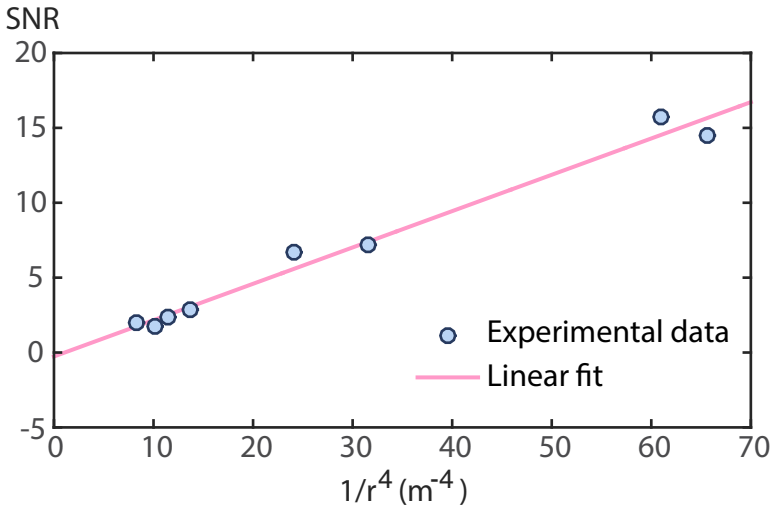

FIGURE 5.16: **Signal-to-noise ratio of experimental data.** The signal-to-noise ratio of our measured data decreases as the target moves away from the field-of-view, following a $1/r^4$ behaviour as expected from the decrease in intensity of a spherical wave as a function of distance in $1/r^2$. In our best measurement, at a signal-to-noise ratio around 15, we detect 750 photons coming from the hidden target, corresponding to a detection rate of 0.07 photons per frame. The lowest signal, giving a SNR of 2, corresponds to a detection rate of about 0.01 photons per frame coming from the hidden target.

In the farthest position of our set of measurements, we obtain a signal-to-noise ratio of $\sim$2 for a target distance of $|\vec{r_o} - \vec{r_l}| \sim |\vec{r_o} - \vec{r_i}| = 60$ cm. To detect targets located at larger distances, some improvement to the SNR can be made with modifications to both the detection setup and the hardware used:

1. We use in our experiment a target that has a 10x30 cm surface. The signal detected by the SPAD camera is proportional to the surface of the target, so that for human-size hidden target, the signal should increase by a factor ~ 30. If the target is a car, this improvement can be potentially higher.

2. Depending on the surface used to scatter the laser pulses and detect the back-scattered signal, the signal may vary significantly, given the scattering properties of the surface. In our experiment, we use a white cardboard as a floor - a darker material may scatter less light, leading to a decreased signal. However, since the scattering is not isotropic from a scattering surface, but is preferential in a direction normal to the surface, having access to a different surface such as a wall rather than a floor might significantly improve the signal detected. This comment is only a qualitative assessment on the different parameters that may play a role, and a thorough study of different surfaces would be relevant in the development of a detection system destined to be used in real-life scenarios.

3. The active area of the SPAD camera is also limited in efficiency to about 6.5%. Single-photon avalanche diodes working at our operating wavelength have been reported with efficiencies of the order of 25-30$\%$ [105, 106]. Optimising the sensitivity of the SPAD detector to our operating wavelength could therefore allow a gain of 5x in detected signal.

4. The SPAD array detection efficiency is now limited by its low fill-factor ($\sim 2\%$). With arrays where the counting electronics is vertically integrated to the SPAD detector in each pixel, or using microlens array, fill factors up to $\sim 80\%$ could be achieved, increasing the efficiency of the detector by a factor 40x.

5. As discussed in section 3.3.4, the dynamic range and SNR ratio of the SPAD camera increases with the number of acquired frame. In the experiments reported above, we restricted ourselves to acquiring signals over 10 000 frames, to keep the data acquisition time low (3 seconds). The read-out rate is currently limited by the USB2.0 connection between the SPAD camera and the computer, as mentioned in section 3.3.5. Work is currently underway to incorporate a USB3.0 connection to the SPAD camera [61], which would allow us to acquire more frames in the same period of time, given we reduce the exposure time of the frames. It is reasonable to expect a gain of about a factor 10 using an increased number of frames in each acquisition.

Taking into account all these factors, an improved version of our system aiming at detecting human-size targets could reach a factor $\sim$60 000 in SNR improvement, meaning that it could detect targets at about 10 meters away from the field of view. The ultimate limit in target distance is also dependent on the surroundings, as the SPAD camera can only detect one photon per frame: if the target is located in an empty space with no other significant signals coming from the environment, a bigger portion of that 1 photon/frame

signal can be taken up by the signal coming from the target. However, if the room in which the target is hidden is very cluttered, the distance at which the target can be detected may be reduced as the signal from the target will be less significant with respect to the background signal from the environment. In any case, given this analysis, it is reasonable to expect a non-line-of-sight tracking system based on the method described here to locate human-size targets a few meters away from the field of view of the camera, which is sufficient to be an asset in real-life situations.

The other relevant dimension in a non-line-of-sight tracking system is the stand-off distance r_c between the camera and the surface it images beyond the edge of the obstacle (in our case, the floor). When the laser pulses are directed at the surface, there is technically no loss of intensity as a function of distance as the pulses are then collimated. However, the scattered intensity of the signal travelling from the imaged surface to the camera is expected to decrease in $1/r_c^2$. This loss of signal can however be compensated by imaging a larger area on the scattering surface. If the dimension d of a dxd field of view increases linearly with r_c, the area imaged would increase with r_c^2, compensating perfectly for the loss of signal due to the decreasing intensity of the scattered wave. Depending on the imaging geometry, the stand-off distance between the camera and the edge of the obstacle can therefore significantly be increased. In section 5.4, we will discuss experiments performed with single-pixel detectors that reached a stand-off distance of 50 m while still being able to detect hidden people around a corner about 2 m away from the corner. Before detailing this experiment, we will first cover some measurements done at larger distances to detect hidden people using the SPAD camera.

5.3.4 Scaling to larger dimensions

A first test to locate a person hidden from view was done using the same type of imaging geometry than the one used in the proof-of-principle experiment, but at larger distances. The SPAD camera was looking down at the floor from a height of 1 m, imaging an area on the floor 2.1 meter away of roughly 20x54 cm dimension. A person was positioned around the corner from an obstacle blocking the view of the camera (in this case, an optical table). The geometry of the imaging system and the positions of the target are shown in figure 5.17. In this measurement, we also investigated the possibility of taking measurements with lights on: some lights above the optical tables in the laboratory were open, and we placed a 10 nm bandwidth filter at the operating wavelength (780 nm) in front of the SPAD detector. The noise level was slightly raised compared with a situation where all lights were closed, but it did not prevent us from acquiring and isolating a signal from the target.

We observed that, in this geometry, acquiring 10 000 frames was not sufficient so we acquired data over 100 000 frames of 300 μs exposure time, for an acquisition time of

30 seconds. This is too slow for a real-time acquisition and retrieval of the hidden object position, but we nevertheless investigated the signals that can be detected from a person hidden around the obstacle. The data was processed by first subtracting a background and applying a median filter in the spatial dimension. A Gaussian was then fitted to each histogram;. an example of a Gaussian fit is shown in figure 5.18a. Unfortunately, the Gaussian fitting does not work as rapidly here than it did in the small-scale experiment, due to the fluctuations in the data and the spread of the signal; it takes about 6 seconds to fit all histograms of one data set. Again, this is too slow for a real-time implementation of a non-line-of-sight tracking system. The retrieved probability densities for five different positions of the person, assuming a height of $z_o = 1$ m, are shown in figure 5.17. The

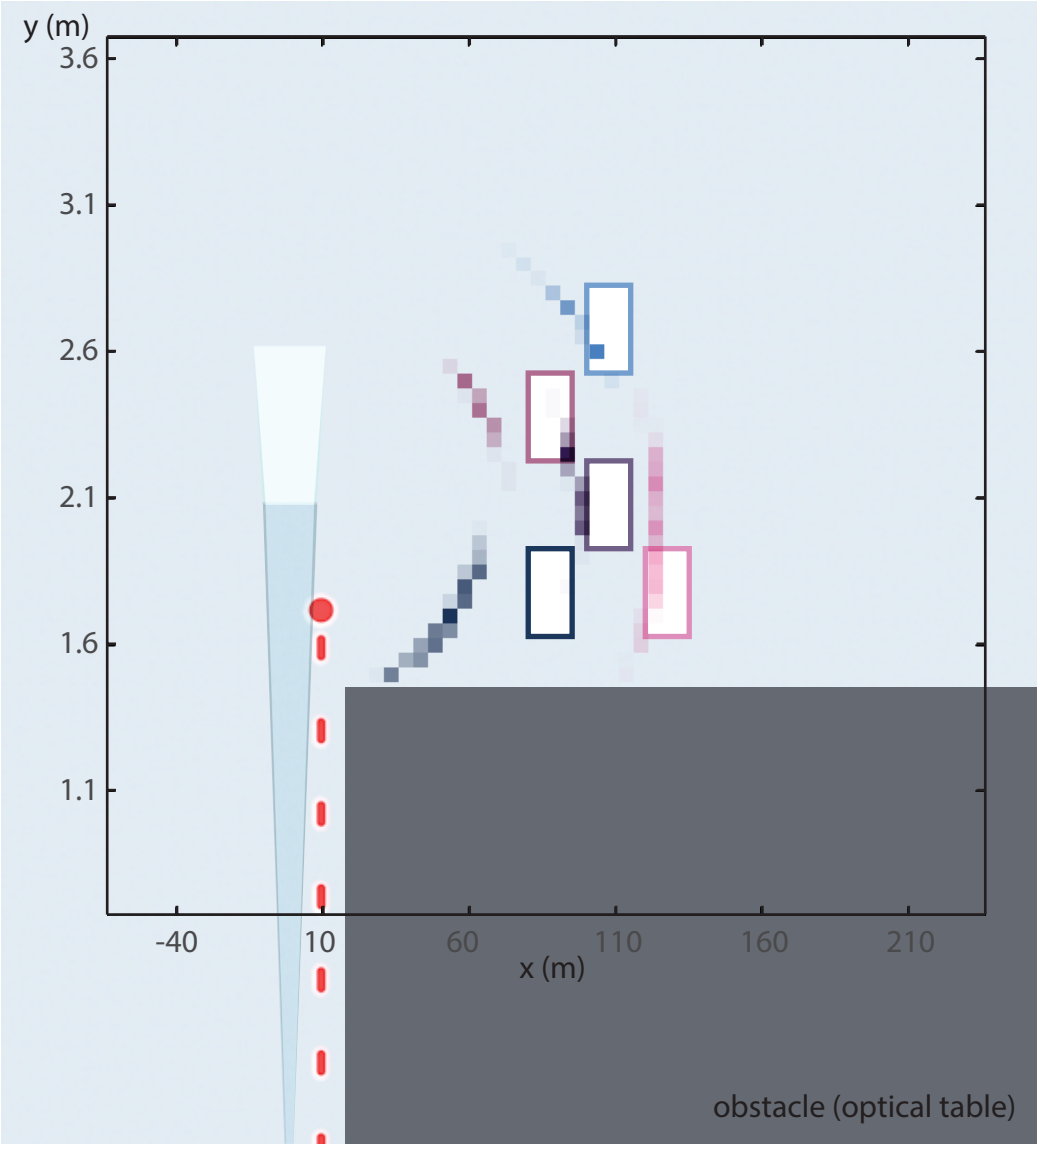

FIGURE 5.17: **Locating a person hidden from view.** A measurement was performed with the SPAD camera looking down at the floor 2.1 meters away, during a 30 second acquisition time. The recorded data was processed to find the signal time of arrival, and the position of the hidden person was retrieved with the algorithm detailed in section 5.2 using a 60x60 search grid with a pixel pitch of 5 cm. The probability densities are shown and the actual position of the hidden person is indicated by a rectangle of the corresponding contour colour.

method gives satisfactory results for targets located about 1 meter away from the field of view.

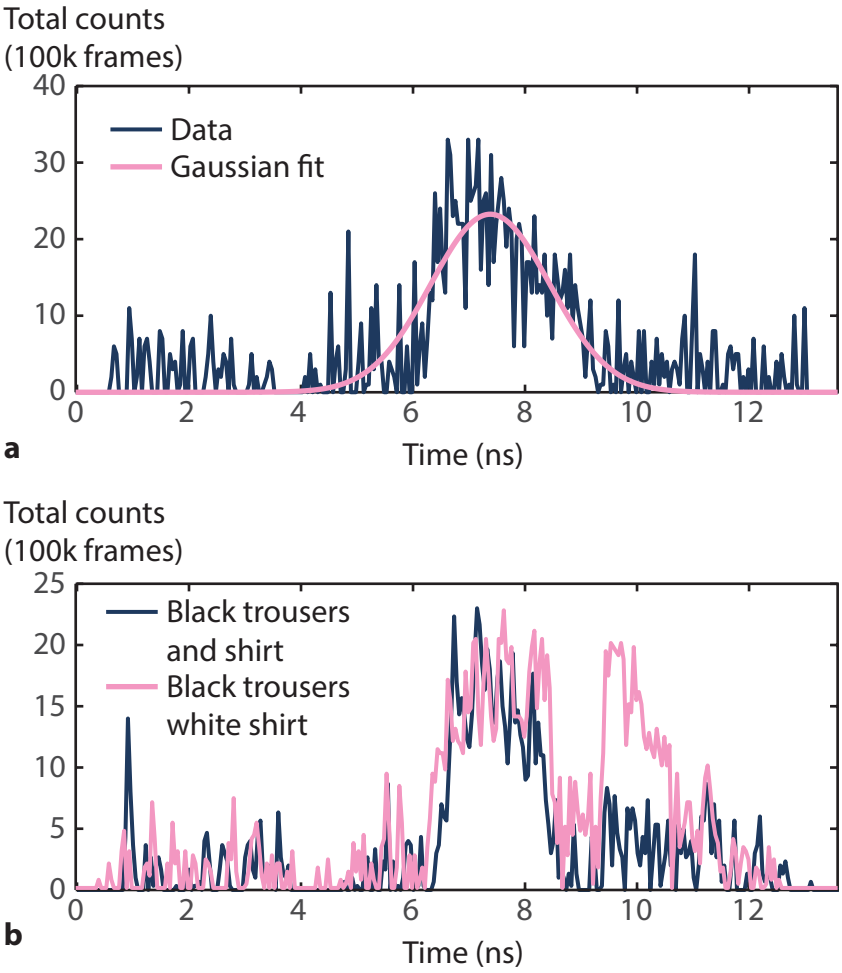

FIGURE 5.18: **Histograms recorded from a hidden person.** a) An example histogram corresponding to the target position ($x = 1.1$ m, $y = 2.1$ m) is shown after data processing, as well as its Gaussian fit. b) We also observe changes in the recorded histograms depending on the clothes worn by the target, suggesting a way to retrieve additional information about the hidden target or to distinguish multiple targets in one scene depending on their signature signal.

