



Université catholique de Louvain
Faculté des Sciences Appliquées
LABORATOIRE DE TÉLÉCOMMUNICATIONS
ET
TÉLÉDÉTECTION

B - 1348 Louvain-la-Neuve

Belgique

Multidimensional Photogrammetry of Short-Lived Events

Damien Douxchamps

*Thèse présentée en vue de l'obtention du grade de
Docteur en Sciences Appliquées*

Examination Committee:

Benoît MACQ (UCL/TELE) - *Supervisor*
Michel VERLEYSSEN (UCL/DICE)
Christophe DE VLEESCHOUWER (UCL/TELE)
Laurence MUTUEL (Thales Avionics, France)
Hervé CAPART (National Taiwan University, Taiwan)
Luc VANDENDORPE (UCL/TELE) - *President*

October 2004

Acknowledgements

Along my years spent at the TELE lab I have come across the path of a lot of people who have helped and guided me towards the final goal of it all, which is now in your hands. In order of appearance, I must first thank my advisor Benoît Macq for believing in my sometimes original ideas (which is not so far from crazy ideas in my case), and giving me the time to realize them. Also from the early days, Hervé Capart has always been a precious source of ideas and his help and enthusiasm always came even in the most difficult situations. Together with Christophe Craeye, he has been advisor for my master degree, which has led to the second chapter of this thesis. I would also like to thank Dimitri Devriendt for the velocimetric analysis and Wei-Jay Ni for the experiences used for the error validation.

Chapter 3 was elaborated as a part of the IWAKE project¹ which has been remarkably conducted by Laurence Mutuel, always ready to help, advise and communicate her enthusiasm. The IWAKE consortium was composed of Thales Avionics (France), Airbus (Germany), the DLR (Germany), LISA Laser (Germany), the University of Hamburg (Germany), the NLR (Netherlands), Faunhofer IOF Jena (Germany) and UCL TERM (Belgium), all of which should be thanked for their indispensable participation to the project. At UCL, as graduate student and colleague, Yannick Verschueren has produced great work on wake vortices detection and signal processing.

At last, the art.live project² and its team provided the funding, hardware and friendship that finally result in the more experimental work presented in Appendix A. The team spun off to create the Alterface company

¹The IWAKE project is funded by the European Commission under contract number G4RD-CT-2002-00778, project number GRD1-2001-40176. The early spectral analysis is based on the data from the M-FLAME project, also funded by the European Commission under contract number BRPR96-182.

²The art.live project is funded by the European Commission under grant number IST-1999-10942.

which has always been supportive. Many thanks to Xavier, Benoît, Xavier (another one), David, Kily, Alok, Laurence, Xavier (yet another one), Luc and Stéphane. Among the team a special thanks go to Toshi, as a colleague and a friend, and also for offering me a good reason to start to *write*. I'm also grateful to Christophe De Vleeschouwer for seeing the world in 2.73 dimensions with me.

Finally, a few names do not really fit in the big picture but have to be mentioned nonetheless: Benoît Spinewine of Civil Engineering, Yasutoshi Onishi for graciously lending several cameras and of course the lab technicians Robert Nicolas and Christian Kinon who spend a considerable amount of time to create the necessary frankenware for my experiments.

Contents

Introduction	1
1 Traffic monitoring with line-scan cameras	3
1.1 Introduction	3
1.2 Principle and feasibility	5
1.2.1 Accuracy, line rate and baseline	5
1.2.2 The effects of acceleration	9
1.2.3 Exposure with natural outdoor lighting	10
1.3 Calibration	11
1.3.1 Geometry of the camera rig	12
1.3.2 Large pixel problem	15
1.3.3 Calibration setup	15
1.3.4 Calibration procedure	18
1.4 Experiments	20
1.4.1 Horizontal and vertical setups	20
1.4.2 Camera synchronisation	21
1.5 Image processing for line-scan sequences	23
1.5.1 Detection of the boundaries of the vehicle	23
1.5.2 Feature detection	25
1.5.3 Motion estimation	26
1.5.4 Linear regression	31
1.5.5 Other measurements	31
1.6 Results and discussion	34
1.7 Conclusion	35
2 Three-dimensional surface reconstruction from sparse data	47
2.1 Introduction	47
2.2 Experimental setup and imaging configuration	50
2.3 Particle imaging algorithms	53

2.3.1	Particles detection and positioning	53
2.3.2	Particle tracking	53
2.3.3	Imaging geometry and camera calibration	55
2.4	Stereoscopic reconstruction	56
2.4.1	Ray intersection and matching	56
2.4.2	A posteriori synchronisation	58
2.4.3	Bayesian surface interpolation	59
2.4.4	Balance of noise and attenuations errors	64
2.5	Velocimetry	66
2.6	Results and discussion	68
2.7	Technique validation	75
2.8	Conclusions	78
3	Dense measurements of wind velocities	81
3.1	Introduction	81
3.2	Wake vortex detection with a Doppler lidar	83
3.2.1	Laser wavelength considerations	83
3.2.2	Pulsed lidar operation	84
3.2.3	Geometric configuration of the ground tests	86
3.3	Classic spectral estimation	86
3.3.1	Spectrum based	88
3.3.2	Correlation based - ARMA	93
3.4	Model-based spectral estimation	97
3.4.1	Single signal primitive	97
3.4.2	Two signal primitives	99
3.4.3	Two signal primitives without noise spectrum	109
3.4.4	Physical interpretation	110
3.4.5	Convergence issues	112
3.5	Spectral estimation of on-board signals	112
3.5.1	Aircraft installation	113
3.5.2	Signal processing adaptation	113
3.5.3	Shakedown tests	120
3.5.4	Flight tests	125
3.5.5	Airborne results and discussion	125
3.6	Conclusions	127
	Conclusion	139

A	3D matching with perspective distortions	141
A.1	Introduction	141
A.2	Local planarity constrain	142
A.3	Implementation	144
A.4	Results and discussion	145
A.5	Conclusions	147
B	Flight Logs	149
	Bibliography	157

Introduction

This thesis presents a collection of photogrammetry techniques aimed at the analysis of short-lived events. The small time scale on which these phenomena happen makes them difficult to capture not only from the point of view of hardware but also because one can not obtain several measurements of events that are not repeatable. Three subjects are presented in ascending order of complexity. A simple unidimensional image sensor for solids opens the discussion followed by a set of video cameras for liquid flows measurements and finally a full three-dimensional data set of the atmosphere is analysed. A study of perspective distortions in three dimensional image matching which did not fit in this scope is presented in Appendix A.

In the first chapter we detail a new technique for road traffic monitoring and law enforcement. The current trend today is to use different sensors for these applications because they seemingly appear to follow very different goals. The approach for road monitoring is to use roughly placed and oriented low-cost cameras (often installed years ago) which provide images that are difficult to interpret. As the results are not accurate enough for law enforcement, this approach imposes to use one sensor for each task: for example one camera and one radar. The underlying idea behind our approach is at the opposite of what is commonly done today: we choose to use a line-scan camera setup that has strong geometrical constraints. This efficiently displaces the complexity of the image analysis into a more complex but one-time off-line calibration and placement. The resulting degenerated view of the vehicles can then be interpreted in real-time with simple image processing algorithms and a great stability and accuracy. It is also able to provide a large number of measurements: speed, acceleration, length, height and more, making this solution ideal for both traffic monitoring and law enforcement applications.

The second chapter presents imaging methods developed to characterise the oscillatory free surface of rapid flows. The techniques are ap-

plied to torrential currents over sediment antidunes with the aim of obtaining high-resolution relief maps of the free surface topography. Two measurement principles are tested, both based on a new approach of imaging of floating tracers dispersed on the rapidly flowing surface. The first technique involves direct stereoscopic measurements while the second technique is indirect and exploits a Bernoulli relation to derive surface elevations from the horizontal velocity field acquired using a single camera. Relief maps obtained for various bedform patterns are presented allowing comparison between the two techniques. The stereoscopic approach is also applied to a known surface to verify the soundness of the error analysis.

At last, chapter three presents the signal processing unit of a Doppler lidar system aimed at the detection of aircraft wake vortices. The wakes are usually detected by looking for the Doppler shift and the broadening of the spectrum of the return signal. Three common techniques have been tested to that effect: a classic spectral estimation, the Auto-Regressive-Moving-Average (ARMA) which works directly on time data and a Gaussian spectral fit. All approaches fail to a certain degree to provide an adequate detection. A new more appropriate model which consists in two primitives is introduced and applied to ground and airborne data. The results are significantly better and allow the detection of the wake in both cases.

Chapter 1

Traffic monitoring with line-scan cameras

1.1 Introduction

As the traffic density increases over the years, our society looks for ways to control and optimise the flow of vehicles on land, air (see Chapter 3) or sea. The ground traffic has traditionally more attention from the population as it is directly concerned by it, and also poses the largest problems due to the uncoordinated individual behaviours of the drivers. To that effect, public institutions are looking for tools to increase the throughput of roads not only by using appropriate traffic management but also by making the roads safer, all of which requires a good level of law enforcement. A large number of tools has therefore been developed in that perspective, from the infamous speed radar to surveillance cameras, traffic displays showing alternate routes and warnings, radio bulletins, Internet,... Most of these systems designed several years ago are still poorly automated, which given the critical situation of the road traffic is starting to be an insufficient approach. Public authorities are looking more and more for multipurpose sensors due to the large road network to instrument and the resulting investments that will be necessary to upgrade obsolete traffic monitoring systems. In that perspective, this chapter presents a new sensor based on line-scan cameras that is able to measure several parameters of the traffic, among others the vehicle speed, acceleration, length and height. Moreover, the target accuracy of the proposed approach is less than 1%, which is better than the standard 3% of current speed radars.

The systems currently in use roughly split in two categories: speed

measurement for law enforcement and video cameras for traffic information. For speed measurements the most popular and commonly used device is the radio-frequency (RF) radar operating in the tens of GHz range. These devices are based on the Doppler effect and compare the frequency f_o of the emitted signal to the frequency f_r of the echo which has been modified according to the vehicle speed component v that is aligned with the line-of-sight of the radar:

$$f_r = f_o \frac{c}{c - v} \quad (1.1)$$

where c is the speed of light. While the equation above is rather simple, the practical realisation of a radar is difficult because v is much smaller than c and the evaluation of the speed therefore requires an expensive heterodyning system. The operating conditions of the radar are also limited because rain, fog and other hydrometeores lead to inaccuracies as the typical radar wavelength (millimetric waves) is close to the size of water drops. This further conducts to parasitic reflections, a higher attenuation and thus a reduced range and accuracy. Some other factors can also have an important influence on the estimated speed like the vehicle acceleration and the proper calibration of the radar. A more portable version of the radar is the lidar, which also uses the Doppler effect but with a laser instead of an RF source. The lidar used for measuring the speed of vehicles have a limited range and are used to aim at hard targets which results in a smaller package that can be operated by a single person like a hand gun. However like the RF version they only serve one purpose.

Midway between speed measurement and traffic information systems, magnetic loops are very popular for counting vehicle and measuring their speed. They can be set up easily for temporarily instrumenting a problematic crossroad but permanent installations require to close the road temporarily as the loop must be integrated in the road, which quickly become unpractical when it comes to instrumenting a highway.

Information systems mainly use a classic video camera placed over the highway to detect traffic jams, dangerous behaviour and accidents. Although the accuracy of such device for measuring speed is very limited, it is possible to use it to detect simple events like traffic jams or accidents [1] [2]. Most traffic analysis systems are actually trying to replace the operator-camera combination in order make the whole chain fully automatic. This leads, among others, to the ability of building road statistics and vehicle classification at a great computational cost and for a limited accuracy that is insufficient for law enforcement applications [3] [4].

An interesting approach which is close to our line-scan idea is the one using laser profilers. It consists in two vertical lasers placed one or two meters apart along the highway, typically under a bridge. Each laser is measuring the distance to the closest object with a time-of-flight approach, yielding two profiles of the vehicle that can be matched to find the time used by the vehicle to go from one laser to the other, and thus to measure its speed. The major problem of this approach is that the laser only scans the road at one point on each lane. A vehicle can thus pass between the lasers and avoid detection. The cost of fast and accurate laser profilers would also hinder their wide deployment.

Following the idea of the laser profilers, we imagined a new multi-purpose system comprising two line-scan cameras to scan the road (fig. 1.1a). As vehicles go by they are scanned by the cameras and we can calculate the time $T_2 - T_1$ needed by a vehicle to go from the first to the second camera. Because the line-rate of these cameras is very high (typically over 4000 lines/sec) we can expect an accuracy that is better than the laser profilers, even if we are using a vision system. A novel calibration of the cameras has also been designed to accommodate our new sensor. We will start by laying out the principles behind our new line-scan approach and proving its feasibility. The expected accuracy claimed here will require a specially-designed calibration procedure which is then detailed. After a brief description of the experiments, a section reviews the image processing tools used to extract all information from the line sequences. The results of this process are then presented and discussed.

1.2 Principle and feasibility

1.2.1 Accuracy, line rate and baseline

The setup presented on fig. 1.1a consists in two line-scan cameras separated by a distance b , called the baseline. The field of view of each camera is a vertical plane, orthogonal to both the road plane and the direction of the traffic. The beginning and end of the scan of a vehicle by each camera correspond to four key instants $T1...T4$, represented on fig. 1.1b. The resulting line-scan images are shown on fig. 1.1c. Note that the vehicles will appear 'stretched' or 'compressed' depending on their speed. Given this simple geometry, the camera baseline b , the speed s of the vehicle and the time $\Delta T = T_2 - T_1$ needed to go from one camera to the other are related

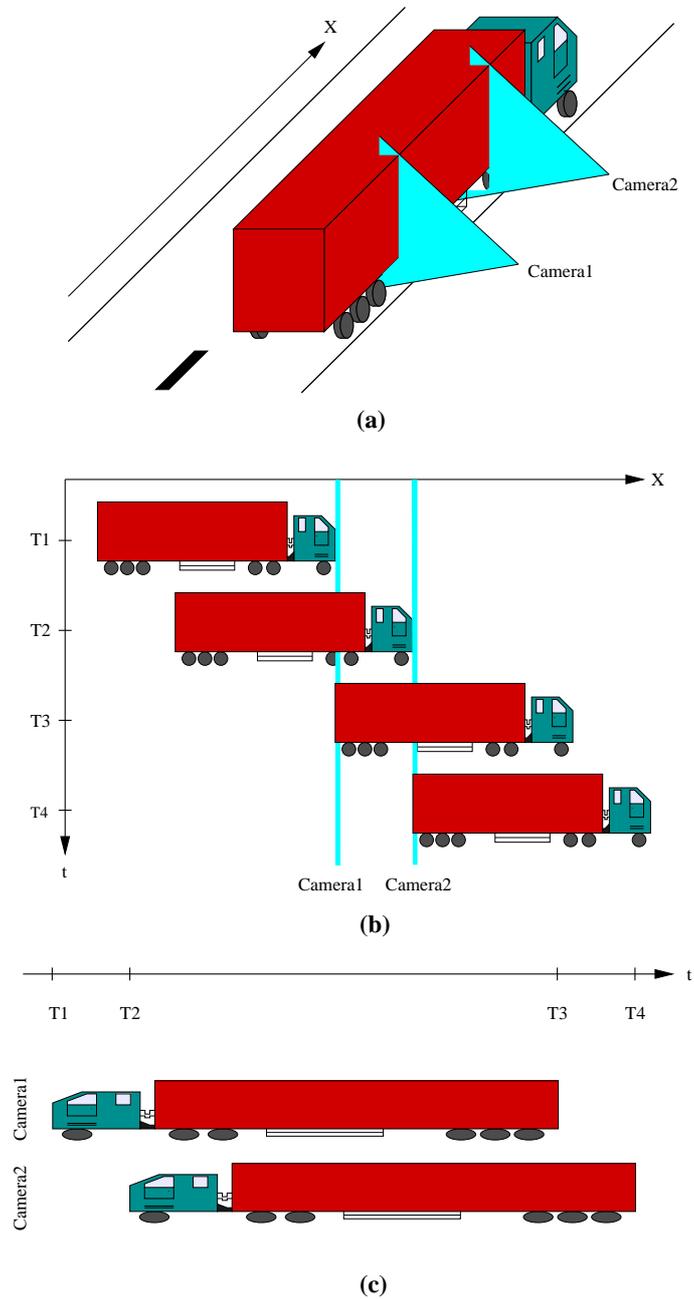


Figure 1.1: Principles of the line-scan speed measurement system: (a) three dimensional view of the setup, with the scanning planes in blue; (b) side view of the setup at entering/leaving times $T_1 \dots T_4$; (c) the two image scans of the vehicle

with the following expression:

$$\frac{b}{\Delta T} = s \quad (1.2)$$

Because the camera samples the scene at discrete times, ΔT is a multiple of the sampling period of the camera, represented here by the line rate f_l :

$$\frac{b}{n/f_l} = s \quad (1.3)$$

where n is the number of lines of delay between the appearance of the vehicle on each line-scan sequence. Using this simple expression we can already verify that the parameters involved are within the capabilities of current hardware. Let us suppose as a first approximation that the error on the detection of points in the image is one pixel. Since two indexes are used in the calculation of ΔT , the error on its evaluation is $2/f_l$, and at least $2n$ lines will be needed for a relative error of $1/n$ on the speed. The relative error ϵ_s on the vehicle speed can be estimated by a function of the camera distance, the sampling rate and the vehicle speed:

$$\epsilon_s = \frac{\Delta s}{s} = \frac{\frac{bf}{n} - \frac{bf}{n+2}}{\frac{bf}{n}} \quad (1.4)$$

which can be reduced to

$$\epsilon_s = \frac{2s}{bf + 2s} \quad (1.5)$$

Fig. 1.2 shows the required baseline and line rate for four accuracy targets. The required line rate is 6600Hz for a typical baseline of one meter and an accuracy of 1% at 120km/h, which is within the capability of line-scan cameras. An interesting fact to consider is that, all other things being equal, the error will be lower at a lower speeds because the number of lines composing a vehicle will be larger. For example, if one sets the error to be 1% at 120km/h it will become 0.42% at 50km/h and 1.5% at 180km/h. The line-scan approach has *relative* error that is lower at lower speeds. It is a serious advantage for traffic analysis because it allows to enforce lower speed limits while keeping a reasonable relative error at high speed. This contrasts with Doppler systems which have a constant *absolute* error, making them poor candidates for low speed measurement.

Note also that the error estimation described above is only valid for speed measurements which require the location of two image points. A measurement with four image points like the length or the acceleration will have an accuracy lowered by a factor of two.

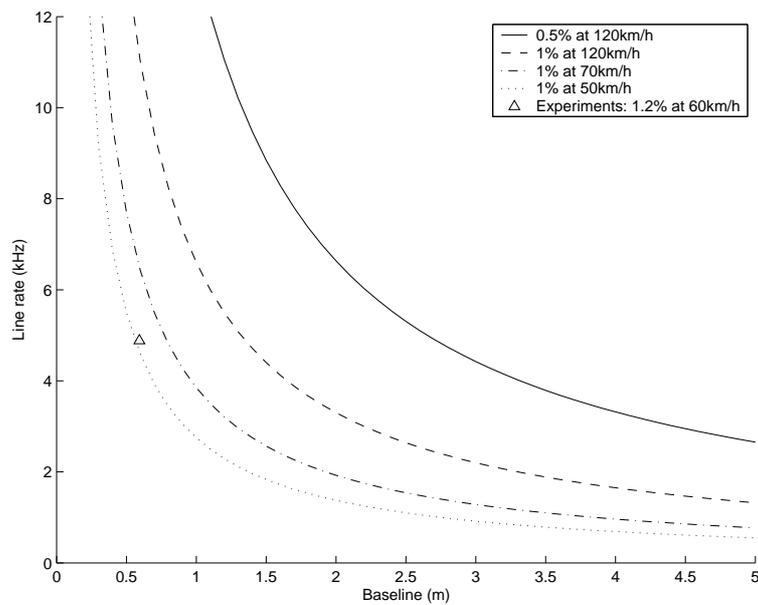


Figure 1.2: Speed error as a function of baseline and line rate: accuracy targets of 0.5% and 1% for different speeds. The operational point of the tests is shown as a triangle.

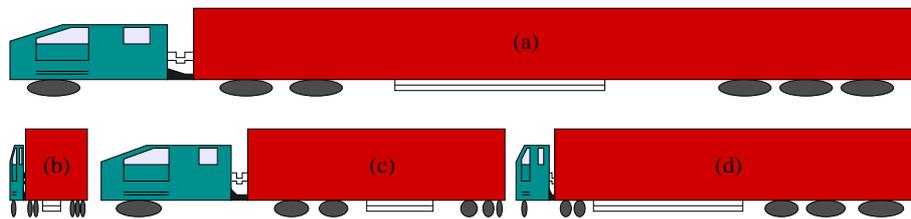


Figure 1.3: Deformations in the observed sequences due to speed and acceleration: (a) slow, (b) fast, (c) accelerating and (d) decelerating vehicles

1.2.2 The effects of acceleration

The speed of the vehicle might not be constant during its pass in front of the cameras, resulting in deformed vehicles in the line-scan images (fig. 1.3). If we exclude the case of traffic jams where the speed will vary greatly, we can limit ourselves to a first order approximation of the speed which is then described only by its initial value v_0 and the acceleration a of the vehicle. The position x of the vehicle between the two cameras can then be described as

$$x = v_0 t + \frac{at^2}{2} \quad (1.6)$$

Isolating the time t in this equation we obtain

$$t = \frac{-v_0 + \sqrt{v_0^2 + 2ax}}{a} \quad (1.7)$$

The error $\epsilon_{s,a}$ on the vehicle top speed due to its acceleration is then equal to half the extreme speed difference observed between the two cameras. The speed difference being equal to at we obtain the following expression for $\epsilon_{s,a}$:

$$\epsilon_{s,a} = \frac{-v_0 + \sqrt{v_0^2 + 2ax}}{2} \quad (1.8)$$

Let us take a few examples of accelerations and see how this affects the estimated speed. A normal car needs about 60m to stop from a speed of 120km/h, resulting in a deceleration of -6.4m/s^2 . A Porsche 911 with carbon disk brakes has been selected as the best braking vehicle, achieving a full stop from 200km/h in 5.2s, or a deceleration of -10.7m/s^2 . An accelerating vehicle will not necessarily produce less errors. The best accelerating vehicle is the Gillet Vertigo, which can go from 0 to 100km/h in 3.27s. This corresponds to a constant acceleration of 8.5m/s^2 . Fig. 1.4 shows the speed error that would result from accelerations of a normal car and of the two sport cars mentioned above. The error on the maximum speed of the vehicle will always be negative since the baseline acts like a low-pass filter on all the speeds taken by the vehicle between the two cameras. This means that we will always underestimate the maximum speed, thereby providing conservative measures. Another interesting note is that the error on the speed due to acceleration is larger for lower speeds, which is against the claim that we laid before that the line-scan approach works well under those circumstances. We will however show later that this claim still holds because we can estimate the acceleration of the vehicle from its speed profile.

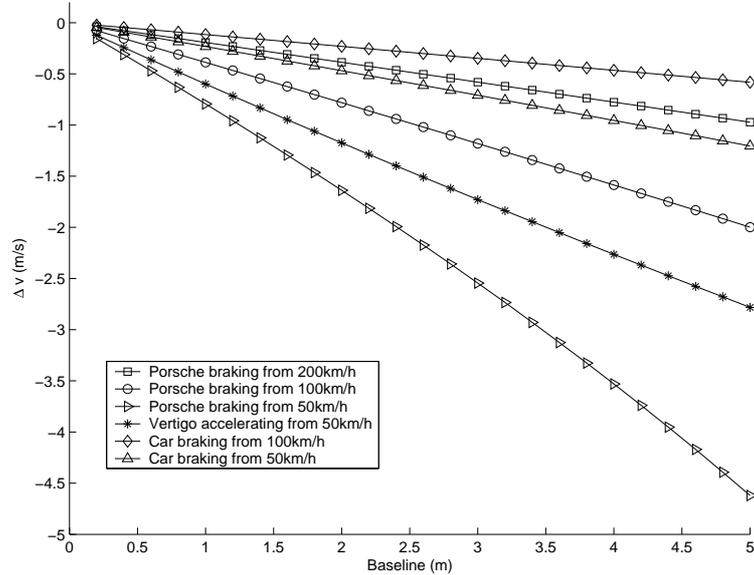


Figure 1.4: Errors due to acceleration as a function of the baseline

1.2.3 Exposure with natural outdoor lighting

A major consideration when operating outdoor is the capability of the camera to properly expose a scanned sequence given the high line-rate required. The problem becomes evident when one notices that a classic video camera needs about $1/100$ s to expose a scene while our line rate gives us a maximum exposure time that is one or two orders of magnitude below. This limitation is the reason why line scan cameras are often limited to industrial indoor use on a small range with a strong artificial lighting. We have here the opposite situation: we work outdoor, with ambient lighting only (for stealth and safety reasons) and on range of up to 8 meters (the width of two lanes, see 1.5.5). Since the light can not be controlled we are left with two camera parameters to adjust: the lens aperture and the sensor sensitivity. The lens aperture is limited just like area-scan cameras: few lenses go above $f/1.2$. This leaves the higher sensitivity as the only solution to work outdoor. The sensitivity being proportional to the number of photons captured by a unit of time, it is also proportional to the sensitive area of one pixel. For a standard video camera the pixel surface is typically $6\mu\text{m} \times 6\mu\text{m} = 36\mu\text{m}^2$, but a line-scan sensor can have much larger pixels because their size is only limited in the dimension in

which the line extends. The size of the pixel in the other dimension can be much larger, for example $6\mu\text{m} \times 125\mu\text{m} = 750\mu\text{m}^2$ for the sensor used in our tests, which is roughly 20 times more than a classic camera pixel. This will let us use a shutter speed 20 times faster, and thus to operate at line-rates over 1kHz.

As we will see later, these large pixels are not without problems: 4 meters away from the camera the field of view of the scanning plane has a width of 5cm which reduces the optical resolution of the system.

1.3 Calibration

The calibration of any camera system is one of the most important step if one wants to obtain high accuracy metric measurements. Unlike the calibration of area-scan cameras, we can not locate points in the image and fit those points into a camera model to derive its parameters: the line-scan cameras contain only one line which is not sufficient to describe the two-dimensional deformations that occur on the image plane. Also, in the case of the area-scan camera the deformations modelled by the calibration will be used to perform operations on neighbour pixels (like image rectification) which is something not possible on a uni-dimensional image. For these reasons we must develop other techniques to calibrate the camera rig.

The main assumption of the linear-camera system is that we have two parallel image planes that are perpendicular to the speed vector of the vehicles (fig. 1.1), ensuring an effective baseline that does not change if the object is closer or farther from the cameras. This is more important than the absolute locations of the cameras since the speed measurement rely on an accurate knowledge of the baseline to convert an image offset to a speed as shown by (1.3). For this reason the calibration will be aimed at properly positioning and orienting the cameras relative to each other. Our calibration procedure will thus aim at placing the cameras accurately and then perform a simple calibration of other parameters of minor influence. On the contrary, working with area-scan cameras would rather lead us to place the cameras approximately and then estimate their calibration parameters accurately.

1.3.1 Geometry of the camera rig

The different geometric parameters of the line-scan camera rig are presented on fig. 1.5. The description of these parameters will be split between positions and angles among which the latter are the most important. A distinction between relative and mean angles or distances is made and the influence of the error on each parameter on the final results is estimated.

Camera position

It is necessary to know the baseline distance b between the two scanning planes in order to determine the vehicle speed with (1.3). The accuracy of this measure will be directly reflected on the speed estimation and this distance must be calculated in the best possible way. The mean position of the camera rig is not as important:

- A displacement Δx of the rig long the road axis will have no effect on our measures.
- Placing the rig in a higher or lower position (Δz) does not affect the measures either, but a lower vantage point is recommended to avoid the shadow issue which overestimates the length of the vehicles (see Section 1.5.1).
- At last, the distance Δy from the camera to the vehicle is a simple matter of field of view: the cameras should be far enough to see the whole vehicle, and not too far to keep a good resolution. Note that if the vehicle is too close from the camera the system will still operate but some measurements like the vehicle height will not be possible. There will also be a risk that too few features will be seen.

Camera orientation

The camera orientation angles are also defined on fig. 1.5. Let us describe each one of them one by one.

- **Mean yaw** $(\gamma_1 + \gamma_2)/2$: The angle of the rig with the road axis must be kept close to zero. If the angle is significantly different (which is the same as having the speed vector of a vehicle not aligned with the

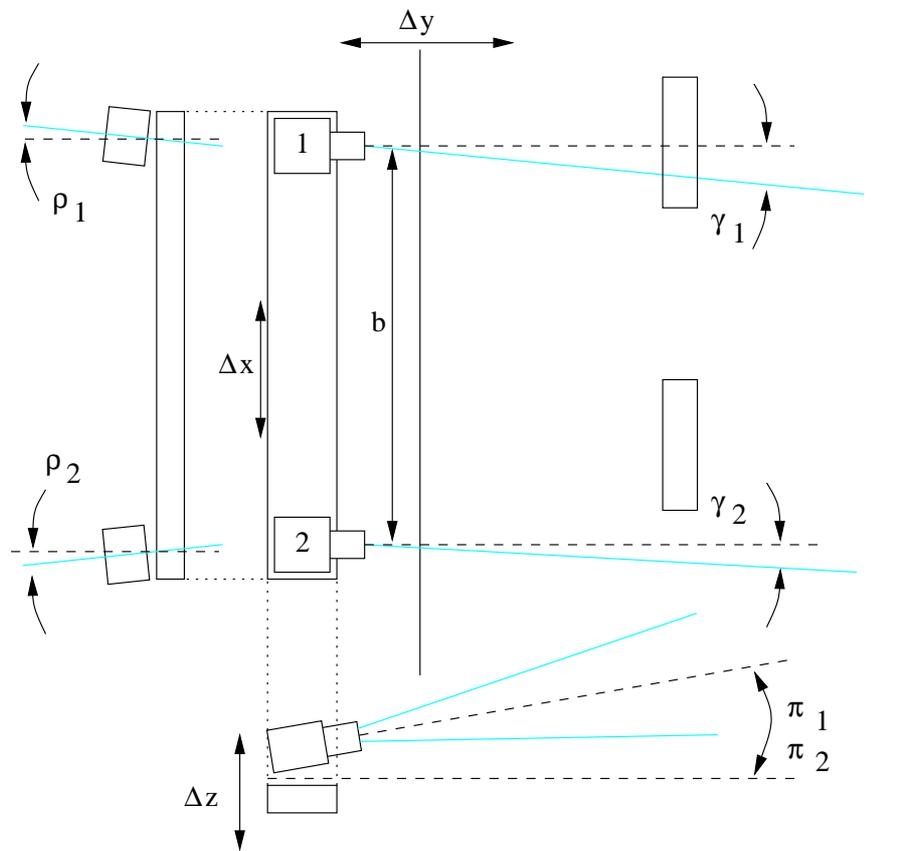


Figure 1.5: The different angles involved in the calibration process: pitch π_i (side view), roll ρ_i (rear view) and yaw γ_i (top view)

road axis), the error on the baseline $\epsilon_{b,1}$ will be

$$\epsilon_{b,1} = b - \frac{b}{\cos(\gamma)} \quad (1.9)$$

and this will inevitably translates into errors in (1.3). It will also lead to inaccuracies regarding the vehicle length since the vehicle width or height will influence the results.

- **Relative yaw** ($\gamma_1 - \gamma_2$): This is the parameter which has the most serious influence on the measurements: if the camera planes are not parallel the baseline will vary with the vehicle distance to the camera rig:

$$\epsilon_{b,2} = \sin(\gamma_2 - \gamma_1) d \quad (1.10)$$

For instance, a baseline of one meter with a relative yaw of one degree will become only .95m at a distance of 3 meters, leading to an error of 5%. As the distance d is unknown we can not compensate that effect in post-processing. A special care will therefore be taken during the calibration to avoid any relative yaw.

In general, the sensor will not be aligned with lens focal point and the actual camera plane will not be perpendicular to the lens axis. This parameter is another horizontal angle that adds with the camera yaw mentioned above. Both will be calibrated together.

- **Mean pitch** $((\pi_1 + \pi_2)/2)$: The vertical axis of the camera can be placed to get the desired point of view. The pitch will depend on the application and can take the typical values of zero degrees (side-viewing) or 90 degrees (top viewing). It will not introduce errors in the measures.
- **Relative pitch** $(\pi_1 - \pi_2)$: If both cameras do not have the same angle a vertical offset will be seen on the scans. This offset can be considered fixed and will be compensated in software (see section 1.6).
- **Mean roll** $((\rho_1 + \rho_2)/2)$: The camera rig must be parallel to the ground. If not, the vehicle will be sampled by slanted planes. It will have two effects: the speed will not be constant on a image column and the length of the vehicle will suffer overestimation.
- **Relative roll** $(\rho_1 - \rho_2)$: Finally, the camera relative roll will also introduce inconsistencies in speed estimation because the baseline will

not be constant with both the distance d or the vertical position in the scanned line. It is thus also an important parameter to calibrate. The maximum error on the baseline introduced by this parameter is:

$$\epsilon_{b,3} = \sin(\rho_1 - \rho_2)H \quad (1.11)$$

where H is the height of the largest vehicle (around 4m).

1.3.2 Large pixel problem

The large size of the pixels which is necessary to operate in natural light conditions is the source of a important limitation of the system. Using the simple geometry of the camera it is easy to determine which horizontal size H_s the pixels will represent at a given distance d from a camera equipped with a lens of focal length f :

$$H_s = \frac{d}{f}125\mu m \quad (1.12)$$

For a distance d from the camera of roughly 4 meters and a focal length f of 8mm, the actual pixel width is 5cm. This optical filtering of the scene reduces the true time-resolution of the line-scan image of a vehicle. We can define a boundary for the sampling frequency f_l in such way that the motion of the vehicle at top speed s_{max} in front of a camera will not be more than half the width of the scanning plane:

$$f_{l,max} = \frac{2s_{max}f}{d125\mu m} \quad (1.13)$$

For the experiments performed this translates into a maximum useful sampling frequency f_l of 4080Hz. This problem could be solved by using cylindrical lenses with a larger aperture which contrary to spherical lenses will not enlarge the pixel width. However, these optics are expensive and do not exist in the wide angle versions which is what we need to capture the full height of a truck two meters away. Nevertheless, we will show that it is possible to obtain good results with off-the-shelf lenses if we keep this limitation in mind.

1.3.3 Calibration setup

The specific setup built to calibrate the line-scan camera pair is shown on fig. 1.6. It consists in a flat base $0.8m \times 2.4m$ with four vertical and parallel panels that create two channels, one for each camera. This large size

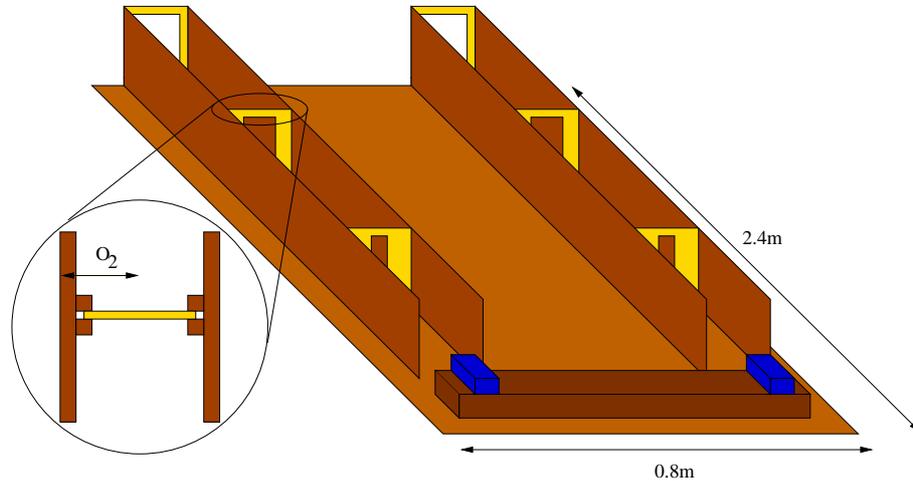


Figure 1.6: Calibration setup schematic showing the two channels and the sliding slots

is necessary for a calibration setup to be able to pinpoint small angular changes. The cameras are placed on a strong aluminium profile at one end of the channels. The profile is equipped with adjustable platforms (fig. 1.8) which can be moved to adjust the yaw and roll of the camera. Each adjustment is performed by screws to obtain a precise and stable setting. Three vertical slots are placed along each channel at distances of 400mm, 1300mm and 2200mm from the camera. The slots have been drilled using an accurate 2-axis Printed Circuit Board (PCB) drilling machine to match the width of the camera scanning volume at each distance. Since the lens is 8.5mm and the pixel width is $125\mu\text{m}$, (1.12) gives us the widths of 5.88mm, 19.12mm and 32.35mm respectively at distances of 400mm, 1300mm and 2200mm. The calibration being performed indoors, a powerful lighting must be used to achieve a proper exposure. The light source should also be diffused to generate a homogeneous luminance over the whole slot width. To this effect two 300W halogen bulbs were placed behind a double layer of translucent plastic panel. The variations in signal amplitude when the slots are moved can be very small (less than .1 intensity value in 8bit, or 0.04%). The room where the calibration setup is performed had all its windows opacified to avoid parasitic variations at that level of lighting sensitivity. The power supply of the lighting is chosen to be DC for the same reason. Fluorescent tubes have presented a much too large variation in intensity, as expected, but AC halogen were also too flickering.

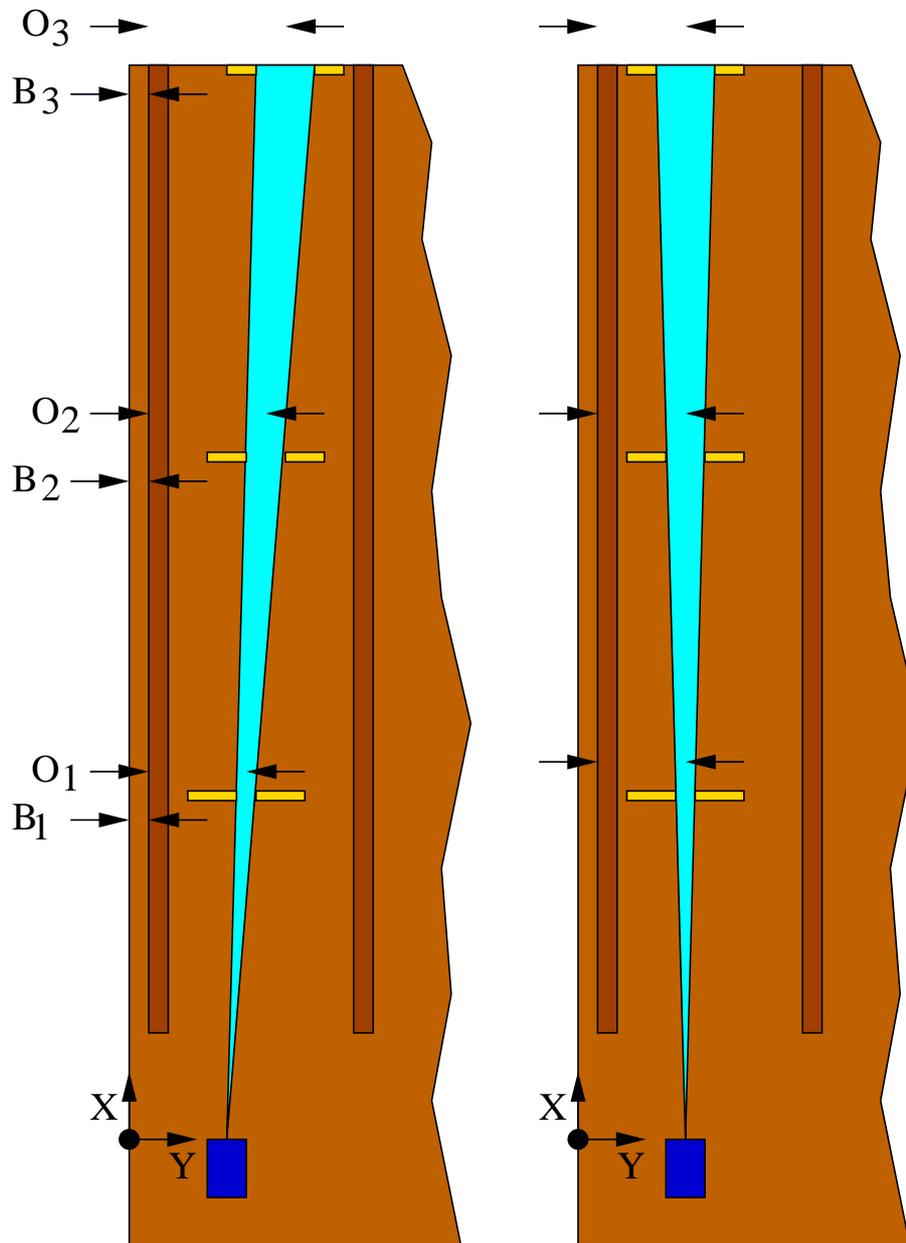


Figure 1.7: The different measurements performed on the calibration rig: offsets to the plank side B_i and slot offsets O_i



Figure 1.8: The camera pan/roll bases. The linear shape of the CCD is visible in the camera.

1.3.4 Calibration procedure

At first only the farthest slot is placed in the channel. The camera is focused on the last slot to have a maximal contrast in the measurements for this slot which, being wider, has a lower response to offset changes. In order to stabilise the electronics, lamp and camera are left alone for twenty minutes before measurements are made. The camera platform is then adjusted to obtain the maximal luminance for the useful portion of the line corresponding to the slot (only 100 pixels on the total 1024 for the channel height of 300mm). This is done by first setting the yaw (also called heading) and then the roll of the platform. One or two iterations on these two parameters are enough to obtain stable results as the angles involved are very small.

Once the camera is aligned the last slot is removed and each of the three slots is placed one by one in the channel. The slots are displaced across the channel width and the global luminance of the line is measured for a discrete set of offsets. The offset corresponding to the maximum of luminance gives the position of the viewing volume. The depth gauge of a calliper is used to measure the offsets with an accuracy of 0.01mm by locking it to the desired value and plunging it into a lateral hole to push

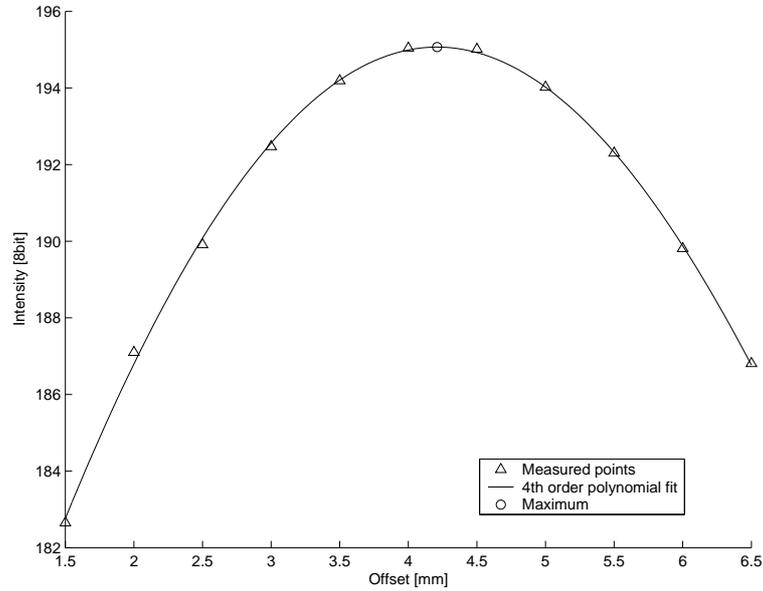


Figure 1.9: Optimisation of the slot offset with a 4th order polynomial

the slot plate at the desired location.

A large number of sampled lines (typically 1000) are averaged in order to stabilise the luminance measurement and obtain an accuracy of 0.01/255. The luminance obtained for each offset is shown on fig. 1.9. A fourth order polynomial is fit to the points obtained to find the true optimum and the procedure is repeated for each slot. The three offsets are then fit with a line corresponding to the line of sight of the camera. The independent term of the interpolated line is the necessary offset at which all three slots must be set to obtain a viewing volume that is parallel to the channel axis (fig. 1.6), while target for angular coefficient is zero. The camera must then be adjusted to aim through this new slot setup. Once the aim is good, another calibration step is performed to ensure that the camera is aligned properly. This is necessary because the positioning of the camera is not as accurate as the slot alignment. The final results for the two cameras are shown on table 1.2.

The same process is applied to the second camera. Both cameras are then aligned with the edge of the base plank and both camera scanning volumes are therefore parallel (fig. 1.7). Once the cameras are calibrated,

Slot	Offset from outer wall B_i	Offset from border O_i
1, right	81.96	18.90
2, right	82.00	18.59
3, right	81.98	18.65
1, left	81.36	18.23
2, left	81.66	18.29
3, left	81.76	18.48

Table 1.1: *Offset values*

	slot 1	slot 2	slot 3	interpolated line
Left camera	104.59	104.54	104.46	$-7.222 \times 10^{-5}x + 104.624$
Right camera	103.05	103.04	103.14	$5.000 \times 10^{-5}x + 103.012$

Table 1.2: *Final offsets $O_i + B_i$ and the resulting interpolated line of sight*

the distance between the cameras can be estimated as

$$\hat{b} = P_w - X_1 - X_2 \quad (1.14)$$

where P_w is the width of the whole calibration setup and X_i are the calibrated offsets. This translates in our case into $800.65 - 104.62 - 103.01 = 593.02\text{mm}$ which given the accuracy of the calibration can be rounded to 593mm .

1.4 Experiments

1.4.1 Horizontal and vertical setups

The line-scan monitoring rig can be used in two different basic configurations: with the camera on the roadside or with the camera under a bridge, looking downward (see fig. 1.10). The lateral view is interesting if classification of the vehicles is required. It is also easier to setup and allows the estimation of the vehicle height which is important at the entrance of tunnels, bridges,... The lateral view of a vehicle also boosts more features than its top view, providing more matching opportunities because logos and window frames are mostly located on the vehicle side. A typical sequence and the detail of its first vehicle are presented on fig. 1.11 and 1.12.

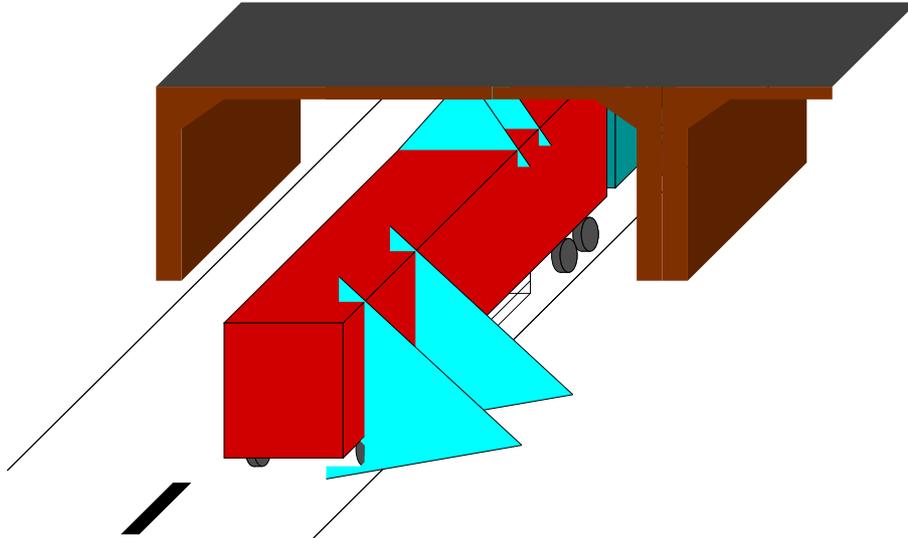


Figure 1.10: *Lateral and overhead setups*

Note the very large image size: the image stretches on 15000 pixels horizontally, for a little less than three seconds of recording at a line rate of 4882.8125Hz. Among the problems of this approach we can note that it is more prone to theft, weather and vandalism. The background is also in general non-uniform and it can be difficult to handle the case of passing vehicles in the case where several lanes are present. This can also be considered an advantage as only one rig is necessary to monitor two or three lanes simultaneously.

The overhead view provides a better platform with its uniform background (i.e. the road surface). It is also hidden and out of human reach, not to mention the better protection against water, dirt,... The drawback is that one will not be able to estimate the height of the vehicle. This approach has not been tested for it requires a much more complex installation.

1.4.2 Camera synchronisation

In order to estimate the delay nT_l used to estimate the speed in (1.3) the cameras must be synchronised so that 1) the cameras start exposing a line at the same time and 2) both cameras start the acquisition at the same time. The cameras we use are based on the IEEE1394 bus, which does not allow



Figure 1.11: An excerpt of a typical sequence: camera 1 (top) and camera 2 (bottom)



Figure 1.12: Detail of the first vehicle seen on fig. 1.11

synchronisation so that an external trigger must be used to synchronise the cameras. Although an external circuit was designed, the cameras do not respond well to this trigger and their erratic behaviour prevented us from using an external signal for triggering. As a result the line will not be acquired at the exact same time. This difference will be typically T_l but can not be known accurately without modifying the camera firmware and will therefore results in an error of T_l on the speed estimations.

Even if the synchronisation is not possible with an uncertainty under T_l , it is mandatory to obtain at least an accuracy of T_l for (1.5) to remain valid. To that effect an external strobe was used to synchronise the sequences: the flash hits both sensors at the same time and its position on the sequences will show the actual delay between the cameras. For reasons proper to the IEEE1394 bus and also to the implementation of the IEEE1394 system under Linux, the delay between the flash imprint on the sequences is -17 lines or 3.48ms . Fig. 1.13 shows the flash as it appears on both line sequences.

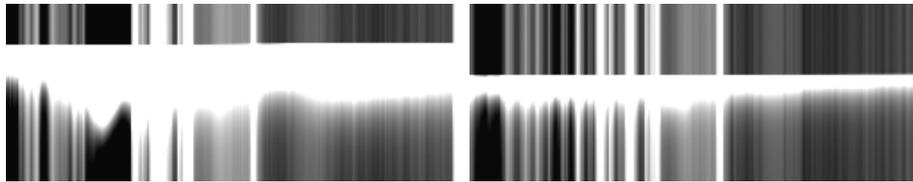


Figure 1.13: *The flash as seen by the line-scan cameras: camera 1 (left) and camera 2 (right). The flash in the first camera is located 17 lines before the one in camera 2.*

1.5 Image processing for line-scan sequences

The image processing required to analyse the sequences is described below. Due to the high data rate which can reach 80 Mbit/s for 10kHz sampling rate, the real-time goal will only be reached on current hardware with simple algorithms. The complex calibration performed above will pay-off in this situation: the degenerated view of the scene is easily interpreted by simple real-time algorithms and the scope of the metric informations inferred from the image measurements is broad: speed, acceleration, length, height, inter-vehicle distances,... In order to obtain these characteristics for the vehicles we only need to obtain a limited set of image measurements: all information will come from the position of boundaries of the vehicles in the scanned sequence and from a more accurate displacement measure.

1.5.1 Detection of the boundaries of the vehicle

The detection of the vehicles in the sequence of lines is performed with a background comparison. A reference background line is obtained regularly as the average of a bundle of typically a hundred lines. As can be seen on figures 1.11 or 1.12 the background is very stable across the whole vehicle length. Indeed, on the contrary to real images where such large averaging of data is not possible, line-scan cameras operate at such a high frequency that the 'large' average on 100 lines has a temporal length of only a few hundredth of a second. The detection will therefore not be subject to problem related to illumination changes or other slow background changes because such events happen on a much larger time scale. Given the simple operating conditions, the edge of the vehicle can be found by comparing the current line with the background line: a vehicle is detected

as long as the difference between the two lines is over a certain threshold. A temporal filter of ten lines is used to avoid detecting spurious objects or to have vehicles split in several parts.

The only significant inconvenient with this approach comes from the shadows that are included in the bounding box of the vehicle (fig. 1.14). In extreme cases, this could result in several vehicles being merged due to their overlapping shadows at dusk or dawn. Detecting the vehicles only on a fraction of the height of the image that is not polluted by shadows improves the results at the expense of stability. It is also possible to use colour processing to detect shadows but colour cameras are less sensible and line rate will have to be lowered. The best solution to this problem of shadows is to place the camera close to the ground to avoid them in the largest fraction of the scanned image but this will also limit the accuracy of lateral distance and height measures. Shadows have been ignored in what follows, leading to an overestimation of the vehicle length.

From the vehicle boundaries in the two images we can obtain a first estimation of the vehicle speed S_v , length L_v and acceleration A_v using (1.3):

$$S_v = \frac{b}{(B_f^2 - B_f^1)/f_l} \quad (1.15)$$

$$L_v = S_v \frac{B_r^1 - B_f^1}{f_l} \quad (1.16)$$

$$A_v = \frac{\frac{b}{(B_r^2 - B_r^1)/f_l} - \frac{b}{(B_f^2 - B_f^1)/f_l}}{(B_r^1 - b_f^1)/f_l} \quad (1.17)$$

where B_f^i and B_r^i are respectively the front and rear line indexes of a vehicle in camera i . Two other estimates can be obtained using $(B_r^2 - B_r^1)$ in (1.15) and $(B_r^2 - B_f^2)$ in (1.16). An example of detected boundaries and the resulting first estimates are given on fig. 1.14 and table 1.3. The corresponding metric speed S_v is 63.3km/h, which is below the 70km/h displayed in the vehicle but still plausible given the tolerances of such meter. The estimated length is 4.66m while the true value is 4.39m, showing the importance of the shadow overestimation. A manual inspection of the more recognisable wheelbase yields 2.64m, much closer to the true value of 2.61m. These early measurements are used to adapt the next processing steps to the highly variable image size that can be observed for a vehicle depending on its speed. Since two estimates are obtainable for the length



Figure 1.14: The detected boundaries of the vehicle. Overestimation of the length due to shadow is clearly visible in the front of the vehicle.

	Estimation 1	Estimation 2
Front line number	38416	38597
Rear line number	39714	39893
Length	4.7m	4.7m
Speed	63.6km/h	64.3km/h
Acceleration	0.73m/s ²	

Table 1.3: First measurements using only the vehicle boundaries

and the speed we can roughly estimate their coherency and act accordingly, for example by not processing the vehicle if the first estimates are too far apart.

1.5.2 Feature detection

The feature detection is used to isolate portions of the vehicle that are useful to determine its speed in order to accelerate the block-matching processing (see section 1.5.3). As the speed is unidirectional, the matching process used to detect the offsets between the images will be performed only in the horizontal direction. For this reason, the horizontal variance will be the sole responsible of a good match, and is the only one to be taken into account. The feature detection is based on a classic Moravec operator [5] [6], which has been modified to detect only horizontal features. For every image point $I(i, j)$ with i and j respectively the horizontal and vertical



Figure 1.15: The feature map for the vehicle of fig.1.14: dark pixels correspond to featured areas

indexes, the feature content $F(i, j)$ will be

$$F(i, j) = \frac{\sum_{p=-k/2}^{k/2} \|I(i, j) - I(i + p, j)\|}{k + 1} \quad (1.18)$$

where k is a typical feature size. Note that (1.18) is nothing more than a gradient averaged on a certain width k . Since the feature size depends on the speed of the vehicles we use the initial speed estimation to adapt k according to a fixed feature size of 20cm:

$$k = \frac{0.2}{S_v T_l} \quad (1.19)$$

A typical feature map is presented on fig. 1.15, where dark pixels represent a high horizontal variance. Since k can be quite large (often around 100), (1.18) is only calculated for a limited set of offsets p to speed-up the computation of the feature map.

1.5.3 Motion estimation

Although a first speed estimation has been obtained in 1.5.1, it is not accurate and relies on unstable boundaries that are subject to shadows. In order to obtain higher accuracy measurements a Block Matching Algorithm (BMA) is used to compute the optical flow between the two scans of the vehicle. In general, the BMA looks for the image offset $(a(i, j), b(i, j))$ that minimises the difference of an image block of size $(2m + 1) \times (2n + 1)$ around a point (i, j) in image one with a block of the same size in image two:

$$(a(i, j), b(i, j)) = \arg \min_{p, q} \Delta (I_1(i, j), I_2(i + p, j + q)) \quad (1.20)$$

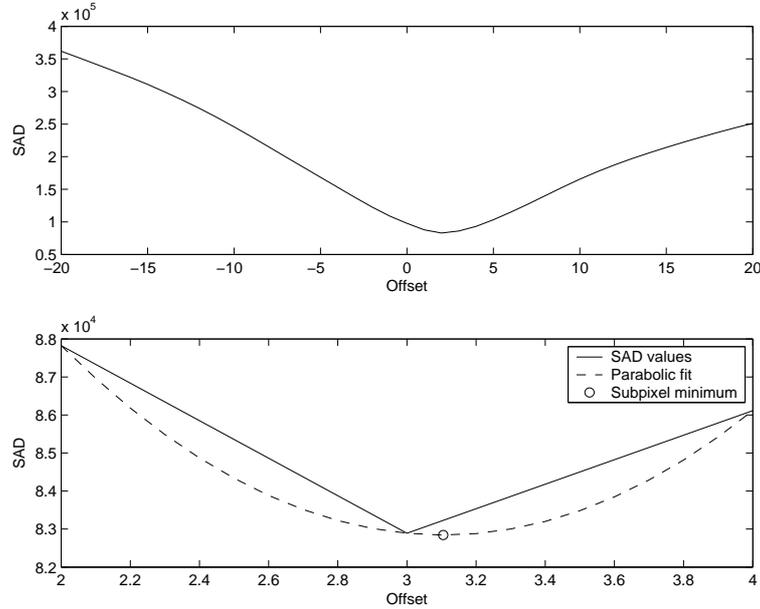


Figure 1.16: The SAD profile across the search radius. A parabolic fit is performed near the minimum to derive sub-pixel accuracy

where Δ is an operator to compare images I_1 and I_2 around (i, j) . In this particular application the vehicle motion is unidirectional so that the search for a matching block can be performed along a single direction too, improving the speed of the algorithm:

$$a(i) = \arg \min_p \Delta (I_1(i, j), I_2(i + p, j)) \quad (1.21)$$

Among the several operators available for Δ we choose to use the Sum of Absolute Differences (SAD) for its good performance and low complexity [7] [8]. (1.21) is then written as

$$a(i) = \arg \min_p \sum_{v=-m}^m \sum_{w=-n}^n |I_1(i + v, j + w) - I_2(i + v + p, j + w)| \quad (1.22)$$

Fig. 1.16 shows the values of the SAD with respect to offsets p for a typical image block. The metric speed $s(i)$ is obtained in obtained from the offsets $a(i)$ using (1.15).

Like the feature detection, block sizes are determined by the size obtained by the first simple evaluation of the speed. Due to the assumption

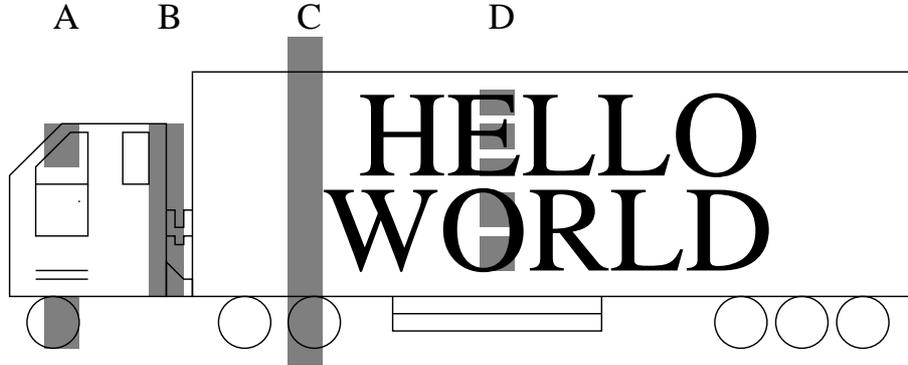


Figure 1.17: Different block shapes use for matching: A, B and C only cover areas with enough horizontal variation. D spans across the whole height of the scan sequence without taking the features into account

of unidirectional motion, the blocks can fill the whole height of the image of the vehicle and only their width must be adapted. As the blocks may become quite large, an improvement on processing speed is possible if only the 'featured' lines are taken into account during the matching process. Fig. 1.17 illustrates which lines would be selected as valid in a synthetic image.

Once the best correspondence is found with (1.22), a classic sub-pixel computation is performed to refine the first motion estimate. The difference δ_a from the integer offset a is the zero of the parabolic curve passing by $(0, SAD(a))$, $(-1, SAD(a-1))$ and $(1, SAD(a+1))$:

$$\delta_a(i) = \frac{SAD(a-1) - SAD(a+1)}{2(SAD(a+1) + SAD(a-1) - 2SAD(a))} \quad (1.23)$$

The resulting speed profile is presented on fig. 1.18. Regions that do not contain enough horizontal variance do not have their speed calculated which results in a broken graph. An interesting artefact can be seen on fig. 1.18: the speed is varying greatly between lines 200 and 300. Correlating the speed profile with the image of the vehicle, one realizes that the variation takes place at the location of the front wheel. Indeed, wheels do not have the translation motion that was an assumption of the line-scan camera system. If the wheel is plain (it does not contain features) then the matching can be performed because the global motion of the wheel is still translational. However, wheels presenting features create an imprint on the scanned image that is not a simple horizontal dilation or con-

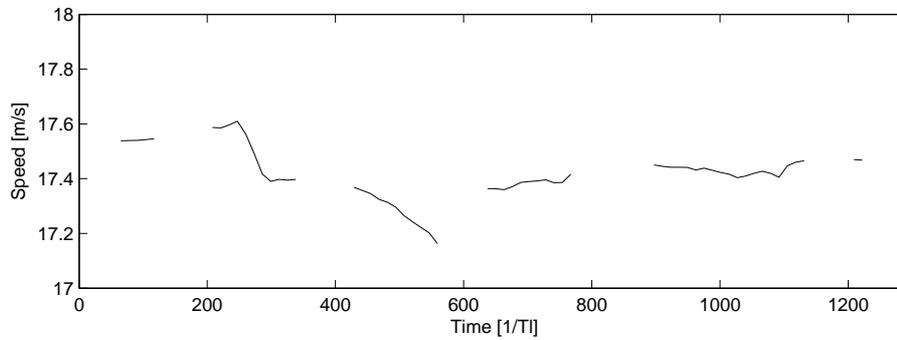


Figure 1.18: *A typical speed profile*

traction because the speed is higher in the upper wheel, and zero in the lower wheel (further assuming no slip). Fig. 1.19a shows the wheel detail of fig. 1.18 and the associated speed profile. The lower speed in the lower part of the wheel is clearly visible on fig. 1.19d ('stretched' venting hole), the inverse being also true in the higher part (small or 'compressed' features). More importantly, the profile measured will depend on which angular difference exists between the wheels as they are acquired by camera 1 and camera 2. If the wheels are captured without angular difference their imprint will be the same on both scanned images and the speed will be correct, as seen on fig. 1.19b. In general the rear wheel does not show such a clear variation because the motion estimation can lock on other more significant features in the same column.

At last, to eliminate as much as bad matches as possible the only measures that will be sent to the linear regression step will be the ones with

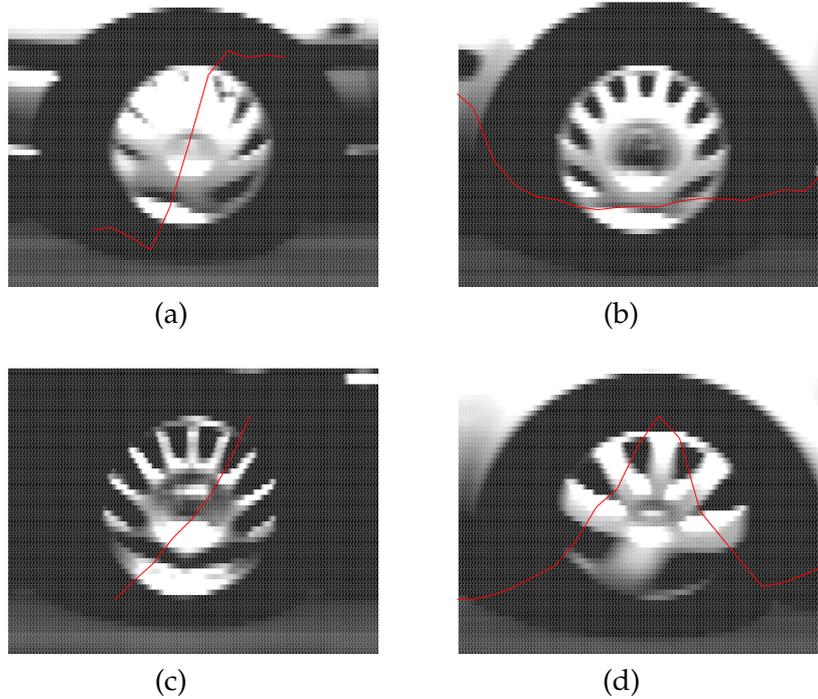


Figure 1.19: A closer look on the speed profile around a few wheel

enough quality. The evaluation of the quality Q is as follows:

$$Q = \frac{\sum_{i=-1}^1 |SAD(a+i) - SAD(a)|}{\sum_{i=-1}^1 SAD(a+i)} \quad (1.24)$$

and has been threshold to 0.01.

While the block matching performs reasonably well with line-scan sequences, other techniques such as feature matching would have to be tested. The SUSAN operator [9] [10] comes in mind and might avoid some artefacts of the BMA.

1.5.4 Linear regression

The linear regression serves three purposes: stabilise the estimations, extract the acceleration and group all speed estimates into a single scalar form. To that effect the speed profile is fitted with a linear approximation which parameters are the speed and acceleration of the vehicle. As detailed in section 1.2.2, the vehicle speed cannot change much during the very small time interval it is seen and we can suppose that the acceleration, if any, will be constant. Suppose that the speed profile consists in r valid estimations of the speed $v(i)$, each obtained at the relative time $t(i) = i/f_l$. The over-determined linear system used to find the speed S_v^f and the acceleration A_v is then:

$$\begin{pmatrix} 1 & t(0) \\ \vdots & \vdots \\ 1 & t(r) \end{pmatrix} \begin{pmatrix} S_v^f \\ A_v \end{pmatrix} = \begin{pmatrix} v(1) \\ \vdots \\ v(r) \end{pmatrix} \quad (1.25)$$

which is nothing else than an expression of the simple formula

$$v = v_0 + at \quad (1.26)$$

The linear system can be solved with a Singular Value Decomposition (SVD) technique [11] and the results of this fit are presented on fig. 1.22 to 1.37. As (1.26) shows, the estimated speed S_v^f is the speed of the front of the vehicle, hence the f superscript. The mean speed is obtained by averaging the speed at the front and at the rear of the vehicle:

$$S_v = \frac{2S_v^f + \frac{B_r^1 - B_f^1}{f_l} A_v}{2} \quad (1.27)$$

1.5.5 Other measurements

Besides the speed, acceleration and length measurements discussed above, a line-scan system can provide other extra information like the lateral position of the vehicle on the road, its height or the inter-vehicle distance. The extraction of the vehicle colour will not be discussed here as it requires colour sensors. The following discussion is based on a side-scanning camera. In the case of a overhead-scanning camera the same measurements can be performed except for the height which is replaced by a measure of width obtained in a similar fashion.

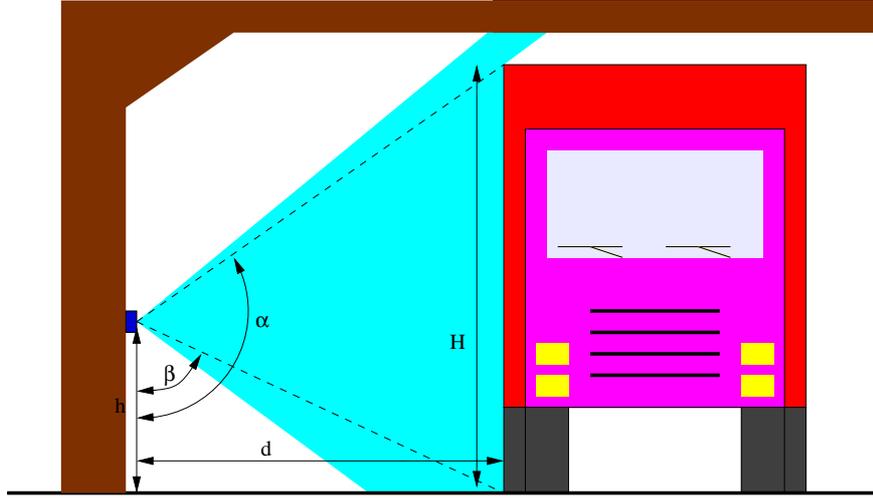


Figure 1.20: Measuring the lateral position and the height of vehicles

Lateral position

The lateral position of the vehicle can be extracted to provide an information about which traffic lane the vehicle is driving on, thereby offering a full monitoring for several lanes with one camera rig. This measure is inferred from the position P_b of the vehicle lowest point in the image which is used to estimate the angle β of the line of sight (fig. 1.20):

$$\beta = \text{atan} \left(\frac{(P_b - S_x/2) P_x}{f} \right) + \frac{\pi}{2} \quad (1.28)$$

where P_x is the pixel size (e.g. $6\mu\text{m}$), S_x is the size of the sensor (e.g. 256 pixels) and f is the camera focal length. The addition of $\pi/2$ reflects the horizontal orientation of the camera, but this angle can be adapted if the camera is mounted otherwise. From β we obtain the vehicle distance d immediately:

$$d = \tan(\beta) h \quad (1.29)$$

where h is the altitude of the camera rig with respect to the ground. Suppose the image measurements have an accuracy of one pixel. From (1.28) and (1.29) we can estimate the accuracy of the estimation of the lateral position to be around 8cm for the following set of parameters: $h=1\text{m}$, $P_x=6\mu\text{m}$, $P_b=50$, $S_x=256$ and $f=8\text{mm}$. The measure of lateral position of the vehicle can also be used to automatically compensate for a

lateral speed: if P_b is different on each camera the vertical offset of the block matching process can be adapted and, in extreme cases, a full two-dimensional block matching can be performed.

Height

Once the lateral position of the vehicle is calculated the measure of its height follows a similar pattern: from the vehicle upper boundary in the image P_h we calculate α

$$\alpha = \text{atan} \left(\frac{(P_h - S_x/2) P_x}{f} \right) + \frac{\pi}{2} \quad (1.30)$$

The vehicle height is then

$$H = h + \frac{d}{\tan \left(\alpha - \frac{\pi}{2} \right)} \quad (1.31)$$

Inter-vehicle distance

The enforcement of the minimum inter-vehicle distance is quite difficult today especially in the extreme cases where vehicles follow each other so closely that they will appear as one vehicle for most monitoring systems. This typically happens with area-scan cameras but line-scan cameras are able to differentiate the vehicles and measure their distance thanks to their high speed and proper orientation with respect to the road.

Suppose that two vehicles have been detected, one after the other. The delay between their appearance on the first camera is

$$\Delta T^1 = T_l (B_{f,2}^1 - B_{r,1}^1) \quad (1.32)$$

where the second subscript indicates the vehicle index. Their speed $S_{v,1}$ and $S_{v,2}$ have been calculated, which leads to the following estimation of their mean distance ΔD^1 :

$$\Delta D^1 = \Delta T^1 \left(\frac{S_{v,1} + S_{v,2}}{2} \right) \quad (1.33)$$

A second distance ΔD^2 can be obtained by substituting the delay information from the second camera ΔT^2 to ΔT^1 :

$$\Delta T^2 = T_l (B_{f,2}^2 - B_{r,1}^2) \quad (1.34)$$

If a difference exists between $S_{v,1}$ and $S_{v,2}$ it will affect the measure of the distances that will also differ. As for other measurements, this double checking can be used to obtain a level of confidence.

1.6 Results and discussion

The results for a few selected vehicles are shown on fig. 1.22 to fig. 1.37. Table 1.4 summarises some important measurements performed on the vehicles. Browsing through the speed profiles, the first striking observation is that the speed almost always shows a decreasing value on the windshield of the car, as on fig. 1.33 between lines 400 and 550. The variation is much beyond the accuracy of 1.2% claimed before in Section 1.4.1. These errors can actually be observed elsewhere but there is in general no other features than the windshield frame around the windshield itself so that it is most remarkable there.

To find the reason to this problem we have to go back to the descriptions of the experiments. As the light level is limited in outdoor conditions, pixels have been binned down to the small resolution of 256 pixels per line. This effectively increases the sensor sensitivity by a factor 4 compared to a resolution of 1024, thereby allowing outdoor operation. However, this also decreases the vertical resolution down to a level where matching slanted features becomes a problem: within a single vertical pixel shift the difference in the horizontal position of the edge location changes too much. Fig. 1.21 shows for example that a vertical shift of $1/5$ pixel leads to an error of one full pixel horizontally. The slightest vertical misalignment of the cameras will therefore create an unacceptable shift. To avoid this problem the experiments should be performed with scanned images that have a height/width ratio not too far from the reality (around $1/3$ for cars). In the experiments described here the ratio is close to $1/12$ and the full 1024 pixel height is thus necessary, instead of the limited resolution of 256 used for illumination purposes. Faster vehicles would not pose the same problem as their effective length on the scanned sequence is shorter but all other things being equal, the theoretical accuracy for the measure of their speed is also lower. This brings us to conclude that the accuracy of rapid vehicles is close to the theoretical values presented before, but that slower vehicles will not have a significantly better estimation if the vertical resolution is not sufficient. It would also be possible to align the sensors vertically during the calibration but a much more rigid calibration system would be required to reach an accuracy of $1/4$ pixel.

Even with this problem in mind, the results of table 1.4 have some interesting aspects. We first note that the repeatability of the measures is very good: four cars have a true length close to 3.84m and all measured values are 4.04m. The test car, which is 4.39m long, was also measured with consistent results: all test involving a constant speed yield the same

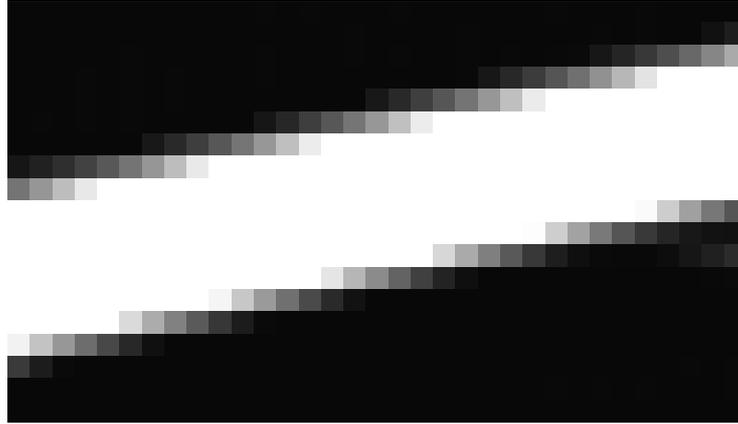


Figure 1.21: A closer view on the windshield frame showing the large offset between the lines

length of 4.63. Due to the higher error in the estimation of the acceleration, the results from braking tests are not as accurate, but still within 1%. From those few measurements we can conclude that the repeatability of the measures is currently only limited by the estimation of the acceleration.

1.7 Conclusion

We have presented a new system for traffic monitoring and law enforcement. The approach we took of using a geometrically constrained camera setup is paying off because the degenerated view of the scene makes it much easier to interpret the vehicles recordings. The results of the first tests are promising and with a few identified modifications could reach the theoretical accuracy, like a higher sensor resolution, a better camera position to avoid shadows or a feature-based matching. The system has several advantages over currently existing systems. First it is able to provide much more than speed: acceleration, length, height, lateral position and inter-vehicle distances can be measured with a single sensor, and a visual output is available. The measures are also more accurate than existing techniques and can even be verified by hand for the purpose of law enforcement. At last, the proposed solution can easily be implemented in real-time with current hardware.

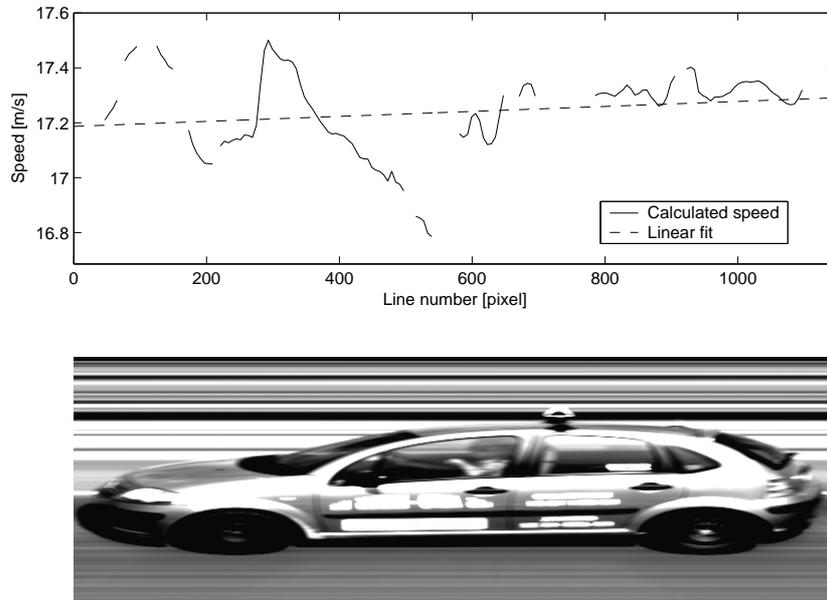


Figure 1.22: Results for car number 2123

However, the line-scan solution poses several practical problems. First it requires a difficult calibration. The accuracy required for this step also means that the camera rig must be very stable to maintain the advertised capabilities of the system. Specifically, meteorological conditions like excessive wind, ice or heat could result in a deformation or a relative motion of the cameras which will in turn yield errors in the speed estimation. For this environmental reason the camera rig must be enclosed in a properly isolated casing which will limit the maximal practical baseline which can be used. On top of that it is possible that other disturbances will have to be tackled: for example if one considers placing the system under a bridge on which the passage trucks might create significant vibrations. Secondly, at the sensor level, its limited optical resolution (with off-the-shelf lenses) and exposure issues like the limited dynamics of the sensor and the pixel size limit the maximum theoretical accuracy that can be reached. These practical issues will have to be overcome before a commercial exploitation of this line-scan traffic analysis system can be envisioned.