Interestingly, we can observe that the signal detected in the histograms (figure 5.18a) is much wider than the one detected in the proof-of-principle experiment where the target used was much smaller, which confirms that the width of the signal increases with the area of the hidden target. We also observe a clearly different temporal distribution of the target signal depending on the clothes worn by the hidden person. In figure 5.18b, we show a histogram acquired while the person was standing a meter away from the field-of-view wearing black trousers and a white shirt, and a histogram acquired at the same target location, but with the person wearing black trousers and a black shirt. It is interesting to note that we can observe these differences in the temporal structures of the signal, which could potentially be exploited for retrieving some three-dimensional information about the target or to help distinguish the targets if several are present in the scene.

A different geometry using the SPAD camera was also investigated, where we used a wall as a scattering surface rather than the floor. In this experiment, the camera was directed towards a wall placed 5 meters away, imaging a 37x37 cm portion of the wall at about a 1.3 meter height from the floor. The laser pulses were directed on a point 31 cm below the field of view. In this configuration, data also needed to be acquired over 100 000 frames of 300 μs, leading to a 30 seconds acquisition time. Since this is not sufficient to achieve our goal of real-time tracking, we limited ourselves to study here only the target distances at which a signal can be detected compared with the imaging geometry using the floor as a scattering surface, rather than concentrating our efforts to perform a thorough position retrieval experiment. Interestingly, we found that a signal could still be detected while the camera was farther away than in previous measurements ($r_c = 5$ m), and for a hidden person placed up to 3 meters away from the wall. The distance at which a signal could be detected depends on the exact location of the person with respect to the wall due to the fact that the scattering of the laser pulses at the wall is not isotropic but is scattered preferentially at a normal to the surface. In figure 5.19, we show the signals recorded by one pixel of the camera for the target positioned at three different positions. When the target is close to the wall, a signal can be acquired for a target distance up to 1.5 meter. If the target is distanced perpendicularly to the wall rather than along it, a target at a 3 meter distance produces a significant back-scattered signal. A third intermediary

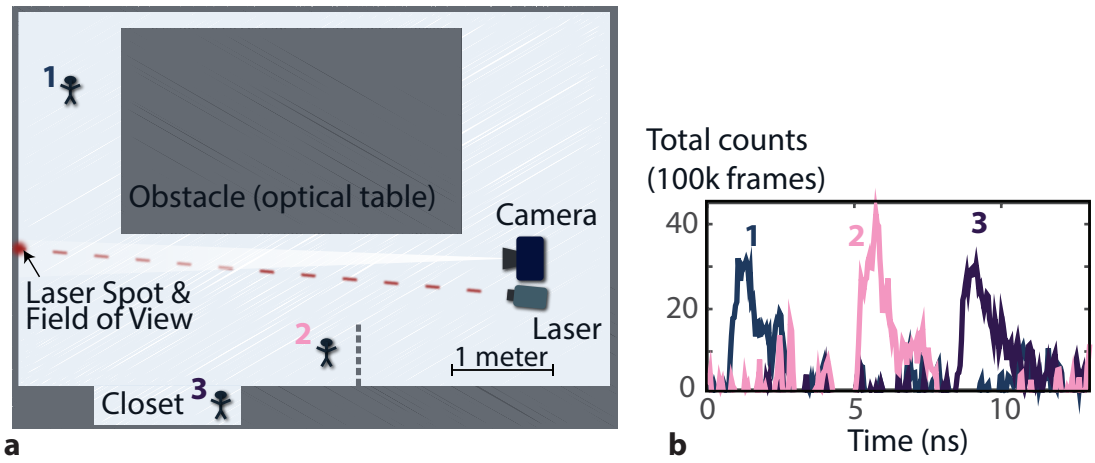

FIGURE 5.19: **Detecting a signal from a person hidden a few meters away from the camera.** a) We detect a signal over a 30 second acquisition period for a person hidden at three different positions in the laboratory. At position 1, the person is placed 1.5 meters away from the field of view, along the wall, and is hidden from the camera by an optical table. At position 2, the person is not physically hidden from the camera, but is out of the field of view in a position corresponding to a situation where a wall would be placed between the camera and the person, as illustrated by the dashed line. At this position, the person receives more of the scattered light from the wall, so it generates a good signal even at a distance of 3 meters from the wall. The last position corresponds to the person physically hiding in an open-door closet, 2.5 meters away from the field of view. b) We observe in the example histograms from one pixel of the SPAD camera that the level of signal is equivalent at all these positions, as the signal depends not only on the total distance between the target and the field of view, but also on its position with respect to the scattering intensity distribution of the laser pulses from the wall.

position showed the detection was possible for a target placed 2.5 meter away, for a person hidden in a closet (open door) in the laboratory. The results shown in figure 5.19 show that we can detect a signal from a person located to a distance up to 8 meters away from the camera (3 meters from the wall, plus 5 meters to the camera).

It is clear from these results that the current version of the SPAD camera is not efficient enough to be used for real-time tracking of hidden people at a few meters distance. Implementing some of the improvements discussed in section 5.3.3 will greatly improve the practicality of our non-line-of-sight tracking system exploiting the SPAD camera as the detector. In the meantime, we investigated an alternative solution: using a number of single-pixel single-photon avalanche detectors rather than a SPAD array. The advantages of this second approach is that we can reach 100% fill factor and have a lot of flexibility in placing the points imaged by the pixels far enough away to obtain an acceptable precision in the retrieval even for targets located far away from the scattering wall. The next section will therefore detail the measurements performed with single SPADs as the detectors used in our non-line-of-sight tracking system.

5.4 Single pixel implementation

To evaluate the potential of using single-pixel SPAD detectors to capture a signal from an object hidden from view, we performed a first experiment at a scale similar to that presented in section 5.3. The experimental setup is shown in figure 5.20: a room is created with black cardboards to a size of 102x102x76 cm. There is an opening in the front wall, mimicking a window into a room at a reduced scale, and the laser and detectors stand at a distance of 55 cm from the window. An object is placed into the room outside the line of sight of the detection system. Similarly to the method described above, we send ultrashort laser pulses on a scattering surface in line of sight of the detector: here, we direct the pulses through the window and onto a white surface on the back wall, 157 cm away from the detectors. Rather than imaging a squared portion of the wall with the SPAD camera, we use three single-pixel SPAD detectors (Excelitas SPCM-AQRH) and image a spot on the wall with each of them using one-inch collection optics and105 μm core optical fibres. The three pixels and the laser spots are spread over 24 cm horizontally and 3 cm vertically, as illustrated in figure 5.20. The three SPAD detectors are connected to an external TCSPC module (time-correlated single-photon counting), not shown in the image (Picoquant HydraHarp 400), which builds up histograms for each pixels with time bins of 64 ps and a jitter of 350 ps, characteristic of the SPAD detectors we use.

The use of single pixel detectors provides a significant gain in detection efficiency compared to the MF32 SPAD camera:

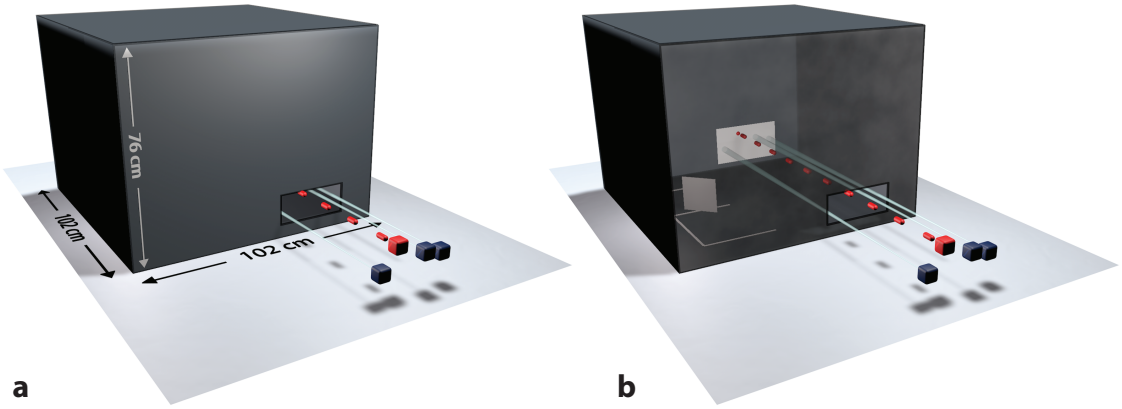

FIGURE 5.20: **Detecting hidden object with single-pixel SPAD detectors.** a) We re-create a situation, at reduced scale, where an object is hidden in a room (102x102x76 cm) and the only optical access in the room is through a window (28x12 cm). b) The same room is shown with the front wall transparent to see inside. We direct laser pulses at the back wall onto a white piece of paper. We image three points around the laser spot using collimators on optical fibres coupled to three separate single-photon-avalanche diodes (SPAD), and record histograms at those three points to collect the signal back-scattered by the hidden target.

1. The SPADs used are optimised for detection at 700 nm with an efficiency of 70%, the efficiency being still at 65% at a 780 nm wavelength.

2. The active area of the SPADs has a diameter of 180μm, and light is coupled to the detector with a 105μm fibre, giving us a significantly larger detection area compared to the SPAD array active areas.

3. The TCSPC module can operate at a count rate up to 12.5 million counts per second, which allows us to detect and time-tag many more photons in one second than with the SPAD camera, which is currently limited to a read-out rate of 3 kHz.

Given these improvements in detection efficiency, we can significantly lower the laser power used to send a signal in the hidden scene and detect a signal scattered back by the hidden object. We therefore use as a source a 780 nm pulsed laser diode emitting 50 ps pulses at 80 MHz with an average power of 2 mW.

A first experiment was performed by placing the hidden object, a 15x15 cm white screen, at 11 different positions in the hidden scene along a pre-defined E-shaped track, visible in figure 5.20b. Data was acquired by the SPAD detectors for one second. To process the three acquired histograms, we first subtract a background acquired without the object present in the environment - a strategy based on computing the median of multiple acquisitions could alternatively be used, as to not rely on any controlled access to the scene. In contrast to the data processing performed on the SPAD camera data, we do not use here any spatial correlation between the pixels to enhance the signal, but still perform a processing step in the temporal domain. We apply on each histogram a soft Savitzky-Golay filter, which fits successive sub-sets of the histograms with a low-order

polynomial in order to smooth the signals. We use here a 6^{th} order polynomial fit on subsets of 24 bins. An example of signals acquired for a given target position is shown in figure 5.21a after background subtraction and in figure 5.21b after the Savitzky-Golay filtering.

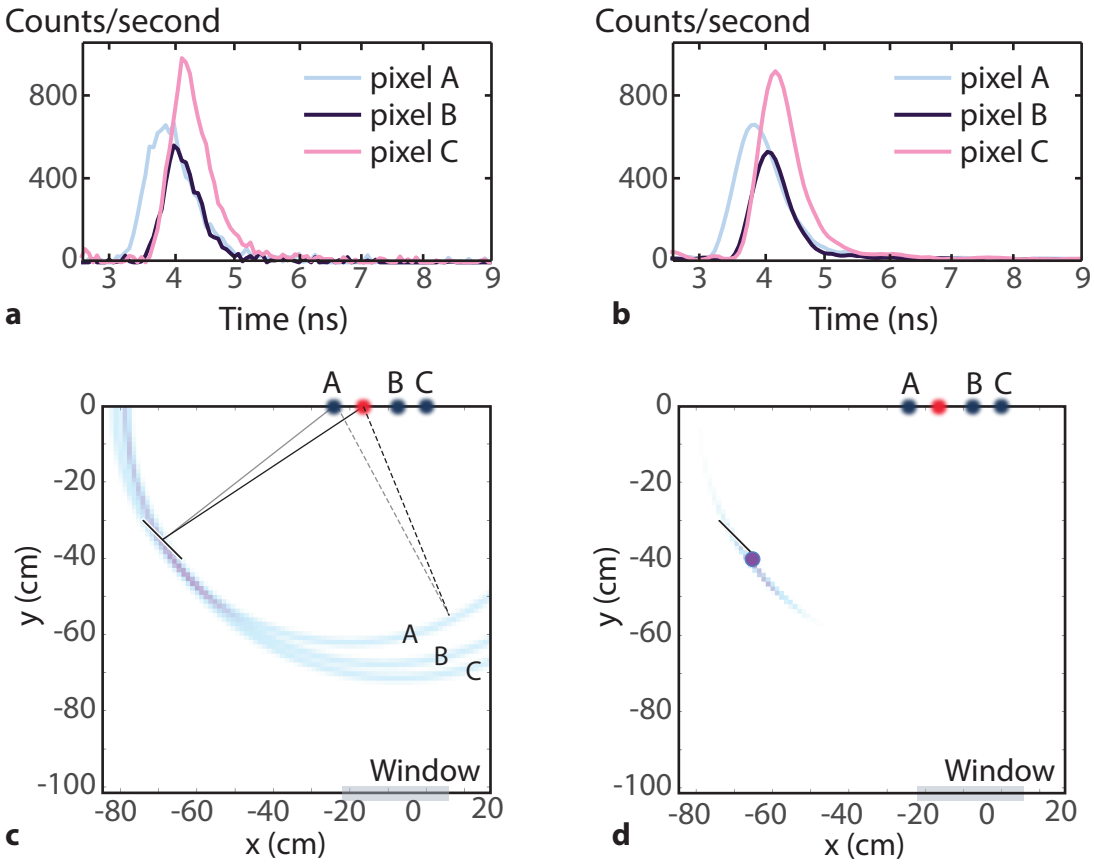

FIGURE 5.21: **Analysing the single pixel detectors data.** a) We first subtract a background to the measured data to isolate the signal coming from the hidden target. Here, the signal is clearly visible on all three pixels (A,B and C as indicated in (c) and (d)), and arrives at different times due to the geometry of the measurement setup. This detected signal is equivalent to that measured by individual pixels of the SPAD camera, imaging different points in the field of view. We collect here between a few hundred thousands and a million counts per second on the SPAD detectors, including 2000 counts coming from the hidden target within a 1 second acquisition time. b) We then smooth temporally each histogram and find the peak position of the signal, which we use as $\langle t \rangle_i$ in our reconstruction (equation 5.4). c) A probability density $P_i^{\text{ellipse}}(\vec{r}_o)$ is calculated for each of the three pixels. We show here the three ellipses overlaid, and see they intersect at the position of the object. The uncertainty σ_{t_i} used for this reconstruction was taken as the pixel pitch of the reconstruction to be the limited factor on the uncertainty (1 cm pixel grid). d) The multiplication of the three ellipses gives the joint probability density of finding the hidden object at position $\vec{r}_o$. The x and y coordinates are given with respect to the position of pixel B, taken arbitrarily as a reference position.

We then proceed to retrieve the hidden target position with the algorithm detailed in 5.2. We calculate the probability density ellipses $P_i^{\text{ellipse}}(\vec{r}_o)$ for each of the three pixels. An example of those three ellipses overlaid is shown in figure 5.21c. Here, since the signals acquired all show a good signal-to-noise ratio, and it is easier to monitor the signals

over three pixels than over 1024 pixels to make sure of this, we simply multiply together the probability density ellipses $P_i^{\text{ellipse}}(\vec{r}_o)$, without adding a uniform probability density to avoid contributions from bad measurements. We therefore obtain a joint probability density similarly to what we found in the experiments detailed in section 5.3 - an example of a retrieved $P_i(\vec{r}_o)$ is shown in figure 5.21d.

In this experiment, we calculate the probability densities over a search space covering a wide area in front of the scattering wall of 3×1.5 m, divided into 256×128 pixels. We also limit our retrieval to a 2D plane above the floor, using an educated guess on the height of the target, in order to simplify and accelerate the retrieval process. The results are shown in figure 5.22 where we show each probability densities for the eleven positions, the mean position calculated from these probability densities and the actual target position. We find a very good agreement between the retrieved position and the actual location of the target, with a precision of ±7 cm in x and ±6 cm in y on average. This uncertainty is consistent with the values expected while considering a lower number of pixels for a same spread of

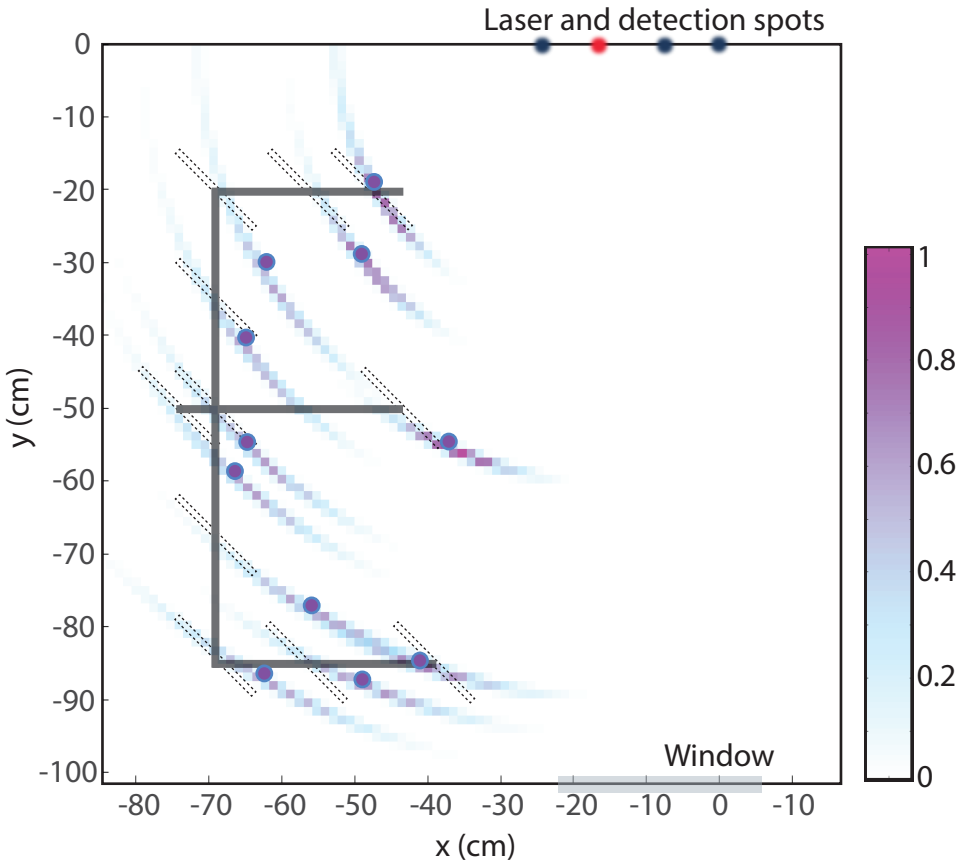

FIGURE 5.22: **Experimental results for locating a hidden object with single pixel SPADs.** We place the target at eleven distinct positions – the dashed rectangles show the actual positions and orientations of the target. The retrieved probability densities are shown, each normalised to a peak of 1 for better visualisation, as well as the position of their weighted mean. The experimental results show good agreement with the actual target positions, and we can locate a hidden target up to one meter away from the imaged portion of the wall, which is itself located 1.5 meters away from the imaging system.

the detection area in the imaging plane (see figure 5.13).

The data processing on all three pixels and retrieval algorithm takes only half a second. Given this very fast acquisition and processing time, the whole process was implemented to be performed directly on the computer as the data was acquired, to provide information about the hidden target in real-time. Having developed a real-time non-line-of-sight tracking system, we performed an acquisition and retrieval as the object was moving in the hidden space at a speed of 5.7 cm/s, updating the retrieved position every 1.5 seconds (1 second acquisition time + 0.5 seconds processing time). This measurement resulted in a movie where we can follow the motion of the target as its retrieved position is updated simultaneously to its motion (see appendix A). In figure 5.23, we show three frames from that movie, where the retrieved position is given as well as a frame from a video of the moving target, taken at a separate time with the front wall removed and the lights open to show the reference position and speed of the target. This result demonstrates that single-pixel SPAD detectors can be used to create a real-time non-line-of-sight tracking system, using only a 2 mW laser for tracking objects at small scales, located about 2.5 meters away from the detection system. The uncertainty is slightly increased with respect with using

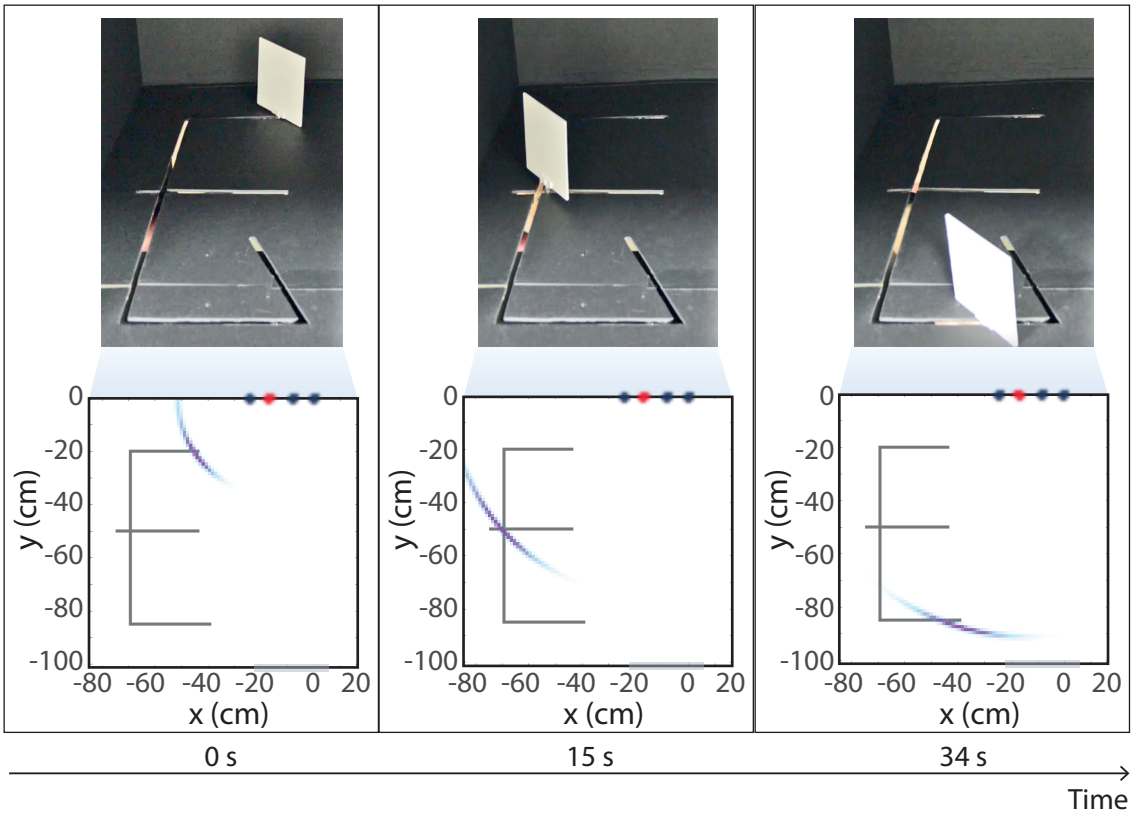

FIGURE 5.23: **Experimental results for tracking a hidden object with single pixel SPADs.** The data acquisition takes only 1 second and the processsing 0.5 seconds, so that our system is capable of acquiring and processing data in real-time as the target moves in the hidden space, actualising the retrieved position every 1.5 seconds. We show here the retrieved position at three different times, next to frames of a video that was acquired to show the motion of the target, moving along an E-shaped track at about 5.7 cm/s.

the SPAD camera, and the form factor of the detection system is larger (the three single pixel SPAD detectors and the separate TCSPC module occupy much more space than the very compact SPAD camera), but this implementation allows us to perform real-time acquisition and processing at larger distances with lower laser power, due to the higher detection efficiency reached with the single pixel detectors. These results are detailed in a manuscript under preparation by Chan $et\ al.$ [107].

5.4.1 Non-line-of-sight detection at large stand-off distances

To further explore the potential of this implementation, we set out to perform an experiment at a much larger scale, that is relevant in real-life scenarios to detect hidden people at large stand-off distances. The experimental conditions are shown in figure 5.24: we placed our detection system at one end of a corridor in our building, looking down at a screen 53 meters away, at the other end of a corridor. A person was positioned around the corner from the edge of the corridor, hidden from view.

Our detection system is here too based on the use of single-pixel SPAD detectors. In order to test the method at a wavelength in an eye-safe region of the electromagnetic spectrum, we use a 1532 nm laser (100 ps pulse emitted at 40 MHZ with an average power of 400 mW). The laser beam reaches the screen at the end of the corridor with a diameter of about 5 cm. The use of this wavelength outside the visible spectrum also allows us to block more of the ambient light by using the appropriate filter - we use here a highpass filter that blocks light below 1500 nm, as this is what we had available; using a bandpass filter of narrow bandwidth at the operating wavelength would be ideal. To detect a signal at this wavelength, we use an InGaAs/InP single-photon avalanche diode (ID Quantique ID230) combined with a TCSPC module (Picoquant PicoHarp 300) that provides histograms of detected photon arrival times with time bins of 4 ps and an impulse response function of 250 ps at a maximum count rate of 500 kHZ. The detection efficiency is about 10% and light is coupled onto the detector using a fibre of 62.5μm core and numerical aperture of 0.22.

To collect as much signal with the SPAD detector as possible, one wants to optimise two surface areas: the area over which the imaging system can collect light (diameter of the imaging system aperture), and the area that is imaged onto the screen by the SPAD detector. We here use a Schmidt-Cassegrain telescope of 20 cm diameter to couple light into the fibre, and image a spot size of about 3 cm diameter on the wall positioned 53 m away, onto the fibre core. These large collection and imaging diameters compensate for the loss in intensity as we move the detector away from the scattering wall, as discussed in section 5.3.2. Ideally, one wants to use imaging optics with a numerical aperture that matches the aperture of the fibre (see appendix D). The telescope used here has a smaller aperture than the fibre (f/10), so we use a second lens of 1 inch diameter and 30 mm

focal length to create a secondary image of the fibre before the telescope and gain a factor ∼4 in magnification. The highpass filter is positioned right at the output of the fibre as