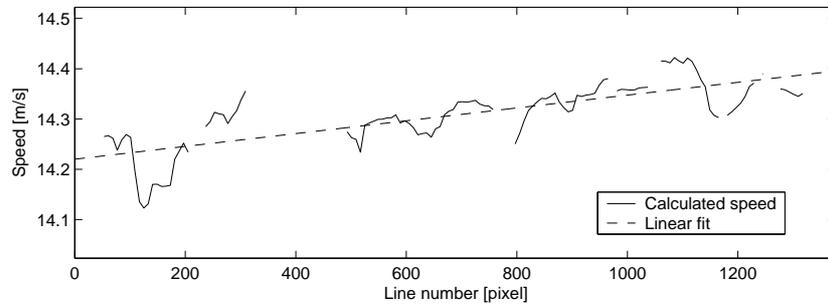


Figure 1.23: Results for car number 6676

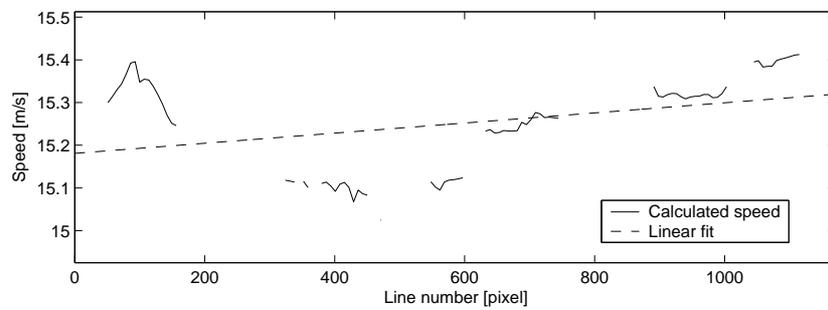


Figure 1.24: Results for car number 11874

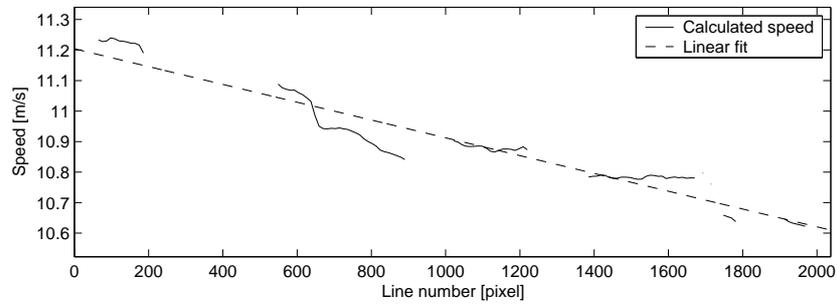


Figure 1.25: Results for car number 25774

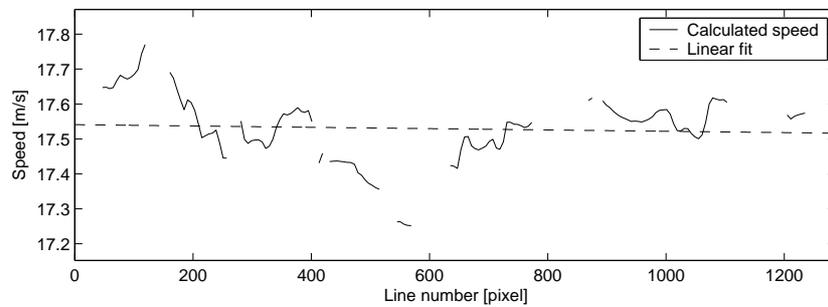


Figure 1.26: Results for car number 27896

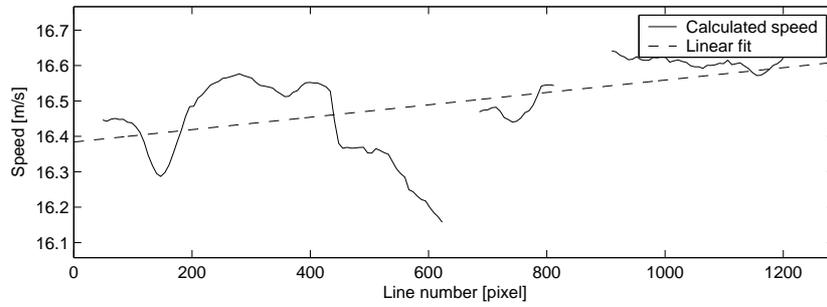


Figure 1.27: Results for car number 37050

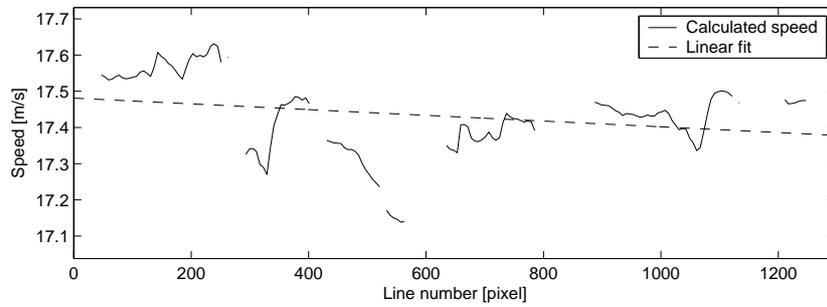


Figure 1.28: Results for car number 38416

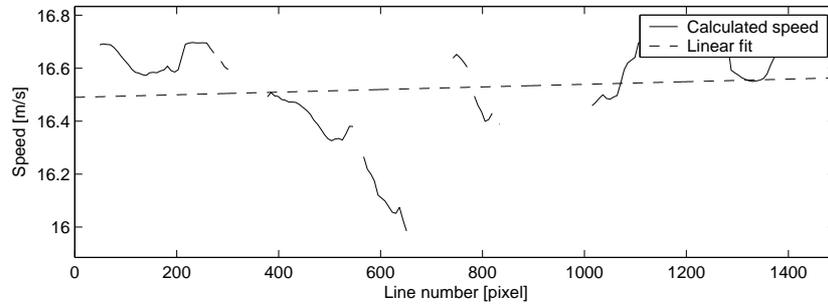


Figure 1.29: Results for car number 50580

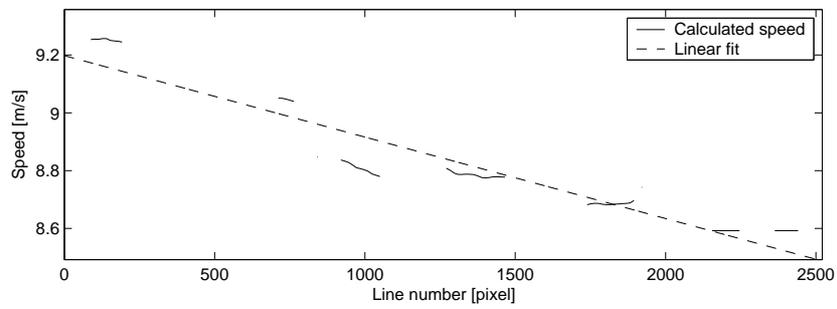


Figure 1.30: Results for car number 61176

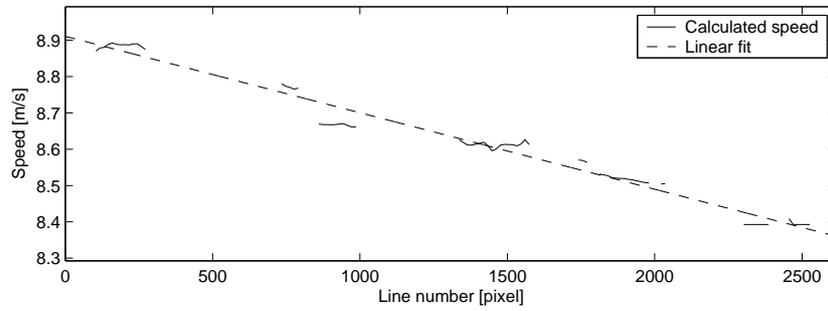


Figure 1.31: Results for car number 64991

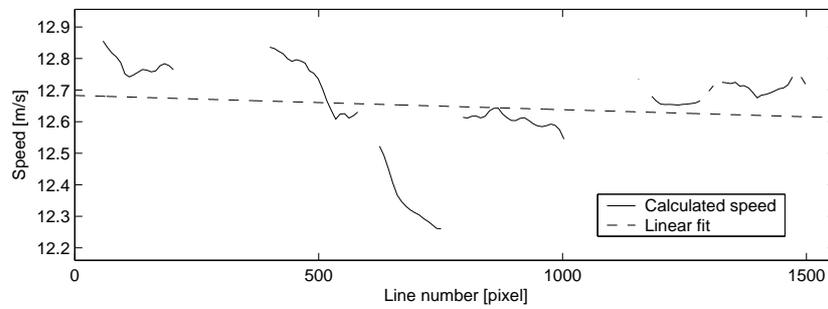


Figure 1.32: Results for car number 68407

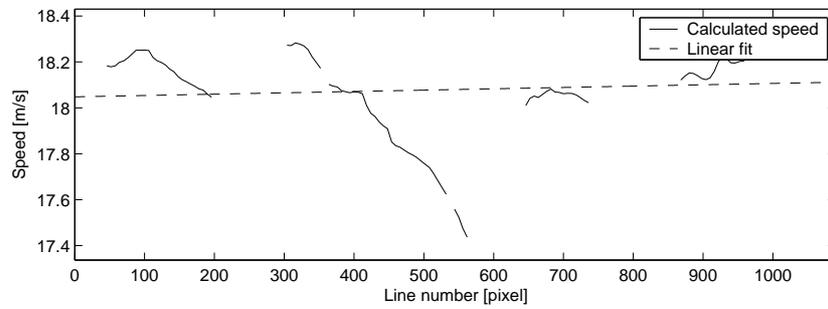


Figure 1.33: Results for car number 74948

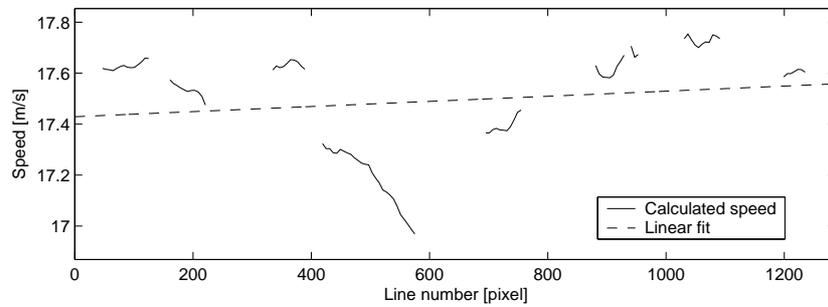


Figure 1.34: Results for car number 78769

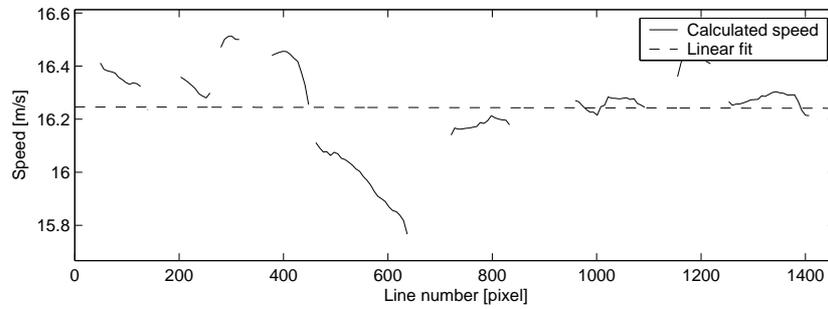


Figure 1.35: Results for car number 89378

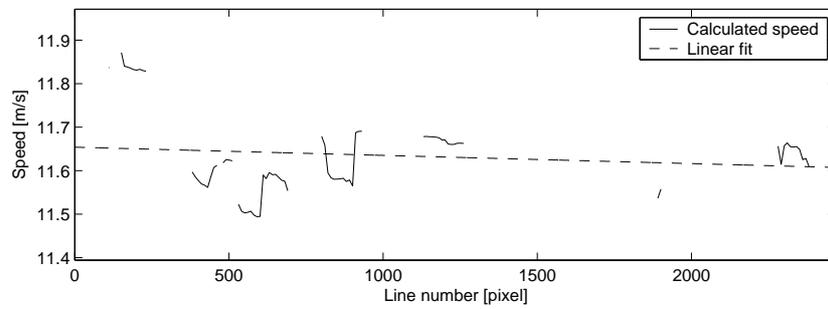


Figure 1.36: Results for car number 94881

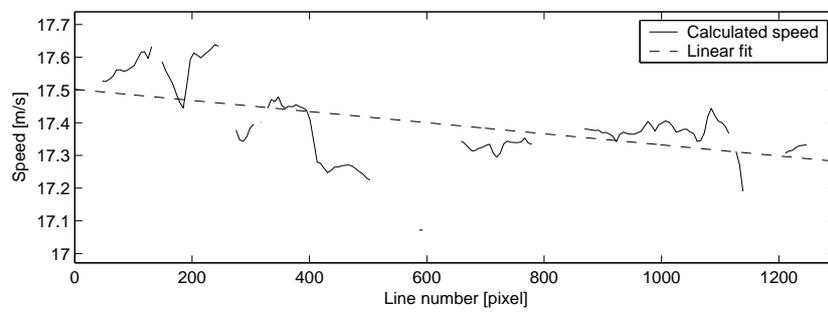


Figure 1.37: Results for car number 97647

Car ID	First estimations			Final estimations			Counter speed <i>km/h</i>	True length <i>m</i>	Rel. fit error %
	Speed <i>km/h</i>	Accel. <i>m/s²</i>	Length <i>m</i>	Speed <i>km/h</i>	Accel. <i>m/s²</i>	Length <i>m</i>			
2123	63.0	0.45	4.10	62.1	0.44	4.04		3.84	0.63
6676	51.9	0.00	4.06	51.5	0.62	4.04		3.84	0.22
11874	55.3	-1.02	3.69	54.7	0.58	3.65		3.54	0.54
25774	38.8	-0.58	4.50	39.3	-1.42	4.56	decel	4.39	0.35
27896	63.4	-1.22	4.65	63.1	-0.09	4.63	70	4.39	0.41
37050	58.9	-2.10	4.32	59.4	0.85	4.36		4.15	0.49
38416	63.2	-0.80	4.66	62.8	-0.39	4.63	70	4.39	0.46
50580	60.3	0.00	5.12	59.5	0.24	5.05		4.73	0.82
61176	31.7	-1.24	4.57	31.8	-1.38	4.59	decel	4.39	0.60
64991	31.1	-0.96	4.65	31.1	-1.03	4.64	decel	4.39	0.23
68407	46.1	0.00	4.09	45.5	-0.22	4.04		3.82	0.81
74948	65.8	0.52	4.08	65.1	0.29	4.03		3.83	0.79
78769	62.8	-1.59	4.61	63.0	0.49	4.62			1.00
89378	58.9	-0.62	4.89	58.5	-0.02	4.85		4.51	0.76
94881	42.5	-0.67	5.97	41.9	-0.09	5.89		5.40	0.63
97647	62.8	-1.58	4.64	62.6	-0.83	4.63	70	4.39	0.42

Table 1.4: Results for a few cars

Chapter 2

Three-dimensional surface reconstruction from sparse data

2.1 Introduction

Free surface measurements constitute an important challenge in experimental fluid mechanics. The free surface is often of analysed because of the key role it plays in the flow phenomena under consideration [12] [13] [14]. Moreover, it may be the only flow feature which is readily accessible to measurements by imaging or remote sensing methods [15] [16] [17]. Both circumstances are encountered in the present study of antidune flows.

Antidunes are trains of bed waves [18] [19] [20] [21] which appear when rapid, shallow currents flow over coarse granular material, and are characterised by in-phase coupling between the oscillatory sediment bed and the water free surface. As illustrated on fig. 2.1, they can be observed both in the field [22] [23] and in the laboratory [24]. Along with other types of fluvial bedforms encountered in modern flow environments and in ancient sediment deposits, they are of interest to geomorphologists and hydraulic engineers [25], for example to determine the proper height of dikes to prevent floods. Antidunes are notable in particular for their evanescent character: they vanish rapidly once the flow wanes, and leave few lasting traces aside from bedding and grain sorting effects. As a result, their geometrical configuration is best studied when the flow is active.

Long- or short-crested, arranged in regular arrays or in narrow trains of peaks and troughs, antidunes come in a variety of patterns. In the present study, a characterisation of such patterns is sought through mea-

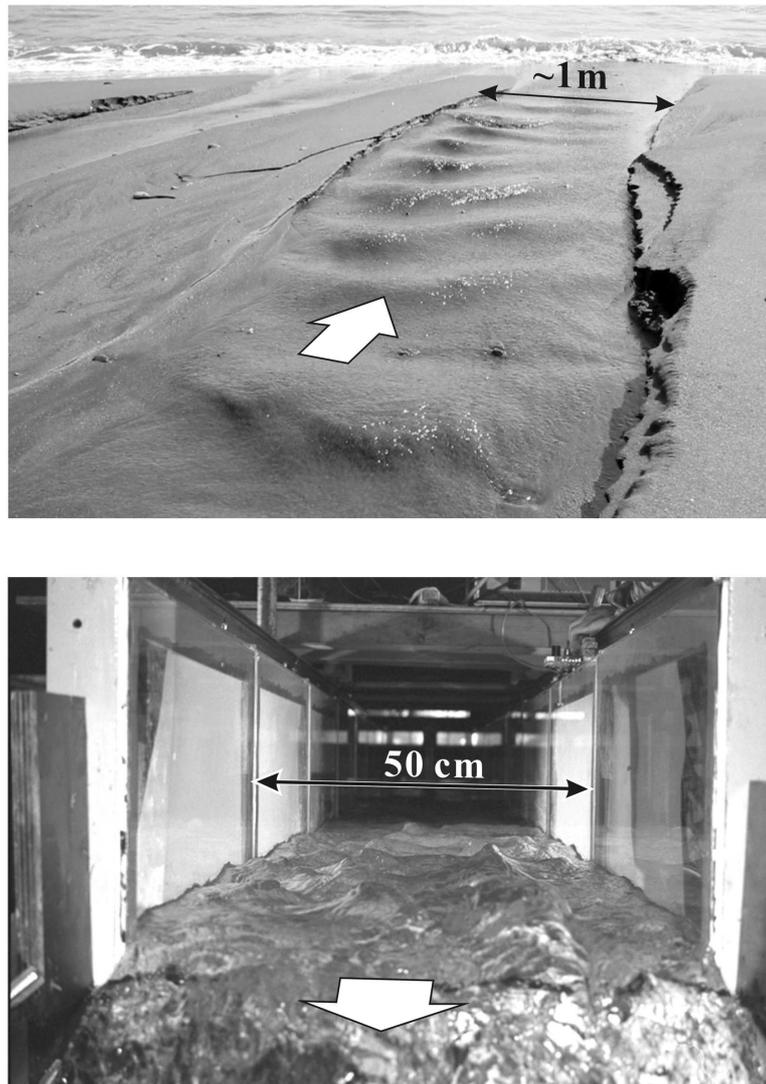


Figure 2.1: Antidune flows on a beach of Eastern Taiwan (top) and in the Louvain laboratory flume (bottom). Courtesy of H. Capart and B. Spinewine

measurements of the water free surface topography. Visual access to the underlying sediment bed surface is impeded by the rough and shallow character of the flowing water layer, hence no direct measurement of the bottom topography is possible. Since the two surfaces are locked in phase with each other, however, the water surface topography provides an indirect image of the bedform pattern.

Measurements of water surface topography most often involve point sensors, placed in multi-sensor arrays at fixed locations or scanned across the surface [12] [26]. Sensors used to measure water elevation include resistive or capacitive gauges, pressure transducers and acoustic beams. Resistance gauges are inapplicable in the present case because of the high sensitivity of antidune flows to intruding objects. More generally, the spatial and temporal resolution of point sensors is limited by the number of available devices or the time required to scan a sensor across the field of interest. This makes them unsuitable for the transient (on a time scale of a few seconds), spatially varied surfaces of flows over evolving bedforms.

Imaging techniques in contrast permit non-intrusive, fast whole-field measurements. Imaging approaches that have been used to measure free surface topographies include shadowgraph [27] [14], shape-from-shading [13], and reflection-based methods [17] [28]. These techniques, however, are restricted to smooth surfaces and small slopes whereas the present application features rough and highly varied water surfaces. Two alternatives are explored in the present work. Both are based on the imaging of floating tracers dispersed on the flowing water, but they involve different reconstruction principles. It is to the best of our knowledge the first time that height maps of free surfaces are derived from particle imaging measurements. The first technique is a classical stereoscopic technique based on the matching of particles on image pairs acquired from two cameras. The second is an original velocimetric technique [29] requiring only a single vertical camera: the horizontal velocity field is first acquired by particle tracking velocimetry, then converted into a vertical elevation map by exploiting a Bernoulli principle derived from the fluid mechanics of the water free surface. Floating tracer particles are required for both methods because small-scale features of the water free surface itself are unsuitable for both tracking in time and matching between views. This is due to the evolving nature of the short wavelength features and to the specular character of the reflection at the water surface [13].

The parallel development of two distinct techniques was motivated by the following considerations. First, it makes possible to weigh the

respective advantages of the methods and evaluate their applicability to more complex situations. Second, an assessment of the validity of both techniques can be obtained by comparing their results. To permit cross-validation, the two methods are applied to the same experiments, while each relies on its own separate camera footage and data analysis pipeline. Also, special care is taken to estimate the errors incurred at each stage of the reconstruction processes. Preliminary results of the present research effort were reported in [30], [31] and [32].

This chapter is organised as follows. In the next section, the experimental set-up and camera systems used for the measurements are presented. Then, basic particle imaging algorithms used as building blocks by the two reconstruction methods are outlined. The next two sections are then devoted to a detailed presentation of each of the two reconstruction methods. This includes special procedures for the estimation and filtering of measurement errors at each step of the analysis. Finally, results are shown, compared and discussed in the last section, followed by a verification of the error estimations which precedes the overall conclusions of the work.

2.2 Experimental setup and imaging configuration

Experiments are conducted in a hydraulic flume having the following dimensions: length = 6m; width = 50 cm; side wall height = 50 cm. The flume is tilted to obtain a bottom slope of 1%. A 5 cm deep layer of loose sediment covers the flume bottom and is replenished during the flow by an upstream silo. Sand of nearly uniform size distribution is used as sediment material, and has the following characteristics: mean grain diameter = 1.65mm; density = 2615kg/m³. For the proof-of-concept experiments presented hereafter, the water inflow was not tightly controlled but rather loosely operated in order to produce a variety of patterns. Water discharges in the flume ranged from 15 to 35l/s, yielding flow depths of 4 to 6 cm and Froude numbers¹ in the range $Fr = 1.4$ to 2. For such high Froude numbers, the water free surface responds strongly to the underlying oscillatory topography of the bedforms, up to the point of breaking at the wave crests when antidunes are fully developed.

Starting from a plane sand bed, the antidunes emerge as longitudinal

¹The Froude number can be seen as the ratio between buoyancy and inertial forces. Turbulent flows will have a high Froude number because inertial forces are dominant. On the contrary, laminar flows will have a small Froude number.

trains of crests and troughs initiating from downstream but stationary in phase. The antidunes are observed to respond to transient changes in the flow rate (both increase and decrease) by temporarily growing in amplitude. Amplitude responses and gradual shifts in pattern occur on a time scale of tens of seconds. By contrast, on the shorter time scale corresponding to the image acquisition (of the order of 2s), the hydrodynamics can be assumed to be quasi-steady and this is exploited hereafter to derive a single surface from each measurement sequence. The flow is however observed to exhibit small unsteady pulses, and this physical source of noise may slightly affect the results concurrently with measurement errors.

The measurement section is placed some 2m upstream of the flume outlet. Moments before image acquisition, a uniform dispersion of floats is dropped onto the mean surface by means of staggered metal meshes. The tracer particles are white wooden pearls 9mm in diameter. Image sequences are obtained using digital cameras placed above the flow. The velocimetric and stereoscopic methods require rather different image acquisition systems, hence some care is necessary in order to operate them simultaneously and observe the same scene.

Two commercial digital cameras (miniDV, PAL, 25 frames per second) are used for the stereoscopic measurements. These cameras offer good image resolution (768 by 576 pixels) but cannot be synchronised with each other during the acquisition. To avoid motion-stereo ambiguity, it will thus be necessary to synchronise them *a posteriori* using an interpolation procedure (see below). The stereo cameras are placed above the flow with oblique optical axes contained in a vertical plane parallel to the direction of flow (see fig. 2.2). For the velocimetric measurements, on the other hand, a high frame rate camera (250 frames per second; resolution of 239 by 192 pixels) is placed directly over the flow with a nearly vertical optical axis. Due to the high frame rate, this camera requires strong lighting, obtained with four 2 KW light sources. Such powerful lighting saturates the commercial cameras even at maximum shutter speed, and these have to be fitted with dimming filters.

After positioning, the viewpoints of all three cameras are determined by placing a calibration target in their shared viewing volume. This is essential for stereo reconstruction and allows the results of the two methods to be obtained in the same three-dimensional referential. For ease of reference in the next sections, we adopt labels $V = (L; R; T)$ to identify the three viewpoints. Labels L and R designate the left and right oblique cameras used for stereoscopic reconstruction, while label T refers to the top

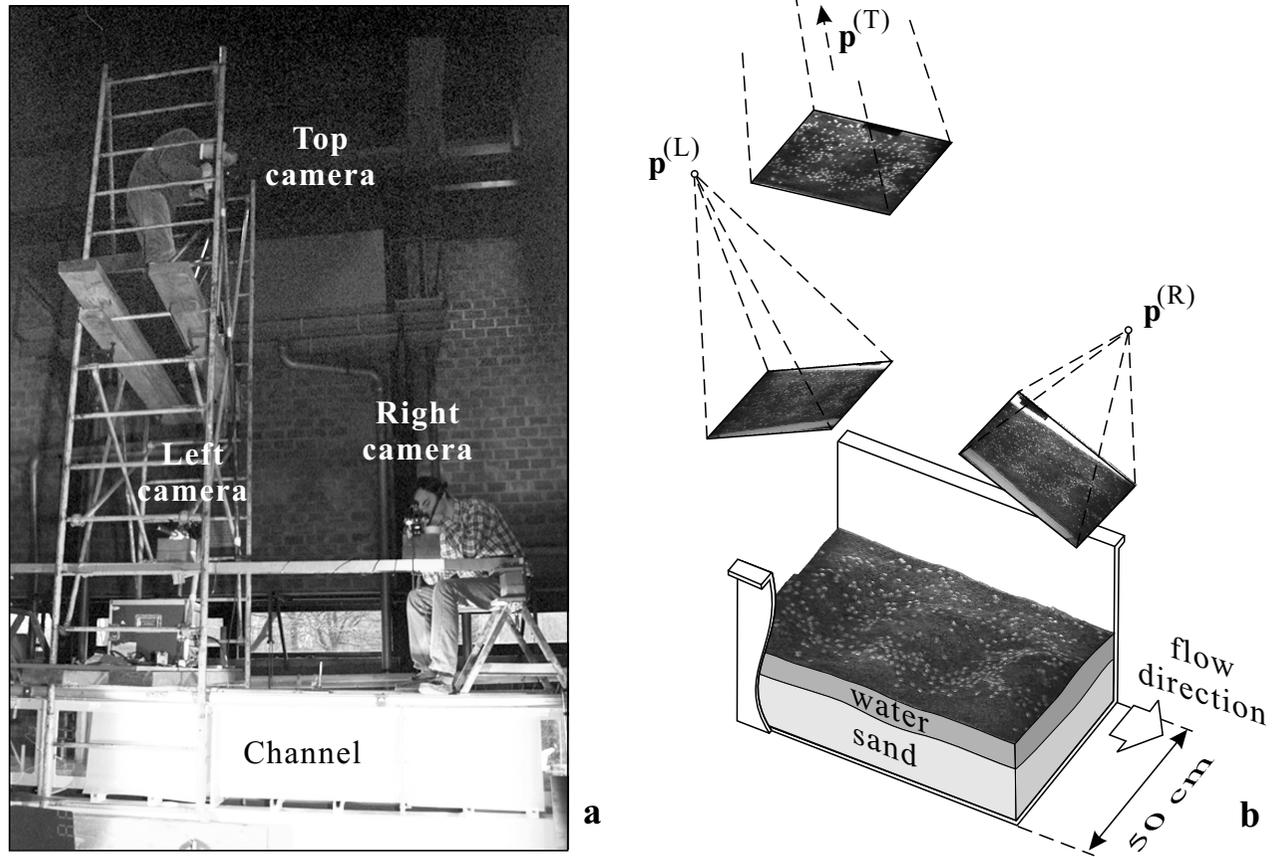


Figure 2.2: Imaging configuration

camera used for velocimetric reconstruction. The camera configuration shown on fig. 2.2b is reconstructed based on actual calibrated viewpoints.

2.3 Particle imaging algorithms

2.3.1 Particles detection and positioning

Consider image sequences acquired under any of the three viewpoints $V = L, R, \text{ or } T$. For both reconstruction techniques, the first step of the analysis is the localisation of particle centres on individual images. One seeks to extract for each viewpoint V a set of positions $\{\mathcal{R}_{i,m}^{(V)}\}$ where $\mathcal{R}_{i,m}^{(V)} = (X_{im}^{(V)}, Y_{im}^{(V)})$ denotes the position of the i -th particle visible on the m -th image of the sequence. The white particles show up on the images as bright spots against a dark, relatively noisy background (see Section 2.2). Such bright spots are first highlighted by convoluting the image with a laplacian-of-Gaussian filter [33]. Local maxima of the highlighted images are then sought. The position of each maximum is finally refined to sub-pixel accuracy by way of a quadratic interpolation surface (for more details, see [34]). Illustrating these steps, a filtered image with detected particle positions is shown on fig. 2.3a. The expected root mean square accuracy on the X and Y image coordinates obtained with such a procedure is of the order of 0.1 to 0.25 pixels [35] [34].

2.3.2 Particle tracking

Particle tracking is required in the present work for both reconstruction methods. For velocimetric reconstruction, particle tracking is used to acquire the horizontal velocity field and constitutes an essential component of the method. For stereoscopic reconstruction, tracking is not an essential component, but is required in the present work because of the use of non-synchronised cameras. A posteriori synchronisation is needed to avoid stereo-motion ambiguity, and is achieved by interpolating particle positions along their tracked trajectories.

The particle tracking step aims at establishing correspondence between particles seen on one image frame and particles seen on the next. Since particles are identical, the information available is limited to the two sets of particle positions $\{\mathcal{R}_{i,m}^{(V)}\}$ and $\{\mathcal{R}_{j,m+1}^{(V)}\}$ obtained for the two successive frames m and $m + 1$. The tracking problem is then one of finding

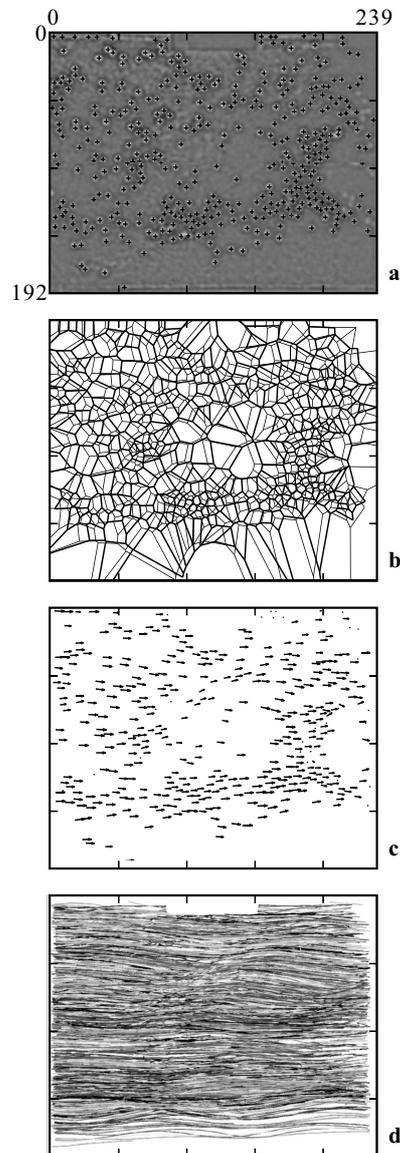


Figure 2.3: Particle positioning and tracking operations applied to a sequence acquired with the top camera: (a) filtered image and particle positions (+); (b) Voronoi tessellations constructed on particle positions identified on one frame (thin lines) and the next (thick lines); (c) displacement vectors obtained by matching the shapes of the Voronoi cells [34]; (d) particle trajectories obtained by concatenating displacements over the whole sequence.

the pairing $j(i)$ associating each particle j on frame $m + 1$ to its most reasonable counterpart i on the previous frame m .

In the present work, this operation is performed using the Voronoï algorithm of Capart et al. [34]. This algorithm tracks local patterns of neighbouring particles using Voronoï cells as match templates. Voronoï cells (see e.g. [36]) are polygons surrounding each particle, encompassing the region of the plane closer to that particle than to any other. Qualitatively, the Voronoï match algorithm pairs particles which are characterised on two successive images by similar Voronoï cells. This measure of similarity is obtained by comparing the relative position of the neighbours using polar coordinates. The procedure is illustrated on fig. 2.3: Voronoï diagrams constructed on two successive frames are shown on panel 2.3b and the corresponding displacement vectors are plotted on panel 2.3c. As illustrated on fig. 2.3d, displacements can finally be concatenated to extract particle trajectories over the entire image sequence. For certain image pairs, a small number of mismatches are obtained. They are filtered out using the procedures detailed in [34].

2.3.3 Imaging geometry and camera calibration

One must now be able to relate the image coordinates $\mathcal{R}^{(V)} = (X^{(V)}, Y^{(V)})$ associated with each camera viewpoint V to a system of 3D world coordinates $r = (x, y, z)$ measured in a shared reference frame. The required transformation can be modelled as a central projection followed by an affine transformation in the image plane [37] [38]. This amounts to specifying for each viewpoint V a matrix $A^{(V)}$ and a vector $b^{(V)}$ allowing image coordinates $(X^{(V)}, Y^{(V)})$ to be computed from a known position $r = (x, y, z)$ according to

$$\alpha \begin{pmatrix} X^{(V)} \\ Y^{(V)} \\ 1 \end{pmatrix} = A^{(V)} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + b^{(V)} \quad (2.1)$$

Conversely, a point having image coordinates under viewpoint V belongs to a 3D ray defined by the parametric equation

$$r^{(V)}(\alpha) = p^{(V)} + \alpha q^{(V)} \quad (2.2)$$

where scalar α is the free parameter. Vectors p and q denote the position of the projection centre and the ray direction, respectively, and are given by

$$p^{(V)}(\alpha) = - \left(A^{(V)} \right)^{-1} b^{(V)} \quad (2.3)$$

$$q^{(V)}(\alpha) = \left(A^{(V)} \right)^{-1} \begin{pmatrix} X^{(V)} \\ Y^{(V)} \\ 1 \end{pmatrix} \quad (2.4)$$

A least-squares camera calibration procedure is needed to obtain matrix $A^{(V)}$ and vector $b^{(V)}$ for each viewpoint, as explained in more details in [39]. In the present experiments, each of the viewpoints $V = L, R,$ and T is calibrated using some 200 points of known world and image coordinates. They are acquired from snapshots of a target placed within the viewing volume before or after the experimental runs.

2.4 Stereoscopic reconstruction

The first reconstruction method is a feature-based stereoscopic matching procedure used in digital photogrammetry [40] and applied here to the floating tracer particles. The core of the method lies in the matching of rays associated with the particle images picked up by the two oblique cameras. This results in a three-dimensional cloud of points, from which the desired water free-surface is recovered by surface fitting. The procedure developed in the present context includes two original aspects: a posteriori camera synchronisation needed to avoid stereo-motion ambiguity and maximum-likelihood surface fitting used to filter out stereoscopic mismatches. These various steps are now detailed.

2.4.1 Ray intersection and matching

Consider two candidates for a stereoscopic match characterised by coordinates $\mathcal{R}_i^{(L)}$ and $\mathcal{R}_j^{(R)}$ on images seen by the left and right cameras ($V = L, R$). The corresponding rays are given by parametric equations (see (2.2))

$$r_i^{(L)}(\alpha) = p_i^{(L)} + \alpha q_i^{(L)} \quad (2.5)$$

$$r_j^{(R)}(\alpha) = p_j^{(R)} + \alpha q_j^{(R)} \quad (2.6)$$

With reference to fig. 2.4, the points on the left and right rays corresponding to the minimum distance between the two rays are parameterised by the values α, β which are solutions to the system

$$\begin{pmatrix} q_i^{(L)} \cdot q_i^{(L)} & -q_i^{(L)} \cdot q_j^{(R)} \\ q_j^{(R)} \cdot q_i^{(L)} & -q_j^{(R)} \cdot q_j^{(R)} \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} q_i^{(L)} \cdot (p_j^{(R)} - p_i^{(L)}) \\ q_j^{(R)} \cdot (p_j^{(R)} - p_i^{(L)}) \end{pmatrix} \quad (2.7)$$

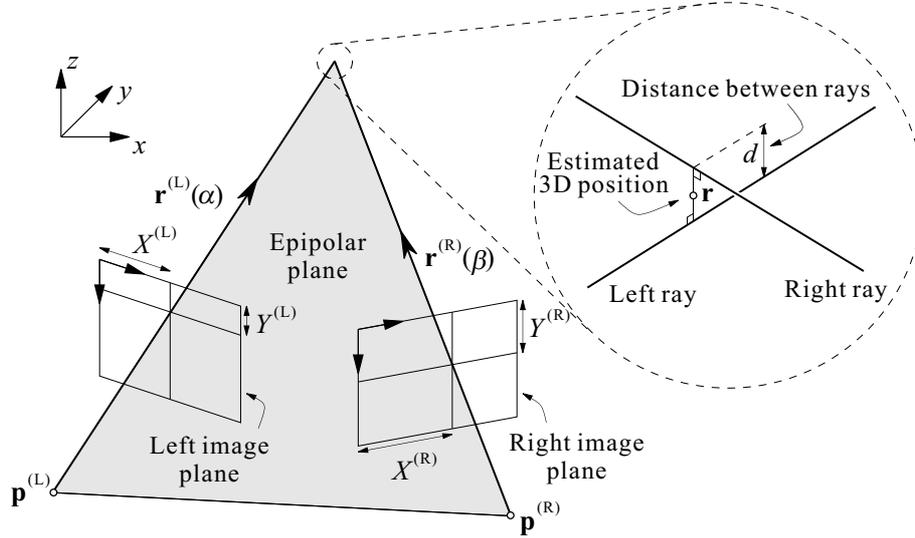


Figure 2.4: Intersection of two rays associated with the left and right camera viewpoints.

where the dot denotes the usual dot product. The midpoint $r_{i,j}$ and distance $d_{i,j}$ between the two points of closest encounter are then given by

$$r_{i,j} = \frac{1}{2} \left(r_i^{(L)}(\alpha) + r_j^{(R)}(\beta) \right) \quad (2.8)$$

$$d_{i,j} = \left\| r_j^{(R)}(\beta) + r_i^{(L)}(\alpha) \right\| \quad (2.9)$$

Due to the limited accuracy of the camera calibration and image plane measurements, rays corresponding to the same physical point will not perfectly intersect. Midpoint $r_{i,j}$ constitutes an approximate intersection, while the distance of closest encounter $d_{i,j}$ provides an indicator of the quality of the approximation.

A procedure to select stereo pairs follows. Let $\{\mathcal{R}_i^{(L)}\}$ and $\{\mathcal{R}_j^{(R)}\}$ designate the two sets of particle positions picked up at the same time t by the left and right cameras. The task is to find the pairing $j(i)$ most likely to match together rays corresponding to one and the same physical particle. The stereo matching problem is in this regard identical to the velocimetric tracking task sketched in Section 2.5 and documented in greater detail in [34]. It can be cast as an optimisation problem, whereby one seeks to

minimise the objective function

$$\sum_i d_{i,j(i)} \quad (2.10)$$

under the constraint that a given particle image can participate in only one stereo pair. This bijective graph optimisation problem can be solved approximately using the Vogel algorithm: consider for each ray the best and second best match, then construct a reasonable global optimum by picking ray pairs in the order of maximum difference between first and second best choices. The set of 3D positions is finally obtained from (2.8) as $\{r_{i,j(i)}\}$. The matching and intersection of rays associated with particle images thus yield the world coordinates of the physical particles. Note that the Voronoï-based matching used for particle tracking in the velocimetric technique could not be used for stereoscopic analysis because the high-speed camera position forces the two stereo cameras to be widely separated. This will result in large perspective difference between the two stereo views, and therefore in very different Voronoï cells. A possible improvement of the stereo technique could result from using a shorter baseline stereo setup and exploiting the particle patterns (or their projective invariant characteristics).

2.4.2 A posteriori synchronisation

In the case of a dynamic scene, an essential requirement of the procedure sketched above is that two images used for stereo matching be acquired at the same time t . Otherwise, the displacements of the particles between the two different times will register an artificial change in three-dimensional position [41]. The commercial cameras available for the present work, however, did not permit synchronised video acquisition. As a result, sets $\{\mathcal{R}_{i,m}^{(L)}\}$ and $\{\mathcal{R}_{i,m}^{(R)}\}$ are sampled at separate times $\{t_m\}$ and $\{t_n\}$. To avoid stereo-motion ambiguity, positions on the right images must first be resampled at times $\{t_m\}$ at which positions on the left images are available.

This is done in two steps. First, a LED display clock placed in the view of both stereo cameras makes it possible to determine the index n such that $t_n \leq t_m < t_{n+1}$. Resampled positions $\mathcal{R}_j^{(R)}(t_m)$ are then obtained by linear interpolation along particle trajectories, i.e.

$$\mathcal{R}_j^{(R)}(t_m) = \mathcal{R}_j^{(R)}(t_n + \tau) = \mathcal{R}_{j,n}^{(R)} + \frac{\tau}{\Delta t} \left(\mathcal{R}_{k(j),n+1}^{(R)} - \mathcal{R}_{j,n}^{(R)} \right) \quad (2.11)$$

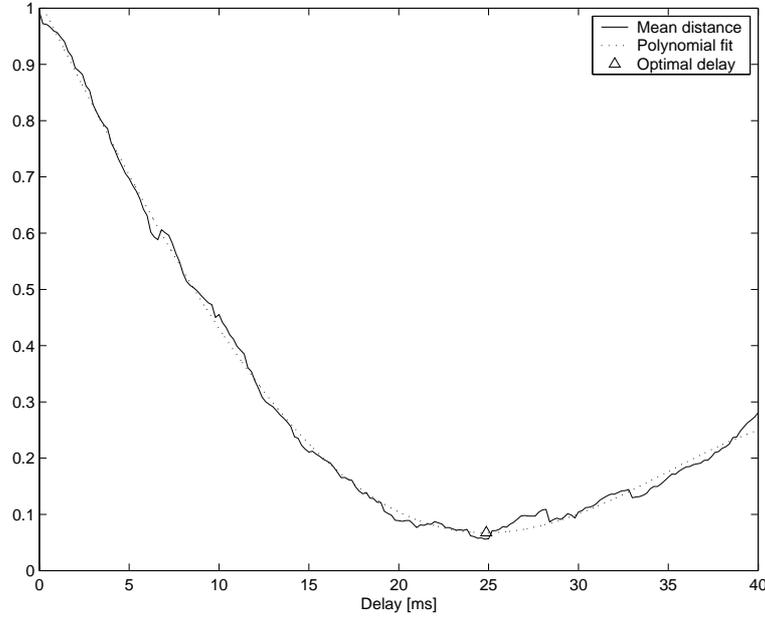


Figure 2.5: Variation of the mean distance between rays with the camera desynchronisation delay

where velocimetric matching $k(j)$ derives from the tracking procedure outlined in section (2.3.2), and where $0 \leq \tau < \Delta t$. At this point, the lag time τ is still unknown. The correct estimate of τ , however, should yield the best stereoscopic fit. It can thus be found by searching for the value τ which minimises $\frac{1}{N} \sum d_{i,j(i)}$ where $d_{i,j(i)}$ is the ray distance between stereo matches and N is the number of matched pairs. Since $d_{i,j(i)}$ depends on the particle matching the latter process will have to be carried out for each different τ . The variation of the resulting mean distance between rays is shown on fig. 2.5. Since the stereo cameras operate at the same frequency, the inter-camera lag time τ is a constant which only needs to be estimated once for each acquired sequence.

2.4.3 Bayesian surface interpolation

The above matching procedure results in sets of points $\{r_{i,m}\}$ positioned within the viewing volume at time instants $\{t_m\}$. Since in the present application the flow is assumed to be quasi-steady over the duration of each video acquisition, the different subsets $\{r_{i,m}\}$ can further be pooled into

one single set $\{r_k\}$. A continuous water free surface elevation $\eta(x, y)$ remains to be extracted from this cloud of points. If the matching process were perfectly accurate, all points would belong to the desired surface and interpolation would suffice. In practice, however, various sources of error conspire to offset the points around the true surface, and some form of fitting is required.