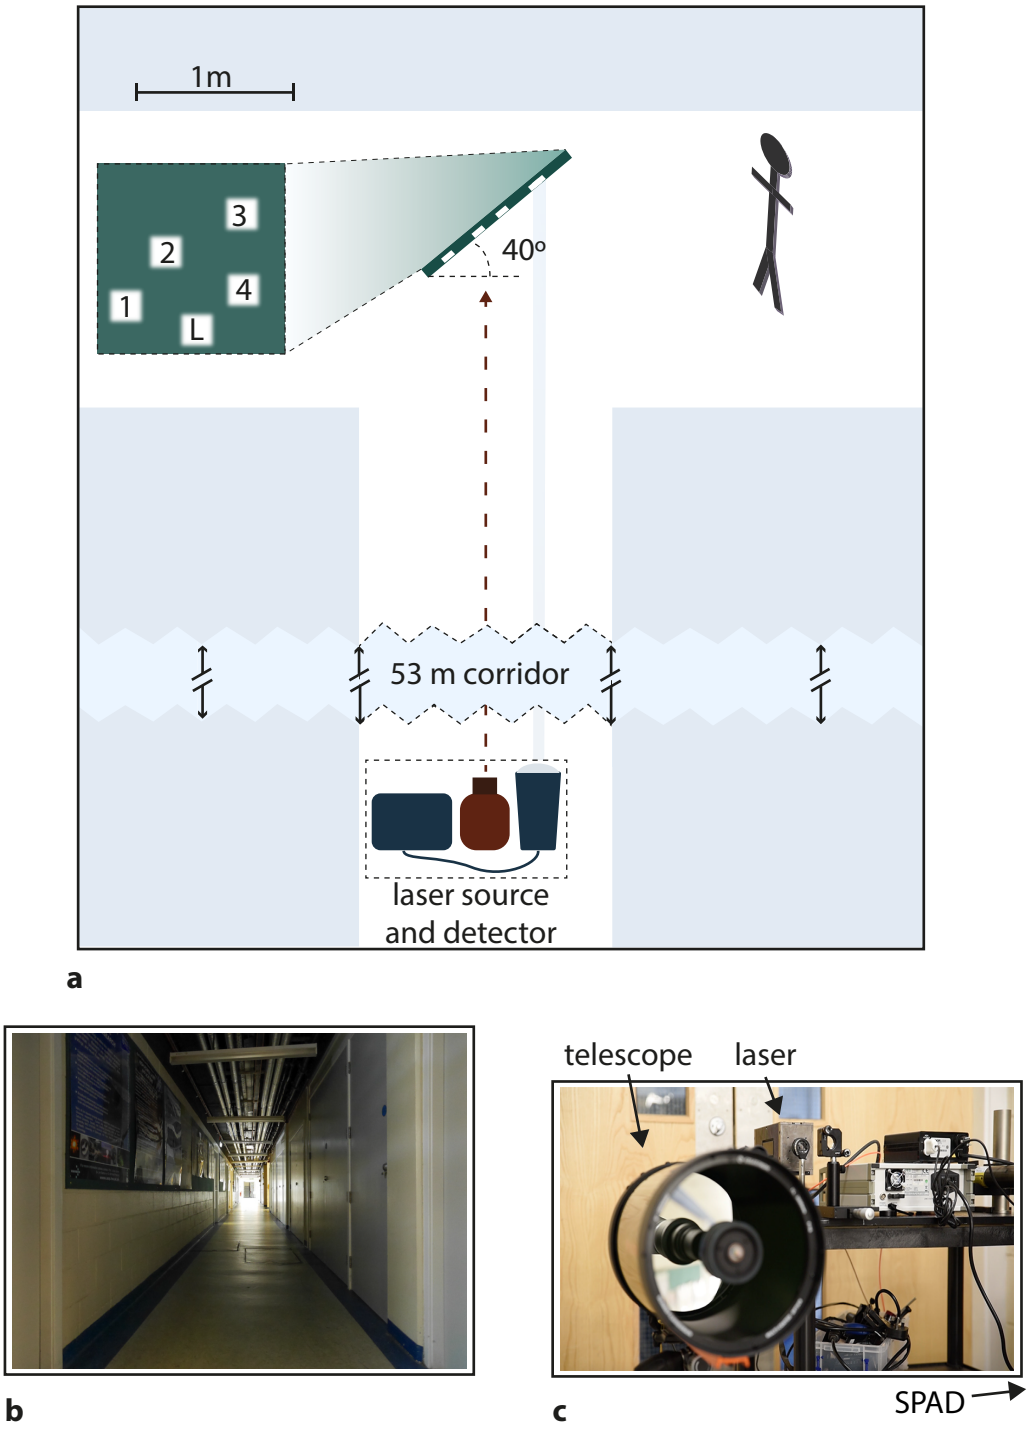

FIGURE 5.24: **Tracking hidden people down a corridor.** a) We install our detection system at the end of a corridor and place a screen at a 40° angle at the other end, to use as a scattering surface to send light around the corner and collect a signal coming back. A person is hiding around the corner at the end of the corridor. The SPAD detector images four spots consecutively, spread over a 75x75 cm area on the screen. White scattering surfaces are positioned on the screen at the laser spot and pixel positions to enhance the signal. b) This picture, taken at the end of the corridor where the laser and detector were installed, shows the length of the corridor and the ambient light conditions. c) We use a telescope to collect light and couple it onto a single-pixel SPAD detector, optimised at the operating laser wavelength of 1532 nm.

to remove as much signal from ambient light as possible. The lighting conditions in the corridor when the experiments were performed is shown in the picture in figure 5.24b; the main contribution to ambient light was the sunlight coming through large windows at the far part of the corridor. This was creating a noise level of about a thousand counts per second in the recorded histograms.

In a first measurement, a person was positioned at different distances from the edge of the corridor in the hidden scene. Data was acquired with the SPAD detector for one second at each of the four pixel positions, successively. In figure 5.25a, we show an example of the four histograms acquired at a given position of the hidden person. As in previous experiments, we see that the time of arrival of the signal, calibrated to correspond to the

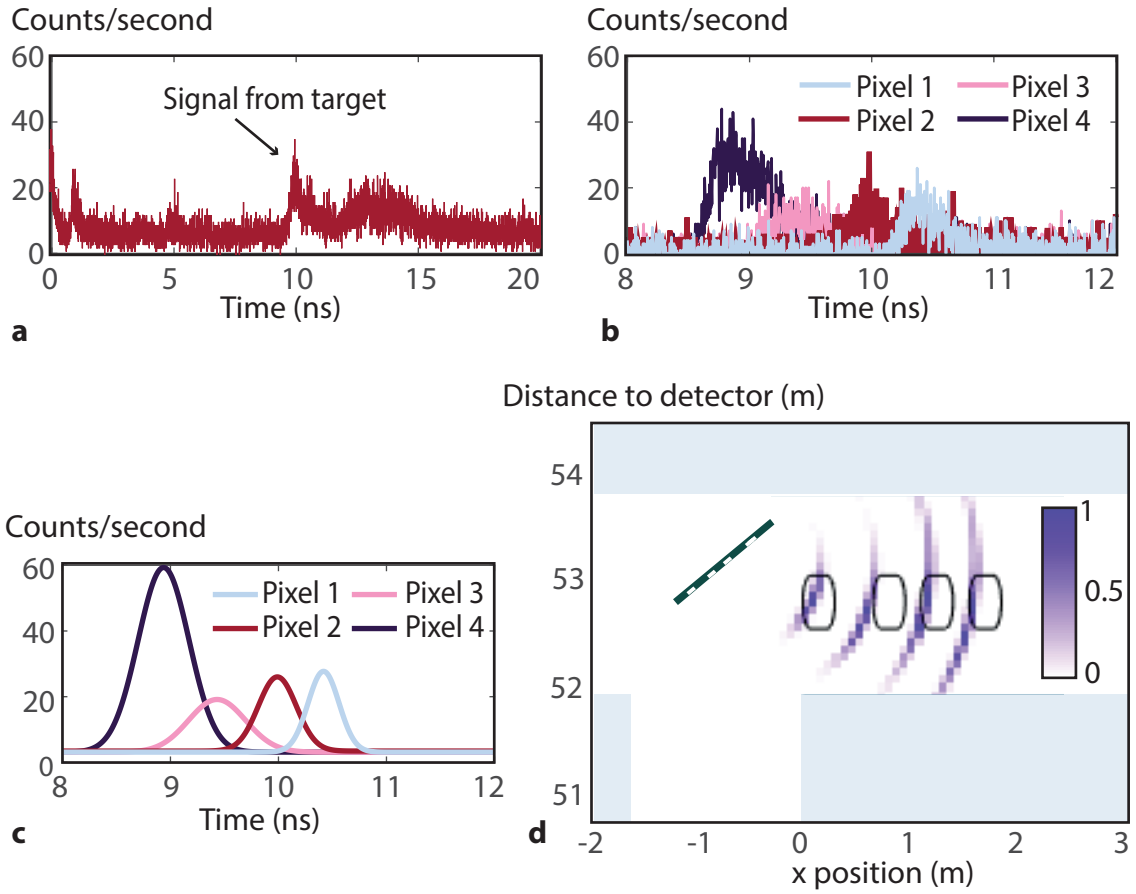

FIGURE 5.25: **Position retrieval of a person hiding around a corner 53 m away from the detector.** a) Example of a raw histogram acquired at pixel 2, which shows a signal from the target as well as from the other elements in the environment. b) Example of the four histograms acquired at the four pixel positions indicated in figure 5.24b, after background subtraction. We acquired data over 1 second and collect on average 60 000 counts in a second, including about 2000 counts coming from the target signal. c) We fit a Gaussian to these histograms to recover the time of arrival of the signal and its standard deviation, which we use in our position retrieval algorithm. d) We retrieve the position of a person hiding up to 1.8 meters away from the edge of the corridor and 53 meters away from the detector, using a search space of about 9x2 m with a resolution of 7.1 cm, on a plane at $z_o = 1.45$ m. The rounded rectangles show the actual positions of the person – our retrieval is in good agreement with the ground truth.

back and forth travel between the laser, the object and the pixel, changes as a function of the pixel position on the screen. We apply a background subtraction to the recorded data and fit a Gaussian to the histograms (figure 5.25c) to retrieve the arrival time $\langle t \rangle_i$ and its uncertainty σ_{t_i}. The position is retrieved by calculating the probability ellipses $P_i^{\text{ellipse}}(\vec{r}_o)$ for each pixel and multiplying them to obtain the joint probability density $P(\vec{r}_o)$. The probability densities obtained for four different positions of the hidden person are shown in figure 5.25d. We see that we can reliably retrieve the location of a person located up to 1.8 meters away from the edge of the corridor, with a precision always better than 50 cm (better for closer positions).

To test the ability of this system to detect multiple hidden targets, we acquired data with two people hidden at two distinct positions, separated by about one meter, as shown

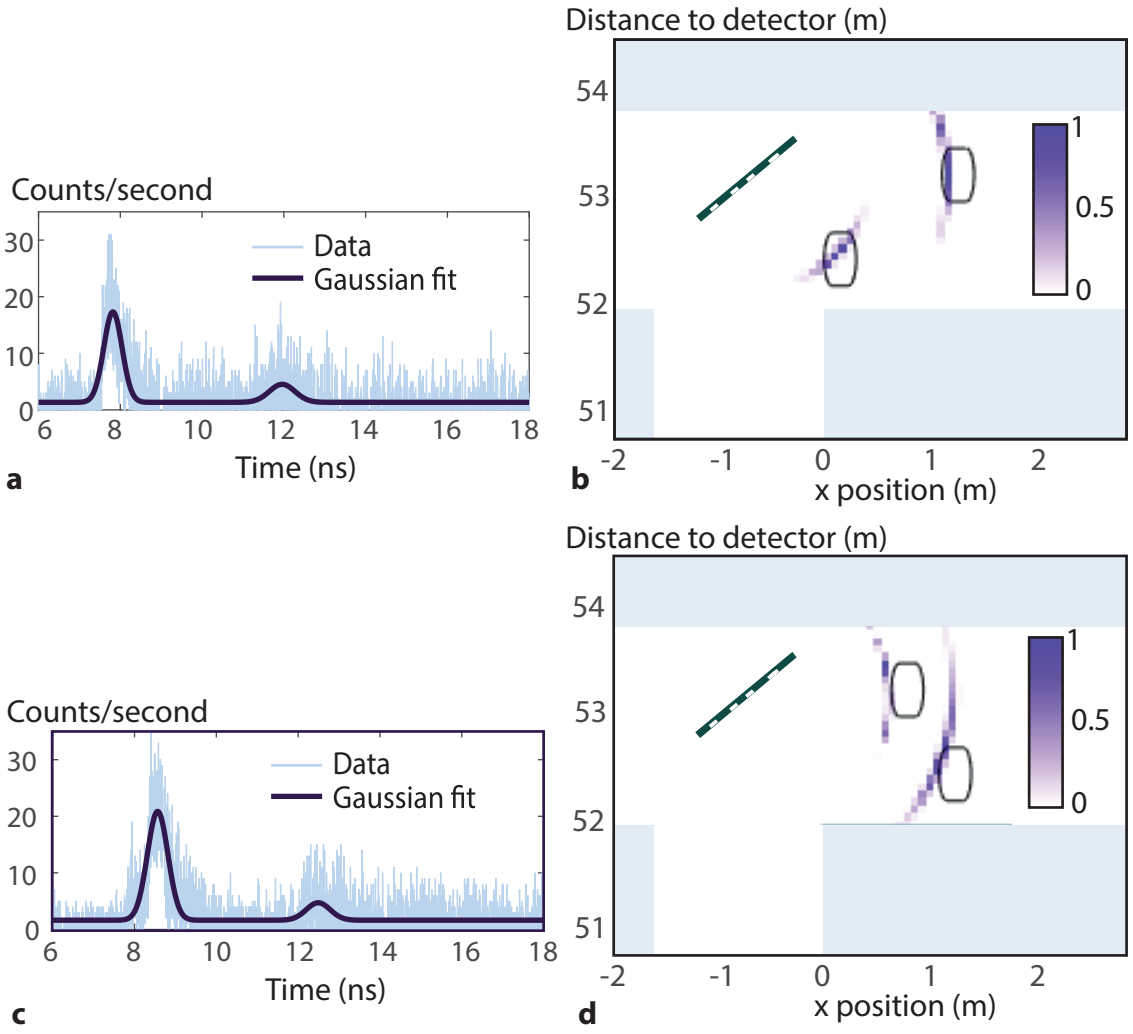

FIGURE 5.26: **Locating two people hidden around a corner 53 m away from the detector.** a) and c) In the background subtracted histograms recorded with two people positioned in the hidden scene, we can distinguish two signals coming from the two different targets. We perform a Gaussian fitting including two Gaussian peaks in the fit to retrieve the time of arrival of both targets. We can then use the position retrieval algorithm on the two targets separately to retrieve their position, shown in (b) and (d). We find a very good agreement between the retrievals and the actual positions of the two people, separated by less than a meter from each other.

in figure 5.26. As discussed in section 5.3.3, the requirement to locate multiple targets is to be able to distinguish their signal in our measurement. Here, as illustrated in figure 5.26a and 5.26c, we see two distinct signals in the recorded histograms after background subtraction, originating from the two people hidden around the corner. We can identify the signal from each target in all four histograms and compute the probability densities for each target. In figure 5.26b and 5.26d, we see that we can reliably locate two hidden targets with similar precision to the location of a unique target, and obtain good agreement with the actual target position.

We also find, similarly to the characteristics in the data discussed in figure 5.18, that we obtain different temporal signatures for the two different person that we are trying to locate. In figure 5.27, we show a histogram recorded at pixel 4 for the two people standing at the nearest position to the wall, one at a time. We find that they emit a signal at the same time of arrival but with a slightly different temporal distribution. One person was dressed with white trousers and a white shirt, and the corresponding signal is spread over a large temporal distribution and shows a double peak structure. The other person was dressed with jeans and a dark shirt, and the corresponding signal is both narrower and stronger. As discussed previously, this can be exploited to identify different targets and gain additional information about their shape and dimension. A deeper analysis of how this information can be optimally exploited, although it is outside the scope of this thesis, would be of great interest to evaluate how much information can be gathered with a limited number of pixels for detection.

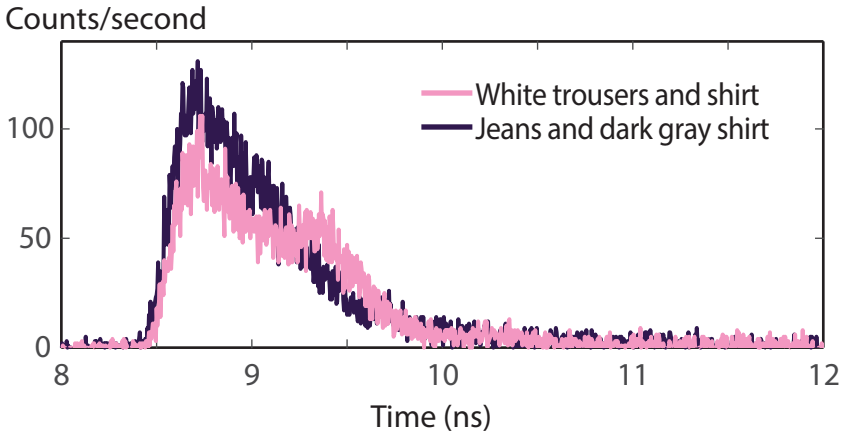

FIGURE 5.27: **Detected signal from two distinct people.** The temporal distributions associated with each person hiding around the corner show clear differences between the two targets, in both structure, strength of signal and temporal width.

We also performed a measurement with the scattering screen at a angle of $0°$, i.e. perpendicular to the incidence of the laser and the imaging system of the SPAD detector, as shown in figure 5.28. Due to the fact that the scattering from the wall is not perfectly isotropic and is preferential in a direction normal to the surface, such a geometry reduces

the signal reaching the target and coming back into the imaging spot of the SPAD detector. We therefore here integrated a signal over 10 seconds rather than 1 second, and limited our measurement to only two pixel positions as illustrated in figure 5.28, on either side of the laser spot and distanced by 70 cm. We then investigate the possibility of locating two people that are hidden on either side of the edges of the corridor. A recorded histogram on one pixel is shown in figure 5.28a, and by analysing the data in a similar fashion than we analysed the data used to obtain the results in figure 5.26, we can retrieve a probability density for each target. Here, since we restricted our measurement to two pixels, the uncertainty is much larger, so we show a thresholded probability density in figure 5.28b. Although this measurement shows a large uncertainty, it demonstrates that we can retrieve the positions of multiple targets even when they are located in different regions of the hidden scene. These results at long range are the object of a submitted manuscript by Chan *et al.* [108].

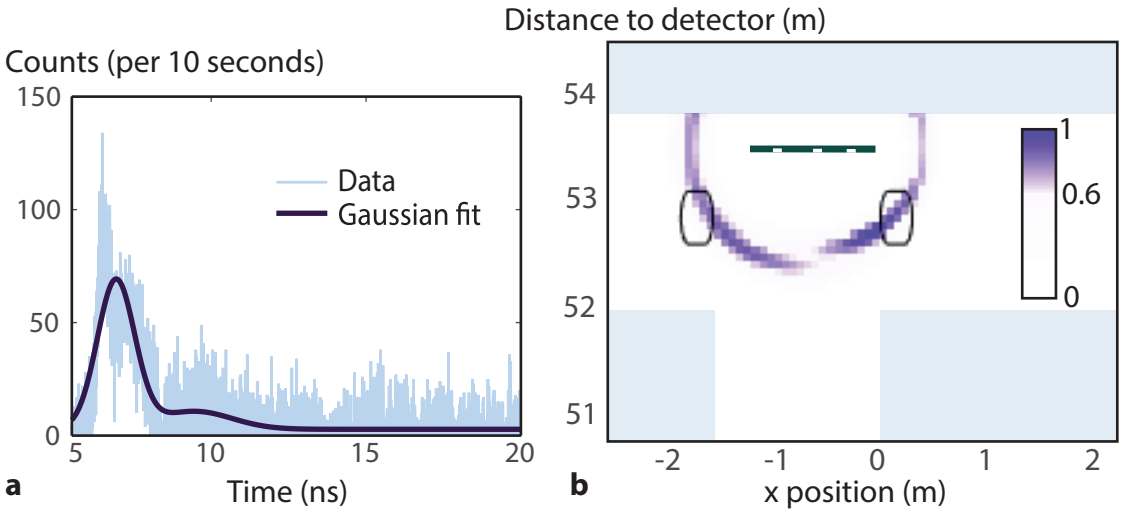

FIGURE 5.28: **Locating two people hiding on two different sides of the scattering wall.** a) An example of one of the two recorded histograms and its Gaussian fit shows the distinct signals coming from the two targets, where the signal coming from the target located on the same side of the laser spot than the pixel is much stronger. The relative intensities from the two targets are reversed for the second pixel position. b) We retrieve the positions of the two targets using our position retrieval algorithm and, although the result shows a large uncertainty due to the fact that we only use two pixel positions, we retrieve a position for the two targets on either side of the scattering wall that is consistent with the ground truth.

In this chapter, we have detailed significant contributions of this thesis to the field of non-line-of-sight detection. We first demonstrated a system capable of not only locating hidden objects, but also tracking their motion as the SPAD camera used for detection allows us to gather enough spatio-temporal information about the back-scattering signal from a hidden target to recover its position in only 3 seconds. We therefore showed a way to retrieve information about the position of a hidden target in real-time. To extend this method, the SPAD camera in its current version is unfortunately limited in both efficiency, due mainly to its low fill factor, and in read-out rate, due to its slow connection to

the computer. We therefore implemented a version of our detection method which relies only on a few single-pixel SPAD detectors rather than an array of them. The increased efficiency and photon counting rate were exploited to locate the position of people hiding around a corner at a very large stand-off distance (50 meters). This experiment greatly extended the range at which non-line-of-sight detection is commonly done to scales that are relevant for applications at human scales. This work lays the basis on which further work can be done to develop a system that can be used in real-life scenarios, in applications such as autonomous driving and search and rescue missions.

Chapter 6 – Conclusion

We have discussed in this thesis several contributions made to the state-of-the art in the field of imaging light in motion, and its use in detecting objects hidden from view. We have presented a new technique to capture light in motion, which relies on the use of a camera composed of a 32x32 array of single-photon avalanche diodes, developed by the CMOS Sensors and Systems group at University of Edinburgh, which can detect single photons and tag their arrival time with a resolution on the order of 100 picoseconds.

This allowed us to perform light-in-flight measurements of very low light intensity phenomena, which were not possible to observe and study directly with previously developed light-in-flight measurement methods. In particular, we showed the first recording of the propagation of laser pulses in air by detecting light scattered by air molecules in the path of the laser. We also detailed the qualitative and quantitative study of various phenomena, including laser-induced plasma, propagation in fibres and slow light in a rubidium cell. Our measurement system is compact and easy-to-use, making it practical to implement in a variety of scenarios and travelling with the detector itself, which we did for some experiments in collaboration with research groups outside the UK. The SPAD camera, on the other hand, does present some limitations. One of them is its limited spatial resolution, at 32x32 pixels. We have compensated for the low resolution images acquired either by software, by performing an interpolation on the data to create better visualisations of the events recorded, or with hardware, by scanning the camera either by the full width of its field of view, or by sub-pixel steps in order to acquire higher resolution images, of up to 160x160 pixels. It is likely that the next generations of SPAD arrays will incorporate more pixels than the current versions available. An alternative method which could be used in situations where one wants to perform a light-in-flight recording with higher spatial resolution would be to use another type of detector with high sensitivity and high temporal resolution, such as an ICCD camera. We have shown that this type of camera, although it is not as practical as the SPAD camera due to its higher cost and form factor, can record the propagation of laser pulses through air with a resolution of about 200 ps.

Another limitation of the presented method is its inherent requirement for averaging the signal over many laser pulses for a given event under study. There has been a lot of interest in the ultrafast imaging community for single-shot measurements and many methods have been developed to record events unfolding at the speed of light in a single pulse acquisition. With the SPAD camera, we are limited to studying repetitive events since we can only detect one photon per pixel per event. An interesting avenue to investigate would be to combine light-in-flight measurements done with the SPAD camera with techniques relying on the capture of only one photon per pixel in a measurement, such as first-photon imaging [109], to significantly lower the number of pulses over which we

need to integrate a signal to capture light in motion.

Another interesting route to study using the SPAD camera and its ability to record light in motion is the potential of using the spatio-temporal information it provides to look through a thick scattering medium and infer information about its internal structure or any hidden object in the medium. In this thesis, we presented some very preliminary results in this direction, but further work needs to be done to assess how much information can be retrieved from the data, and in what scenarios - medium dimensions and scattering properties, imaging geometries - this tool could be most useful.

The main application of capturing light in motion developed in this thesis is the ability to capture scattered signal from an object hidden from view. We have designed a measurement method relying on an ultrafast laser to send light around an obstacle and on a SPAD camera to detect a signal back-scattered by a hidden object, and developed a retrieval algorithm that allows us to infer the position of the object within a few seconds acquisition and processing time. This method allowed us to demonstrate that it is possible to track the motion of objects hidden from view using the data recorded with the SPAD camera. This demonstration was done in a small scale experiment, where a 30 cm high target was hidden about a meter away from the camera. Its position was recovered as it was moving at a speed up to 5.3 cm/s, with a precision of about 20% the target dimension in x and y. The scaling of this technique to bigger dimensions (larger targets farther away) was tested with the SPAD camera, but some factors of the currently available detector now limits the possibility to scale up this technique: the detection efficiency of our SPAD array is only 6.5% at the operating wavelength which, combined to the low-fill factor of the camera ($\sim 2\%$) limits our overall efficiency to 0.1%. Additionally, the count rate is limited by the connection between the camera and the computer at 3 kHz.

To overcome these current limitations, we investigated the use of three single-pixel SPAD detectors rather than the 32x32 pixel SPAD array in our technique. The detectors used in a first experiment have 65% detection efficiency at 780 nm, a 100% fill factor and a larger detection area and count rate. We performed an experiment at small scale and demonstrated that we can reliably track the position of a hidden target about 2 meters away from the detectors using a 2 mW laser. In this experiment, the acquisition was done in 1 second and the data processing and position retrieval only took 0.5 seconds. These two steps were implemented in a real-time tracking software, that updates the position of the hidden target every 1.5 s. The uncertainty on the measurement was slightly increased compared to the results obtained with the SPAD camera, but as we discussed, this can be compensated by a wider repartition of the spots imaged by the single pixel detectors.

Following the promising results obtained with the single-pixel implementation of our method, we performed an experiment at a much larger scale, to investigate a situation that

is closer to real-life applications of a non-line-of-sight tracking system. We performed measurements down a corridor, with the laser and SPAD pointing at a wall 53 meters away, to detect people hiding from view around the edge of the corridor. We showed that, using a 400 mW laser at 1532 nm, operating in an eye-safe region of the optical spectrum, we could accurately find the position of one or two people hidden from view with a precision of about 50 cm for a person located 1.8 meters away from the wall. This is the closest demonstration so far to a non-line-of-sight tracking system working in real-life scenarios – the measurement was done here with successive acquisition using only a single SPAD detector and changing its position on the wall, but a real-time system could be developed using multiple single-pixel detectors.

Our detection method provides uniquely the information about the hidden object position in a 2D plane. Although this information is sometimes sufficient, it would be interesting to investigate how much information can be acquired about the hidden scene and target given the data that we record. We have highlighted that our measurements already contain some information about the target, as differently dressed people show different temporal structure in their detected signal. A route for further development would be to develop ways to analyse this type of structure in order to gain additional information about the scene. The current techniques that can image scenes around a corner rely on the acquisition of more information than what we have here (more pixels and laser positions), but it would be very interesting to assess how much information about a 3D structure can be retrieved with low number of measurements, while keeping the real-time acquisition and processing capabilities of the system.

A non-line-of-sight tracking device would be a real asset in a variety of real-life scenarios. For example, it could be used to assess a room or area dangerous or impossible to access, to evaluate if people are in the room, how many people and where they are – this would provide valuable information, for example, in search and rescue missions, or hostage situations, to assess a situation from afar and gain as much information as possible without optical access to the area. It could also be installed on cars to detect incoming vehicles from around a blind corner and avoid accidents. These ultimate goals still require a lot of work to achieve, but this thesis laid the groundwork of a possible route to get there. We have showed that measurements of non-line-of-sight target position was possible on a person hiding 1.8 meters away from a corner, with the detection system placed 53 meters away from the edge of the obstacle. This result was obtained by using single-pixel SPAD detectors rather than the SPAD camera, given its current limitation. However, the SPAD camera provides the advantage of a very compact packaging which includes the SPAD detectors and the TCSPC module - the next generation of SPAD arrays promise to reach higher detection efficiency, fill factors and read-out rates: they would therefore be a tool of choice to perform non-line-of-sight tracking in real-world scenarios.

Appendix A – Light-in-flight and tracking videos

The videos referred to in this thesis have been made available online at the following links:

1. Video 1: Light-in-flight measurement of laser pulses propagating in air
   https://youtu.be/QUiEBgzB3ow
2. Video 2: Light-in-flight measurement of a femtosecond laser pulse
   https://youtu.be/iLN4liCKON8
3. Video 3: Light-in-flight measurement with an ICCD camera
   https://youtu.be/eLUT_4K6a7Q
4. Video 4: Light-in-flight measurement of a laser-induced plasma
   https://youtu.be/ajvsLo3H-TA
5. Video 5: Light-in-flight measurement of laser pulses propagating in an optical fibre
   https://youtu.be/fMurCQc9NhU
6. Video 6: Looking through scattering media: light transmitted through a hand and forearm
   https://youtu.be/q52NWWJC3Rk
7. Video 7: Tracking a moving hidden target with a SPAD camera
   https://youtu.be/2WMRDhBXbQw
8. Video 8: Tracking a moving hidden target with single pixel SPADs
   https://youtu.be/2m-pSl4NCbc

Appendix B – Increasing the spatial resolution of the SPAD camera

In order to increase the spatial resolution of the SPAD camera, we mount the camera on a 2D translation stage to move its position and acquire data at different positions. Two strategies have been considered in different experiments: moving the camera so that the field of view is shifted by the whole width of the field of view, and shifting the camera by a distance less than the width of the image of a single pixel. In the first case, different data acquisitions can be stitched together side by side to form a bigger image. In the second case, shifting the camera has the effect of shifting the active area of a single pixel over the whole area of that pixel. Since the active area has a diameter of 8 μm and the whole pixel has a width of 50 μm, we perform scans where the camera is shifted on a 5x5 grid of positions, that leads the active area to cover a 5x5 grid on the pixel area. We can therefore recover a larger image by stitching the data acquisitions side by side, pixel by pixel, as the active area on each pixel at different positions do not overlap with each other, as shown in figure B.1a.