In the present case, adequate fitting requires careful consideration of the structure of the measurement errors. The measured points set $\{r_k\}$ includes two distinct subsets: (i) points resulting from correctly matched rays (i.e. they do indeed correspond to one and the same physical particle) but positioned with a limited accuracy; (ii) points resulting from mismatched rays. It is reasonable to expect that points of the first subset will cluster around the physical surface with an error distribution that is approximately Gaussian, due both to position inaccuracy and to physical fluctuations. Mismatches, however, create a very different kind of background noise.

With reference to fig. 2.6a let us suppose for the sake of argument an idealised, symmetrical viewing geometry with focal points of the projection sent to infinity, and a uniform distribution of particles on a wavy surface. Geometrically, mismatches correspond to the establishment of a wrong correspondence between two rays that approximately belong to the same epipolar plane (i.e. the plane containing the point of intersection of the two rays and the focal points of both projections). Among the family of such rays, each one is equally likely to produce a mismatch, hence mismatched points will be distributed randomly in the viewing volume rather than cluster around the physical surface. In an actual imaging geometry, the reasoning should still hold in an approximate way, hence the structure of the point distribution is expected to combine: 1) correct matches corrupted by Gaussian noise; 2) mismatches uniformly distributed in the viewing volume. In the presence of a high proportion of mismatches, highly biased results can be obtained if this structure is overlooked.

Rather than the whole set, let us consider the points belonging to a prismatic bin of dimensions $(\Delta x, \Delta y)$ around position (x, y) . Let ω_1 designate the correctly matched state in which a particle can find itself and ω_2 designate the mismatched state. $P(\omega_1)$ and $P(\omega_2)$ are the corresponding probabilities. Based on the above model, the point elevations $\{z_k\}$ are randomly distributed with probability density function given by

$$p(z|\Theta_1, \Theta_2) = p(z|\omega_1, \Theta_1)P(\omega_1) + p(z|\omega_2, \Theta_2)P(\omega_2) \quad (2.12)$$

where Θ_1, Θ_2 designate the parameter vectors which characterise the dis-

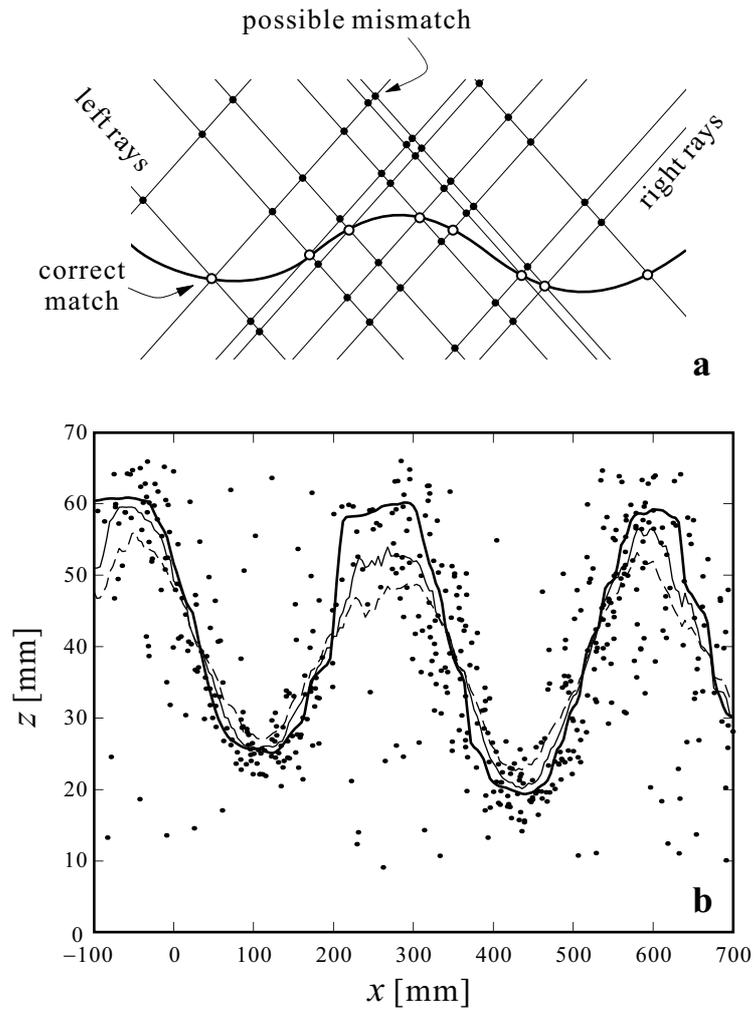


Figure 2.6: Free surface reconstruction in the presence of stereo mismatches: (a) conceptual sketch; (b) comparison of mean (dashed line), median (thin line) and maximum-likelihood reconstructions (thick line) of the free surface based on stereo data points (dots).

tribution of each class. This is a so-called mixture density, in which conditional densities $p(z|\omega_1, \Theta_1)$, $p(z|\omega_2, \Theta_2)$ are called component densities and probabilities $P(\omega_1)$, $P(\omega_2)$ are called the mixing parameters [42]. For the normal and uniform distributions postulated above, we have

$$p(z, |\omega_1, \Theta_1) = \frac{1}{\sigma_\eta \sqrt{2\pi}} \exp\left(-\frac{(z - \eta)^2}{2\sigma_\eta^2}\right) \quad (2.13)$$

$$p(z|\omega_2, \Theta_2) = (z_b - z_a)^{-1} \quad (2.14)$$

hence parameter vectors are given by $\Theta_1 = (\eta, \sigma_\eta)$ and $\Theta_2 = (z_a, z_b)$, where η , σ_η are respectively the mean and standard deviations of the distribution of correct matches and z_a, z_b are the (possibly truncated) vertical limits of the viewing volume at location (x, y) . Symbol η has been used to designate the mean of the correct matches because it designates the surface elevation that is ultimately sought.

The problem of finding the surface elevation has thus been cast into one of finding an optimal estimate for parameter η of a component density. Since it is not known a priori which points are correct matches and which are mismatches, the task at hand belongs to the category of unsupervised learning problems [42]. The elevation boundaries z_a and z_b are known and by definition $P(\omega_1) + P(\omega_2) = 1$. Three parameters thus remain to be determined: η , σ_η and $P(\omega_1)$. In a manner analogous to the case of normal mixtures treated by Duda and Hart in [42], maximum-likelihood estimates $\hat{\eta}$, $\hat{\sigma}_\eta$ and $\hat{P}(\omega_1)$ must satisfy

$$\hat{P}(\omega_1) = \frac{1}{n} \sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta) \quad (2.15)$$

$$\hat{\eta} = \frac{\sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta) z_k}{\sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta)} \quad (2.16)$$

$$\hat{\sigma}_\eta^2 = \frac{\sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta) (z_k - \hat{\eta})^2}{\sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta)} \quad (2.17)$$

where by Bayes' rule

$$\hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta) = \frac{p(z_k|\omega_1, \hat{\eta}, \hat{\sigma}_\eta) \hat{P}(\omega_1)}{p(z_k|\omega_1, \hat{\eta}, \hat{\sigma}_\eta) \hat{P}(\omega_1) + p(z_k|\omega_2, z_a, z_b) (1 - \hat{P}(\omega_1))} \quad (2.18)$$

and where $p(z_k|\omega_1, \hat{\eta}, \hat{\sigma}_\eta)$ and $p(z_k|\omega_2, z_a, z_b)$ are given by distributions (2.13) and (2.14). Despite the involved expressions, estimates (2.15)-(2.17) can be interpreted rather simply as a frequency ratio, a sample mean and a sample variance respectively but with samples weighed according to the conditional probability that they belong to the class of correctly matched particles. A solution of the above set of implicit equations can be obtained by iterations. Initial estimates for $\hat{\eta}$, $\hat{\sigma}_\eta$ and $\hat{P}(\omega_1)$ can first be used to evaluate using (2.18). Equations (2.15)-(2.17) can then be used to update the estimates and so on. If a quality measure or a cost is available for each match it is possible to add a weight parameter W_k to (2.18) and (2.15), resulting in the following more generic expressions:

$$\hat{P}(\omega_1) = \frac{\sum_{k=1}^n \hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta)}{\sum_{k=1}^n W_k} \quad (2.19)$$

$$\hat{P}(\omega_1|z_k, \hat{\eta}, \hat{\sigma}_\eta) = \frac{p(z_k|\omega_1, \hat{\eta}, \hat{\sigma}_\eta) \hat{P}(\omega_1)}{p(z_k|\omega_1, \hat{\eta}, \hat{\sigma}_\eta) \hat{P}(\omega_1) + p(z_k|\omega_2, z_a, z_b) (1 - \hat{P}(\omega_1))} W_k \quad (2.20)$$

In our case the weight can include the inter-ray distance of the match d_k and the distance r_k of the reconstructed three-dimensional point to the lattice point (x, y) :

$$W_k = \frac{1}{r_k d_k} \quad (2.21)$$

This enhancement did not provide better results because the inter-ray distance is not a stable indicator of the quality of the match. It could however be used for applications where a meaningful quality measure is available.

While convergence to a single solution is not guaranteed in all cases, the procedure was found to perform adequately in this application, provided reasonable initial estimates are chosen. For the present experimental conditions and procedures, the converged estimates for $\hat{\sigma}_\eta$ and $\hat{P}(\omega_1)$ are respectively $\approx 5\text{mm}$ and ≈ 0.6 . The latter number is especially significant: about 40 percent of the data points can be assessed to result from mismatches! This is due to the combination of limited image resolution with a dense dispersion of tracer particles on the free surface, creating many opportunities for epipolar ambiguity. The high likelihood of mismatches underscores the importance of filtering out their effect on the final measurements.

Fig. 2.6b presents results for a sample longitudinal slice of the measurements of run 3. Estimates obtained by binning the values, then taking the local mean or median rather than applying the above maximum-likelihood estimate are also shown. Although the median is slightly less affected, both these simpler estimates are seen to be significantly biased by the presence of uniformly distributed mismatches in addition to the more classical Gaussian scatter. The maximum-likelihood Bayesian estimate exhibits a better behaviour, disregarding the random mismatches and focusing on the densely clustered points. Similar techniques aimed at weighting down outlying measurements in the presence of non-Gaussian noise are used in high-energy physics [43].

2.4.4 Balance of noise and attenuations errors

To carry out the above procedure, individual data points must be assigned to finite size bins. Overlapping square bins are chosen, with bin centres positioned on a regular grid $x_i \times y_j$. These bins have a variable size, each one encompassing a constant number of data points n . An optimal choice of this number n is based on the following considerations. When applied to a spatially distributed set of data corrupted by uncorrelated Gaussian noise, averaging produces two different effects: on the one hand, the Gaussian noise error diminishes in proportion to $1/\sqrt{n}$ if only Gaussian error is present (i.e. there are no stereo mismatches); on the other hand, the coarser the averaging, the higher the attenuation of fine-grained details of the spatial field, inducing an attenuation error which grows with the bin size and number n of binned data. These two effects act at cross-purposes and it is possible to choose a number n which minimises the combined error.

Let $\hat{\eta}^{(n)}(x, y)$ designate the reconstructed surface resulting from assigning n data points to each bin and applying the Bayesian maximum-likelihood estimator of the previous section. An estimate of the residual Gaussian noise error is given by

$$\left\langle \left(\hat{\eta}^{(n)} - \bar{\eta}^{(n)} \right)^2 \right\rangle^{1/2} = \frac{\hat{\sigma}_\eta}{\sqrt{\hat{P}(\omega_1)n}} \quad (2.22)$$

where the angular brackets denote an ensemble average, $\bar{\eta}^{(n)}$ is the averaged surface which would be obtained by bin-averaging in the absence of noise, $\hat{\sigma}_\eta$ is the standard deviation of the Gaussian error as estimated from (2.17), and $\hat{P}(\omega_1)n$ is the number of Gaussian distributed points estimated

from (2.15). Reduced estimate $\hat{P}(\omega_1)n$ is used in (2.22) rather than the full number of points n in order to disregard the non-Gaussian stereoscopic mismatches occurring with probability $(1 - \hat{P}(\omega_1))$.

An estimate of the attenuation error, on the other hand, can be obtained as follows. Let ν designate the density of data points per unit surface obtained after collapsing all data into one single dataset. Then the likely bin size resulting from assigning a constant number of points n to each bin is given by $\Delta x = \Delta y = \sqrt{n/\nu}$. Assume now that the free surface is roughly harmonic, with variations in the x and y directions characterised by wavelengths $\lambda_x \approx \lambda_y \approx [\lambda]$ and amplitude $a \approx [\eta']$, where $[\lambda]$ and $[\eta']$ designate characteristic horizontal and vertical scales of the free surface oscillations. The averaging process can be idealised as a convolution of the harmonic surface with a square top-hat filter of size $\Delta x \times \Delta y$. This results in an attenuated amplitude $\bar{a}^{(n)}$ given by [33]

$$\frac{\bar{a}^{(n)}}{a} = \text{sinc}\left(\pi \frac{\Delta x}{\lambda_x}\right) \text{sinc}\left(\pi \frac{\Delta y}{\lambda_y}\right) \approx \text{sinc}^2\left(\pi \frac{\sqrt{n/\nu}}{[\lambda]}\right) \quad (2.23)$$

where the sinc function is defined as $\text{sinc}(x) = \sin(x)/x$. Developing the sinc function by Taylor around the origin (up to second order), one can estimate the attenuation error due to the spatial averaging as

$$\left\langle \left(\bar{\eta}^{(n)} - \eta \right)^2 \right\rangle^{1/2} = \frac{1}{2} \left(\bar{a}^{(n)} - a \right) \approx \frac{\pi^2}{3} \frac{[\eta']}{\nu[\lambda]^2} n \quad (2.24)$$

where the attenuation error is seen to grow in proportion to the number of binned points n .

Finally, an approximate way of minimising the combined error is to set errors (2.22) and (2.24) equal to each other. For the conditions of the present experiments, this results in an optimal number of points $n \approx 50$, an average bin size $\Delta x = \Delta y \approx 45\text{mm}$ and a residual error on elevation $e_\eta \approx 1\text{mm}$. While not very precise, the procedure provides a valuable guide for the choice of bin size and an estimate of the resulting error. This is especially important for the present experiments because, ultimately, measurements from two highly heterogeneous methods (stereoscopic and velocimetric) are to be compared with each other. As we experienced in the first stages of the present work, the application of different and arbitrary standards in designing the binning procedures for the two methods can lead to over-damping of one set of results and induce significant discrepancies between the two.

2.5 Velocimetry

In contrast with the relatively direct character of the stereoscopic method, the velocimetric technique² is indirect in nature. The technique seeks to reconstruct the surface topography on the basis of measurements of the horizontal velocity field. For that purpose, the method exploits the Bernoulli relation applicable to free surface streamlines. This relation takes a particularly simple form if the flow can be approximated as a superposition of small perturbations upon a rapid mean current. To carry out the procedure, images from the top camera are first used to track particles on the water free surface and estimate their horizontal velocities. The elevation field is then obtained after suitable longitudinal and transverse averaging of the individual velocity vectors.

The following two assumptions are taken to apply locally (on the scale of a few wavelengths) to antidune flows: 1) the time evolution of the loose sediment bed is sufficiently slow that it appears stationary to the rapidly flowing fluid, hence the flow can be taken as quasi-steady; 2) the flow can be considered inviscid (but it does not have to be assumed irrotational). The dynamics of the free surface can then be described by requiring the conservation of mechanical energy along streamlines. Each surface streamline is subject to the following set of ordinary differential equations:

$$\frac{dx}{d\xi} = \frac{u}{\sqrt{u^2 + v^2}}, \quad \frac{dy}{d\xi} = \frac{v}{\sqrt{u^2 + v^2}}, \quad \frac{d\eta}{d\xi} = \frac{w}{\sqrt{u^2 + v^2}} \quad (2.25)$$

$$\frac{d}{d\xi} \left(\frac{1}{2} (u^2 + v^2 + w^2) + g\eta \right) = 0 \quad (2.26)$$

where ξ is a curvilinear coordinate measured along the streamline projection in the horizontal plane (x, y) , $\mathbf{u} = (u, v, w)$ is the particle velocity, η is the free surface elevation and g is the constant of gravity. Equations (2.25) describe the streamline as a parametric curve $(x(\xi), y(\xi), \eta(\xi))$, along which the Bernoulli equation (2.26) is written. Thus if the horizontal velocity field (u, v) of a steady free surface flow is known (and if dissipation can be neglected), then with suitable boundary conditions the set of ordinary differential equation can be solved for the free surface elevation η along each streamline. This constitutes the basis of the velocimetric technique proposed here for the measurement of free surface topography. On

²The velocimetric technique was developed by Hervé Capart and Dimitri Devriendt and is only briefly described here. More details can be found in [31] and [29].

a streamline-by-streamline basis, the above principle can be applied to a wide class of flows beyond the present antidune case.

Before proceeding, it is useful to reduce the system to a dimensionless form. Given a characteristic scale $[u]$ for the velocity, a characteristic wavelength can be defined as [20]

$$[\lambda] = \frac{[u]^2}{g} \quad (2.27)$$

Conversion to dimensionless variables can then be achieved by the transformation:

$$u \rightarrow [u]u, \quad v \rightarrow [v]v, \quad w \rightarrow [w]w, \quad (2.28)$$

$$x \rightarrow [\lambda]x, \quad y \rightarrow [\lambda]y, \quad \eta \rightarrow [\lambda]\eta, \quad \xi \rightarrow [\lambda]\xi, \quad (2.29)$$

resulting in the following set of dimensionless equations:

$$\frac{dx}{d\xi} = \frac{u}{\sqrt{u^2 + v^2}}, \quad \frac{dy}{d\xi} = \frac{v}{\sqrt{u^2 + v^2}}, \quad \frac{d\eta}{d\xi} = \frac{w}{\sqrt{u^2 + v^2}}, \quad (2.30)$$

$$\frac{d}{dx} \left\{ \frac{1}{2} (u^2 + v^2 + w^2) + g\eta \right\} = 0 \quad (2.31)$$

An especially simple situation arises if the free surface can be approximated as small amplitude oscillatory perturbations superposed upon a rapid mean flow. This is the case for the present antidune flows, for which one can further assume that: 1) the mean flow is oriented in the longitudinal x direction; 2) the mean flow is approximately uniform, with only a weak mean velocity gradient along the transverse y direction. In that case, it is possible to decompose velocities into mean and perturbation components and introduce a small parameter ϵ such that

$$u(x, y) = u_0 + u' = u_0(y) + \epsilon u_1(x, y), \quad v(x, y) = v' = \epsilon v_1(x, y) \quad (2.32)$$

where u_0 is the mean current velocity, and $\epsilon = [u']/[u]$ measures the strength of the perturbation velocities u' , v' (having characteristic scale $[u']$) relative to the mean current velocity u (having characteristic scale $[u]$). One can then substitute expansions

$$\eta = \eta_0 + \epsilon \eta_1 + O(\epsilon^2) \quad (2.33)$$

$$\frac{d}{d\xi} = \left(\frac{d}{d\xi} \right)_0 + \epsilon \left(\frac{d}{d\xi} \right)_1 + O(\epsilon^2) \quad (2.34)$$

in (2.30) and (2.31) to obtain, at order 0

$$\eta = \eta_0 \quad (2.35)$$

and at order 1,

$$\eta_1 = -u_0 u_1 \quad (2.36)$$

where it was assumed that transverse variations in mean velocity are at most $O(\epsilon)$. One can finally translate back to dimensional variables and obtain the key result

$$\eta = \eta_0 + \eta' = \eta_0 - \frac{u_0 u'}{g} + \frac{[u]^2}{g} O(\epsilon^2) \quad (2.37)$$

which is particularly simple both to use and to interpret: superimposed upon a mean surface level η_0 , the free surface topography perturbations η' are proportional to the perturbations in longitudinal velocity u' , but opposite in phase. For torrential flows, the water velocity is slowest at the crests and fastest at the troughs, an observation pointed out already by Leonardo Da Vinci in his notebooks several centuries ago [44]. In practice, the ratio of fluctuation velocities to mean velocities $\epsilon = [u']/[u]$ obtained in the present experiments is about 0.1, with mean velocities of about 1 m/s. This leads to errors on the perturbed elevation η' of the order of $(1\text{m/s})^2(0.1)^2/(10\text{m/s}^2) \approx 1\text{mm}$. This lies in the same range as errors arising from other sources, hence there is little point in resorting to more refined approximations of (2.30)-(2.31). Whatever the order of approximation, an obvious limitation of the velocimetric method is that only the topography perturbations are accessible, not the mean elevation. If needed, this has to be obtained by other means. The details of the implementation of the velocimetric method can be found in [45], [31], [32] and [29].

2.6 Results and discussion

Examples of free surface topography measurements obtained using both the stereoscopic and velocimetric methods are presented on fig. 2.7 through 2.9. Taken from a series of 10 runs, the three experimental runs shown (runs 3, 5, and 6) exhibit a range of different patterns: wide-crested rolls for run 3 (fig. 2.7), a zig-zag pattern of staggered peaks and troughs for run 5 (fig. 2.8), and a narrow wave train with sharply-peaked crests on one side of the channel for run 6 (fig. 2.9). The gray-coded surface elevation maps provide full-field pictures of these different patterns. Since the

velocimetric method cannot determine the mean water level η_0 , the latter has been subtracted from the stereo results to obtain for both methods the perturbed topography $\eta'(x, y) = \eta(x, y) - \eta_0$. These maps allow qualitative comparisons between the two methods. Some superficial differences between the stereoscopic and velocimetric results can be noted. First, due to the different binning procedures used, the spatial coverage of the two methods do not fully coincide. Because the corresponding binning assigns an elevation even to areas where few particles were identified, the stereo results extend over the whole domain. The velocimetric results, on the other hand, cover only the restricted domain where particle trajectories were retrieved. Secondly, the surfaces exhibit distinct textures that reflect the peculiarities of each of the two methods. The maximum-likelihood fitting used in the stereo method leads to staircase-like surfaces because the density of the cloud of three-dimensional points can vary abruptly in addition to the matching which is error prone. On the other hand, the velocimetric maps are smoother, but exhibit longitudinal stripes due to the streamlines along which velocities were measured and averaged. These are fine-grained effects, however, which do not affect the overall surface elevation maps.

Qualitatively, the agreement between the two methods is seen to be quite good for all three of the patterns examined. The magnitudes, locations and arrangements of the peaks and troughs are in good correspondence. Local regions of discrepancy are nonetheless present, for instance the spurious crater-like feature exhibited by the stereo results of fig. 2.9a near position $(x, y) = (150, 400)$, but absent from the velocimetric results of fig. 2.9b. This is a region of wave-breaking where the stereo method appears to have broken down, likely due to an insufficient number of successfully matched stereo pairs in that zone. To facilitate precise comparisons, the same results are plotted again in a different format on fig. 2.10 through 2.12. Here profiles $\eta'(x)$ reconstructed by the two methods are plotted together for various positions y across the width of the channel. The stereo profiles (thin lines) and velocimetric profiles (thick lines) are found to be in reasonably close qualitative and quantitative agreement. Results in the central region of the viewing volume are generally better than the results near the edges. Especially poor results are recorded for the stereo profile of fig. 2.12a, which passes through the "crater" zone of fig. 2.9a. For this run (run 6), the stereo cameras were placed lower with respect to the surface than for runs 3 and 5, leading to near-occlusion in the troughs in addition to the difficulties associated with wave-breaking

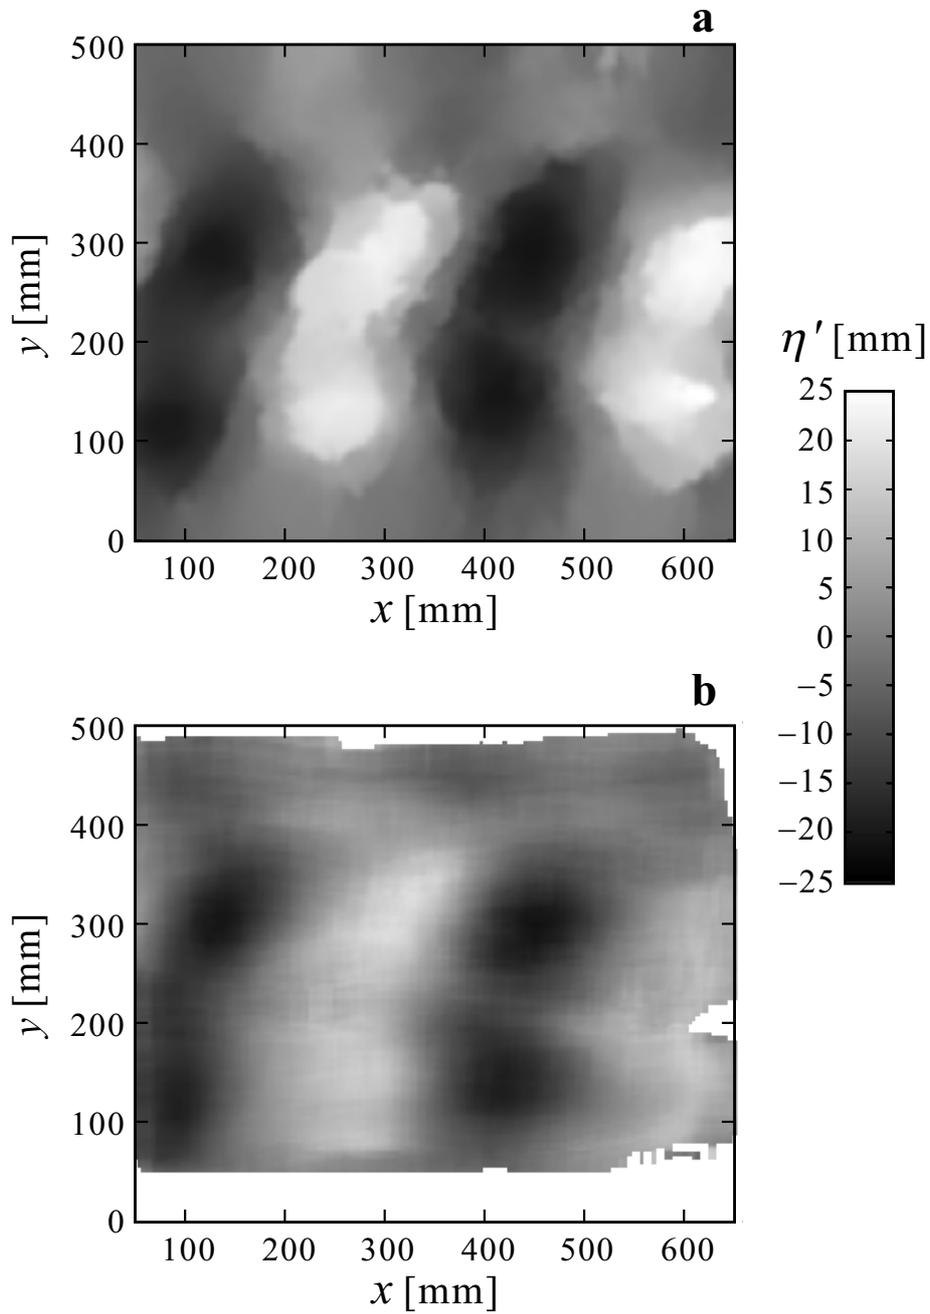


Figure 2.7: Reconstructed free surface topography $\eta'(x,y)$ for run 3 (broad crested rolls): (a) stereoscopic method; (b) velocimetric method

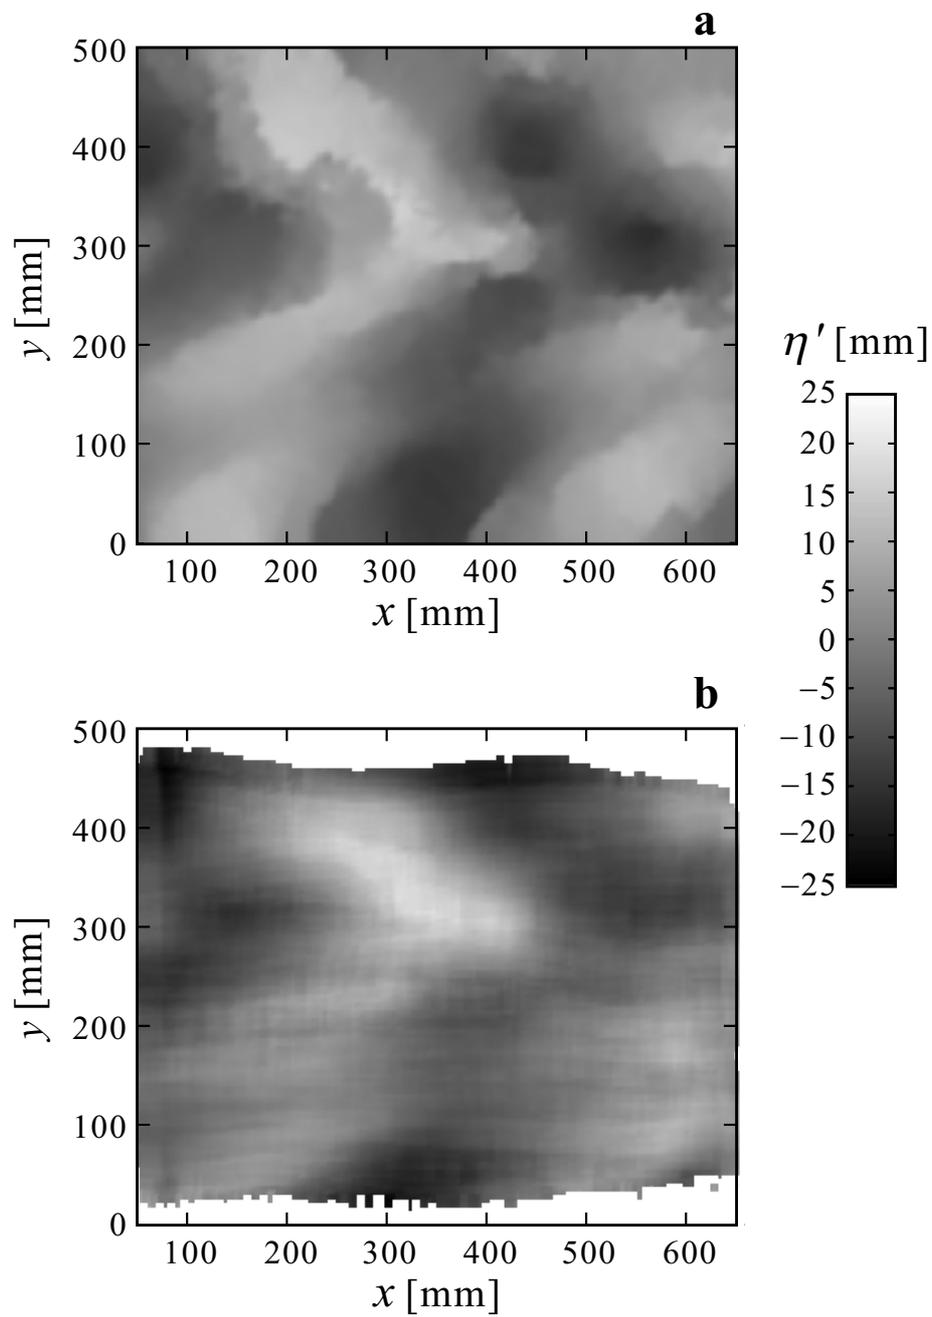


Figure 2.8: Reconstructed free surface topography $\eta'(x, y)$ for run 5 (zig-zag pattern): (a) stereoscopic method; (b) velocimetric method

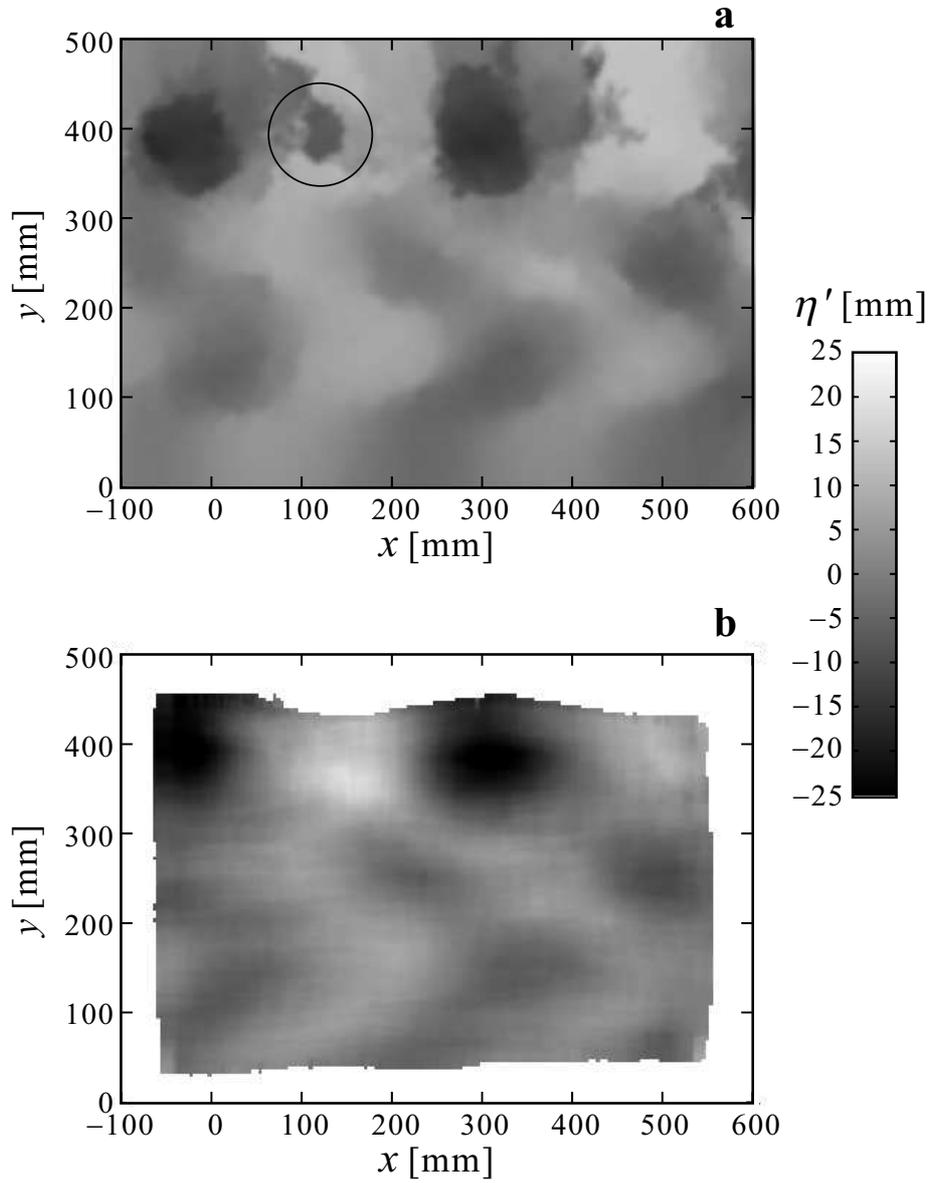


Figure 2.9: Reconstructed free surface topography $\eta'(x, y)$ for run 6 (narrow wave train on one side of the channel): (a) stereoscopic method; (b) velocimetric method. Circled: crater-like region where stereo measurements break down

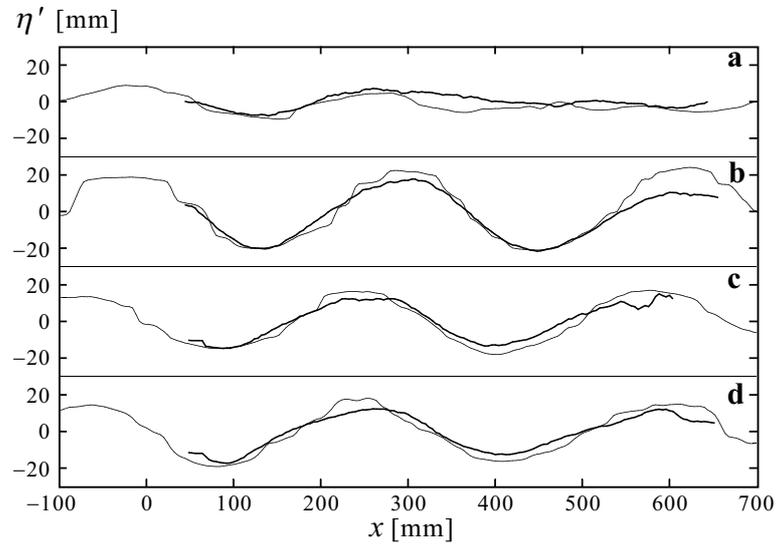


Figure 2.10: Comparison of free surface profiles for run 3 obtained by the stereoscopic (thin lines) and velocimetric (thick lines) methods at selected sections: (a) $y = 400$ mm; (b) $y = 300$ mm; (c) $y = 200$ mm; (d) $y = 100$ mm

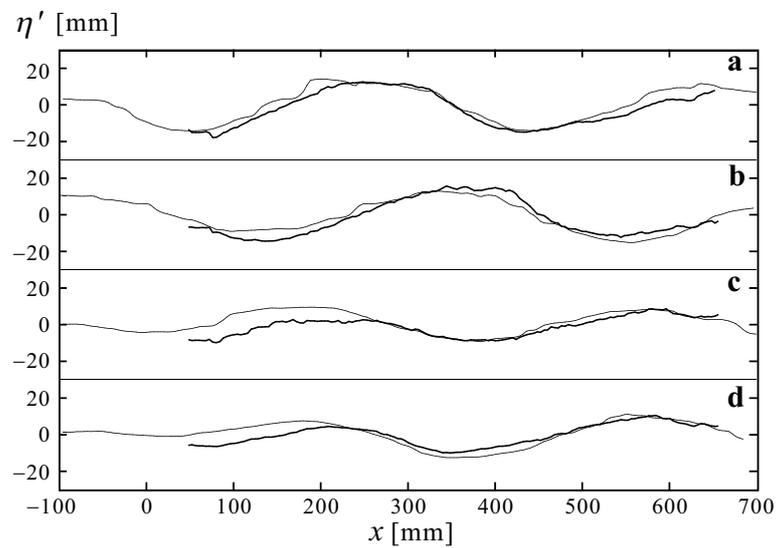


Figure 2.11: Comparison of free surface profiles for run 5 obtained by the stereoscopic (thin lines) and velocimetric (thick lines) methods at selected sections: (a) $y = 400$ mm; (b) $y = 300$ mm; (c) $y = 200$ mm; (d) $y = 100$ mm

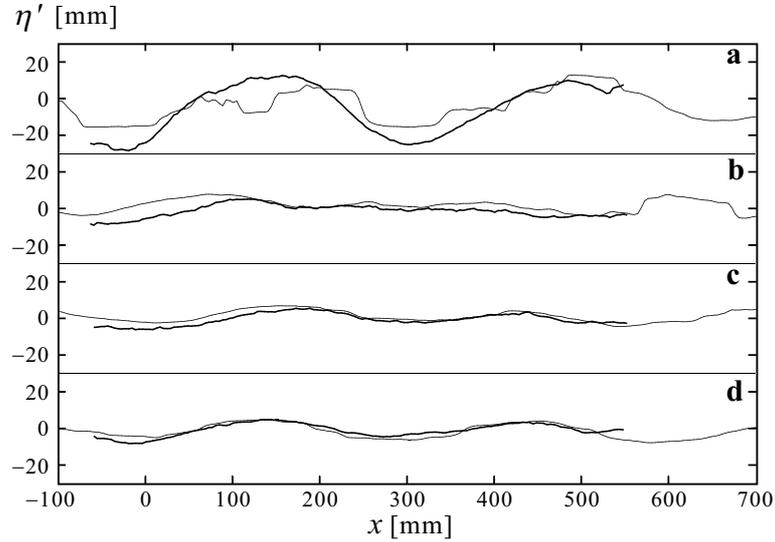


Figure 2.12: Comparison of free surface profiles for run 6 obtained by the stereoscopic (thin lines) and velocimetric (thick lines) methods at selected sections: (a) $y = 400$ mm; (b) $y = 300$ mm; (c) $y = 200$ mm; (d) $y = 100$ mm

at the peaks.

Computation of the root-mean-squared discrepancy $\{(\eta'_{stereo} - \eta'_{vel})^2\}^{1/2}$ between the results of the two methods (averaged over a square domain of $500\text{mm} \times 500\text{mm}$) for the three runs 3, 5, and 6 yields Root Mean Square (RMS) errors of 3.8, 5.1, and 4.7 mm, respectively. Relative to range $\eta'_{max} - \eta'_{min} \approx 40\text{mm}$, this amounts to relative errors of the order of 10 to 15%. These discrepancies appear consistent with the error estimates derived separately for the two techniques in sections 4 and 5. Discrepancies turn out to be roughly twice the predicted error values (of the order of 1-2 mm), an outcome which may be ascribed to three factors: (i) the approximate nature of the error analysis procedures; (ii) model errors (for instance deviations of the image formation process from a perfect projective transformation, or departures of actual streamline dynamics from the Bernoulli equation possibly due to the slightly unsteady character of the actual flows); (iii) the presence of regions where one of the two methods performs particularly poorly, as for instance in the "crater zone" of fig. 2.9a mentioned previously.