To demonstrate the gain in image resolution and quality that can be attained with a sub-pixel scan, we image a simple scene by illuminating it with a diffuse femtosecond beam. We first acquire data at a single camera position to obtain a 32x32 pixel spatial resolution image - in figure B.1b, we show the integrated image over all time bins, after background subtraction. We then perform a second acquisition with the camera scanned over a 5x5 grid, with steps equal to $d_{\text{FOV}}/(32 \times 5)$, where d_{FOV} is the width of the FOV, which we divide by the number of pixels in the field of view and the size of the grid over which we want to scan. The integrated image obtained is shown in figure B.1c. This technique allows us to recover images and frames with 160x160 pixel spatial resolution. A last step we implemented to enhance the quality of the acquired image is to identify the hot pixels, that do not show any signal after background subtraction (black squares in figure B.1c) and apply a matrix completion algorithm to fill in the blank spaces. The pixels were identified using the acquired background as those with a high level of noise (high number of counts in the background). A matrix completion algorithm, provided by Francesco Tonolini, then filled the value of those pixels with a value calculated by computing the discrete cosine transform (DCT) of the image and reducing its values before computing the inverse cosine transform to find a smoothed image: the values of the hot pixels in the original image are then replaced by the values of the same pixels in the smoothed image. After iterations, we obtain an image where the black squares have been filled with a value rendering the resulting image smoother and easier to interpret, as shown in figure B.1d.

These steps can be applied on each frame of an acquisition to enhance the resolution and quality of time-resolved data.

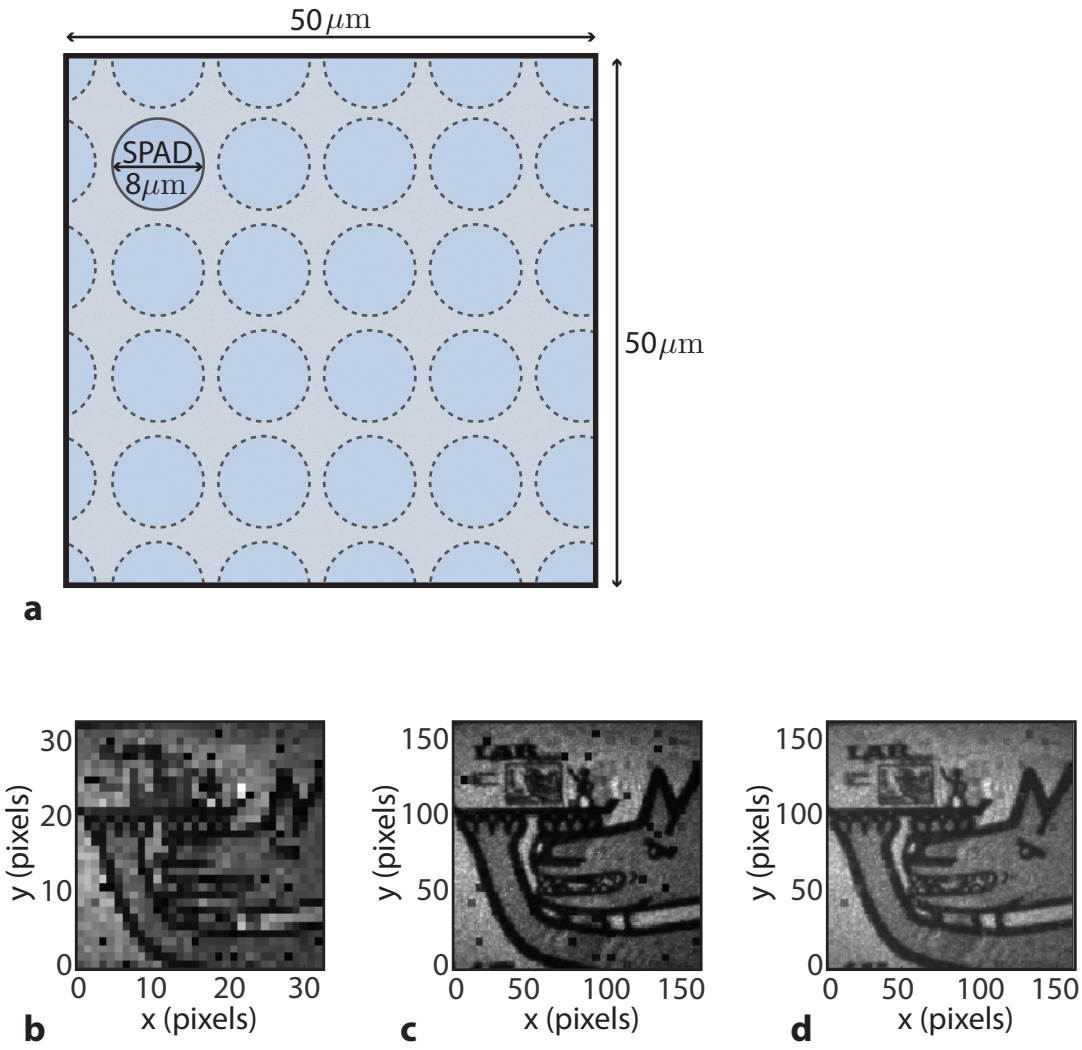

FIGURE B.1: **Enhancing the resolution of images acquired with the SPAD camera.** a) We shift the position of the camera in such a way that the active area of a pixel is shifted within the dimension of this pixel (in the image plane). As a result, the active area covers a 5x5 grid on the pixel area. Data is recorded at each position of that grid. b) An example of integrated image acquired with the SPAD camera only contains 32x32 spatial pixels. The object imaged here is a Thorlabs snack box illuminated by a femtosecond laser source. c) By scanning the camera, we can increase this resolution to 160x160 pixels. d) To get rid of the black squares in the image coming from the hot pixels, we identify these pixels and fill them with a value computed by an iterative algorithm based on smoothing the discrete cosine transform of the image.

Appendix C – Imaging geometry of SPAD camera-based non-line-of-sight tracking system

In the experiments described in section 5.3, we use a SPAD camera imaging a portion of the floor ahead to capture a signal from around a corner. The field of view on the floor has a trapezoidal shape, as illustrated in figure C.1. This field of view can be viewed as a projection of a square field of view in a plane perpendicular to the camera imaging axis onto the floor, which makes an angle θ with the perpendicular plane. To correct the measured time of arrivals for this projected field of view, and to correctly construct the probability densities in the search space, we need to know the geometry of this imaging setup. To do so, we here measure the values shown in figure C.1: the dimensions of the field of view and its distance in y and z to the camera position. In our experiments, we manually measure these distances, but in a more developed implementation of our system, these could be deduced given the angle of the camera and a line-of-sight range detector.

We first find the angle θ between the square field of view axis y and its projection axis y_p on the floor:

$$\theta = \frac{\theta^+ + \theta^-}{2} \tag{C.1}$$

with

$$\tan\theta^- = \frac{\mathrm{y_c}}{\mathrm{z_c}} \tag{C.2}$$

$$\tan\theta^+ = \frac{\mathrm{y_c} + \mathrm{y_{FOV}}}{\mathrm{z_c}} \tag{C.3}$$

This angle can be used to retrieve the positions of each pixel on the axis y_p. Given that 32 pixels are spread uniformly over a field of view of dimension b in the y axis, we have a series of y positions for the pixels of

$$y_i = \frac{\mathrm{b}}{32} i + \frac{\mathrm{b}/32}{2} \tag{C.4}$$

where i varies between 0 and 31 and the second term is added to use the middle of each pixel as its position. We then find the corresponding values in y_p, using the geometrical calculations illustrated in figure C.2:

$$\mathrm{y_{p_i}} = \mathrm{y_1} + \mathrm{y_2} = \mathrm{y}_i cos\theta + \mathrm{z_0}\tan\alpha = \mathrm{y} cos\theta + \mathrm{y}_i \sin\theta \frac{\mathrm{z_c} - \mathrm{z_0}}{\mathrm{y}_c + \mathrm{y_1}} \tag{C.5}$$

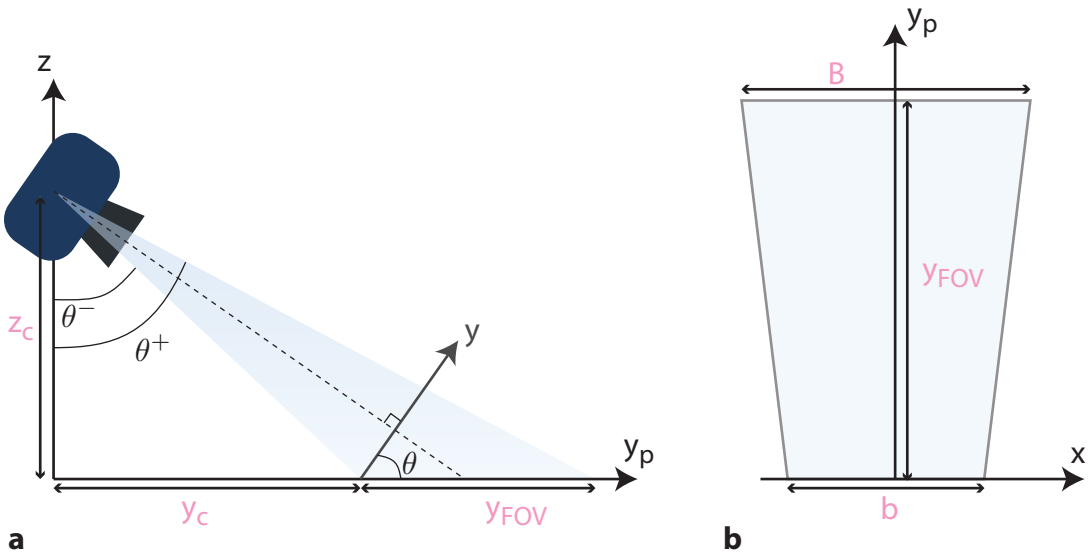

FIGURE C.1: **Imaging geometry.** a) The SPAD camera is imaging down and ahead at a portion of the floor. The field of view on the floor is a projection of a square field of view lying in a perpendicular plane (with axis (x, y)) to the camera imaging axis. b) The field of view therefore has a trapezoidal shape with height y_{FOV}, small base b and large base B. All distances in pink in the figure have been measured in the experimental setup.

For the x positions, we have 32 pixels on each line spread uniformly over a varying x width. We calculate the extra width at each line Δx as illustrated in figure C.2b to obtain the width of the field x_{range_i} at each positions y_{p_i}:

$$\Delta x_i = y_{p_i} \tan\beta = y_{p_i} \frac{(B-b)/2}{\text{y}_{\text{FOV}}} \tag{C.6}$$

$$x_{\text{range}_i} = b + 2\Delta x \tag{C.7}$$

Given the pixel positions on the x axis of x_j spread over a width of b at the bottom of the field of view, with j varying between 0 and 31, we obtain the positions x'_{ji} at different y_p positions of

$$x'_{ji} = x_j \frac{x_{\text{range}_i}}{b} \tag{C.8}$$

The values y_{p_i} and x'_{ji} given by equations C.5 and C.8 are the ones we use as $\vec{r}_i$ in our calculation of the probability density ellipses described in section 5.2.

As mentioned above, we also correct the measured time of arrivals as a function of the extra path taken from the projected field-of-view to the flat field of view represented in figure C.1a, assuming that light coming from all points of the flat field of view travels the same distance to the detector. We subtract this extra time

$$t = c \times \sqrt{y_2^2 + z_0^2 + (x' - x)^2} \tag{C.9}$$

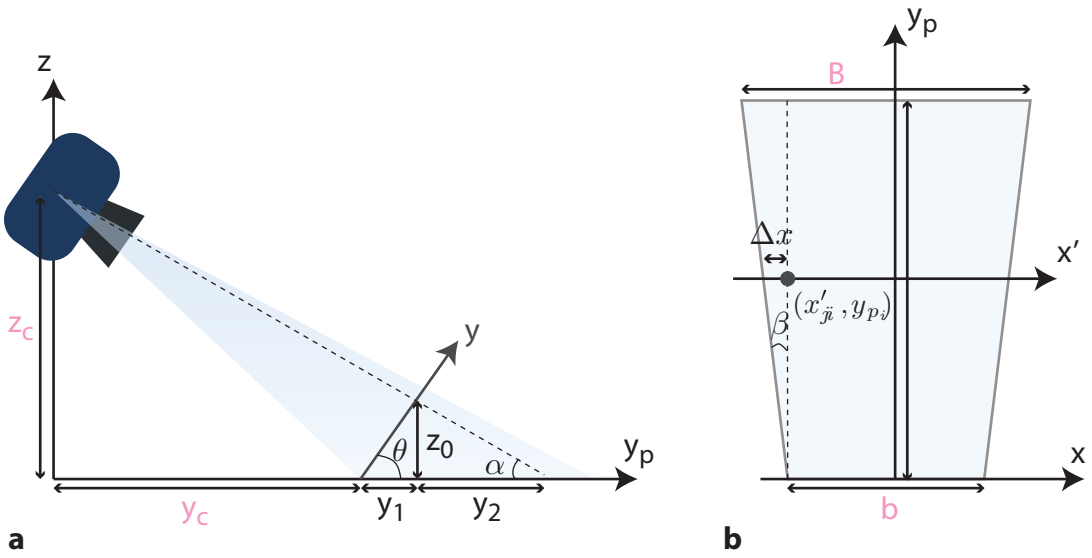

FIGURE C.2: **Calculating the (x, y) positions of the pixels.** a) The positions in y_p are calculated geometrically given the projection of the positions along y. b) The positions x' are calculated as being uniformly spread over a field of view width that varies between b and B depending on the y_p position.

to each pixel measured time of arrival. We also considered a correction that takes into account that all parts of the perpendicular field of view take different times to reach the detector, assuming pinhole-camera-like pathways, but this extra correction does not visibly change our position retrievals.