However, the quality of the comparison is encouraging. Errors are not negligible, but remain within reasonable bounds considering that

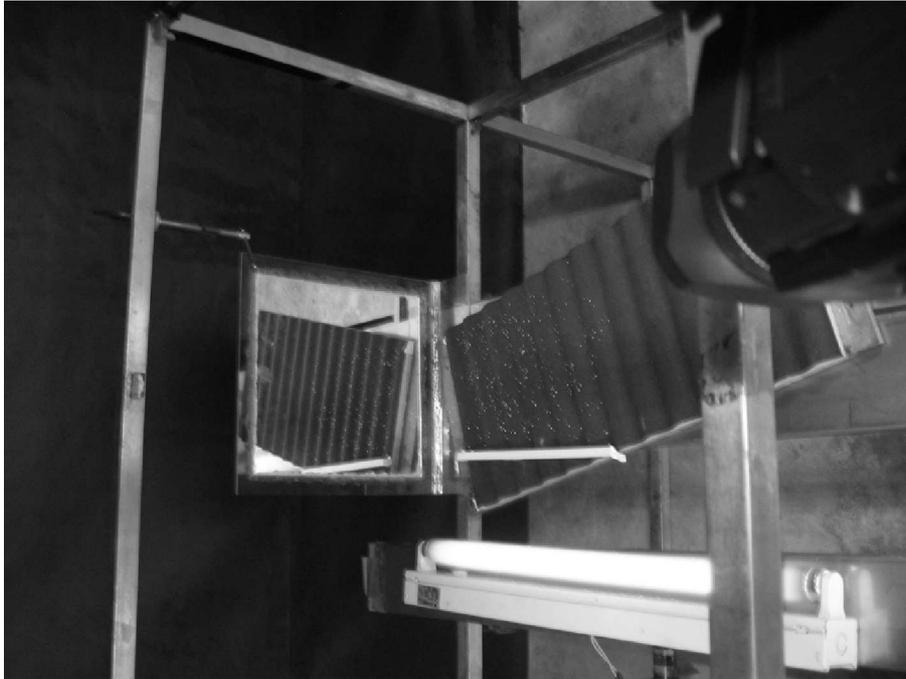
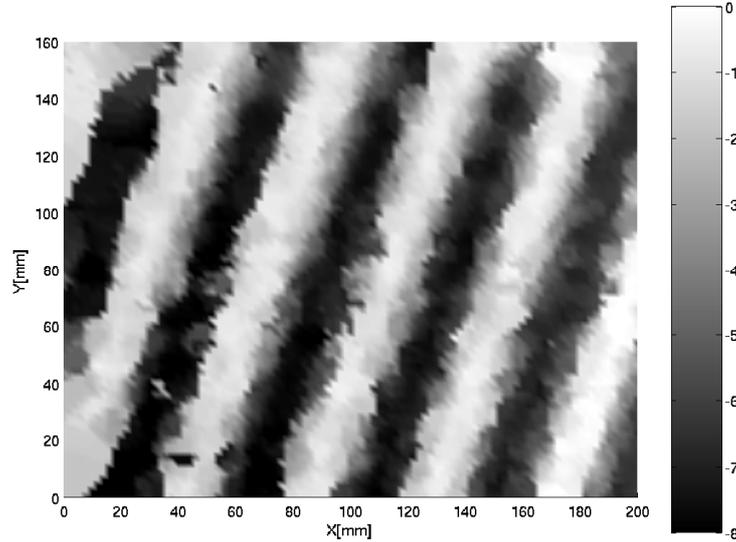


Figure 2.13: *Imaging setup for the verification of the error on $\eta'(x, y)$*

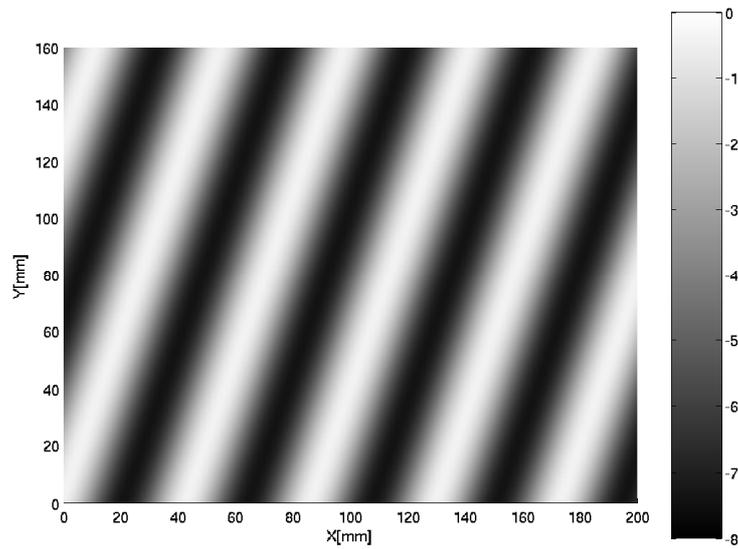
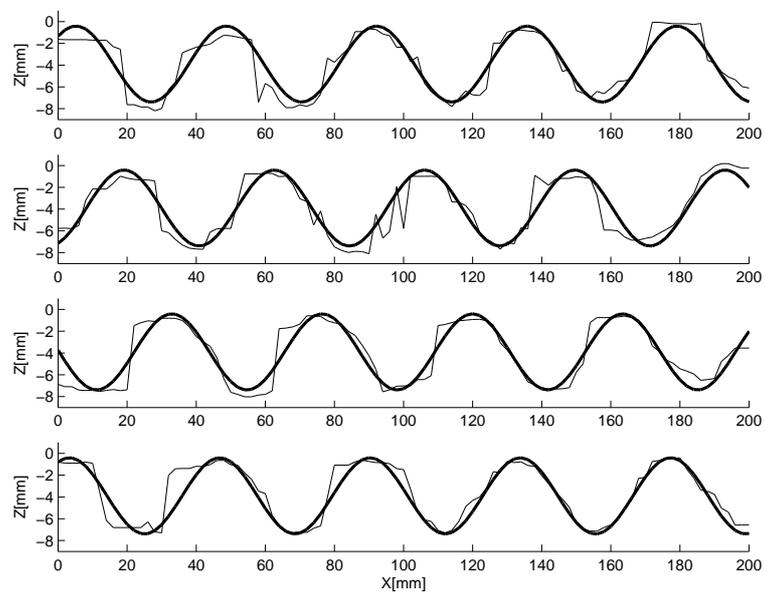
both methods yield whole-field measurements rather than point values. Most importantly, the spatial patterns are well-captured by both methods, vividly depicting a variety of motifs including the broad crested rolls, zigzags, and narrow train of peaks and troughs. Due to the complex way in which light interacts with the rough water surface, these organised patterns could not be so easily grasped by pure visual inspection of the laboratory flows.

2.7 Technique validation

The technique detailed above can be validated by placing particles on a known surface and verifying that the error estimation depicted in Section 2.4.4 corresponds to the true error. To that effect, a surface of a known unidimensional sinusoidal elevation is placed in the field of view of a camera/mirror combination that provides synchronised stereo images (fig. 2.13). The surface is fitted with particles which are lit with black light to enhance the contrast in the images. As for the antidune experi-

Figure 2.14: *Reconstructed surface*

ments, several images of different particle arrangements are taken to obtain a suitable particle density. The surface reconstruction can not be performed with the velocimetric technique as it requires motion while this setup is static and therefore only the stereo algorithms are applied to the raw camera images. The resulting surface is presented on fig. 2.14 and the ground truth surface on fig. 2.15. A comparison of the two surfaces follows on fig. 2.16, where it can be seen that the obtained surface has the same stepping artifacts as the stereo-reconstructed antidunes. The fact that the sequences were synchronised in this case does not seem to improve the results compared to the real tests which bring us to conclude that major issue for the stereo technique is the matching and that the post-synchronisation of the sequences presented in section 2.4.2 is adequate. The estimated error for $\hat{\eta}$ is obtained through the use of (2.22) and (2.24) with the following parameters proper to this experiment: $\nu = 0.15/mm^2$, $[\eta'] = 3.5mm$, $\hat{\sigma}_\eta = 0.91mm$, $\lambda = 41.1mm$ and $\hat{P}(\omega_1) = 0.93$. The estimated number of points n returned is 6.3, which corresponds to a bin size $\Delta x \times \Delta y = 8 \times 8mm$. The RMS error observed is 1.12mm or 16% of the surface peak-to-peak amplitude which, given the rough character of the error estimation, is very close to the 15% estimated for the measurements of antidune flows.

Figure 2.15: *Ground truth surface*Figure 2.16: *Four slices across the surface in $Y=20\text{mm}$, $Y=60\text{mm}$, $Y=100\text{mm}$ and $Y=140\text{mm}$*

2.8 Conclusions

In this chapter, two approaches to the measurement of free surface topography in high Froude number flows were presented and compared. Both rely on the imaging of floating tracers dispersed on the surface, but they resort to different reconstruction principles: on the one hand the stereoscopic matching of rays associated with two camera viewpoints, and on the other hand the indirect estimation of the elevations based on the horizontal velocity field measured using a single camera. The results from the two techniques were compared for flows over antidunes featuring various surface patterns. Discrepancies were found to be limited and generally consistent with error estimates. The obtained relief maps vividly depict the variety of motifs that can evolve as a result of interaction between shallow flows and loose sediment beds.

Drawbacks, advantages and possible improvements can be identified for both techniques. Stereoscopic matching and velocimetric tracking both involve pairing particles on distinct images, but the stereo methods were found to generate a much greater proportion of mismatches (up to 40% of the data points). This stems partly from the fact that the stereo technique pairs isolated particles whereas the velocimetric technique tracks patterns of neighbouring particles. The higher vulnerability of the stereo technique to mismatches further leads to greater accuracy requirements for the camera calibration procedure as well as tighter limits on the density of particles which can be reliably imaged. Camera synchronisation is also a paramount concern, best addressed by synchronising two sensors at the time of image acquisition (rather than a posteriori as in the present work).

A key advantage of the stereo technique is that, in principle, it applies to general flows independently of any specific assumption (occlusion is the only limitation). In practice, the stereo technique requires the flow to vary slowly in order to combine data from multiple image pairs to obtain a surface of sufficient resolution. The velocimetric technique, by contrast, fundamentally depends on certain physical assumptions. The Bernoulli equation requires the flow to be quasi-steady, and its simplified version further assumes moderate perturbations superimposed onto a quasi-uniform mean current. It also requires the mean surface level to be known by other means. When these conditions are met, however, the present results demonstrate that the velocity-based approach works well, and produces reasonable topography measurements. Camera sensors having a better resolution than the relatively limited equipment used

in the present work would also lead to improvements in the accuracy of both approaches.

From a methodological point of view, the present work underscores the importance of characterising the error formation process. Many significant experimental errors cannot be treated as Gaussian random variables of zero mean that will simply average out if measurements can be repeated a sufficient number of times. In the present experiments, non-Gaussian errors which had to be controlled include: 1) stereo mismatches distributed evenly within the viewing volume as a result of epipolar ambiguity; 2) velocity errors due to position noise which cancel each other along particle trajectories; 3) attenuation associated with binning, averaging and filtering; 4) geometric distortions associated with projections on non-planar surfaces. In all these cases, the special character of the errors had to be taken into account in order for their influence on the results to be well balanced. The objective of ultimately comparing two independent sets of measurements (stereo and velocity-based) served as a powerful motivation in striving to get a grip on the errors inherent to both methods.

Provided errors can be controlled, it was found possible to use both stereoscopic and velocimetric techniques to capture the free surface topography of antidune flows, with results that compare favourably with each other. Comparisons with theoretical and computational descriptions of antidune flows will be addressed in future work. Future prospects include the use of synchronised high-resolution cameras and smaller particles in order to obtain more accurate, denser information. Also, further applications of the techniques are contemplated, including measurements of evolving antidune fields at larger scales as well as other free surface flows.

Chapter 3

Dense measurements of wind velocities

3.1 Introduction

Wake vortices are rotating air masses that are generated by aircraft as a result of their bearing pressure. These air flows will induce a rolling moment to a following aircraft entering them, creating a potential hazard especially during take-off and landing phases when the speed induced by a wake vortex is more important than in cruise and the aircraft might not have the resources to recover from such encounter. The investigations on wake vortices started in the 1940s with the work of Russian scientist Belotserkovsky [46] who made it a personal challenge to better understand this phenomenon after the death of Yuri Gagarin in a wake vortex encounter in 1968. The dynamic behaviour of wakes is now well understood and experiments have confirmed the theoretical advances. However, their practical impacts are few and vortices are still a hazard limiting the separation distance of landings and the period of takeoffs. These restrictions are the main factors fixing the boundaries for the throughput of today's airports [47] so that a system allowing a smaller period with the same security would yield more benefits for airports, as well as reduce the need for unpopular new installations. The period between movements is currently fixed by the International Civil Aviation Organisation (ICAO) regulations and depends on the leading and following aircraft weights (implementation of these rules by the Federal Aviation Administration (FAA) can be found in [48] [49]). These regulations should ideally be applied at all times but as they do not take all factors into account and their strict ap-

plication would lead either to a lower movement frequency or to a lower security. Practically, observations shows that these rules are adapted by Air Traffic Control (ATC) personnel to reach an optimum throughput. For example a crosswind moves the vortex away from the trajectory of the following aircraft, leading to a reduced minimal distance between planes. This adaptation is done with the expert judgement of the ATC controllers but it is nonetheless non-optimal and stressful for the operators and pilots. A more systematic approach based on in-situ measurements is therefore sought. In the recent years, ground Doppler Light Detection and Ranging (lidar) systems [50] have been deployed on large airports to detect vortices near the runway threshold [51]. These systems have a very good detection capability because they operate in the transverse plane of the vortex where the wind speed is the highest. However, they operate locally and do not report on the full zone that the aircraft will fly in. To that effect, on-board systems are contemplated to directly warn pilots [52] and serve as one of the inputs to airport-wide wake management.

The on-board integration of a lidar system has been the subject of previous studies, from simple pulsed lidars in [53] or [54] to the more complete system of the NASA AVOSS project [55] [56]. Other investigations include in-situ measurements implying a wake encounter [57] [58]. The European Commission funded IWAKE project which will be described here is however the first to test an operational on-board lidar-based detection of wakes.

This chapter will be organised around the data collected during two test campaigns. The first was a series of preliminary ground tests performed by the MFLAME project in 2000 [59]. The primary goal of this campaign was to study the feasibility of an axial vortex detection. This chapter uses the MFLAME data set for the purpose of testing new algorithms for wake detection and can be considered the following of the work of Keane. A major contribution of this chapter is to introduce and demonstrate the usefulness of a new two-primitive spectral model for IR lidar signals. The data being collected near the ground where particle density is high, the signal-to-noise ratio (SNR) is high and sharp vortex signatures can be extracted.

The second part of this chapter relies on the data collected in June 2004 during the flight tests of the IWAKE project, whose goal is to move the hardware from the ground shelter of MFLAME into an aircraft. Although much of the processing is similar, special care must be taken for three reasons: (i) the aircraft is a moving platform and its position, speed and atti-

tude must be taken into account, (ii) the return signal is much noisier as we are working at a higher altitude of 7000 feet where the particle density is lower, and (iii) we have to adapt the algorithms to the inevitable changes in hardware that occurred since the beginning of MFLAME. The association of the new spectral model with the flight tests of the IWAKE project provides the first proof of the on-board ability to detect wake vortices with an IR lidar.

3.2 Wake vortex detection with a Doppler lidar

The vortex being moving air masses, the basic principle of operation of the vortex detector is the Doppler effect. As the wind speeds inside the vortex are different from their more static environment, a Doppler lidar will be able to differentiate the vortex from the rest of the atmosphere. The Doppler effect in itself describes a change in frequency Δs of a received wave with the relative speed v of the transceiver and the local air mass:

$$\Delta s = \frac{-2v}{\lambda} \quad (3.1)$$

where λ is the wavelength of the emitted signal. Note that this relation is only valid along the axis of propagation joining the receiver and emitter and it will therefore be only possible to measure the component of the air speed vector that is aligned with the line of sight.

3.2.1 Laser wavelength considerations

The wavelength of the emitted light is constrained twofolds: by the size of the particles we wish to observe and by the absorption spectrum of the atmosphere. Regarding particle sizes, the wavelength must be equal or smaller. Since we are working in a clear atmosphere we do not have large particles (like water droplets) so that RF radar is excluded. The particles we can count on are either the atmospheric dust or the air molecules themselves. Atmospheric dust has a particle size distribution that is highly variable, but is mainly composed of particles in the micrometer range, which leads to the use of infrared (IR) waves for their detection ($\lambda=1\rightarrow 10\mu\text{m}$). This approach works well in the lower atmosphere where dust is abundant, but can not be used in the higher atmosphere where large particle are scarce. In that case air molecules are used as tracers and the wavelength required moves to the ultra-violet spectrum, with a wavelength of 300nm

to a few nanometers. One might be tempted to work with the UV solution in the lower atmosphere, but unfortunately the higher water and ozone concentration is attenuating the UV waves much faster than IR waves. For the detection of wake vortices in a landing phase the IR waves is thus the best choice.

Power lasers operating in the IR spectrum do not offer a wide choice of wavelength that are eye-safe and compatible with the absorption spectrum of the lower atmosphere. For the experiments described here a $2.022\mu\text{m}$ diode-pumped Tm:LuAG laser was used as this device is readily available on the market.

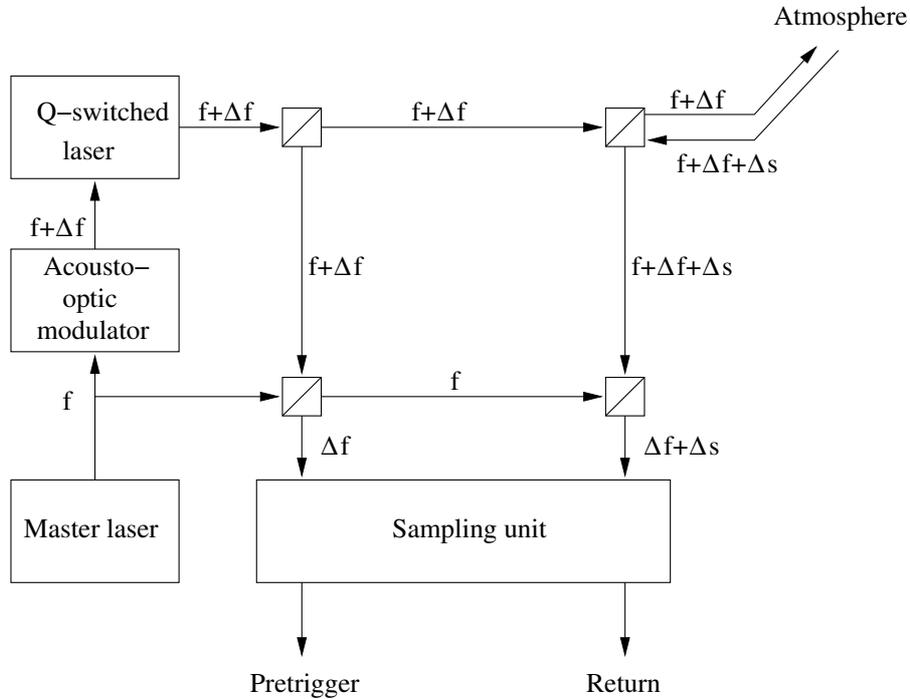
At the wavelength of $2\mu\text{m}$ the vortex airspeeds that we wish to measure are in the order of a few meters per second. This leads to very small relative frequency changes that can not be measured accurately as is because they are several orders of magnitude below the operating frequency of 148THz: given the laser wavelength of $2\mu\text{m}$, a 1m/s wind will be indeed translated to a Doppler shift of 1MHz:

$$\Delta s = \frac{2 \times 1\text{m/s}}{2\mu\text{m}} = 1\text{MHz} \quad (3.2)$$

The signal received is therefore heterodyned at the detector surface with the original laser output so that the signal to be detected will have a mean frequency that is the difference of the two signals. As the wind speed Doppler offset Δs might be positive or negative the resulting heterodyned signal might have a negative frequency which will be observed as a mirrored spectrum. This annoying phenomenon is avoided by shifting the wavelength sent in the atmosphere by a safe amount Δf using an acousto-optic modulator (fig. 3.1). The frequency shift also allows to work at frequencies that are more adequate for electronics than pure baseband. For ground tests where the wind speeds are less than 10m/s a frequency $\Delta f=30\text{MHz}$ was used. In the case of on-board detection the speed of the aircraft (80m/s) must also be taken into account as it will yield a frequency offset of 160MHz . The total offset to be used is the sum of the expected speeds of the aircraft and the wind, or $\Delta f= 200\text{MHz}$.

3.2.2 Pulsed lidar operation

A lidar can be designed for one of two modes of operation: continuous wave (CW) or pulsed. The CW signal can not provide range information: it is focused on one point and all measurements will be performed for that location only. Moreover, an on-board lidar will sometimes aim at the

Figure 3.1: *Optical diagram*

ground and the resulting echo from hard targets will mask all other returns from the atmosphere. Pulsed lidar, on the other hand, is focused to infinity. The time between the pulse emission and the measured return provides the range information. The spatial resolution Δd of the lidar is then limited by the pulse duration t_p :

$$\Delta d = \frac{c}{t_p} \quad (3.3)$$

where c is the speed of light. A pulsed lidar operation is mandatory for the axial wake vortex detection because we wish to find where the vortex is in the atmosphere *volume* in front of the aircraft. As the master laser frequency can slightly vary from shot to shot, it is necessary to evaluate its accurately in order to obtain good wind speed measurements. This measure is performed by the second pretrigger sampling channel shown on fig. 3.1.

3.2.3 Geometric configuration of the ground tests

Contrary to tangential vortex detection which only needs a single angle scan, the axial detection of wakes requires the laser to be swept across the volume containing the vortices. This is achieved by a two-dimensional scan depicted on fig. 3.2 and a proper range gating. The sinusoidal pattern is created by two mirrors (one for the vertical deflection and one for the horizontal deflection) that operate one full sinusoidal scan every 6 seconds, with a total vertical opening of 3° and a horizontal opening of 12° . The mean aim of the scanner is 7° above the horizon. A laser pulse is sent every 2.5ms so that each complete scan will be composed of around 2400 laser shots. As each pulse travels at the speed of light, it is possible to obtain an information about the atmosphere at a specific distance from the lidar by looking at the signal return with a certain delay. For example, the signal received $2\mu\text{s}$ after the pulse has been sent correspond to a distance from the lidar of approximately $3 \cdot 10^8 \times 2 \cdot 10^{-6} / 2 = 300\text{m}$. The return signal can therefore be time-gated to obtain a analysis volume between 800m and 2375m and then digitised at a sampling rate $T_s=128\text{MHz}$. Each resulting time series $S^L(k)$ of shot L is then split in gates of 64 samples, which is roughly equivalent to a spatial range of 75m. To obtain better estimates, the chunks of the previous and next gates are concatenated to the current gate data, which effectively yields 19 range gates of 225m (192 samples), each separated by a distance of 75m. The data available for each range gate g is thus defined as:

$$\begin{cases} S_g^L(k) = S^L(64(g-2) \dots 64(g+1)) \\ g = 2 \dots 19, L = 1 \dots 2400 \end{cases} \quad (3.4)$$

Of the numerous scans measured by the MFLAME project, four will be presented here. These four tests are the same as the ones used by Keane [59] and correspond to the landing of an A340-600, an A340 Beluga Super-Transporter, an A320 and a Fokker 100.

3.3 Classic spectral estimation

As we are looking to estimate the Doppler shift of the return signal, it is natural to base our analysis on the signal power spectrum [60] which is obtained using an Fast Fourier Transform (FFT) on the signal $S_g^L(i)$:

$$P_g^L(k) = \|\text{FFT}(S_g^L(k))\|^2 \quad (3.5)$$

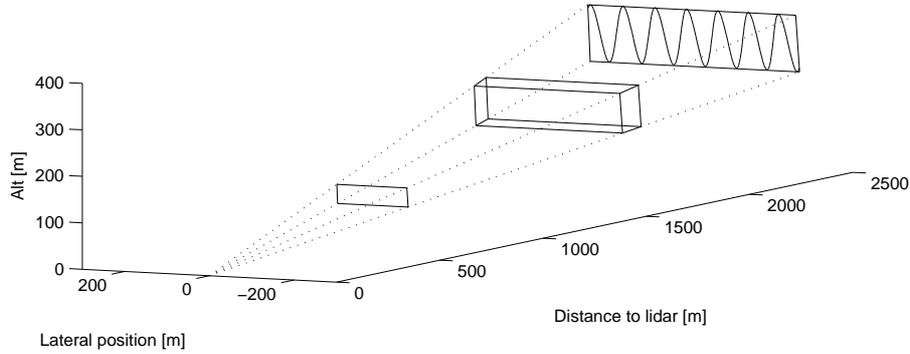


Figure 3.2: *The MFLAME experiments configuration. The first gate is at a distance of 800m, the last one at 2375m. The sinusoidal scanning pattern is shown on the last gate. A typical volume of analysis for a range gate is presented at a distance of 1500m.*

This approach used in the MFLAME project [59] is described below. Another less intuitive possibility is to work on the signal correlation which will be performed directly on the signal $S_g^L(k)$. Among the time-domain techniques available, the Auto-Regressive-Moving-Average (ARMA) has been tested on the MFLAME data.

In [59], Keane not only uses the Doppler shift but also the spectral width of the signal to detect the vortices. Analysing the power spectrum $P_g^L(k)$, the mean returned frequency will reflect the Doppler effect and therefore can be used to provide the mean wind speed in the direction of the line of sight. As every wind speed will yield a different Doppler shift, the variance of the signal is an indication on the broadness of the speed distribution in a certain volume (for a more complete discussion on the broadening effects see [61]). Since vortex are not laminar flows we can expect to detect them not only with the mean speed variations but also by looking at a larger spectrum variances. The estimation of the two spectral moments is the classic procedure for vortex detection: vortex are expected to be turbulent zones (wider peak) with a slightly different airspeed in the core (Doppler shift).

Formally, the signal spectrum mean f and variance σ correspond to the spectral moments of order zero and one. The estimation of these moments is well covered by the literature. The most complete description of the spectral estimation applied to wind measurements is provided by

Doviak and Zrnic [62]. Zrnic also proposes a study of range weighting [63] and a discussion of the particular case of pulse pairs [64], among others. Frehlich details a method for estimating the wind velocities statistics in a stationary atmosphere [65] which essentially uses the same spectral estimation techniques. Dias et al. also discuss some improvements on the spectral estimation of the two first moments [66] [67] [68] in the case of a Gaussian signal in a Gaussian noise. The broad coverage of the subject will allow us to concentrate on the practical aspects of spectral estimation for vortex detection, taking the work of Keane as a starting point.

3.3.1 Spectrum based

The spectrums obtained from (3.5) are subject to local fading and are also very noisy, as can be seen on fig. 3.4. Two architectures can be used in order to obtain better estimates for the spectral moments. The first one is to estimate the set of spectral moments f_g^L and σ_g^L of each individual power spectrum $P_g^L(k)$. At this point the spatial location of the estimates follow a sinusoidal pattern since the scanning is sinusoidal and each spectrum refers to a different Line Of Sight (LOS). In order to obtain regular images suitable for image processing, the spectral moments are resampled on a regular lattice with a radius r that is large enough to filter out the noise in the estimates (fig. 3.3). The resampled spectral moments are then

$$f_g(i, j) = \frac{1}{n} \sum_{L \in r} f_g^L \quad (3.6)$$

$$\sigma_g(i, j) = \frac{1}{n} \sum_{L \in r} \sigma_g^L \quad (3.7)$$

where n is the number of evaluations of the spectral moment that fall within the sampling radius r . This technique is simple and gives good results at reasonable processing cost. However, it has the inconvenience that the noise in the evaluation of the parameters is kept until the last stage of the process. The difficulty to infer stable noiseless spectral moments from a set of unstable noisy spectrums greatly limits the SNR (and thus also the range) at which the system can operate.

The second method is to average before the estimation process all the spectrum which fall within a radius around the lattice point (see for example [69]). One can not simply average the spectrums by summing them point to point because each spectrum is shifted by a different amount by the Doppler effect. The addition would therefore result in a broader peak

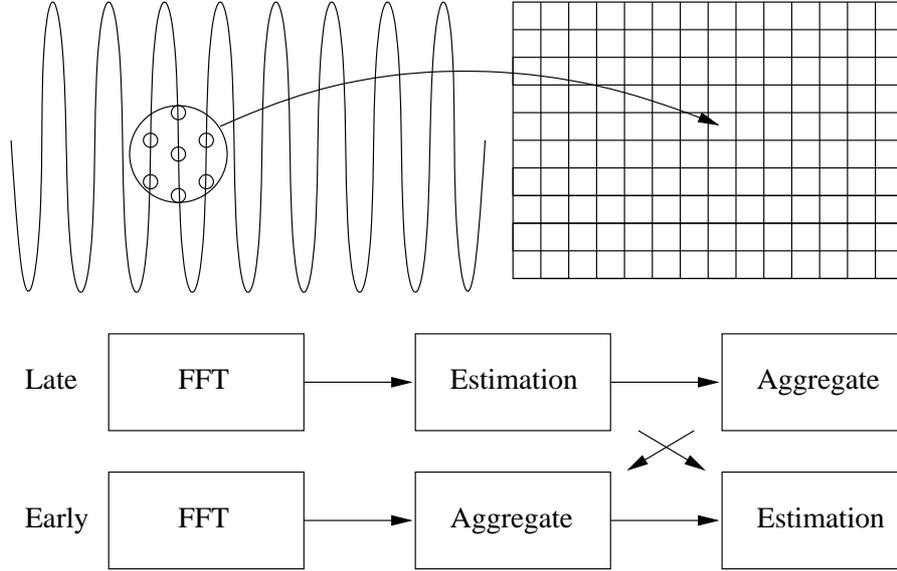


Figure 3.3: Early or late LOS grouping

than what is really observed. To avoid this broadening the spectrums within the sampling radius are shifted to the median position $k_{g,med}(i, j)$ of their peak frequencies k_g^L before averaging them. To further improve the final spectrum the average is weighted with a basic estimation W_g^L of the signal SNR:

$$W_g^L = \frac{P_{g,M}^L - P_{g,m}^L}{P_{g,M}^L} \quad (3.8)$$

where $P_{g,M}^L$ is the power spectrum peak value and $P_{g,m}^L$ is its average value. The resampled spectrum for image point (i, j) is thus expressed as:

$$P_g(i, j, k) = \frac{\sum_{L \in r} W_g^L P_g^L(k - k_g^L + k_{g,med}(i, j))}{\sum_{L \in r} W_g^L} \quad (3.9)$$

A resulting spectrum is presented on fig. 3.5 and it is clear that averaging the spectrums provides a better basis for the estimation of the spectrum moments. This second architecture is however more complex because the filtering is based on merging spectrum together while the first technique was merging the scalar spectral moments. We will show below that this second approach provides very stable spectrum which allows it to be fol-

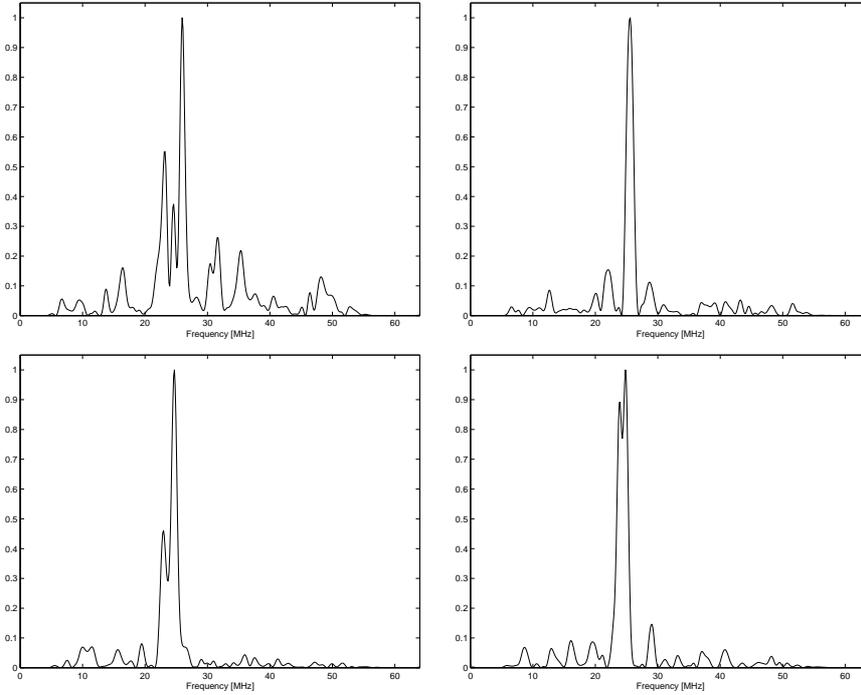


Figure 3.4: *Individual spectrums from a close range gate*

lowed by a more accurate model matching process, which will in turn yield more information about the atmosphere.

With the stable spectrum obtained above it is interesting to extract its noise component before estimating the spectral moments. The power of the reference signal at the detector can be large enough to the point where most noise will come from the shot noise of the reference signal on the detector [70]. This noise has therefore a constant spectrum which can be estimated only once by acquiring data with a direct laser feedback (see fig. 3.1). An example of spectrum is provided on fig. 3.6. This noise spectrum can be removed from the observed spectrum to enhance the extraction of the spectral parameters. It will also be part of the spectrum model detailed in Section 3.4.

The estimation of the mean velocity from the power spectrum is extensively described in [71], [72] and [73]. Techniques that directly work on the signal samples through the use of correlation (like [74]) will be detailed later. The mean velocity of the air mass pierced by the lidar will

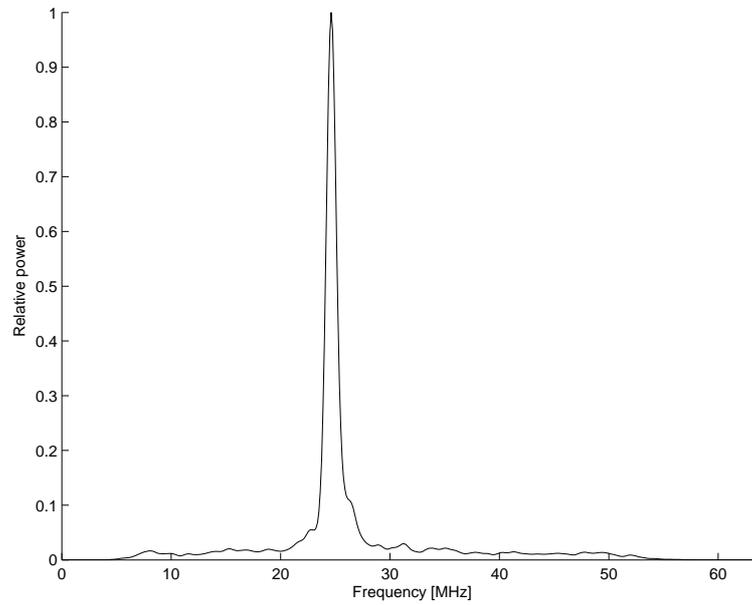


Figure 3.5: Typical signal spectrum obtained with a weighted spectral averaging.

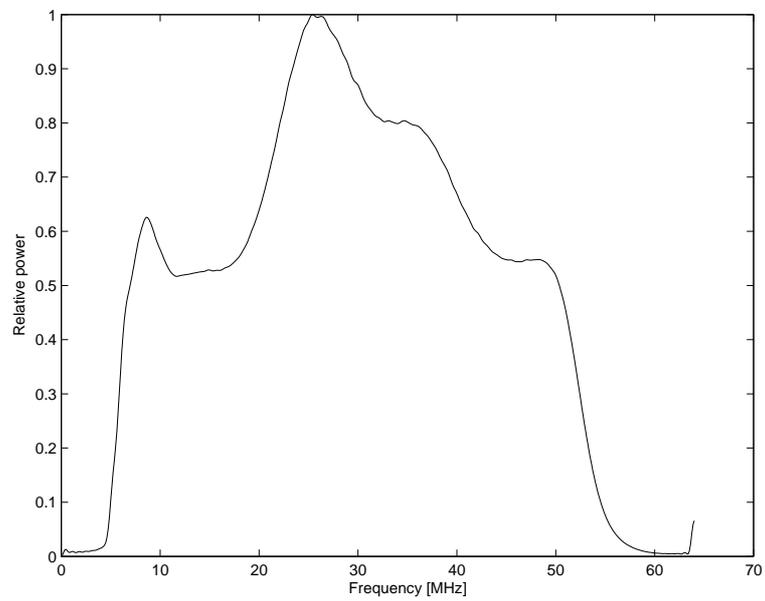


Figure 3.6: Noise spectrum for the ground tests

be reflected on the returned spectrum as the Doppler shift which can be estimated using (3.1)

$$v_g(i, j) = \lambda \frac{f_g(i, j) - f_l(i, j)}{2} \quad (3.10)$$

where $f_g(i, j)$ is frequency of the spectrum peak and $f_l(i, j)$ is the laser frequency, estimated with a direct sampling of the emitted pulse (see above). The peak frequency $f_g(i, j)$ is estimated with more accuracy with a weighted average around the observed maximum

$$f_g(i, j) = \frac{1}{NT_s} \frac{\sum_{-a}^a (k_g(i, j) + a) P_g(i, j, k_g(i, j) + a)}{\sum_{-a}^a P_g(i, j, k_g(i, j) + a)} \quad (3.11)$$

where a is the evaluation radius, $k_g(i, j)$ is the index of the peak frequency in the periodogram, T_s is the sampling frequency and N is the number of points of the Fourier transform.

The estimation of the spectral width is also important and the method described by Keane et al. in [59] for the detection of wakes is actually based only this spectral moment. They have realized that the mean speed deficit in the vortex centre can not be detected with enough stability, as it is both too weak and present in a very thin region. The spectral width estimation is performed on a power spectrum from which the noise spectrum was removed and is then defined as [75] [72]:

$$\sigma_g(i, j) = \sqrt{\frac{\sum_{u=-a}^a \left(f_g(i, j) - \frac{u}{NT_s} \right)^2 P_g(i, j, k_g(i, j) + u)}{\sum_{u=-a}^a P_g(i, j, k_g(i, j) + u)}} \quad (3.12)$$

As Keane notes [59], this method is very sensitive to the width k on which the weighted average is performed and the averaging domain must be fine-tuned in order to give meaningful results. It is nonetheless a computationally efficient approach that gives interesting results for the spectral width. The results published by Keane could not be reached with the approach he described and deteriorate rapidly with the distance to the laser source. Instead of using the late LOS grouping of his technique, we have therefore applied the early LOS grouping to obtain the results of fig. 3.7. As a convention we will always show the nearest gates at the bottom of the

graph and the farthest gates at the top. The range of the measurements (0.7 to 2 m/s) is justified by Keane [59] and will be used in general for spectral width figures.

To further enhance the estimations an automatic adaptation of the evaluation radius a was implemented. The lower boundary a_1 is defined as the smallest index that verifies

$$\begin{cases} P_g(i, j, a_1 + 1) > P_g(i, j, a_1) \\ a_1 < k_g(i, j) \end{cases} \quad (3.13)$$

and the higher boundary of the interval a_2 becomes the highest index verifying a similar constrain:

$$\begin{cases} P_g(i, j, a_2 - 1) > P_g(i, j, k_2) \\ a_2 > k_m \end{cases} \quad (3.14)$$

where k_m is the number of samples in P . The evaluation formula for σ_g is then:

$$\sigma_g(i, j) = \sqrt{\frac{\sum_{u=-a_1}^{a_2} \left(f_g(i, j) - \frac{u}{NT_s} \right)^2 P_g(i, j, k_g(i, j) + u)}{\sum_{u=-a_1}^{a_2} P_g(i, j, k_g(i, j) + u)}} \quad (3.15)$$

The spectral width results shown on fig. 3.8 and 3.9 are significantly better than the original estimation of fig. 3.7. The mean velocity estimates are almost identical to the results of Keane which is not surprising since this measure is easier to obtain.

3.3.2 Correlation based - ARMA

While most techniques rely on the spectrum of the signal to determine its properties, the Auto Regressive - Moving Average (ARMA) methods [76] [77] [78] are directly using the time-domain signal to find the main modes of the spectrum. Several implementations exist for the ARMA. We have here tested the approach of Friedlander et al. [79]. It consists in two separate steps, one for determining the AR parameters which uses the modified Yule-Walker equations (see for example [80]) and a second one to estimate the MA parameters [81]. A similar but iterative ARMA technique has been used for Doppler weather radar by Palmer [82] and later by Keel and Baxa [83].