Appendix D – Optimising the collected signal in a non-line-of-sight tracking system

In our detection system, we use an optical fibre of fixed core diameter d_f and numerical aperture NA_f to couple light onto the single-pixel avalanche diode. Given these parameters, the signal detected will vary given the imaging system used to collect light into the fibre. The signal will depend on three main parameters:

1. The distance r_w between the detector and the scattering wall use to send light and detect a signal around an obstacle. As light scatters approximately as a spherical wave from a diffuse surface, we can consider a loss of signal in $1/r_w^2$ as the distance to the wall increases.

2. The collection optics diameter d_c, as a larger aperture will collect a larger portion of the scattered wave. The detected signal will be proportional to the area of the aperture and increases in d_c^2.

3. The spot size diameter d_w that is imaged on the scattering wall. This is given ideally by the size of the fibre core times the magnification factor M of the optical system $d_w = d_f \times M$. The collected signal increase with the surface of the spot, so is proportional to d_w^2.

We analyse here the effect of the imaging system in a simplistic model, assuming we are given a single lens as an imaging system, in order to investigate the optimal optics to use in the system in terms of focal length and collection aperture. Given a fibre of fixed aperture, the imaging system should then be chosen to match this numerical aperture in order to optimise the detected signal. The fibre core will be positioned in the image plane associated with the object-distance of r_w. All these parameters are illustrated in figure D.1. Light will be accepted in the fibre at a maximum angle given by its NA:

$$\theta = \sin^{-1} \text{NA}_f. \tag{D.1}$$

To fully take advantage of this collection angle, the diameter of the collection optics d_c should be optimised as

$$d_c = 2 \times r_f \tan\theta \tag{D.2}$$

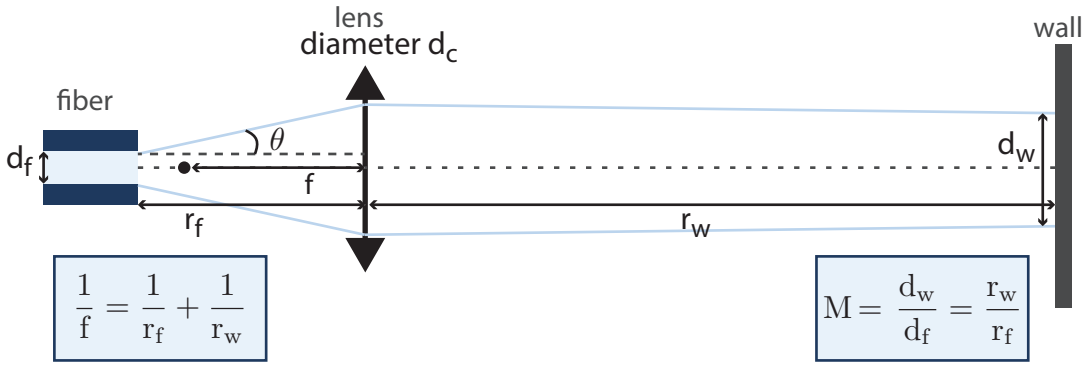

FIGURE D.1: We treat our imaging system as a simple lens following the lens equation to find the expected magnification of a collection optics of given focal length f and the diameter d_c that will optimise the collection into the accepting angle θ of the optical fibre.

where r_f is the distance of the image plane with respect to the lens, where the fibre is located. This distance, given a fixed object distance r_w, depends on the focal length f of the imaging system, following the lens equation:

$$r_f = \frac{1}{\frac{1}{f} - \frac{1}{r_w}} \tag{D.3}$$

We thus find that the optimal collection diameter is dependant on the focal length as:

$$d_c = 2 \times r_f \tan\theta = \frac{2\tan\theta}{\frac{1}{f} - \frac{1}{r_w}} \tag{D.4}$$

On the other hand, we also want to maximise the spot size on the wall which corresponds to the image of the fibre core. As mentioned above, this depends on the magnification of the system, in turn given by:

$$M = \frac{r_w}{r_f} = \frac{r_w}{\frac{1}{\frac{1}{f} - \frac{1}{r_w}}} = \frac{r_w}{f} - 1. \tag{D.5}$$

The imaging spot size diameter d_w therefore also depends on the focal length of the imaging system:

$$d_w = d_f \left(\frac{r_w}{f} - 1 \right). \tag{D.6}$$

To optimise the signal detected given a fixed distance to the wall r_w, we need to optimise the product $d_c^2 \times d_w^2$. By combining equations D.4 and D.6, we find that

$$d_c^2 \times d_w^2 = \left(2\tan\theta d_f r_w\right)^2, \tag{D.7}$$

which does not depend on the focal length f. There is therefore a range of focal lengths that can be used to optimise the signal collection, given that the collection diameter matches the collection aperture of the optical fibre. For a smaller focal length, the optimal collection diameter of the lens will be smaller, but it will produce a larger image

of the fibre core on the wall. This trend is shown in figure D.2a, where the optimal collection diameter d_c is calculated for varying focal length and a fixed distance to the wall of 53 meters, as in our experiment. We also use the fibre parameters used in our experiment for this calculation: NA = 0.22 and $d_f = 62.5\mu$m.

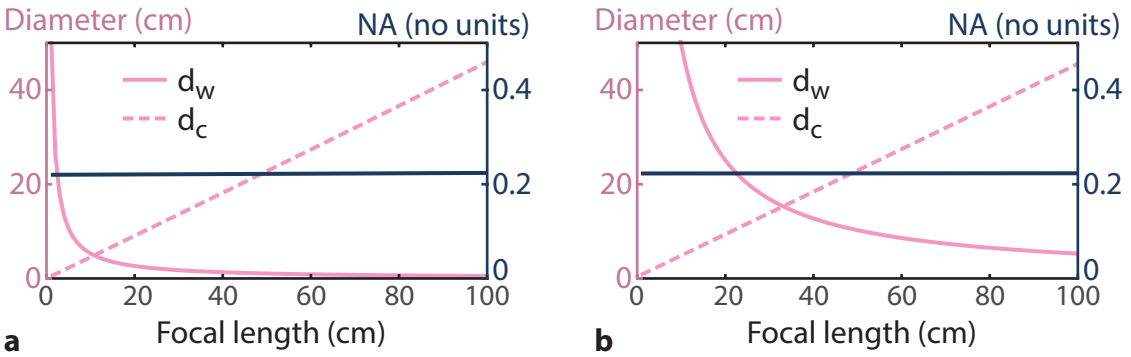

FIGURE D.2: **Optimal collection optics dimensions as a function of focal length.** We calculate that, given an imaging system modeled as an ideal lens, the optimal diameter of the collection optics increases as a function of focal length used to match the numerical aperture of the fibre, while the size of the imaging spot on the wall decreases with focal length. We show the optimal dimensions for a) a wall distance of 53 m and b) a wall distance of 500 m. In both case, the imaging systems need to have a numerical aperture of $\sim$0.22 to match the NA of the fibre, which corresponds to a f-number of 2.2. To obtain a spot size of about 10 cm on the wall, a longer focal length imaging system and a larger collection diameter is needed for a wall placed further away.

The other parameter to be aware of when choosing an imaging element is that the size of the imaged spot on the wall needs to be within the dimension of the temporal resolution: for example, a detected signal with a temporal width of 1 ns corresponds to a dimension of 30 cm. If the imaging spot is larger than this dimension, the signal will be blurred by the fact that it is collected over a large surface. Keeping the size to a surface that will not affect the resolution of the time-of-arrival detection is therefore essential, which sets a lower limit to the focal length that should be used to collect light into the optical fibre. In figure D.2b, we show the same calculations as in figure D.2a, but for a distance to the wall of $r_w = 500$ m. In this situation, a much longer focal length is desirable to keep the imaging spot on the wall to within the temporal resolution of our system.

We also find that the detected signal can increase in r_w^2 when optimising the collection optics, which compensates perfectly for the loss of intensity in $1/r_w^2$. An ideal system should therefore be able to collect a signal from a hidden object no matter the stand-off distance, but practical considerations on the design of imaging systems limit the feasibility up to a certain point: at 500 m, to get a spot size of 10 cm, one needs a 20 cm diameter system with a 50 cm focal length and at 1 km, a 45 cm collection diameter with a 1 meter focal length, which limits in practice the possibility of optimising the collected signal by matching the collection optics to the fibre numerical aperture. To collect as much signal as possible, one can therefore find the imaging optics commercially available (camera

objective, telescope, ...) whose specifications are as close as possible to these optimum values while respecting the size and budget constraints of the system.

Bibliography

[1] The Slow Mo Guys. *The Slow Mo Guys*. https://www.youtube.com/user/theslowmoguys. Accessed: 2016-09-14.

[2] National Geographic. *The World's fastest runner*. http://video.nationalgeographic.com/video/magazine/ngm-cheetah-running. Accessed: 2016-09-14.

[3] E. Jussim & H. Edgerton. *Stopping Time: The Photographs of Harold Edgerton* (Abrams, New York, 1987).

[4] K. Goda, K. K. Tsia & B. Jalali. *Serial time-encoded amplified imaging for real-time observation of fast dynamic phenomena*. Nature **458**, 1145–1149 (2009).

[5] H. Mikami, L. Gao & K. Goda. *Ultrafast optical imaging technology: principles and applications of emerging methods*. Nanophotonics **5**, 98–110 (2016).

[6] N. Abramson. *Light-in-flight recording: high-speed holographic motion pictures of ultrafast phenomena*. Appl. Opt. **22**, 215–232 (1983).

[7] G. Gariepy *et al. Single-photon sensitive light-in-fight imaging*. Nat Commun **6** (2015).

[8] R. Warburton *et al. Observation of laser pulse propagation in optical fibers with a SPAD camera*. submitted to Scientific Reports, in review (2016).

[9] A. Velten *et al. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging*. Nat Commun **3**, 745 (2012).

[10] G. Gariepy, F. Tonolini, R. Henderson, J. Leach & D. Faccio. *Detection and tracking of moving objects hidden from view*. Nat Photon **10**, 23–26 (2016).

[11] J. Richardson *et al. A 32×32 50ps resolution 10 bit time to digital converter array in 130nm CMOS for time correlated imaging*. In *2009 IEEE Custom Integrated Circuits Conference*, 77–80 (2009).

[12] K. Wilson *et al. Slow light in flight imaging*. Submitted to PRA (2016).

[13] B. Heshmat, G. Gariepy, J. Leach, R. Raskar & D. Faccio. *SPAD cameras for biomedical imaging: promis and problems*. In *CLEO: Applications and Technology*, AW4O–3 (Optical Society of America, 2016).

[14] J. A. Giordmaine, P. M. Rentzepis, S. L. Shapiro & K. W. Wecht. *Two photon excitation of fluorescence by picosecond light pulses*. Applied Physics Letters **11**, 216–218 (1967).

[15] M. Duguay & J. Hansen. *Ultrahigh-speed photography of picosecond light pulses*. IEEE Journal of Quantum Electronics **7**, 37–39 (1971).

[16] M. A. Duguay & J. W. Hansen. *An ultrafast light gate*. Applied Physics Letters **15**, 192–194 (1969).

[17] M. A. Duguay & A. T. Mattick. *Ultrahigh Speed Photography of Picosecond Light Pulses and Echoes*. Appl. Opt. **10**, 2162–2170 (1971).

[18] N. Abramson. *Light-in-flight recording by holography*. Opt. Lett. **3**, 121–123 (1978).

[19] N. Abramson. *Time reconstructions in light-in-flight recording by holography*. Appl. Opt. **30**, 1242–1252 (1991).

[20] N. H. Abramson & K. G. Spears. *Single pulse light-in-flight recording by holography*. Appl. Opt. **28**, 1834–1841 (1989).

[21] T. Kubota, K. Komai, M. Yamagiwa & Y. Awatsuji. *Moving picture recording and observation of three-dimensional image of femtosecond light pulse propagation*. Opt. Express **15**, 14348–14354 (2007).

[22] T. Kakue *et al. Moving Picture Recording and Observation of Visible Femtosecond Light Pulse Propagation*. Japanese Journal of Applied Physics **50**, 050205 (2011).

[23] T. Kakue *et al. Digital Light-in-Flight Recording by Holography by Use of a Femtosecond Pulsed Laser*. IEEE Journal of Selected Topics in Quantum Electronics **18**, 479–485 (2012).

[24] A. Velten, E. Lawson, A. Bardagjy, M. Bawendi & R. Raskar. *Slow Art with a Trillion Frames Per Second Camera*. In *ACM SIGGRAPH 2011 Talks*, SIGGRAPH '11, 44:1–44:1 (ACM, New York, NY, USA, 2011).

[25] A. Velten *et al. Femto-photography: Capturing and Visualizing the Propagation of Light*. ACM Trans. Graph. **32**, 44:1–44:8 (2013).

[26] Hamamatsu. *Guide to streak cameras*. https://www.hamamatsu.com/resources/pdf/sys/SHSS0006E_STREAK.pdf. Accessed: 2016-09-14.

[27] L. Gao, J. Liang, C. Li & L. V. Wang. *Single-shot compressed ultrafast photography at one hundred billion frames per second*. Nature **516**, 74–77 (2014).

[28] F. Heide, M. B. Hullin, J. Gregson & W. Heidrich. *Low-budget Transient Imaging Using Photonic Mixer Devices*. ACM Trans. Graph. **32**, 45:1–45:10 (2013).

[29] S. Foix, G. Alenya & C. Torras. *Lock-in Time-of-Flight (ToF) Cameras: A Survey*. IEEE Sensors Journal **11**, 1917–1926 (2011).

[30] A. Kadambi *et al. Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles*. ACM Trans. Graph. **32**, 167:1–167:10 (2013).

[31] NakagawaK. *et al. Sequentially timed all-optical mapping photography (STAMP)*. Nat Photon **8**, 695–700 (2014).

[32] M. Clerici *et al. Observation of image pair creation and annihilation from superluminal scattering sources*. Science Advances **2** (2016).

[33] M. Laurenzis, J. Klein & E. Bacher. *Relativistic effects in imaging of light in flight with arbitrary paths*. Opt. Lett. **41**, 2001–2004 (2016).

[34] J. S. Massa, A. M. Wallace, G. S. Buller, S. J. Fancey & A. C. Walker. *Laser depth measurement based on time-correlated single-photon counting*. Opt. Lett. **22**, 543–545 (1997).

[35] R. H. Hadfield. *Single-photon detectors for optical quantum information applications*. Nat Photon **3**, 696–705 (2009).

[36] E. Charbon, M. Fishburn, R. Walker, R. K. Henderson & C. Niclass. *SPAD-Based Sensors*, 11–38 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013).

[37] E. Charbon. *Single-photon imaging in complementary metal oxide semiconductor processes*. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **372** (2014).

[38] F. Zappa, S. Tisa, A. Tosi & S. Cova. *Principles and features of single-photon avalanche diode arrays*. Sensors and Actuators A: Physical **140**, 103 – 112 (2007).

[39] W. Becker. *Advanced time-correlated single photon counting applications* (Springer, 2015).

[40] M. Gersbach *et al. A parallel 32-by-32 time-to-digital converter array fabricated in a 130 nm imaging CMOS technology*. In *ESSCIRC, 2009. ESSCIRC '09. Proceedings of*, 196–199 (2009).

[41] G. Buller & A. Wallace. *Ranging and Three-Dimensional Imaging Using Time-Correlated Single-Photon Counting and Point-by-Point Acquisition*. IEEE Journal of Selected Topics in Quantum Electronics **13**, 1006–1015 (2007).

[42] J. Lakowicz, H. Szmacinski, K. Nowaczyk & M. Johnson. *Fluorescence lifetime imaging of calcium using Quin-2*. Cell Calcium **13**, 131 – 147 (1992).

[43] R. R. Dunca, A. Bergmann, M. A. Cousin, D. K. Apps & M. J. Shipston. *Multi-dimensional time-correlated single photon counting (TCSPC) fluorescence lifetime imaging microscopy (FLIM) to detect FRET in cells*. Journal of Microscopy **215**, 1–12 (2004).

[44] M. W. Fishburn. *Fundamentals of CMOS single-photon avalanche diodes* (fishburn, 2012).

[45] D. Mosconi, D. Stoppa, L. Pancheri, L. Gonzo & A. Simoni. *CMOS Single-Photon Avalanche Diode Array for Time-Resolved Fluorescence Detection*. In *2006 Proceedings of the 32nd European Solid-State Circuits Conference*, 564–567 (2006).

[46] M. Vitali *et al. A Single-Photon Avalanche Camera for Fluorescence Lifetime Imaging Microscopy and Correlation Spectroscopy*. IEEE Journal of Selected Topics in Quantum Electronics **20**, 344–353 (2014).

[47] D. Bronzi *et al. CMOS SPAD pixels for indirect time-of-flight ranging*. In *IEEE Photonics Conference 2012*, 22–23 (2012).

[48] C. Niclass & E. Charbon. *A single photon detector array with 64 times;64 resolution and millimetric depth accuracy for 3D imaging*. In *ISSCC. 2005 IEEE International Digest of Technical Papers. Solid-State Circuits Conference, 2005.*, 364–604 Vol. 1 (2005).

[49] C. Niclass, A. Rochas, P. A. Besse & E. Charbon. *Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes*. IEEE Journal of Solid-State Circuits **40**, 1847–1854 (2005).

[50] F. Villa *et al. SPAD Smart Pixel for Time-of-Flight and Time-Correlated Single-Photon Counting Measurements*. IEEE Photonics Journal **4**, 795–804 (2012).

[51] F. Villa *et al. CMOS Imager With 1024 SPADs and TDCs for Single-Photon Timing and 3-D Time-of-Flight*. IEEE Journal of Selected Topics in Quantum Electronics **20**, 364–373 (2014).

[52] of University Edinburgh. *Megaframe - CMOS Sensors and Systems Group*. http://www.css.eng.ed.ac.uk/research/megaframe. Accessed: 2016-09-14.

[53] J. A. Richardson, L. A. Grant & R. K. Henderson. *Low Dark Count Single-Photon Avalanche Diode Structure Compatible With Standard Nanometer Scale CMOS Technology*. IEEE Photonics Technology Letters **21**, 1020–1022 (2009).

[54] D.-U. Li *et al. Real-time fluorescence lifetime imaging system with a 32 × 32 0.13μm CMOS low dark-count single-photon avalanche diode array*. Opt. Express **18**, 10257–10269 (2010).

[55] Photon Force. *Photon Force LinkedIn*. https://www.linkedin.com/company/10847471?trk=prof-exp-company-name. Accessed: 2016-09-14.

[56] Princeton Lightwave. *Geiger mode cameras*. http://www.princetonlightwave.com/products/geiger-mode-cameras/. Accessed: 2016-09-14.

[57] Micro Photon Devices. *SPC2*. http://www.micro-photon-devices.com/Products/SPAD-by-Technology/Standard-CMOS/SPC2. Accessed: 2016-09-14.

[58] J. A. Richardson, E. A. G. Webster, L. A. Grant & R. K. Henderson. *Scaleable Single-Photon Avalanche Diode Structures in Nanometer CMOS Technology*. IEEE Transactions on Electron Devices **58**, 2028–2035 (2011).

[59] G. Intermite *et al. Enhancing the fill-factor of CMOS SPAD arrays using microlens integration*. In *SPIE Optics+ Optoelectronics*, 95040J–95040J (International Society for Optics and Photonics, 2015).

[60] Micro Photon Devices. *Understanding SPC2: SNR and DR compared to analogue imagers.* http://www.micro-photon-devices.com/Docs/Application-note/SPC2_SPC2vsCCD.pdf.

[61] N. Krstajić *et al. Improving TCSPC Data Acquisition from CMOS SPAD Arrays*. In *Advanced Microscopy Techniques III*, 879709 (Optical Society of America, 2013).

[62] M. Sneep & W. Ubachs. *Direct measurement of the Rayleigh scattering cross section in various gases*. Journal of Quantitative Spectroscopy and Radiative Transfer **92**, 293 – 310 (2005).

[63] R. W. Boyd. *Nonlinear optics* (Academic press, 2003).

[64] Thorlabs. *Optical mirrors*. https://www.thorlabs.de/navigation.cfm?guide_id=15. Accessed: 2016-09-14.

[65] Andor. *Intensified CCD cameras - The technology behind ICCDs*. http://www.andor.com/learning-academy/intensified-ccd-cameras-the-technology-behind-iccds. Accessed: 2016-09-14.

[66] A. Couairon & A. Mysyrowicz. *Femtosecond filamentation in transparent media*. Physics Reports **441**, 47 – 189 (2007).

[67] H. Xu, A. Azarm, J. Bernhardt, Y. Kamali & S. Chin. *The mechanism of nitrogen fluorescence inside a femtosecond laser filament in air*. Chemical Physics **360**, 171–175 (2009).

[68] Altechna. *Color glass filters transmission curves*. http://www.altechna.com/download/Rutos/Color%20filters.pdf. Accessed: 2016-09-14.

[69] FluorTools. *DecayFit - Fluorescence Decay Analysis*. http://www.fluortools.com/software/decayfit. Accessed: 2016-09-14.

[70] M. Eisaman *et al. Electromagnetically induced transparency with tunable single-photon pulses*. Nature **438**, 837–841 (2005).

[71] N. Akopian, L. Wang, A. Rastelli, O. Schmidt & V. Zwiller. *Hybrid semiconductor-atomic interface: slowing down single photons from a quantum dot*. Nature Photonics **5**, 230–233 (2011).

[72] R. M. Camacho, M. V. Pack, J. C. Howell, A. Schweinsberg & R. W. Boyd. *Wide-bandwidth, tunable, multiple-pulse-width optical delays using slow light in cesium vapor*. Physical review letters **98**, 153601 (2007).

[73] R. M. Camacho, M. V. Pack & J. C. Howell. *Low-distortion slow light using two absorption resonances*. Physical Review A **73**, 063812 (2006).

[74] L. Wang, P. Ho, C. Liu, G. Zhang & R. Alfano. *Ballistic 2-D imaging through scattering walls using an ultrafast optical Kerr gate*. Science **253**, 769–771 (1991).

[75] K. Yoo, Q. Xing & R. Alfano. *Imaging objects hidden in highly scattering media using femtosecond second-harmonic-generation cross-correlation time gating*. Optics letters **16**, 1019–1021 (1991).

[76] E. M. Hillman *et al. Time resolved optical tomography of the human forearm*. Physics in medicine and biology **46**, 1117 (2001).

[77] S. Popoff, G. Lerosey, M. Fink, A. C. Boccara & S. Gigan. *Image transmission through an opaque material*. Nature Communications **1**, 81 (2010).

[78] J. Bertolotti *et al. Non-invasive imaging through opaque scattering layers*. Nature **491**, 232–234 (2012).

[79] O. Katz, P. Heidmann, M. Fink & S. Gigan. *Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations*. Nature photonics **8**, 784–790 (2014).

[80] S. Kang *et al. Imaging deep within a scattering medium using collective accumulation of single-scattered waves*. Nature Photonics **9**, 253–258 (2015).

[81] K. Si, R. Fiolka & M. Cui. *Fluorescence imaging beyond the ballistic regime by ultrasound-pulse-guided digital phase conjugation*. Nature photonics **6**, 657–661 (2012).

[82] G. Satat *et al. Locating and classifying fluorescent tags behind turbid layers using time-resolved inversion*. Nature communications **6** (2015).

[83] D. A. Boas *et al. Imaging the body with diffuse optical tomography*. IEEE signal processing magazine **18**, 57–75 (2001).

[84] A. Gibson, J. Hebden & S. R. Arridge. *Recent advances in diffuse optical imaging*. Physics in medicine and biology **50**, R1 (2005).

[85] A. Torricelli *et al. Time domain functional NIRS imaging for human brain mapping*. Neuroimage **85**, 28–50 (2014).

[86] E. M. Hillman *et al. Time resolved optical tomography of the human forearm*. Physics in medicine and biology **46**, 1117 (2001).

[87] J. M. Pavia, M. Wolf & E. Charbon. *Single-photon avalanche diode imagers applied to near-infrared imaging*. IEEE Journal of Selected Topics in Quantum Electronics **20**, 291–298 (2014).

[88] S. Nag, M. A. Barnes, T. Payment & G. Holladay. *Ultrawideband through-wall radar for detecting the motion of people in real time*. In *AeroSense 2002*, 48–57 (International Society for Optics and Photonics, 2002).

[89] A. Sume *et al. Radar detection of moving targets behind corners*. IEEE Transactions on Geoscience and Remote Sensing **49**, 2259–2267 (2011).

[90] B. Chakraborty *et al. Multipath exploitationwith adaptivewaveform design for tracking in urban terrain*. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, 3894–3897 (IEEE, 2010).

[91] A. G. A. Kirmani. *Femtosecond transient imaging*. Ph.D. thesis, Massachusetts Institute of Technology (2010).

[92] A. Kirmani, T. Hutchison, J. Davis & R. Raskar. *Looking around the corner using ultrafast transient imaging*. International journal of computer vision **95**, 13–28 (2011).

[93] A. Kirmani, T. Hutchison, J. Davis & R. Raskar. *Looking around the corner using ultrafast transient imaging*. International journal of computer vision **95**, 13–28 (2011).

[94] G. A. Kirmani. *Computational time-resolved imaging*. Ph.D. thesis, Massachusetts Institute of Technology (2015).

[95] A. Velten *et al. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging*. Nature communications **3**, 745 (2012).

[96] O. Gupta, T. Willwacher, A. Velten, A. Veeraraghavan & R. Raskar. *Reconstruction of hidden 3D shapes using diffuse reflections*. Optics express **20**, 19096–19108 (2012).

[97] M. Buttafava, J. Zeman, A. Tosi, K. Eliceiri & A. Velten. *Non-line-of-sight imaging using a time-gated single photon avalanche diode*. Optics express **23**, 20997–21011 (2015).

[98] M. Laurenzis, J. Klein, E. Bacher & N. Metzger. *Multiple-return single-photon counting of light in flight and sensing of non-line-of-sight objects at shortwave infrared wavelengths*. Optics letters **40**, 4815–4818 (2015).

[99] F. Heide, L. Xiao, W. Heidrich & M. B. Hullin. *Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3222–3229 (2014).

[100] J. Klein, C. Peters, J. Martín, M. Laurenzis & M. B. Hullin. *Tracking objects outside the line of sight using 2D intensity images*. Scientific Reports **6**, 32491 (2016).

[101] A. P. Mosk, A. Lagendijk, G. Lerosey & M. Fink. *Controlling waves in space and time for imaging and focusing in complex media*. Nature photonics **6**, 283–292 (2012).

[102] B. D. Anderson & J. B. Moore. *Optimal filtering* (Courier Corporation, 2012).

[103] R. Cucchiara, C. Grana, M. Piccardi & A. Prati. *Detecting moving objects, ghosts, and shadows in video streams*. IEEE transactions on pattern analysis and machine intelligence **25**, 1337–1342 (2003).

[104] R. Cutler & L. Davis. *View-based detection and analysis of periodic motion*. In *International Conference on Pattern Recognition*, vol. 14, 495–500 (IEEE Computer Society Press, 1998).

[105] E. A. Webster, L. A. Grant & R. K. Henderson. *A high-performance single-photon avalanche diode in 130-nm CMOS imaging technology*. IEEE Electron Device Letters **33**, 1589–1591 (2012).

[106] S. Mandai, M. W. Fishburn, Y. Maruyama & E. Charbon. *A wide spectral range single-photon avalanche diode fabricated in an advanced 180 nm CMOS technology*. Optics express **20**, 5849–5857 (2012).

[107] S. Chan, R. Warburton, G. Gariepy, J. Leach & D. Faccio. *Real-time tracking of hidden objects with single-pixel detectors*. Manuscript under preparation (2016).

[108] S. Chan, R. Warburton, G. Gariepy, J. Leach & D. Faccio. *Non-line-of-sight tracking of people at long range*. Under review, Optics Express (2016).

[109] A. Kirmani *et al. First-photon imaging*. Science **343**, 58–61 (2014).