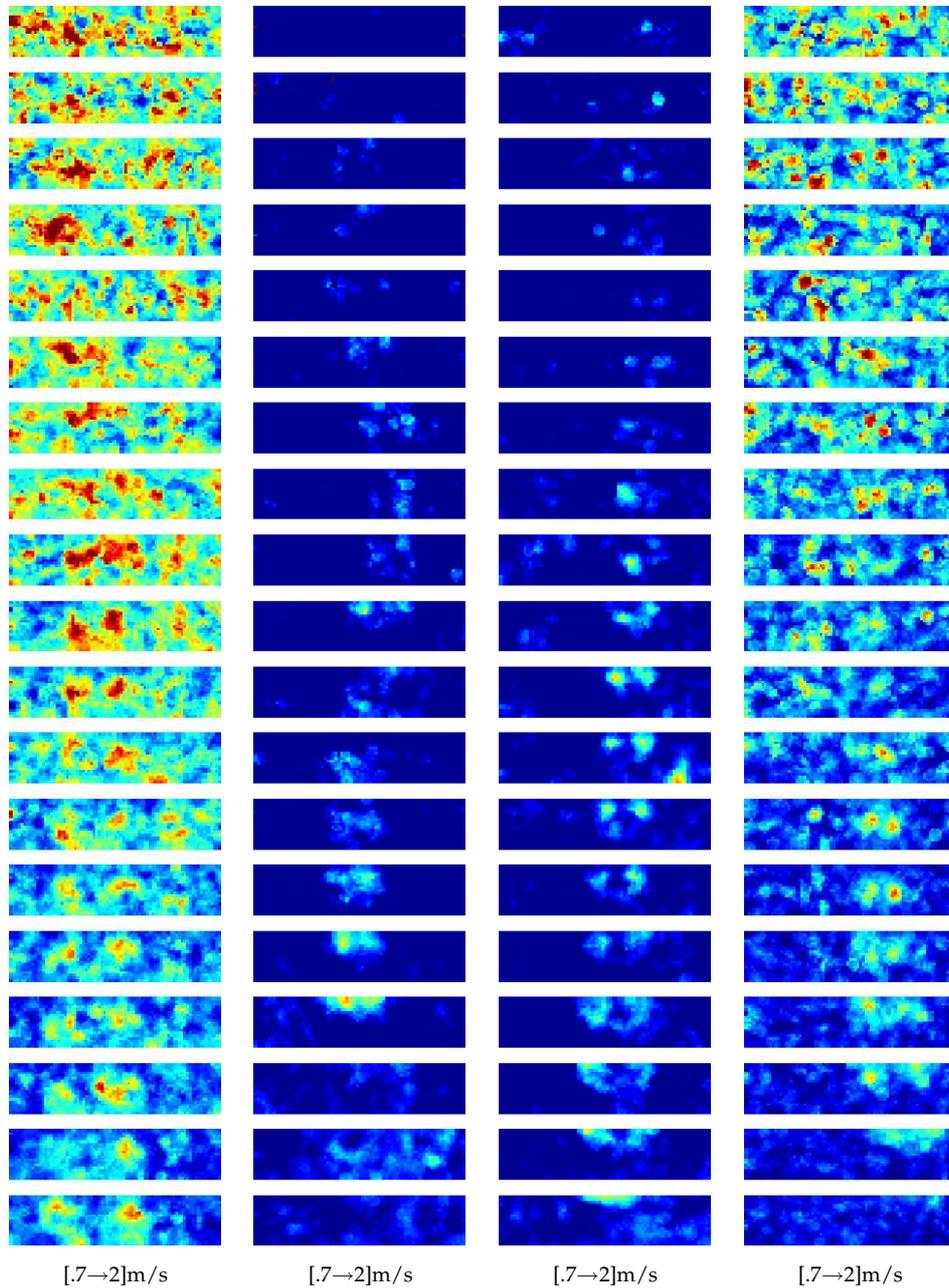


Figure 3.7: Classic estimation of the spectral width without automatic evaluation radius estimation. Left to right: A340, A320, Beluga and Fokker 100 scans. The top gate is the farthest (2300m), the bottom gate is the closest (875m).

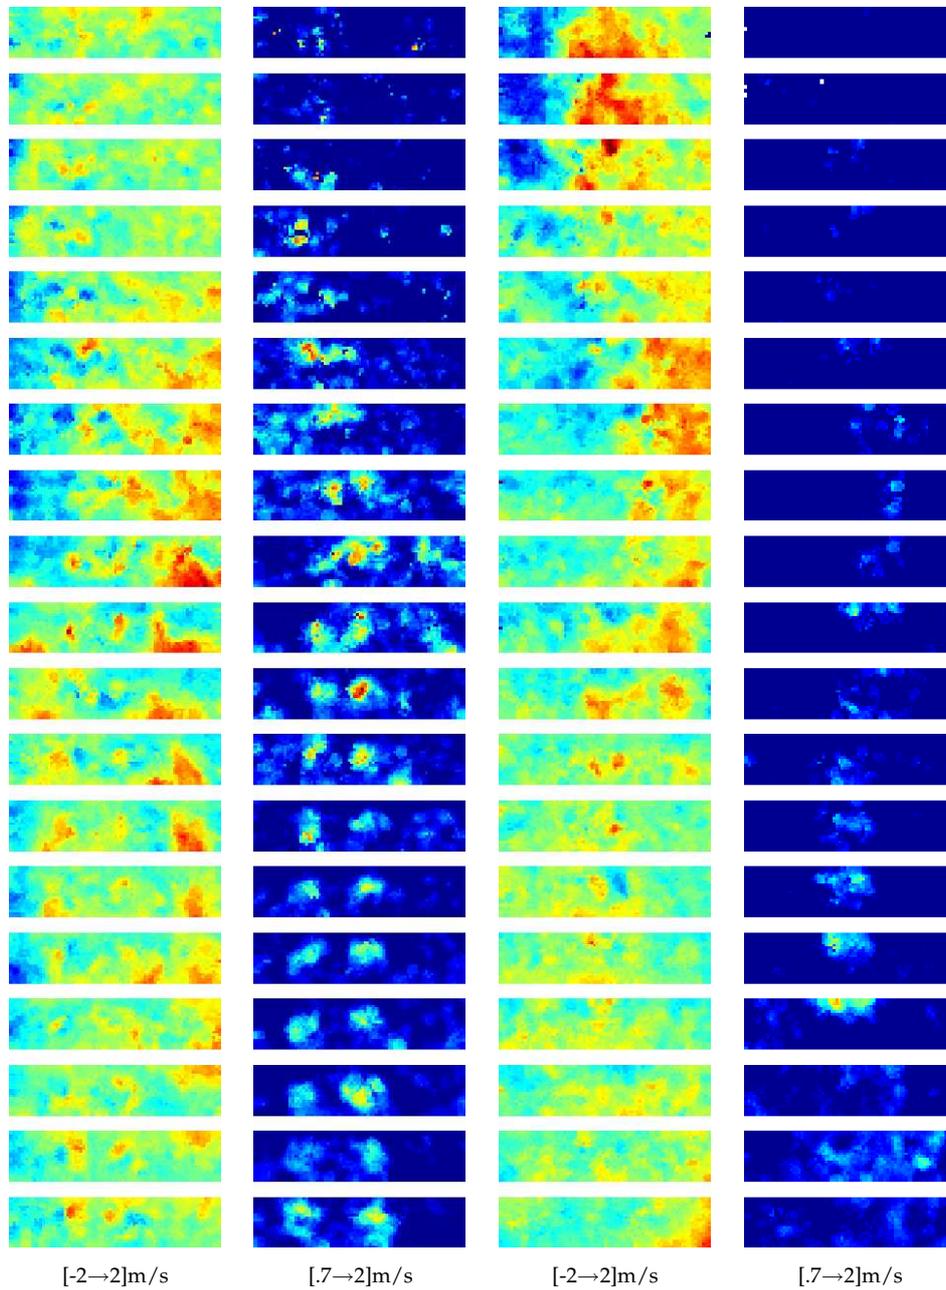


Figure 3.8: Kean's estimations for the A340 and A320 flybys. Left to right: A340 mean velocity, A340 spectral width, A320 mean velocity and A320 spectral width

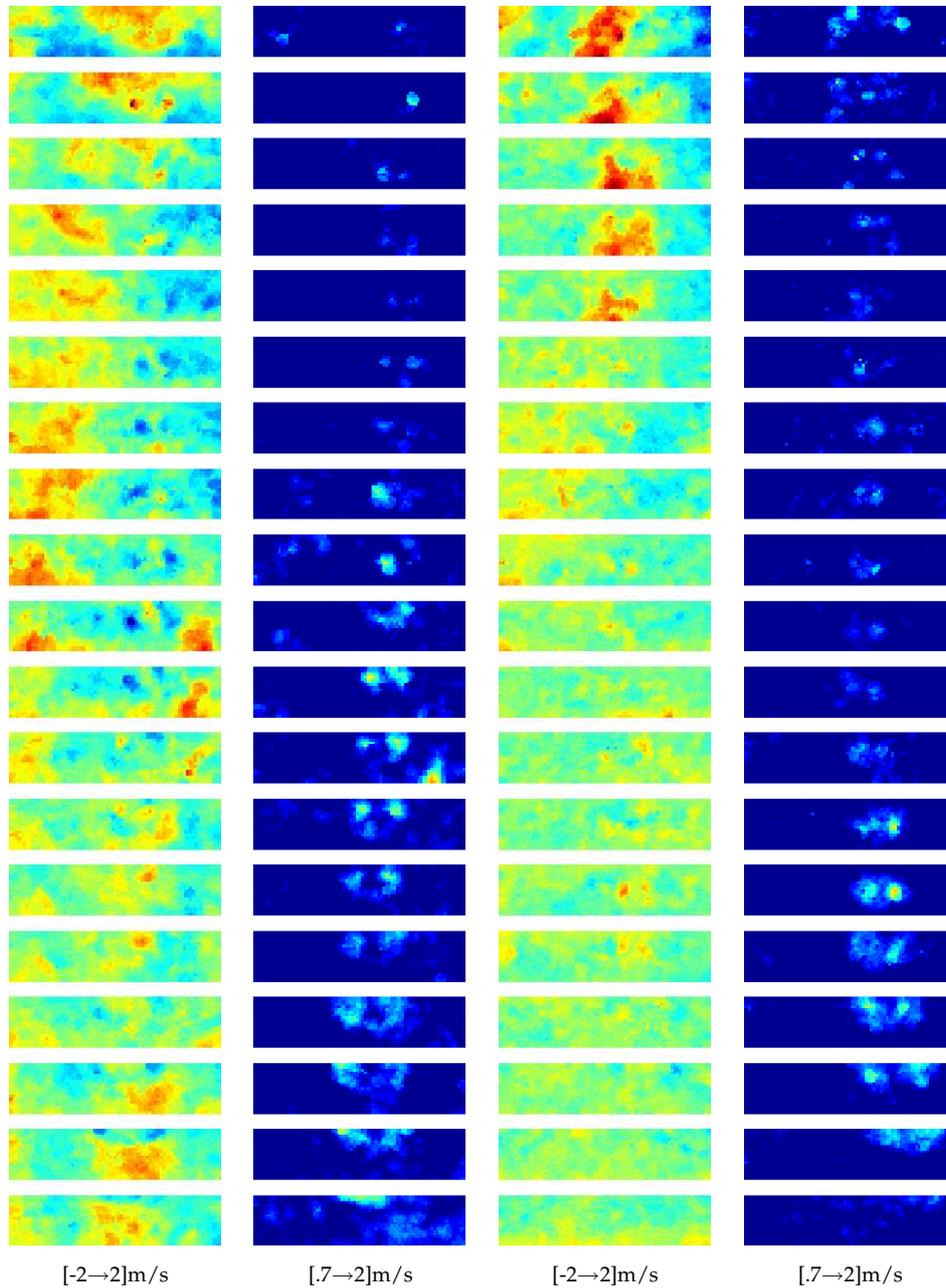


Figure 3.9: *Kean's estimation for the Beluga and Fokker 100 flybys. Left to right: Beluga mean velocity, Beluga spectral width, Fokker 100 mean velocity and Fokker 100 spectral width*

As shown on fig. 3.10, the ARMA provides interesting results for the spectral width estimation but its application is limited in our case because the shape of our spectrum does not fit well with the shape of the modes of the ARMA, thereby requiring a large number of modes to fit the spectrum properly. Indeed, the ARMA is better suited for multiple sharp-peaked spectrums but is not exactly designed to model one Gaussian peak very accurately. This approach is also significantly more time-consuming than the spectral estimation, for results that are better in the near field but degrade rapidly with the distance to the lidar, as can be seen by comparing fig. 3.8-3.9 and fig. 3.10.

3.4 Model-based spectral estimation

Another approach to spectral estimation is to match a spectrum model to the received signal using non-linear fitting techniques [84]. We first present the results of a single primitive fit and then show that a new two primitives model is more appropriate and has a better discriminating capability for wake vortices.

3.4.1 Single signal primitive

A typical spectrum as the one of fig. 3.5 has two identified components in the literature: the returned signal and the detector noise. The signal returned can be modelled as a Gaussian curve (see for example [82] [85] [86] and [84]) with a mean frequency f_1 and a variance σ_1 . For what the detector noise is concerned, its spectrum $N(f)$ is known from specific measurements (see section 3.3.1 above) and only its relative level n must be determined. The simplest model for the signal spectrum is therefore

$$M_1(f) = m e^{-\frac{(f_1-f)^2}{\sigma_1^2}} + nN(f) \quad (3.16)$$

where m is the amplitude of the signal peak. In order to reduce the degree of freedom of the model the spectrum is normalised, leading to the following simplified model:

$$M_1(f) = \frac{e^{-\frac{(f_1-f)^2}{\sigma_1^2}} + nN(f)}{1 + nN(f_1)} \quad (3.17)$$

A few resulting fits are presented on fig. 3.11. The main peak of the spectrum is correctly modelled but the fit is not appropriate at the base of the

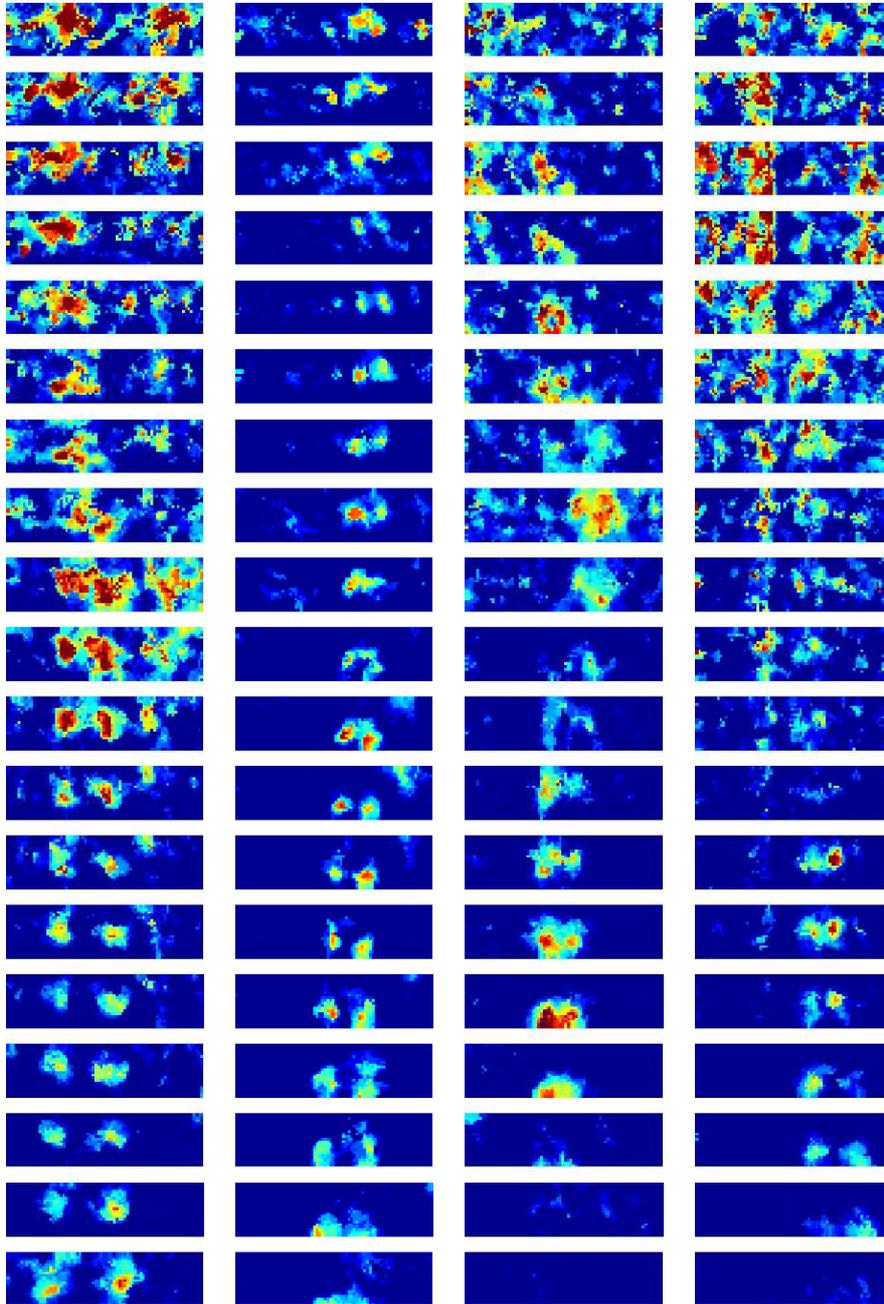


Figure 3.10: ARMA estimation for the four reference tests. Left to right: spectral widths for the A340, Beluga, A320 and Fokker 100. All ranges are $[0.7 \rightarrow 2]$ m/s

peak: not only is the variance insufficient but the spectrum also shows a clear asymmetry that can not be accounted for with the current model. A complete set of results for the four test scans is shown on fig. 3.12. The mean velocity is very close to the results of the classic spectral methods and does not require specific comments. On the other hand, the spectral width differs significantly from the previous estimations with vortex signatures that can barely be seen. This observation might be surprising given that the previous results of Keane showed good vortex signatures in the spectral width field, but it is explained quite easily by observing that all the broadening information is located at the base of the peak where the quality of the fit is poor.

Two more gated data are provided. The first is the relative noise energy E_N defined as

$$E_N = n \int N(f) df \quad (3.18)$$

which is naturally growing with the distance to the lidar source. The noise energy has a small correlation with the vortices presence but nothing significant. The last information provided is the relative residue energy, defined similarly as

$$R = \int \|M_1(f) - P(f)\| df \quad (3.19)$$

Note that there is no metric available for the energy so that the ranges given on the gated figures will not have units. The energies can still be compared: for example the noise energy of fig. 3.12 is (almost) always higher than the residue energy. Observation of fig. 3.12 reveals in fact that the residue is the best way to detect the vortex activity, thereby clearly indicating that the single primitive model is not complete. We can also observe that in farther gates the vortex pattern fades as the noise residue becomes more important than the unmodelled component of the spectrum. We therefore come to the conclusion that the single Gaussian model is not appropriate and that a second component must be used to take into account the broadening of the peak.

3.4.2 Two signal primitives

Given the insufficiency of the single primitive approach we wish to add an additional Gaussian curve to model the base of the spectrum peak. This is a typical approach for UV lidars where two different atmospheric interactions (reflection and molecular interaction) yield widely separated curves.

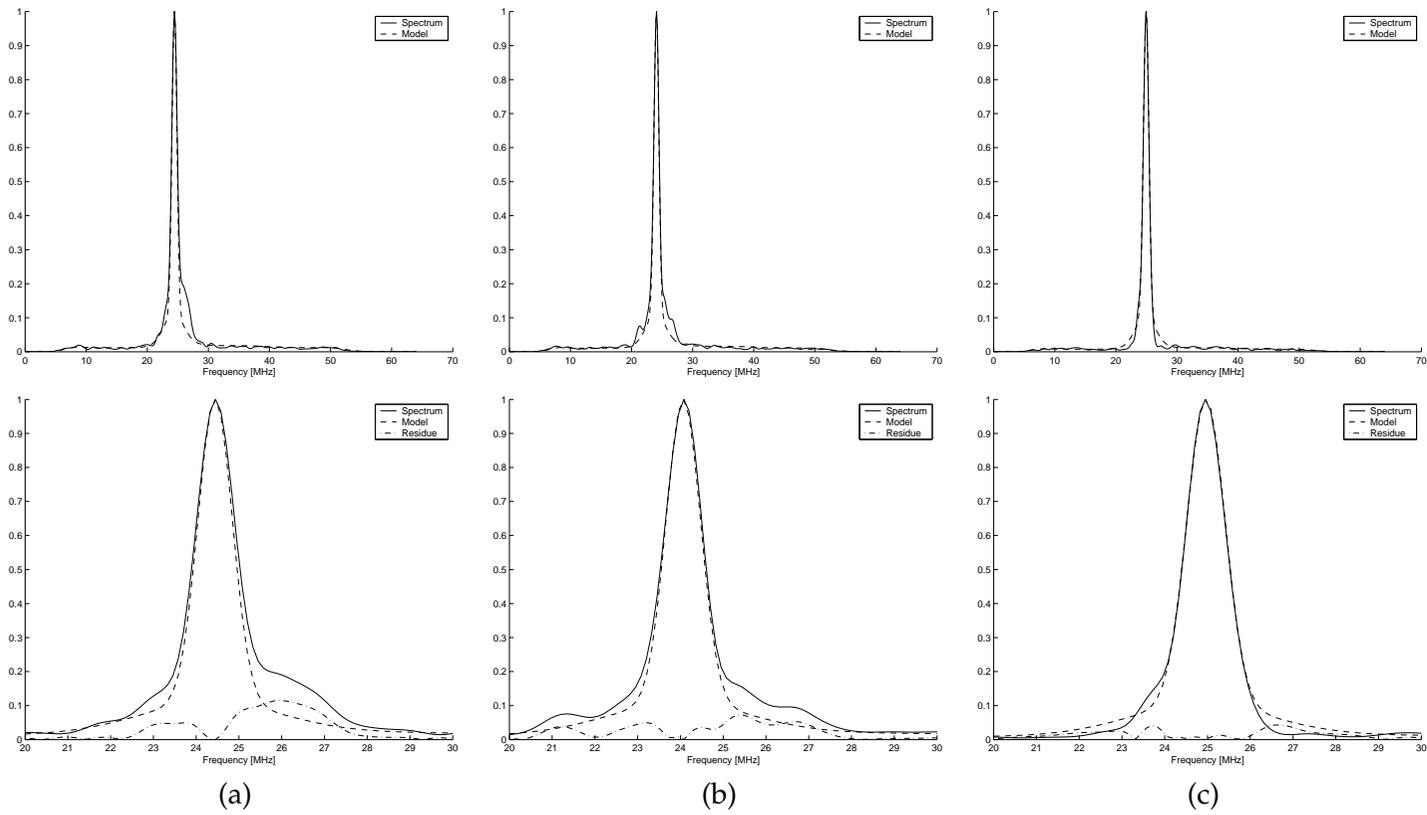


Figure 3.11: Spectral fits with a single signal primitive

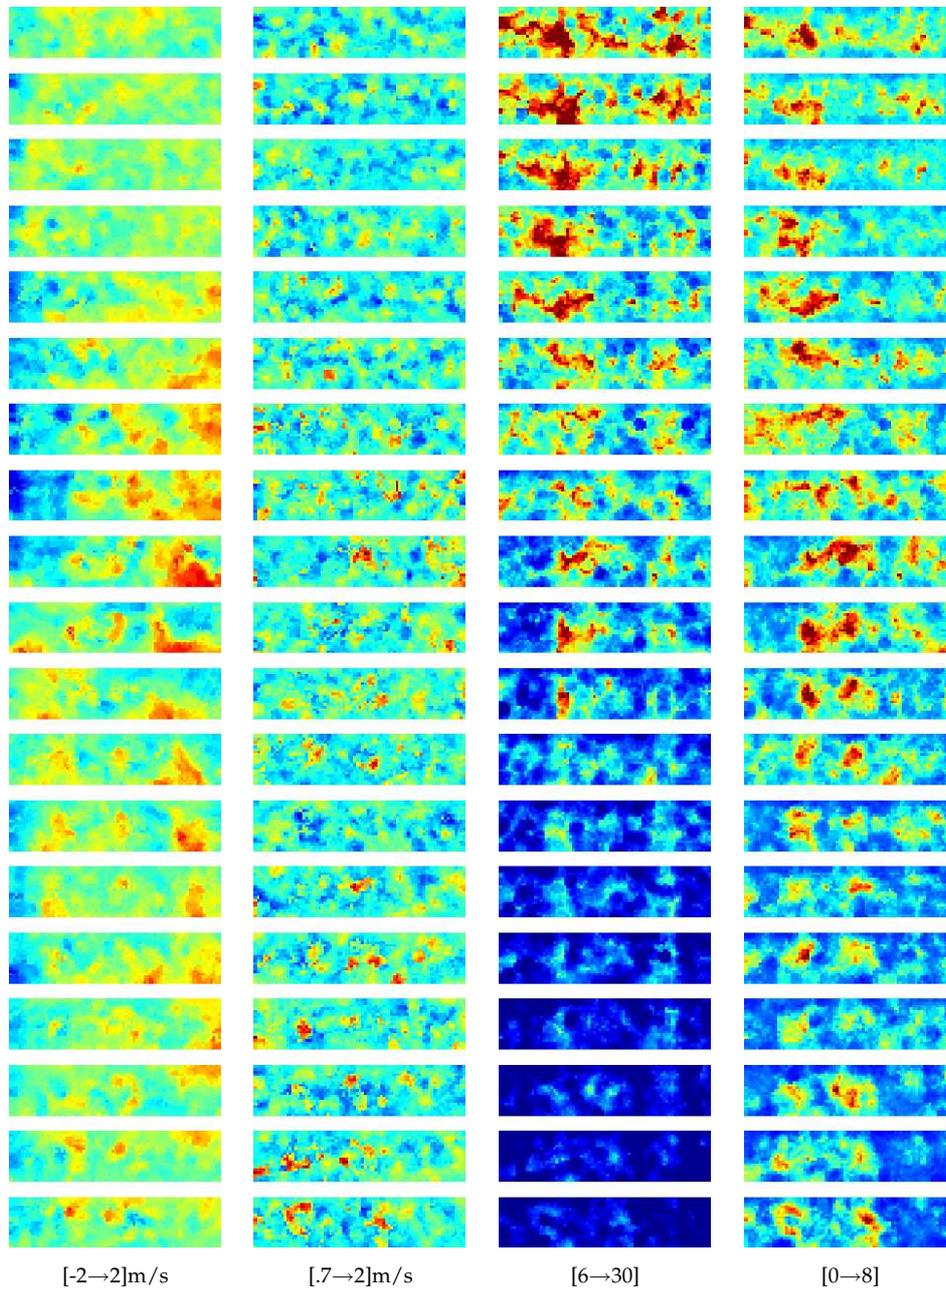


Figure 3.12: Single primitive model parameters: mean velocity f_1 , spectral width σ_1 , noise energy and residue energy

However, the application of a two primitive model to IR lidar measurements is an approach that can be qualified a new to the best of the knowledge of the author.

The secondary Gaussian curve will have a relative amplitude s , a variance σ_2 and a mean frequency f_2 . We call this new component the secondary return and obtain a new model:

$$M_2(f) = me^{-\frac{(f_1-f)^2}{\sigma_1^2}} + se^{-\frac{(f_2-f)^2}{\sigma_2^2}} + nN(f) \quad (3.20)$$

As with the single primitive approach, one degree of freedom is removed by normalising the model to

$$M_2(f) = \frac{e^{-\frac{(f_1-f)^2}{\sigma_1^2}} + se^{-\frac{(f_2-f)^2}{\sigma_2^2}} + nN(f)}{1 + se^{-\frac{(f_2-f_1)^2}{\sigma_2^2}} + nN(f_1)} \quad (3.21)$$

Fig. 3.13 show the results of this two primitives fit on the same spectrums of fig. 3.11. We can see that the base of the main peak and the main peak itself are much better modelled. The residue is very small compared to the original single primitive approach and almost only due to small noise variations. This comforts us in the idea that this model is meaningful and that the spectrum should indeed be modelled by two Gaussian curves of different mean, level and variance.

A full description of the parameters variation is provided on fig. 3.14 thru 3.16. Let us comment each parameter one by one.

- **Mean velocity, main primitive f_1 :** the mean frequency is here represented by the associated wind velocity. This result is almost strictly identical to the single primitive model, which is not surprising since the latter was already able to fit the main peak correctly.
- **Spectral width, main primitive σ_1 :** the spectral width of the main peak, also converted to velocity measures, is completely decorrelated from the vortex signature, showing that the main Gaussian component is not suitable for vortex detection with spectral width measurements.
- **Main peak relative energy E_1 :** since we are working with a normalised spectrum, the energy is directly proportional to the peak variance, leading to similar results for both parameters.

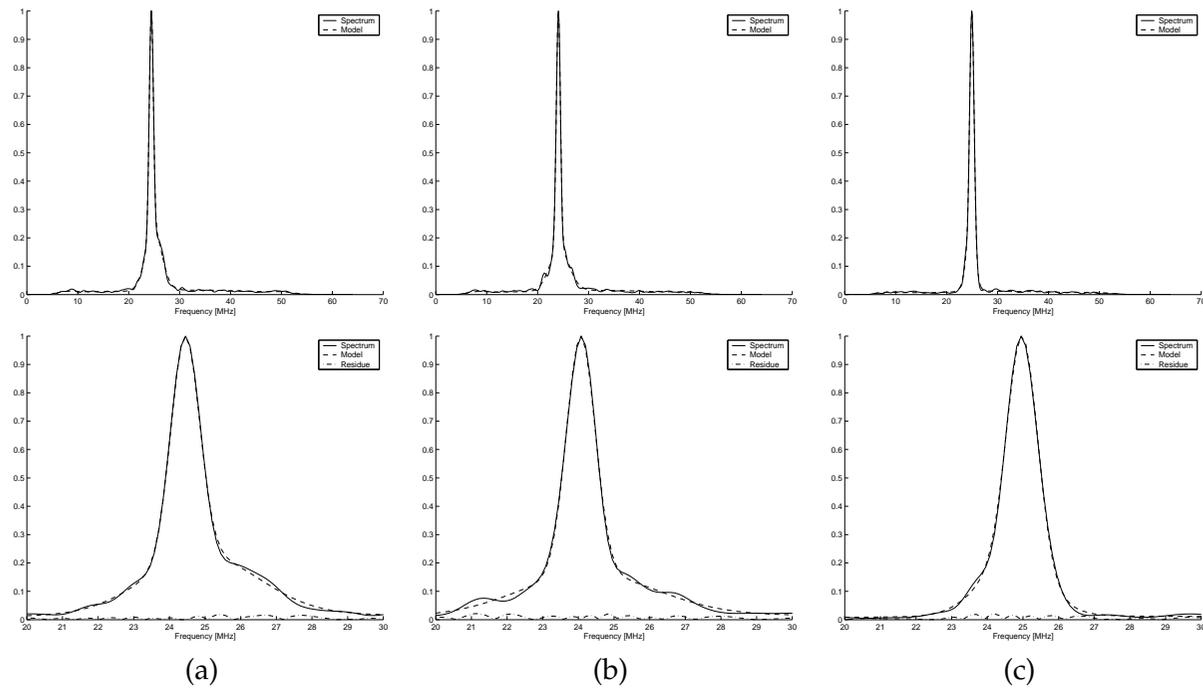


Figure 3.13: Spectral fits with a two signal primitive

- **Relative amplitude, secondary peak s** : the relative amplitude of the secondary return bears only a small correlation with the vortex location. It will be used in the form of the secondary return energy (see below)
- **Mean velocity, secondary primitive f_2** : the mean frequency of the secondary signal is quite similar to the one of the main signal. It contains more noise in the far field because the relative amplitude of the secondary component is close to the amplitude of the noise spectrum variations in this region.
- **Spectral width, secondary primitive σ_2** : the vortex signatures are clearly visible, showing once more that only the secondary return contains valid information about their higher turbidity.
- **Secondary peak relative energy E_2** : the vortex signatures are also vividly detected. The response of this parameters is more binary (large or low values, but few intermediate ones), but has fewer parasite detections than σ_2 in the far field. The fact that this energy correlates with the vortex location indicates that the secondary peak is only present where a vortex is located.
- **Relative noise energy E_N** : as expected, the noise level grows with the distance to the laser source. The two vertical streaks visible across the whole set of gates are not vortices but are due to the difference in the emitted power.
- **Relative signal energy**: the received power is not homogeneous as stated above, and lowers with the distance to the laser. The zones where it is higher correspond to a lower relative noise energy.
- **SNR**: the signal to noise ratio is defined in dB as

$$SNR = 10 \log_{10} \frac{E_1 + E_2}{E_N} \quad (3.22)$$

It reveals that the necessary SNR for a good detection is -10dB which is 10dB lower than the threshold found by Keane in [59].

- **Relative residue energy**: the residue is much lower than on fig. 3.12, confirming that the two primitives approach is better than the single primitive one.

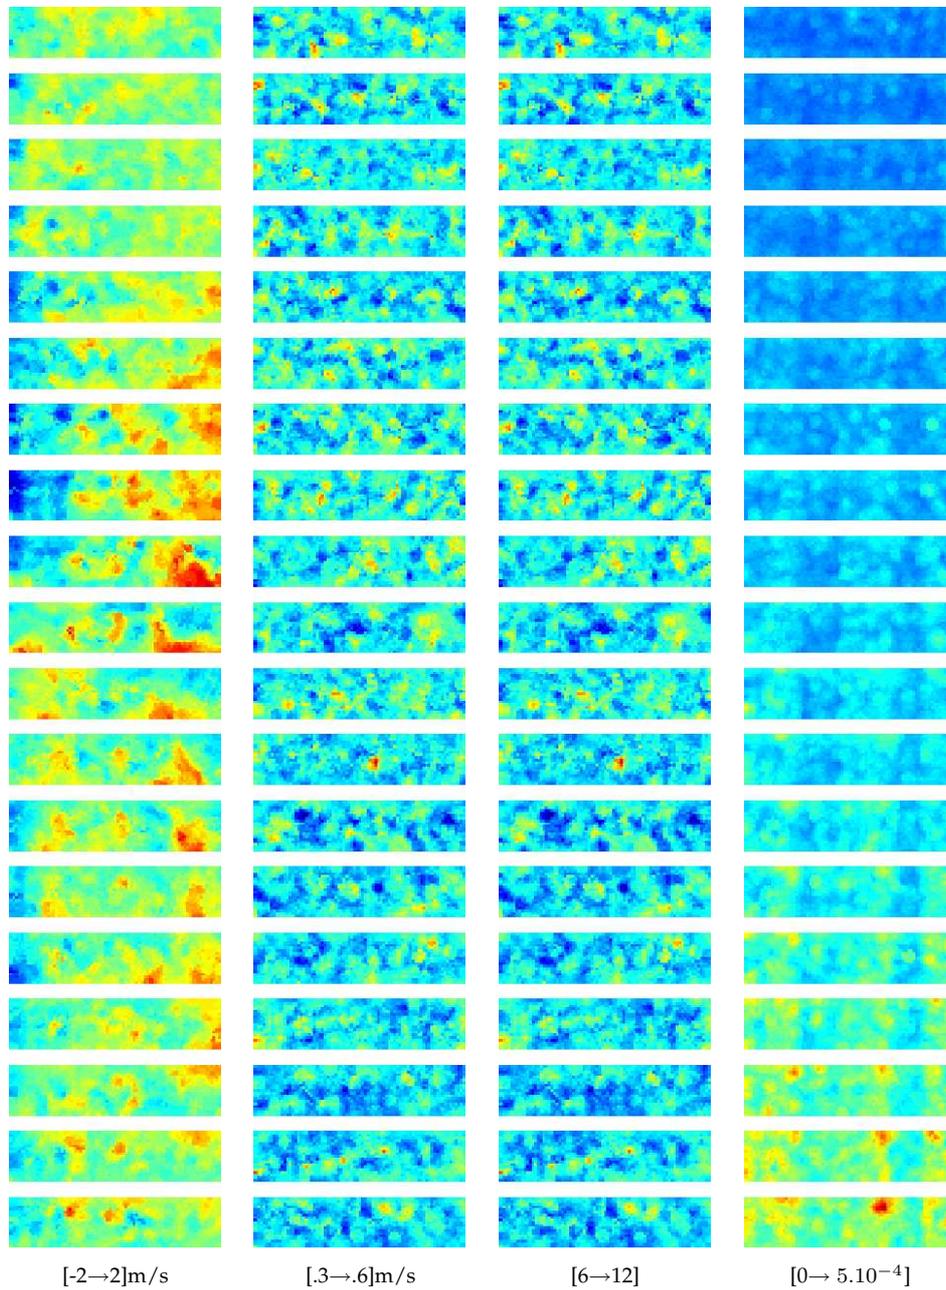


Figure 3.14: Model parameters for the main peak: mean velocity f_1 , spectral width σ_1 , peak energy and global return signal energy

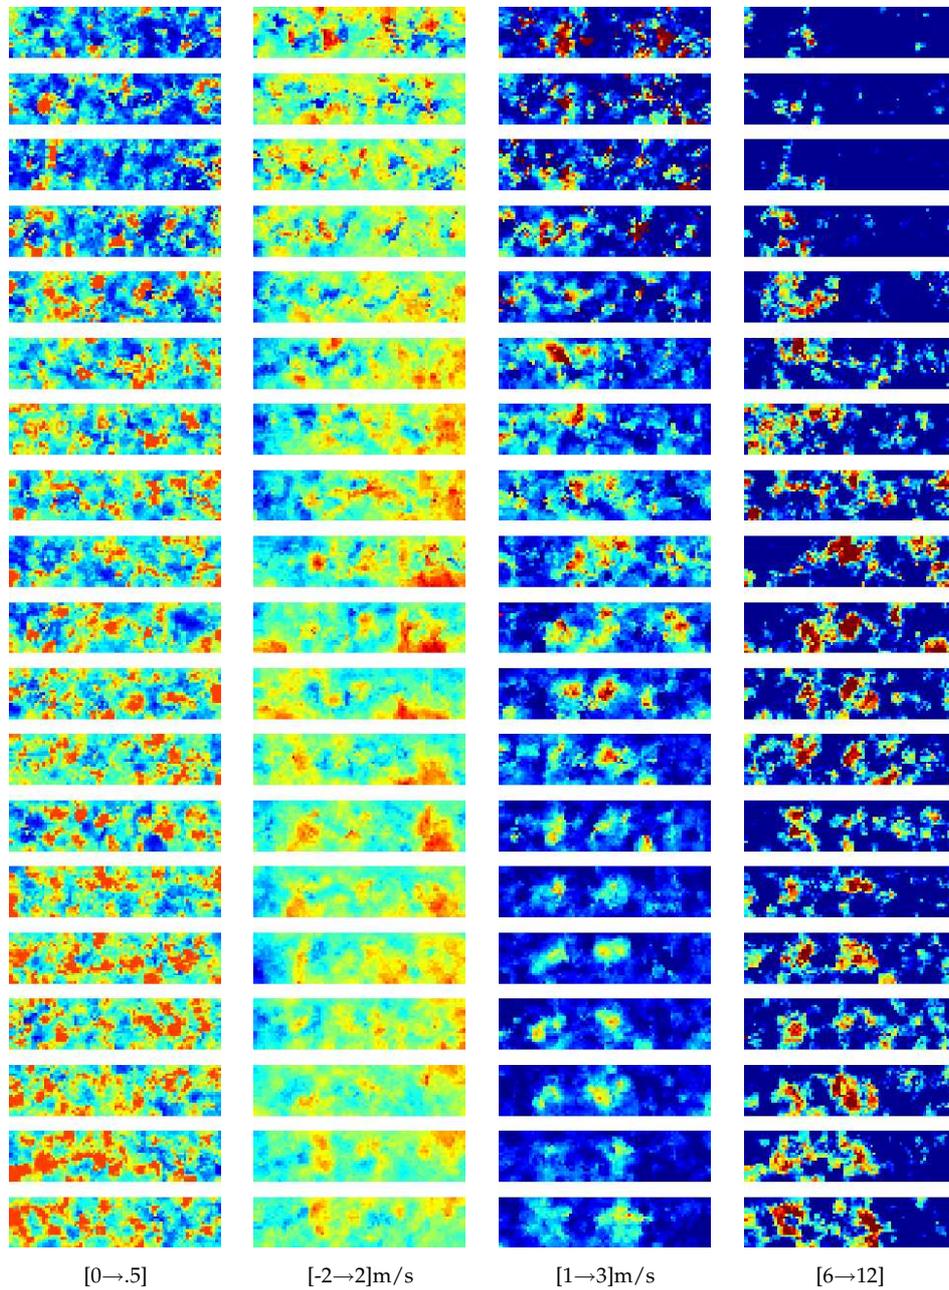


Figure 3.15: Model parameters for the secondary peak: relative amplitude s , mean velocity f_2 , spectral width σ_2 and energy

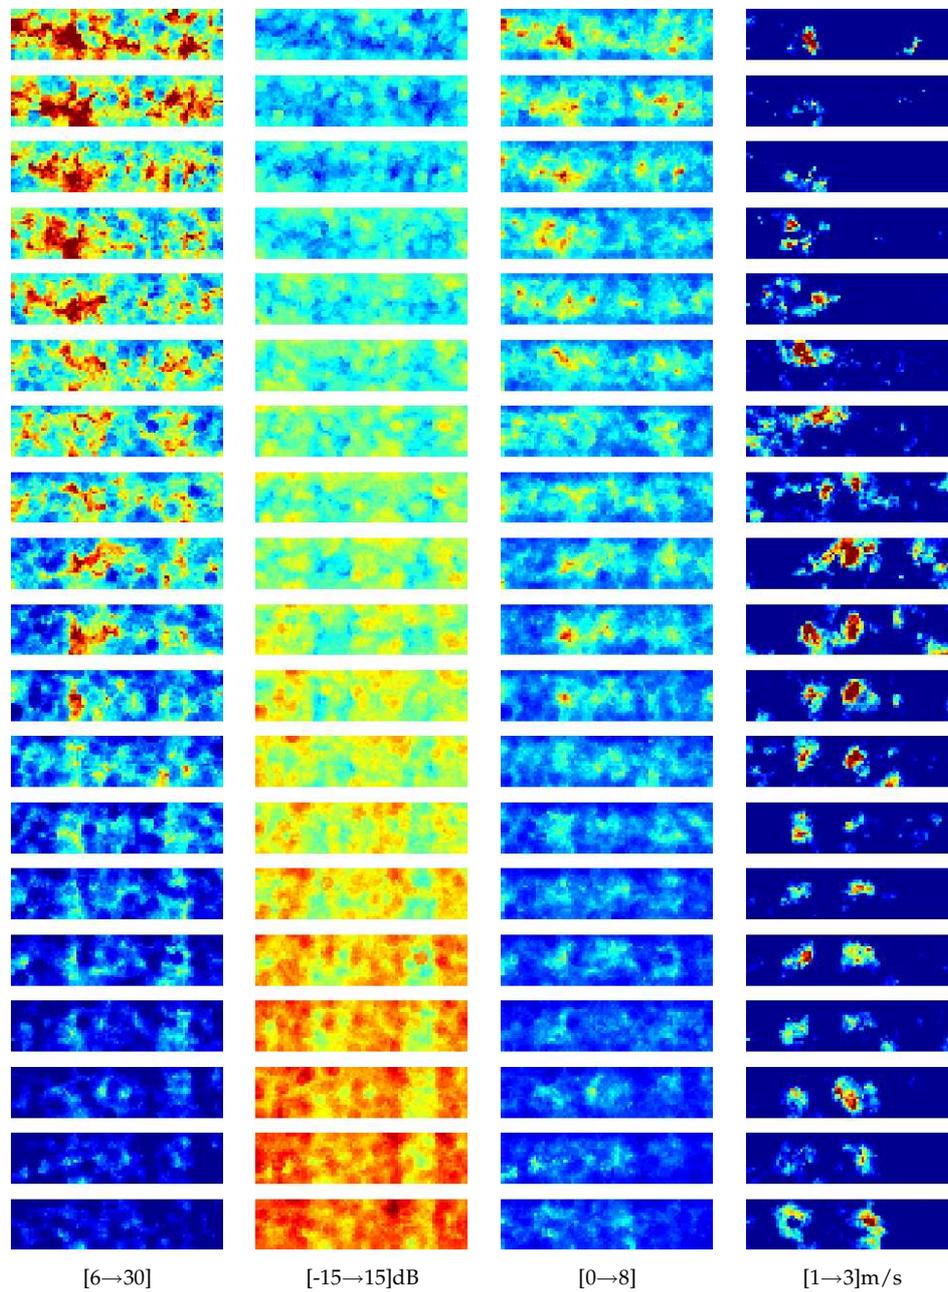


Figure 3.16: Other model parameters: noise energy, SNR, residue and the detection field

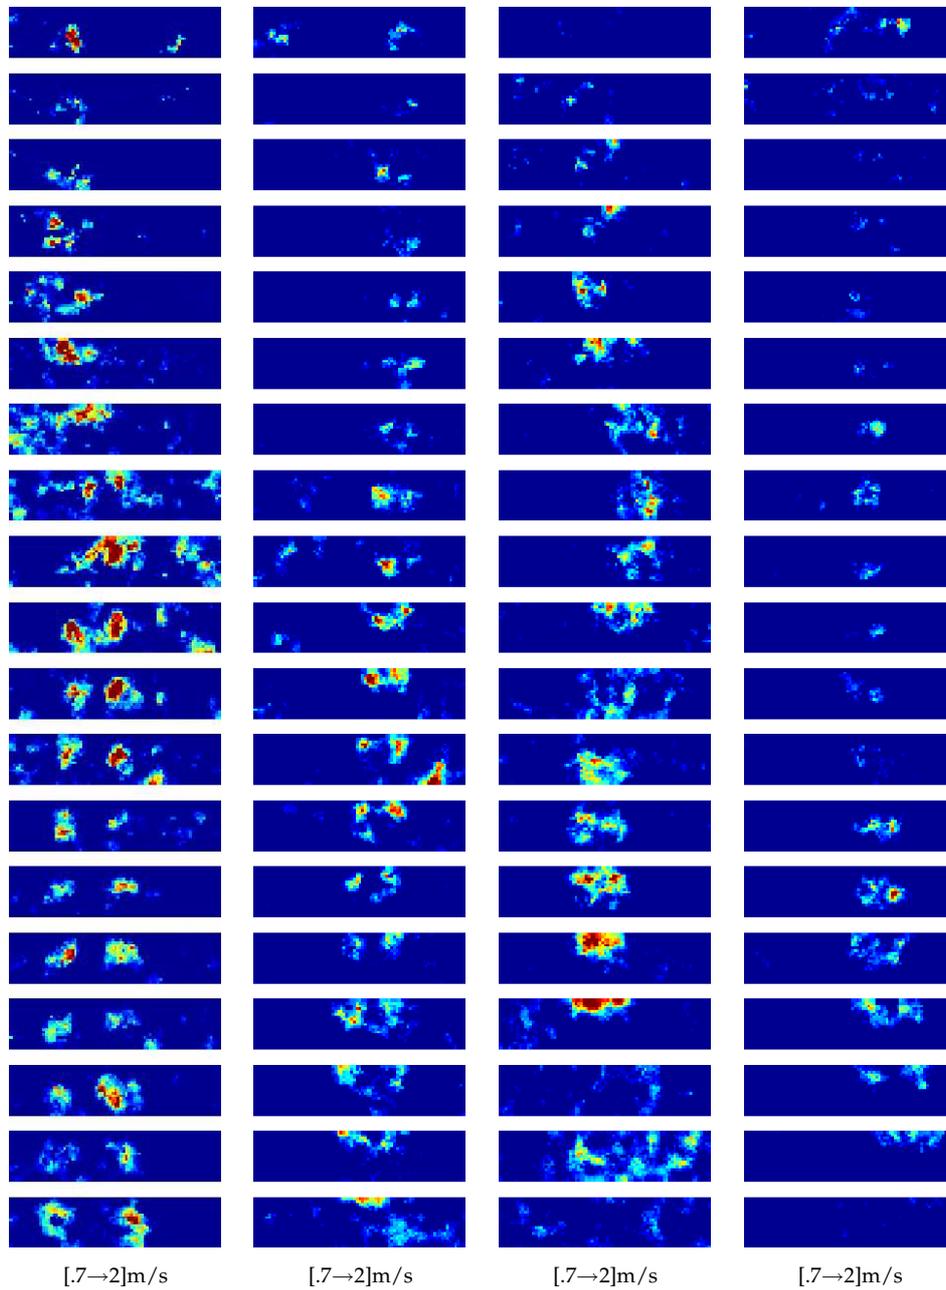


Figure 3.17: Model based estimation with a noise spectrum estimation: the detection field for the A340, Beluga, A320 and Fokker 100 MFLAME tests

From this discussion of the variation of the model parameters we conclude that a good candidate for the detection of wake vortices would be to combine the secondary peak energy and variance into a single detection field $D(i, j)$ ¹

$$D(i, j) = \sigma_2(i, j)E_2(i, j) \quad (3.23)$$

The resulting ranges gates for the four reference tests are shown on fig.3.17. The overall performance of the new detector is better from several aspects. First we note the higher contrast between vortex and atmosphere zones which, together with the lower noise, results in a greater usable range for the lidar. The results for the landing of the Fokker 100 also show that we are able to detect weaker wakes, even in the far field and without any parameter tweaking. The vortices also appear with more contrast than what has been presented on fig. 3.8 and 3.8, showing the superior discriminating power of the two primitives model. At last, a major advantage of our approach is that the SNR is available for each estimation of the detection field. It could be used to weight (3.23) in order to automatically adapt the detection to lower SNR situations that might otherwise trigger unwanted alarms in the cockpit. On a broader scope, the two primitives model provides much more information about the atmosphere and the different available fields might be used to detect other atmospheric hazards like Clear Air Turbulence (CAT), either by using the fields directly or by combining them. A hint on the ability to detect other phenomenon will be discussed at the end of this chapter.

This powerful estimation has the significant drawback of having a huge computational cost as twenty minutes are necessary to obtain a full scan analysis. This is more than two orders of magnitude from a real-time on-board system which can not be built with the currently available off-the-shelf hardware.

3.4.3 Two signal primitives without noise spectrum

In the case where no noise spectrum is available it is still possible to obtain good results by modelling the noise, for example as either a flat spectrum or as a third Gaussian of parameters f_n and σ_n :

$$M_3(f) = \frac{e^{-\frac{(f_1-f)^2}{\sigma_1^2}} + se^{-\frac{(f_2-f)^2}{\sigma_2^2}} + ne^{-\frac{(f_n-f)^2}{\sigma_n^2}}}{1 + se^{-\frac{(f_2-f_1)^2}{\sigma_2^2}} + ne^{-\frac{(f_n-f_1)^2}{\sigma_n^2}}} \quad (3.24)$$

¹From now on we will drop the g subscript denoting the gates to simplify the expressions.

The results shown on fig.3.18 are almost identical to the case where a noise spectrum is used but compared to (3.21) two more parameters must be estimated and the convergence is therefore more difficult. Adding a third Gaussian curve requires in general three times more iterations than if the noise spectrum is used.

3.4.4 Physical interpretation

The exact meaning of the secondary return is not yet known but the results above indicate that some particles within the air mass have a different behaviour. The first obvious observation is that the returned power is lower for the second peak. This can be explained by the lower effective reflectivity (m^2/sr) of the associated particles. Such lower reflectivity can be obtained if the particles are fewer or if they have a smaller size, or both. In our case the particles from the secondary return must have a smaller size because their variance is broader, reflecting that they are more subject to small perturbations. An accurate description of the atmosphere when the tests have been conducted would be necessary to know if the smaller size of the particles is the sole responsible for the lower reflected power or if the particles are also fewer. Another possible explanation is that the lidar wavelength of $2\mu m$ yields the peak response is achieved for particles of $2\mu m$ or larger. The lower return from the secondary particles could mean that the particles are smaller and thus not well detected by the lidar, even if their density is significant.

Concerning the mean velocities f_1 and f_2 , the secondary return follows closely the mean frequency changes of the main peak with some minor differences in some areas. This shows that both peaks are influenced by the mean speed changes generated by the airplane. There is at this time no indication that this change is due to the presence of vortices and the flight tests actually show that variations of the mean air speed are rather due to the the residual air speed behind the aircraft².

A tempting conclusion would be that the secondary primitive is due to the exhaust particles or vapour from the aircraft reactors that are trapped in the vortex. This is backed up by the observation that the secondary particle distribution is only present behind the aircraft and therefor might be generated by it. Flight tests have shown that it was probably not the case because the relative power of the secondary return is not higher when the other particles densities drop at higher altitude.

²Personal communication from Thomas Duquesne, UCL, Thermodynamics Laboratory

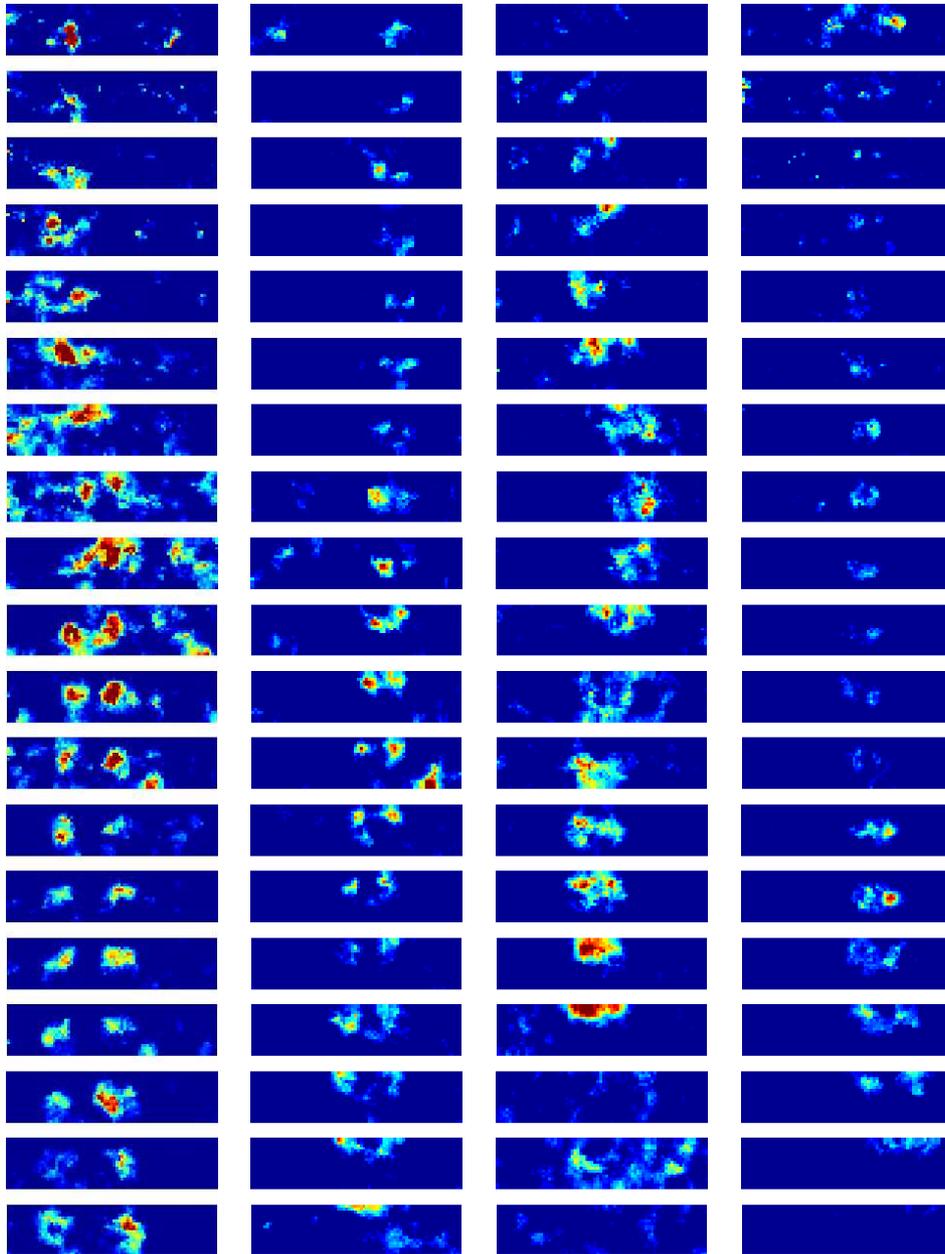


Figure 3.18: Model based estimation with a Gaussian noise: spectral widths and power product for the A340, Beluga, A320 and Fokker 100 MFLAME tests [1→3]m/s

More research will have to be performed in order to derive a definitive conclusion on the origin of the secondary primitive. In particular, experiments on a reduced scale aircraft in a controlled atmosphere might be very valuable but will require very specific short range lidars.

3.4.5 Convergence issues

The issue of convergence must be tackled since fitting the model to the data is done with a non-linear iterative process that is not guaranteed to converge. The non-linear fitting technique used is the Gauss-Newton method with a Levenberg-Marquart type adjustment. A-priori knowledge on the vortex parameters has also been integrated in the form of limited variation domains for each parameter. This limitation is particularly important for the variances of the Gaussian primitives which must have mutually exclusive variation domains in order for the algorithm to converge properly. The values used empirically are $\sigma_1 \in [0.25\text{MHz} \rightarrow 1.25\text{MHz}]$ and $\sigma_2 \in [1.5\text{MHz} \rightarrow 7.5\text{MHz}]$. The result is a convergence that is usually fast, requiring less than 30 iterations. The convergence is not achieved in some cases, especially in the far field which has more variability due to a higher noise level. This is however limited to about 0.5% of the worst-case farthest gate and does not appear on the final results.

In the case where the noise spectrum is not available the convergence is more complicated because the model (3.24) contains two more parameters. The optimisation is therefore split in two parts. First the two Gaussian of the signal are estimated with a fixed noise variance σ_n and mean f_n (only the noise level n is estimated) which is not a major restriction since the noise spectrum does not change significantly (see 3.3.1). Once the signal parameters are obtained σ_n and f_n are set free and the full optimisation is performed. The variation domain for the noise variance is also mutually exclusive with the other variance domains and set to $\sigma_n \in [25\text{MHz} \rightarrow 37.5\text{MHz}]$. With this two-step optimisation the non-convergence after 50 iterations climb up to 2%. However, even with this higher percentage the results are very similar to the ones using the noise spectrum, as fig. 3.18 shows.

3.5 Spectral estimation of on-board signals

After the development of the signal processing algorithms and their test on ground-based lidar data, the next step is to perform flight tests with

the lidar installed in an aircraft. In what follows we will first describe this installation and show its geometric differences with the ground based system. The adaptation of the signal processing is then tackled, followed by a description of the pre-tests (or *shakedown* tests). The flight tests are then depicted and their results shown and discussed.

3.5.1 Aircraft installation

The installation of the lidar in the aircraft is performed with a few constraints in mind. First the equipped aircraft must not encounter the vortex generated by the plane in front of it. As the vortices always move downwards with time the altitude of the following aircraft must be equal or higher than the one of the generating aircraft. This forces the lidar to be pointed downwards which leads to a scanning volume positioned as shown on fig. 3.19. During the flight tests, the measuring plane will follow the generating plane so that the vortices can be detected axially. The second limitation then comes from the aircraft itself which has no opening at its bottom to allow the laser beam to aim straight in the axis of the aircraft. As it is of course excluded to create an additional hole in the fuselage we are left with the existing orifices of the plane, namely the lateral windows. A protruding bubble was therefore designed to protect the optics. For aerodynamic reasons the bubble can not extend too much outside the airplane and the lidar will as a result not aim straight but with a certain angle with the fuselage axis. This horizontal and vertical angles are represented on fig. 3.20 and have a value of 9° .

3.5.2 Signal processing adaptation

From a signal processing point of view, the adaptations necessary for an on-board system are limited to the speed, attitude and position of the plane which must be integrated in our algorithms. Two other aspects related to the operational conditions of the flight tests will be also covered. The first is the new acquisition hardware used for the on-board experiments which required modifications to the software developed for ground lidar. The second is the adaptation of the model-based approach for the lower SNR that is experienced at a higher altitude.

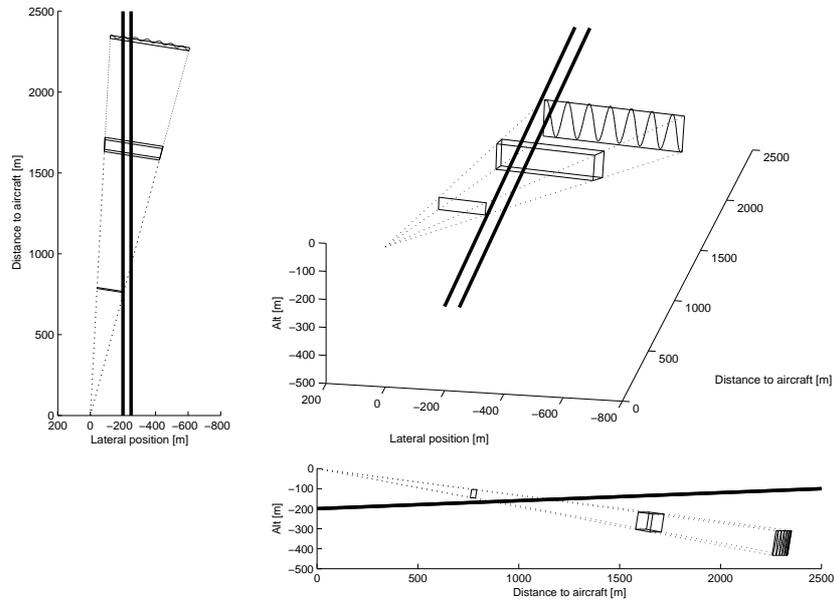


Figure 3.19: Position of the scanned volume relative to the aircraft. The thick lines represent typical wake vortices locations.

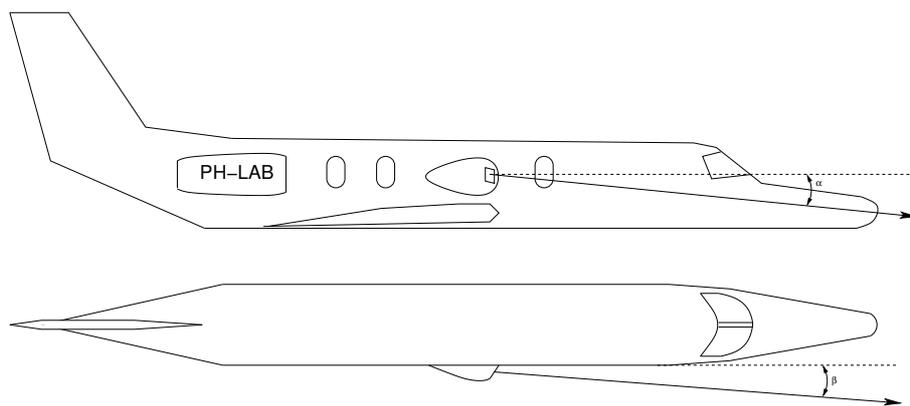


Figure 3.20: Installation of the LIDAR on board the NLR Citation II aircraft

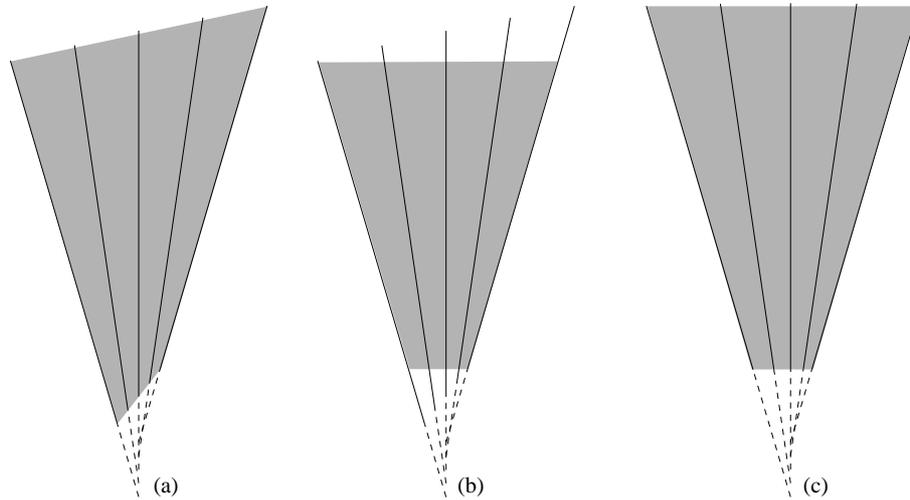


Figure 3.21: *Analysis volume: without compensation, with software compensation and with hardware compensation*

Position

The movement of the laser source means that every shot will depart from a different point in space. Fig. 3.2 is not valid anymore, and a top view of the new scanned volume is shown on fig. 3.21a. Due to the relatively high speed of the plane (80m/s), the difference in position between the first and last shot of a scan will be $80\text{m/s} \times 5\text{s} = 400\text{m}$, which is a significant displacement compared to the LIDAR range of 2400m .

Several possibilities exist to compensate for this motion. One is to fix the first gate distance and modulate the delay of the receiver trigger so that the earlier a shot, the higher the delay applied. This has the advantage that the scanning volume will have a constant range, both for the minimal distance and the maximal one (fig. 3.21c). However, this solution requires modifications to the hardware and is thus not suitable for use within an existing system. Moreover other pose and attitude parameters should be integrated and can not be compensated in this easy fashion. For these reasons a software approximation is used instead, which consists in discarding the first samples of the returned echo (fig. 3.21b). This truncation will lower the maximal range of the lidar for the first shots, but this is not important in our case because this range is much further than the limit allowed by the signal processing.

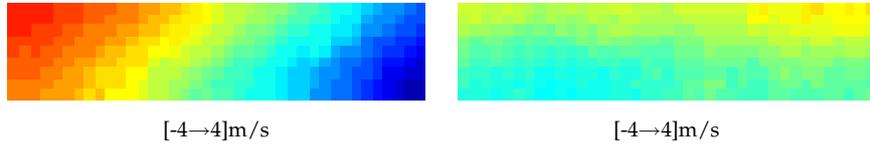


Figure 3.22: The wind speed on a selected gate, before and after application of the aircraft speed compensation

Speed

In a Doppler system the relative speed of the source and the target creates the frequency shift. The speed of the aircraft will therefore have a major impact on the results and must be integrated in the signal processing. This is especially true because both speeds have the same weight in (3.1) and a variation of 1% in the aircraft speed (0.8m/s) will correspond to nearly 25% of the airspeed variations due to a vortex which are within $\pm 2\text{m/s}$. This large influence on the speed measurement requires a special care to be taken during the integration of the speed of the source in our calculations.

The lidar is measuring air speed vectors that are aligned with the direction d^L of its line of sight L and is only influenced by the component of the aircraft speed v^L that is aligned with this direction. Projecting the aircraft speed vector onto the lidar line of sight is a simple matter of three-dimensional geometry and the correction Δv^L to be applied to the Doppler measurements is

$$\Delta v^L = \frac{(v^L, d^L) d^L}{(d^L, d^L)} \quad (3.25)$$

where $(_, _)$ indicates the dot product. Fig. 3.22 illustrates uncompensated and compensated speed measurements. The strong horizontal gradient of fig. 3.22 is due to the fact that the lidar points to the right of the aircraft axis. The span of the speed error on the 12° of the mirror horizontal motion can be estimated as

$$80\text{m/s} \times (\sin(15^\circ) - \sin(3^\circ))/2 \approx 8\text{m/s} \quad (3.26)$$

which corresponds to what is measured on fig. 3.22. A small vertical gradient is still visible on the corrected gate which is either due to mechanical tolerances during the on-board installation of the system or to the limited accuracy of the aircraft attitude information.

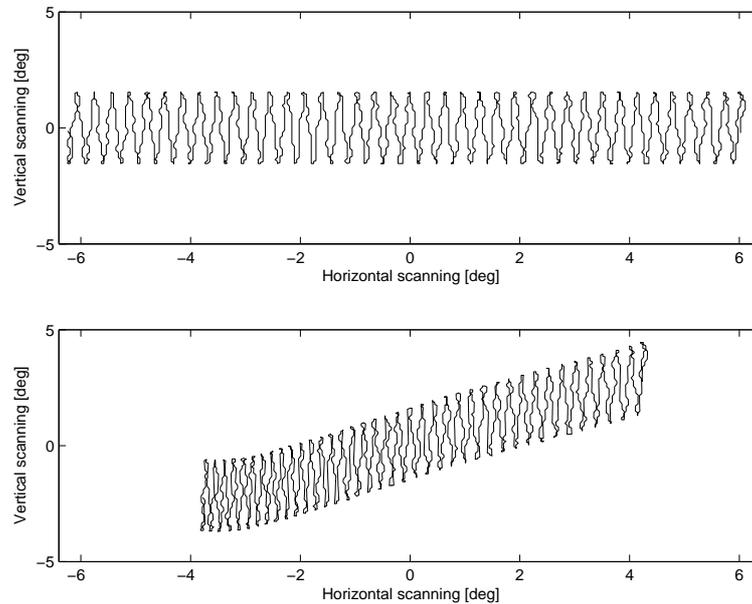


Figure 3.23: *The scanning pattern of the lidar: with only the mirror displacement and with the integration of the aircraft attitude*

Attitude

The aircraft attitude is an important variable that can influence the scanning pattern, as shown on fig. 3.23. Compared to a steady flight, the scanning pattern is compressed horizontally because the plane is turning starboard and follows the movement of the mirror. As the plane turns its roll also changes which is reflected as a tilt of the scanning pattern. Since we are interested only in steady flight conditions these effects will be very limited. They have been integrated in the workflow nonetheless in order to fully verify the integration of all geometric parameters (see section 3.5.3 below). It should be also noted that if one want to aggregate successive scans in time and space it is necessary to take into account the long-term changes in the attitude of the aircraft.

Adaptation to new acquisition hardware

The following analysis is performed on data that has been acquired with another system and this required adaptation of the algorithms designed for the MFLAME acquisition unit. The latter has a sampling frequency of

128MHz and an automatic aircraft speed compensation in the form of a spectral shift in the analog domain but the new unit samples at 500MHz directly without taking into account the speed of the aircraft. This means that the speed has to be accounted for in software. As a consequence a longer processing time is expected because the Fourier transform must be applied to a wider spectrum. The new unit also performs the acquisition continuously and the different scans are therefore not isolated. A specific procedure had to be developed to detect the scans begin and end from the continuous data flow. At last, the stability of the laser output power is not guaranteed and the laser had to be hand-tuned to adapt to the lower air pressure. To this effect the pretrigger pulse was monitored to discard the low-power pulses which contain high level harmonics.

Adaptation to the lower SNR

The aerosol concentrations at the altitude of the flight tests (6000-9000ft) are much lower than on the ground, leading to a lower reflected and received power. As the noise level at the detector is constant, the global SNR of the sampled signal will decrease. Given the SNR observed in the flight test data one must carefully integrate the noise component in the spectrum model. For the first gate close to the observing aircraft (800m) the noise spectrum can be considered flat with respect to the amplitude of the main return peak (fig.3.24a). These spectrums are similar to the ones of the ground based tests. However, as the analysed gate is located further away for the aircraft the SNR lowers rapidly and the larger noise component amplitude can not be considered constant across the spectrum anymore. At a distance of 2360m a typical spectrum looks like the one presented on fig.3.24c. The higher noise level make it impossible to detect the secondary signal component, and is even leading to spurious noise in the spectral width of the main peak.

As explained in section 3.3.1, the shape of the noise spectrum is constant over the scans, letting only an amplitude factor vary. It is thus logical to try to estimate this spectrum and use it in the same way as the ground tests. However, no specific noise-spectrum shots were taken during the IWAKE flight tests. The reason behind this is that the laser is started after take off and that no such tests could be performed in flight. Tests done before take off would not accurately describe the noise spectrum, as the laser will change its characteristics depending on the altitude, vibrations,... For these reasons the noise spectrum is estimated with the last samples of the LOS. These samples correspond to a distance of a little more than

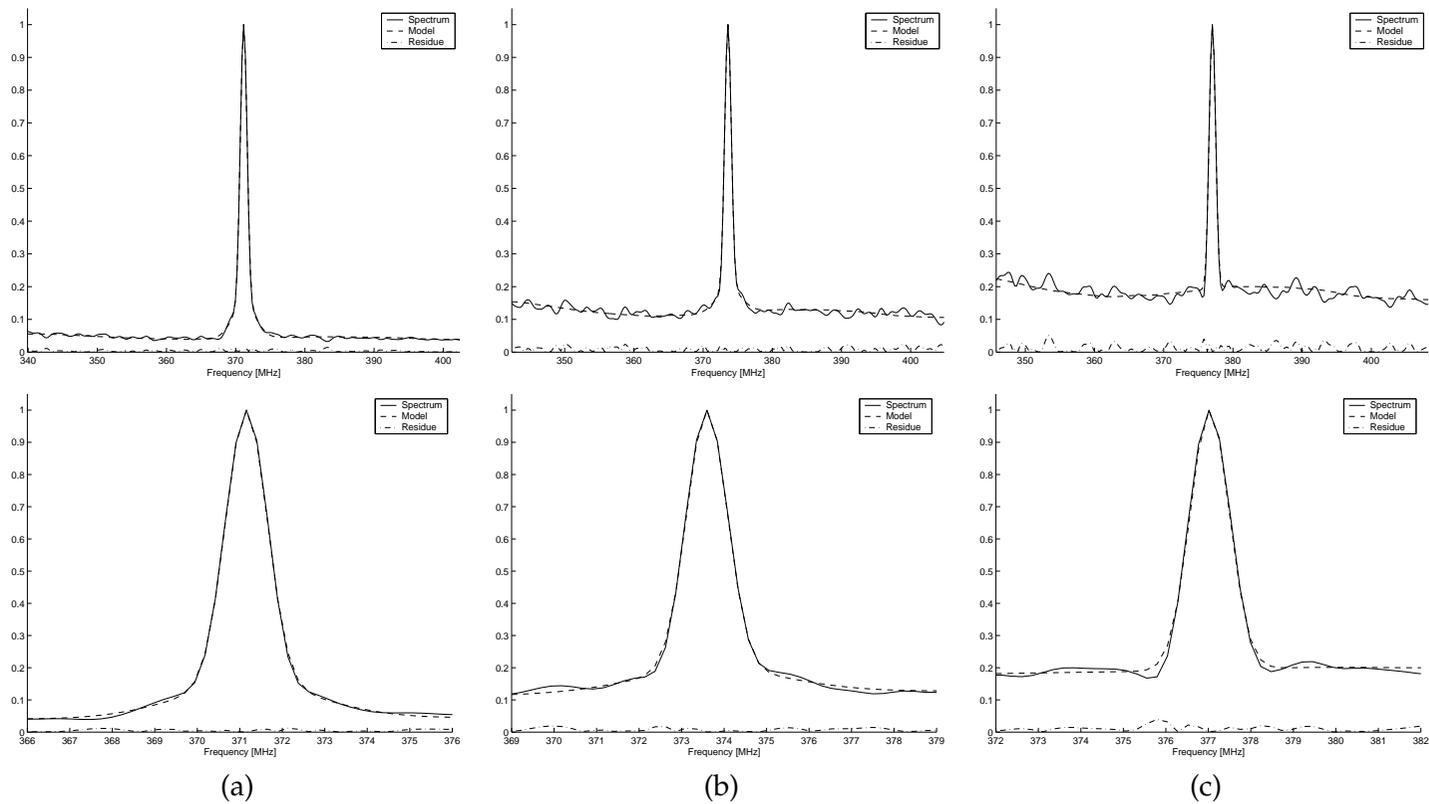


Figure 3.24: Spectral fits of on-board signals with a two-primitives model for range gate 2 (a, 920m), 8 (b, 1640m) and 13 (c, 2360m)

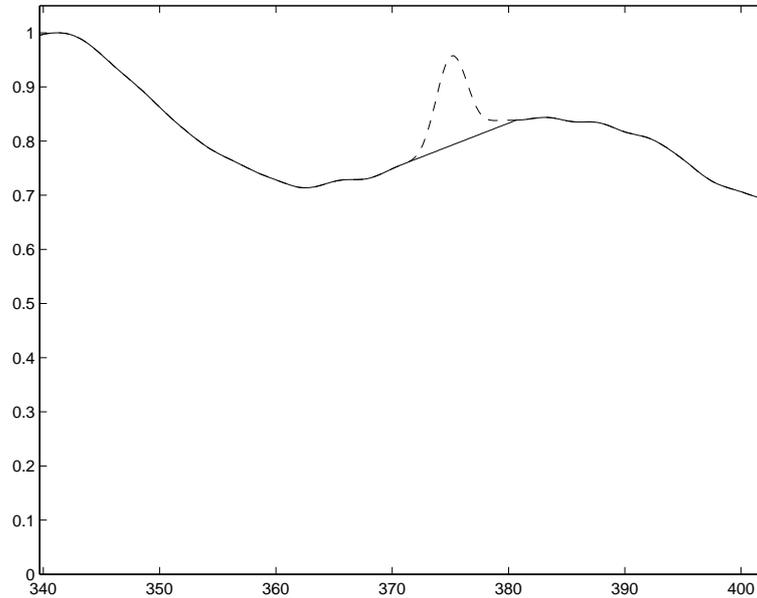


Figure 3.25: *The noise spectrum estimated with the last samples of the line of sight*

5km from the aircraft. This is just outside the range of the lidar and the data should therefore only contain noise. The noise spectrum is estimated for every LOS and then averaged over the whole scan to obtain a stable spectrum. An example of noise spectrum is shown on fig.3.25. The peak around 375MHz is due to a small residual signal found in the far field and is filtered by linear approximation. This approach is more time consuming than the off-line estimation done in the ground-based tests but it has the significant advantage to automatically adapt every six seconds to changes in the spectrum that might occur in flight.

At last, the lower SNR forces us to adapt two parameters specified for the ground tests: the gate length is now set to 120m to assure a better basis for the FFT and the resampling radius r is set to 0.7° to include more shots in the spectrum aggregation. As this radius is significantly higher the resolution of the range data is halved.

3.5.3 Shakedown tests

The proper integration of all geometric aspects was tested during a session of flights prior to the vortex detection tests in Toulouse. The effect of



Figure 3.26: *The Lelystad power plant at the end of the second flyby (52.5744N-5.5295E). Photo courtesy NLR*

the aircraft speed on wind measurement could not be tested as no leading aircraft was present to generate vortices. A test based strictly on distances and angles was devised instead which consists in using the lidar as a ranging device only. The test was carried out in the vicinity of the Schipol Airport (Amsterdam, NL). The target presented on fig. 3.26 is the power plant of Lelystad, located 52.5744N, 5.5295E. The distance from the airplane is inferred from the delay between the emission of the laser pulse and the reception of its reflected wave. This delay must be estimated for every shot from the LOS data. Fig. 3.27 shows a typical LOS signal received. The first 2300 samples are obtained directly from the laser output and consist in the pretrigger signal. The amplifier then switches to high-gain mode and the echo from the atmosphere (or the ground) is measured. As the acquisition runs continuously, the measured delay is simply the difference between the time t_s^L of the emission of the L^{th} laser pulse and the time t_e^L of the reception of its echo. The positions t_s^L and t_e^L of the pulses are obtained by applying a low-pass filter to the absolute value of the received signal

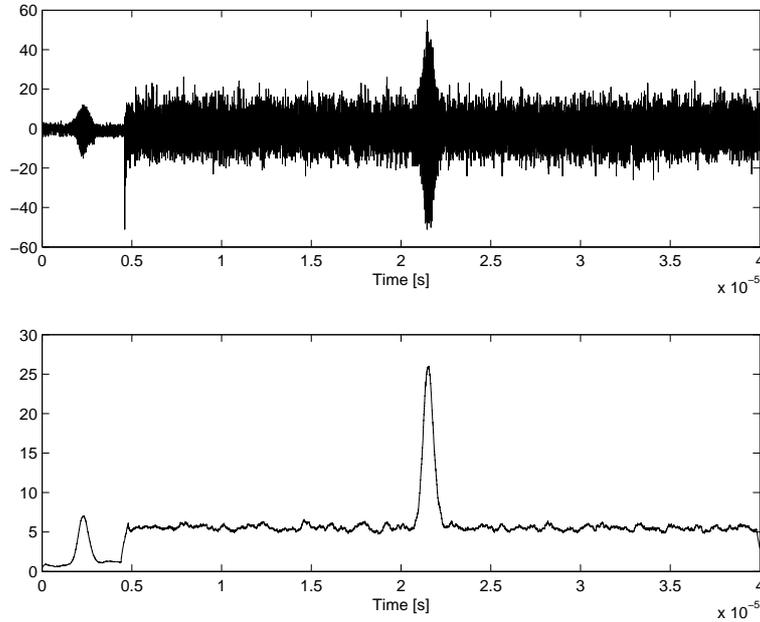


Figure 3.27: Ground signal return and envelope after low-pass filtering

(fig.3.27). The maximum of each peak is then calculated to sub-pixel accuracy with a weighted average, in a fashion similar to the estimation of the mean velocity in (3.11). With the time delay between the two maxima one can immediately obtain the distance D^L from the aircraft along the line of sight L :

$$D^L = \frac{(t_e^L - t_s^L) c}{2} \quad (3.27)$$

Once the distance is measured it must be translated in three-dimensional coordinates. To this effect the aircraft position and attitude as well as the mirrors positions are then used to provide an estimation of the location of the target point in a static earth-fixed referential. The position of the target point in the aircraft referential is first expressed in spherical coordinates $(\rho + \alpha, \gamma + \beta, D)$ where ρ and γ are the mirror angles and α and β are the system angles relative to the aircraft axis (see fig. 3.20). This is translated to Cartesian coordinates in the aircraft referential:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \cos(\gamma + \beta) \cos(\rho + \alpha) d \\ \cos(\gamma + \beta) \sin(\rho + \alpha) d \\ \sin(\gamma + \beta) d \end{pmatrix} \quad (3.28)$$

The Cartesian position $P_a = (x, y, z)$ is then rotated according to the aircraft pitch, roll and yaw angles (θ, ϕ, ψ) obtained from the aircraft flight data. The order of the rotations is the one commonly used in aerospace. This provides the position P_e of the target point in a referential of earth-fixed orientation attached to the aircraft:

$$P_e = (M_\phi M_\theta M_\psi)^{-1} P_a \quad (3.29)$$

where the rotation matrices M_ϕ , M_θ and M_ψ are defined as

$$M_\phi = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi) & \sin(\phi) \\ 0 & -\sin(\phi) & \cos(\phi) \end{pmatrix} \quad (3.30)$$

$$M_\theta = \begin{pmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{pmatrix} \quad (3.31)$$

$$M_\psi = \begin{pmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (3.32)$$

Finally, the GPS position of the plane P_{GPS} is added to P_e to obtain the earth-fixed referential coordinates P_f :

$$P_f = P_e + P_{GPS} \quad (3.33)$$

The two passes of the aircraft over the Lelystad power plant can then be aggregated since the target points P_f are expressed in the same earth-fixed referential. The result are shown on fig. 3.28 and fig. 3.29. The three-dimensional cloud of points shows that the structure of the power plant has recognisable features like the vertical chimneys and the surrounding lower buildings. The square shape of the island is also visible because no echo is measured from the specular reflections on the sea. The errors on the position of the points on fig.3.29 are in the order of 5 meters at a distance of 3 kilometres. This must be compared with the size of a lattice point for a gate at 2375m, which is around 10 meters. Since the error on the position is three times less than the size of a lattice point we can say that these shakedown tests for the integration of the aircraft position are successful.

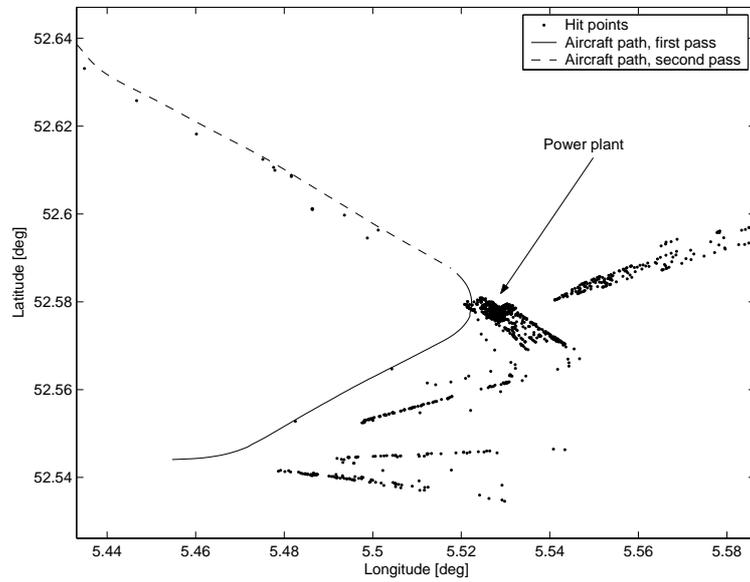


Figure 3.28: Flight paths around Lelystad with the overlay of reconstructed points

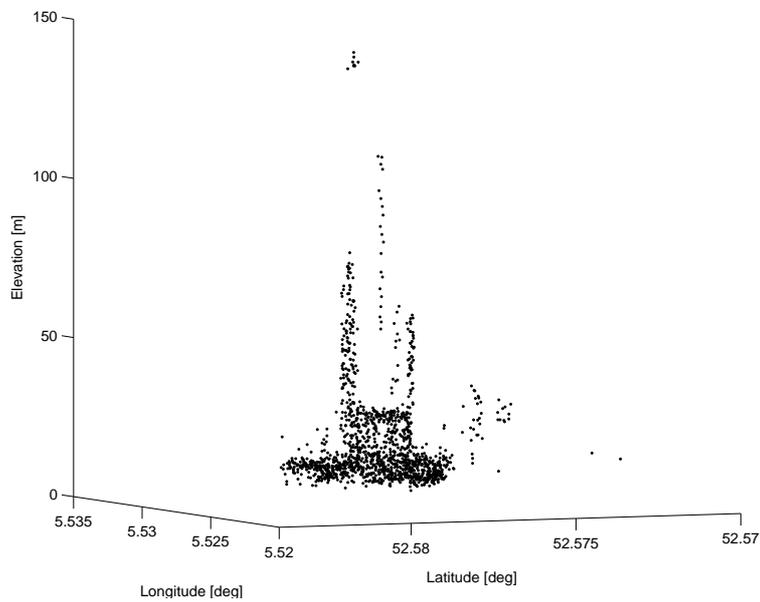


Figure 3.29: Lateral view from the northwest, along the axis of the second pass over the power plant

3.5.4 Flight tests

The final flight tests were conducted in the vicinity of Toulouse-Blagnac airport at the end of June 2004. The tests were split in two flights of one hour and a half with an A340-600 as generating aircraft. The first flight was performed on June 22 at an altitude of 9000 feet (details about the flights can be found in the Appendix B). The strategy to detect the wake vortices was to follow the generating aircraft at a constant relative altitude and track offset. The two parameters were set to maximise the probability that the measuring volume would intersect a vortex. A rapid analysis of the collected data did not provide satisfactory results so that the next flight program was adapted. Instead of keeping a constant relative position between the airplane, the flight of June 26 included a zig-zag route for the following aircraft. This new approach almost guarantees that the vortex will intersect the analysis volume although this also means that the vortex will be visible only for a few scans at each pass. The flight level was also lowered to 6000 feet to increase the aerosols because the rain of the previous day had significantly decreased their concentration in the atmosphere. The results shown in this section come exclusively from this second flight.

3.5.5 Airborne results and discussion

The signal processing algorithms described above have been applied to a selection of scans acquired on the second flight test. From this selection two scans are discussed below. The results of each model parameter is presented for the scan 07-23 (2GB-block number 7, scan 23) on fig. 3.30 and 3.31. Like the ground results, the closer gates are shown below and the distance to the airplane grows by steps of 120m as we progress towards the top of the page. The nearest gate is located at 800m and the farthest at 2120m. Their results are very comparable to the ones of the ground tests but a few differences should be pointed out.

We first note that the parameters are more noisy, which is expected given the lower SNR which is now at best 0dB. As noted in Section 3.4.2 the detection capability quickly degrades when the SNR is lower than -10 dB, which here corresponds to a distance of 1400m or gate 6. The detection field of fig. 3.31 confirms this -10 dB threshold as the first six gates are well rendered but starting from the seventh the results are noisier. Unfortunately the vortex appears in the range from 1200m to 1800m and only three gates (4 to 6) have an acceptable quality. The necessity of a good noise estimation is further illustrated by fig. 3.32 which was obtained us-

ing a simpler flat noise spectrum. The quality of the results clearly degrades much faster and no detection can be performed reliably.

Secondly, two vortices should be detected on the detection field of fig. 3.31 because the gate size being much larger than the vortex separation distance. This is the case for gates 4 to 6, although the strength of the two vortices are not similar. On those three gates we can also see that the vortices are closing as we advance towards gate 6, which actually denotes parallel vortices as the field covered by the scanning mirrors is actually a rectangular based cone. The location of the vortices on the side of the central mean speed deficit area is also coherent. The mean speed deficit is the trail of the fuselage of the moving aircraft and it is logical that the vortices are located on its sides since they have their source on the wingtips. Although the vortex location seems to be correct at this point, the larger sampling radius used for the flight tests data interpretation might have induced the higher spectral width because the aggregation of spectrums was performed on a transition zone where the mean speed has a wider distribution. However, the fact that the vortices have an asymmetric intensity as well as the absence of vortex detection on the upper boundary of the mean speed depression zone show that it is not the case. We can therefore be confident that vortices have been spotted. It should be noted, however, that there is no scientific certainty about this presence at this point and that more research with the flight data, vortex models and meteorological conditions will be necessary to achieve this goal.

The time evolution of the mean speed and the detection field around the scan 07-23 is shown on fig. 3.33 and 3.34. The results are coherent with the previous observations: only the 6 or 7 first gates are exempt of a significant noise background. We can also see that the two separated vortices are only visible on 07-23, thereby showing the limits of the detection.

The results for test scan 09-09 and its time evolution are shown on fig. 3.35 to 3.38. Although the SNR is slightly lower in the first gate, it does not drop as fast as test 07-23 so that in the end one or two more gates have an acceptable detection field. The most striking feature of this test is of course saturation zone in the two last gates. The traits of this zone are a very high SNR, a higher signal power and a low matching residue. The tempting explanation would be that either a cloud, the leading aircraft or a wind shear zone were hit by the laser and returned a much higher power. Looking at the mean wind speed we notice that the vortex trail on gates 5 to 8 indicate that the aircraft is higher than the scanning volume because vortices are moving downwards and start to appear on the top

of the gate. The flight plan confirms that for the last gates the aircraft is too far above to be the source of this anomaly. Another possibility is that a wind shear zone of a higher humidity level was detected by the lidar. The time evolution of the results presented on fig. 3.37 and 3.38 infirm this statement because the perturbation disappears suddenly from one scan to another while other features like the vortex indicate that the motion of the airplane normally yield a displacement of on tenth of a range gate width. If it is a wind shear it must have had either very strong lateral speed of 40m/s or vertical speed of 7m/s. Still, the possibility of a shear zone can not be totally excluded as the flight level was changed during the tests for the very reason that the atmosphere was too perturbed³.

The most plausible explanation remains that the analysis volume has touched the upper surface of the cloud layer. The temporal signal return indeed shows a very sharp peaks at the distance corresponding to gate 13 which can be related to a strong reflection on the cloud layer. More interestingly, we can see that the secondary return accounts for a much higher proportion of the spectrum (fig. 3.36) at the location of the cloud layer. This is an indication that the secondary return might have its origin in the presence of vapour or droplets in the atmosphere. Putting this information back in the perspective of the vortex detection, the fact that the secondary return is only present at the vortex locations is an indication that we could actually be tracking the vapour exhaust behind the aircraft, as suggested earlier. It is at this time the most plausible explanation for the presence of the secondary return.

At last, looking again at the different model parameters we notice that the only one which is not subject to a dramatic change in this zone is the spectral width of the secondary peak σ_s . This kind of information could also be used to detect other atmospheric hazards like wind shears.

3.6 Conclusions

In this chapter we have presented the spectral estimation of lidar signals in the goal of detecting wake vortices in an axial fashion. Several methods and variations have been tested on the data collected during both ground and on-board test campaigns. The current techniques found in the literature fail to detect the vortices reliably but the new approach of dual-primitive spectral modelling which is introduced succeeds to this task. It has two major advantages over the previous approaches. First it allows a

³from the flight log of Stephan Rahm, DLR

correct detection at lower SNR of -10dB, which proved to be mandatory for the analysis of the on-board signals. Secondly, the technique provides more information about the atmosphere and could be used to detect other events like wind shears or clouds, although this affirmation will require a more in depth analysis. On the other hand, the proposed approach has the very significant drawback of requiring a large computation power which might be difficult to integrate on-board. In this perspective, the new FIDELIO project⁴ will concentrate on hardware implementation of the algorithm presented here. The LASEF project⁵ will also profit from the latter and will further investigate the results of this research for ground-based fibre lidar applications.

⁴Accepted European Project proposition 12008 of call FP6-2003-AERO-1

⁵Belgian Walloon Region project, Direction Générale des Technologies de la Recherche et de l'Énergie, 2003

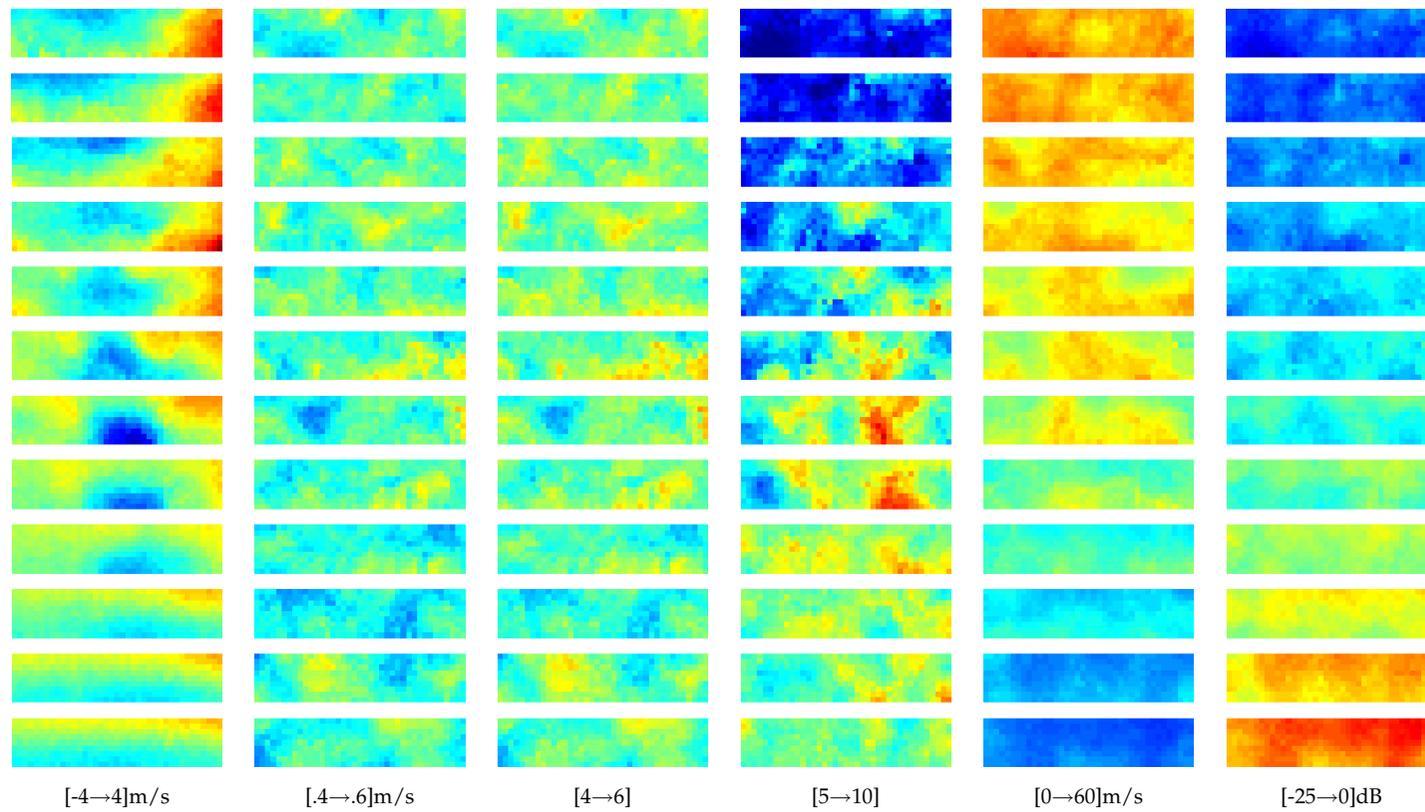


Figure 3.30: *I-Wake flight test: scan 07-23 with measured noise. Parameters shown are the main peak mean velocity, spectral width and energy, the global signal energy, the noise energy and the SNR. Note that the SNR is 10dB below the one found in the MFLAME ground tests (fig. 3.16).*

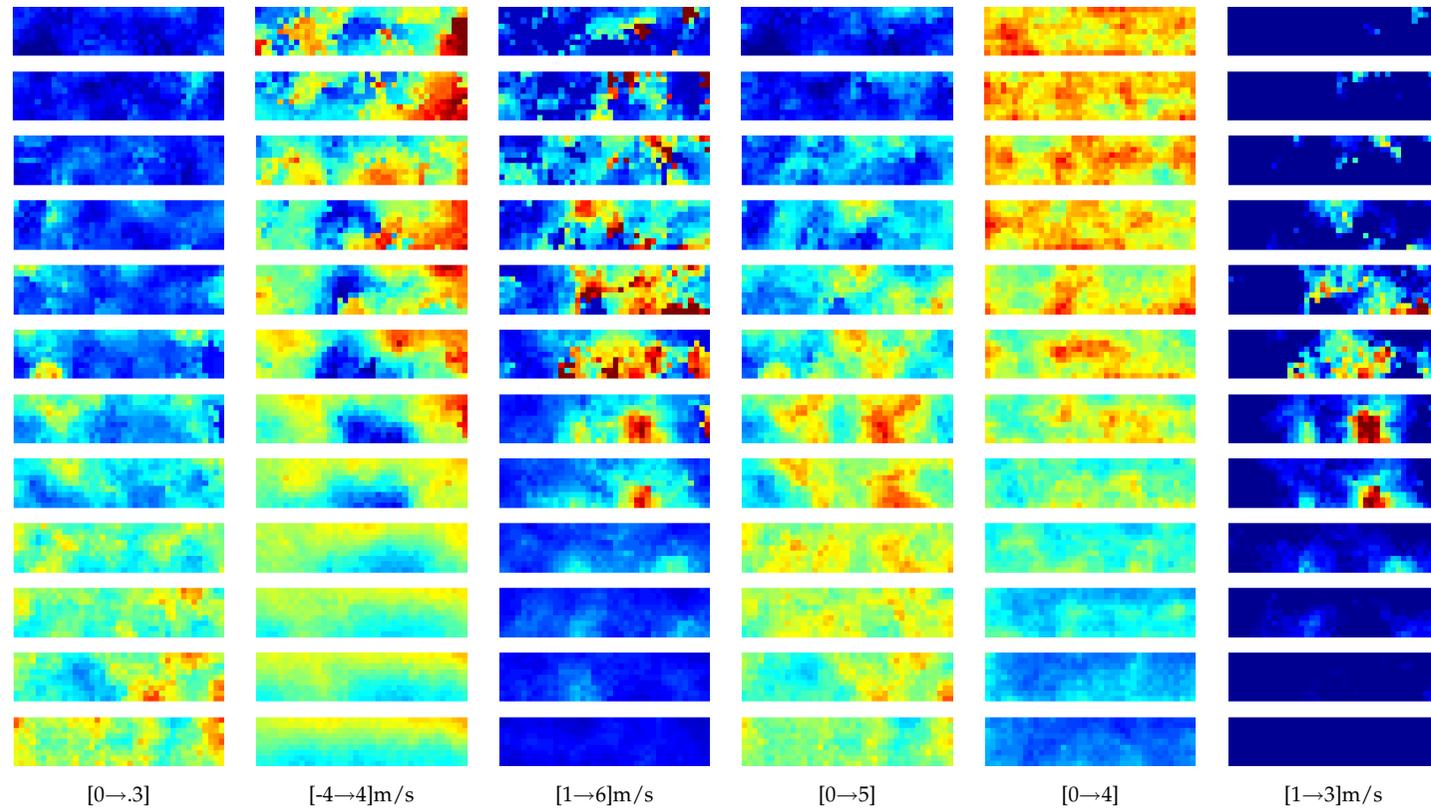


Figure 3.31: *I-Wake* flight test: scan 07-23 with measured noise. Parameters shown are the secondary peak level, mean velocity, spectral width and energy, the residual energy and the detection field

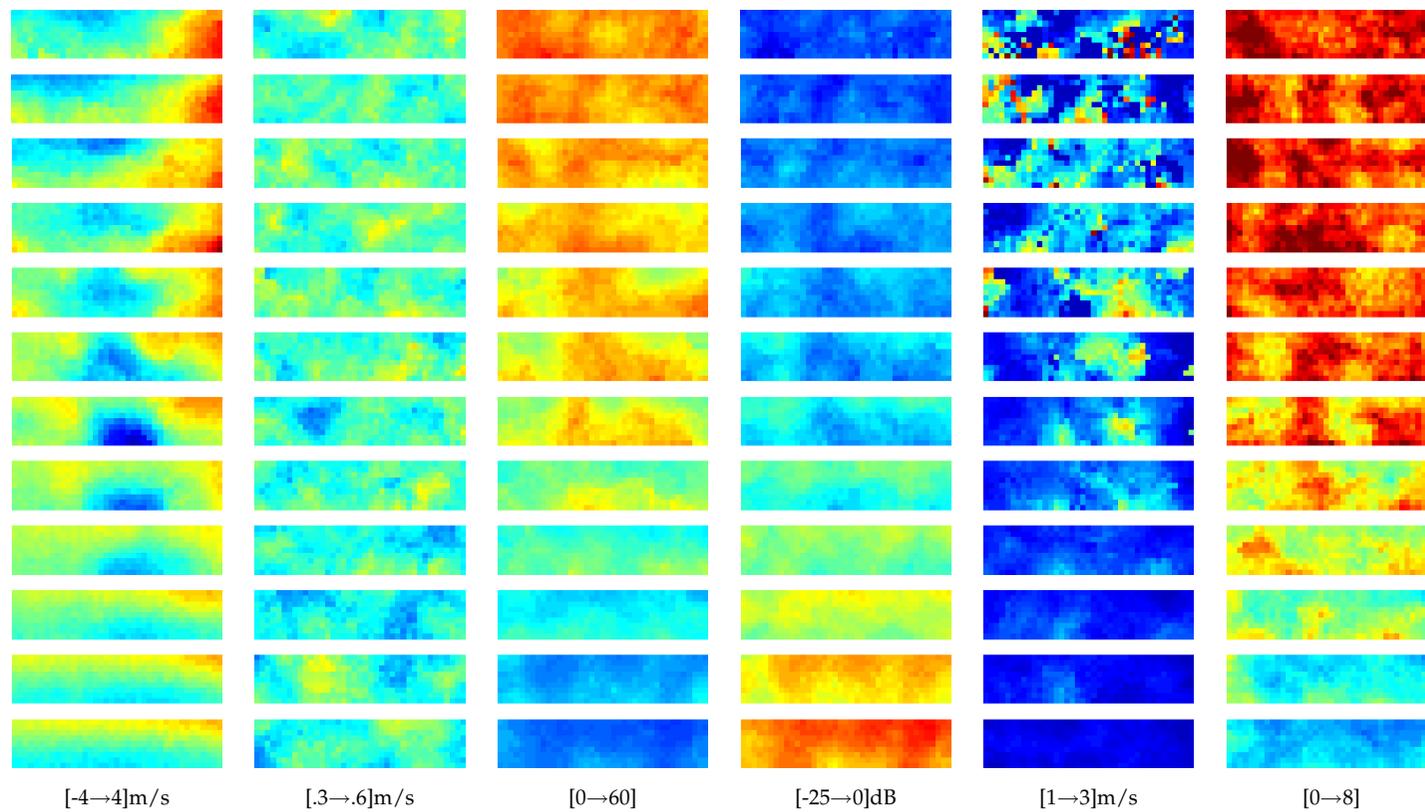


Figure 3.32: *I-Wake flight test: scan 07-23, flat noise for some parameters: main peak mean velocity, main peak spectral width, noise energy, SNR, secondary primitive spectral width and residue energy.*

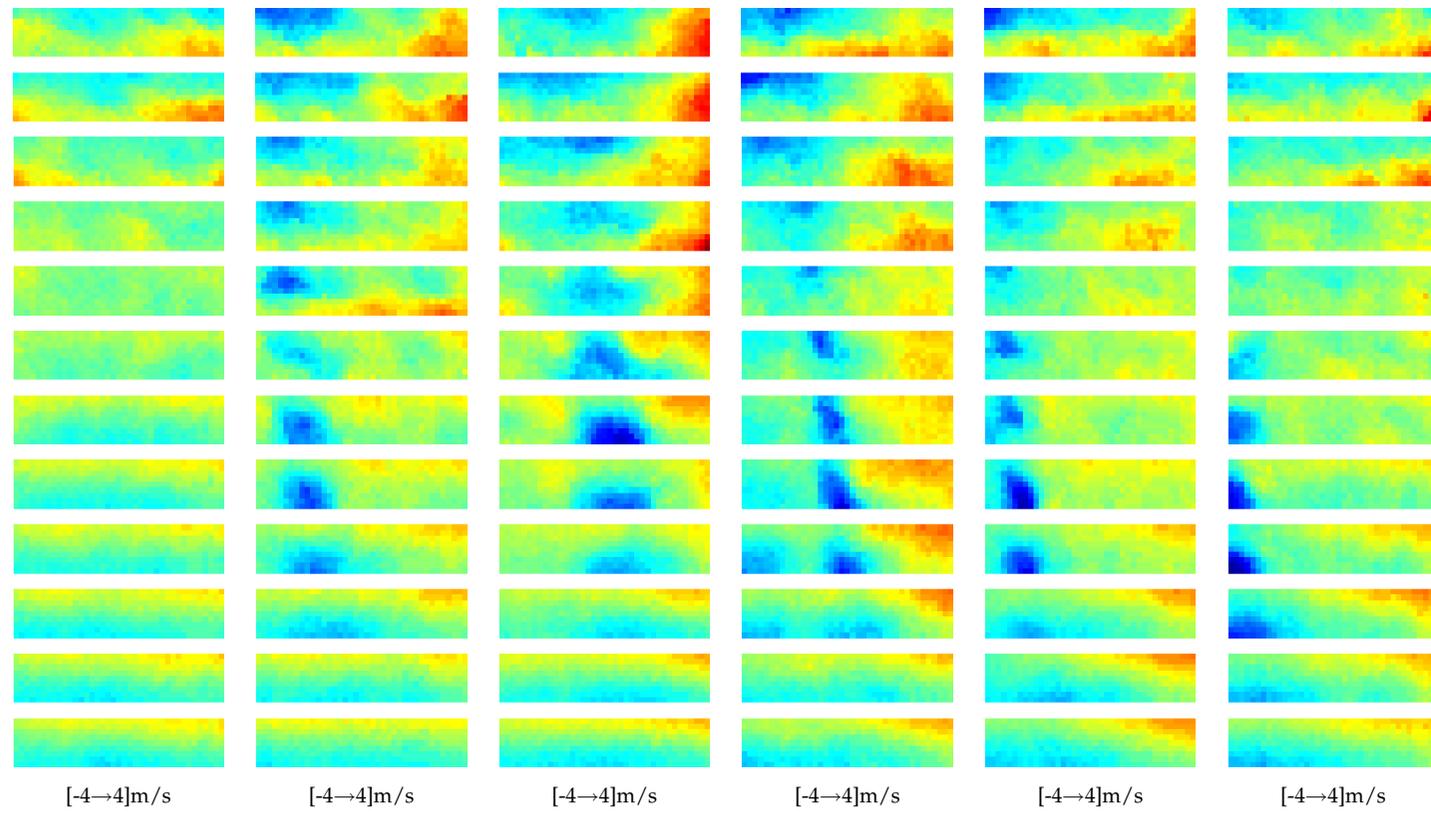


Figure 3.33: *I-Wake flight test: the mean velocity of the main primitive for scans scans 07-21 to 07-27.*

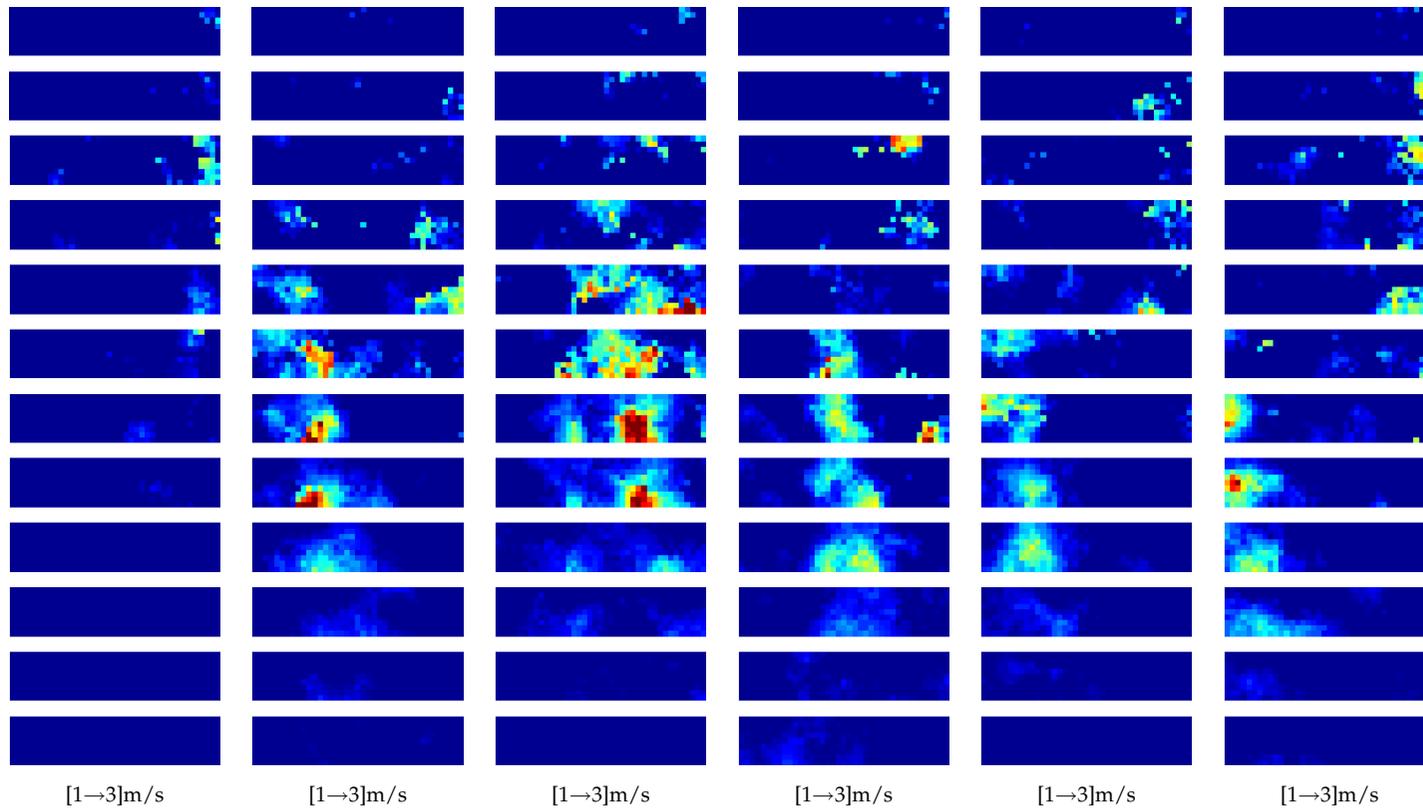


Figure 3.34: *I-Wake flight test: the detection field for scans 07-21 to 07-26.*

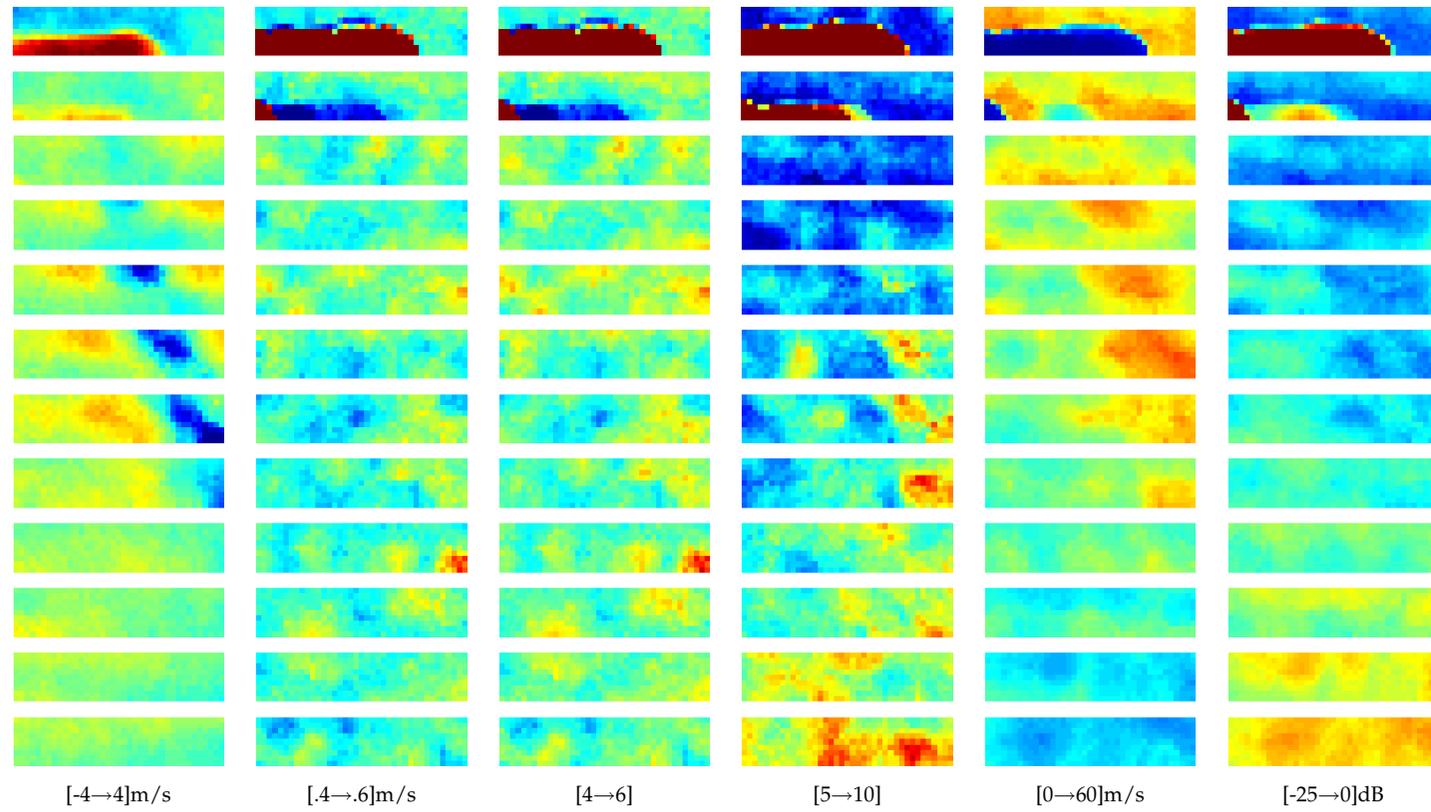


Figure 3.35: I-Wake flight test: scan 09-09, measured noise. Parameters shown are the main peak mean velocity, spectral width and energy, the global signal energy, the noise energy and the SNR

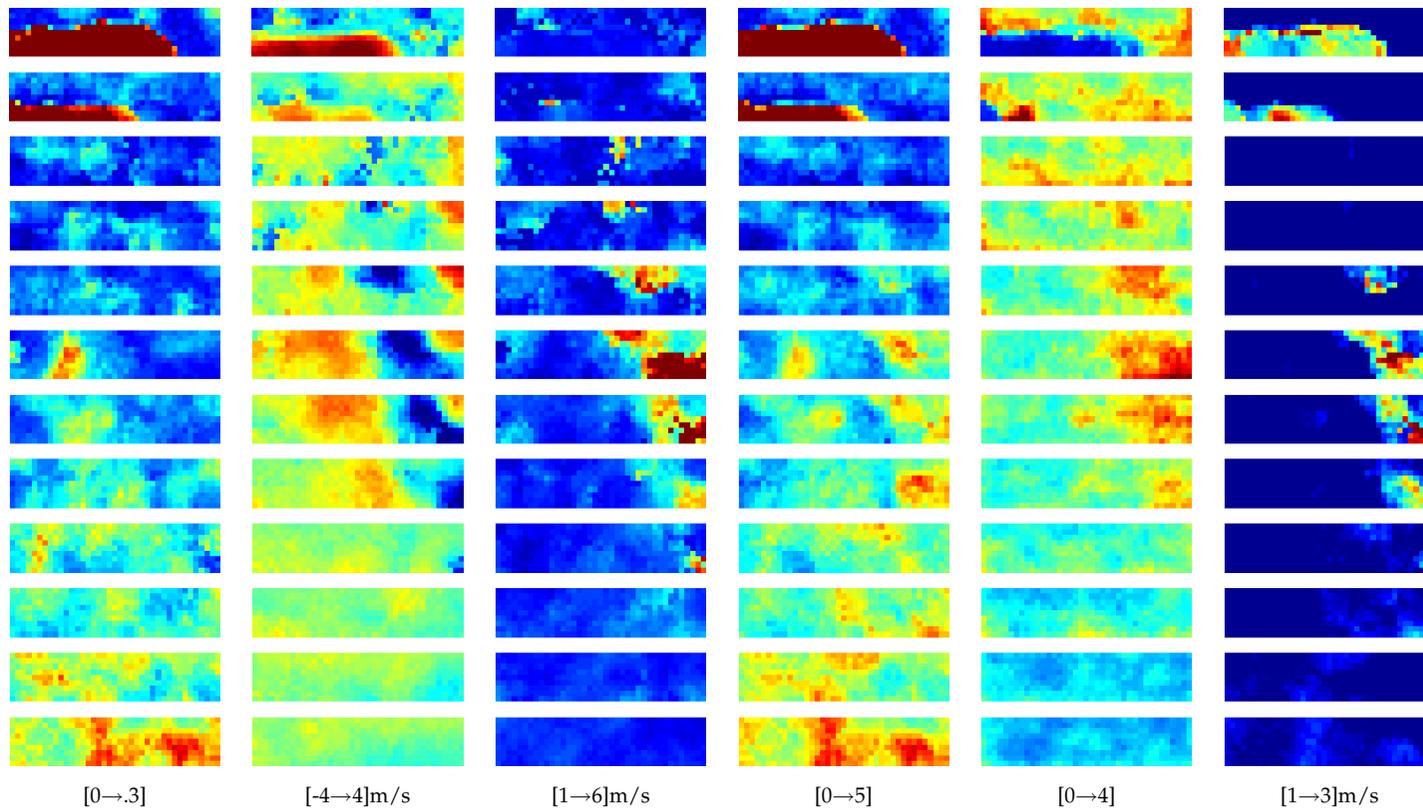


Figure 3.36: *I-Wake flight test: scan 09-09, measured noise. Parameters shown are the secondary peak level, mean velocity, spectral width and energy, the residual energy and the detection field*

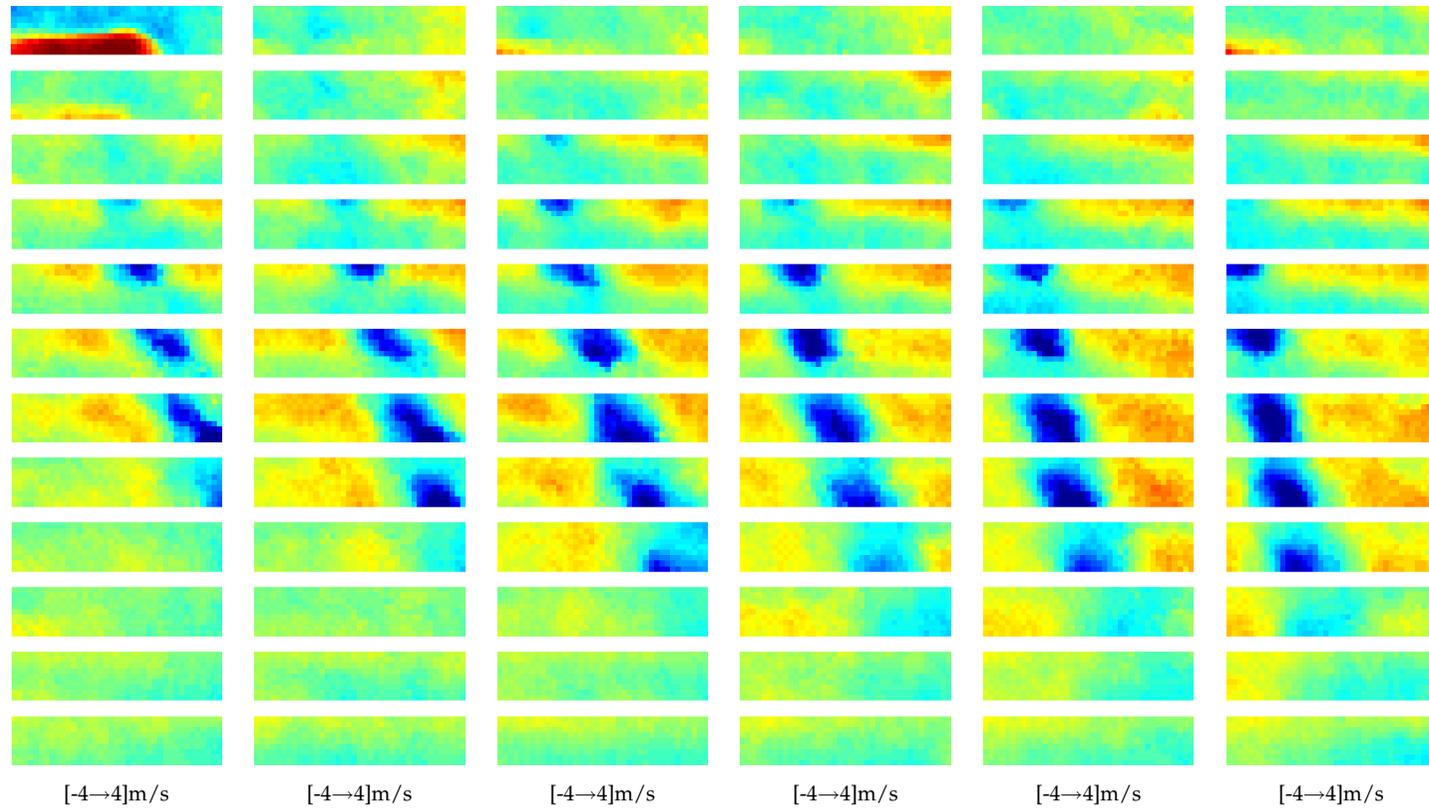


Figure 3.37: *I-Wake flight test: the mean velocity of the main primitive for scans scans 09-09 to 09-14.*

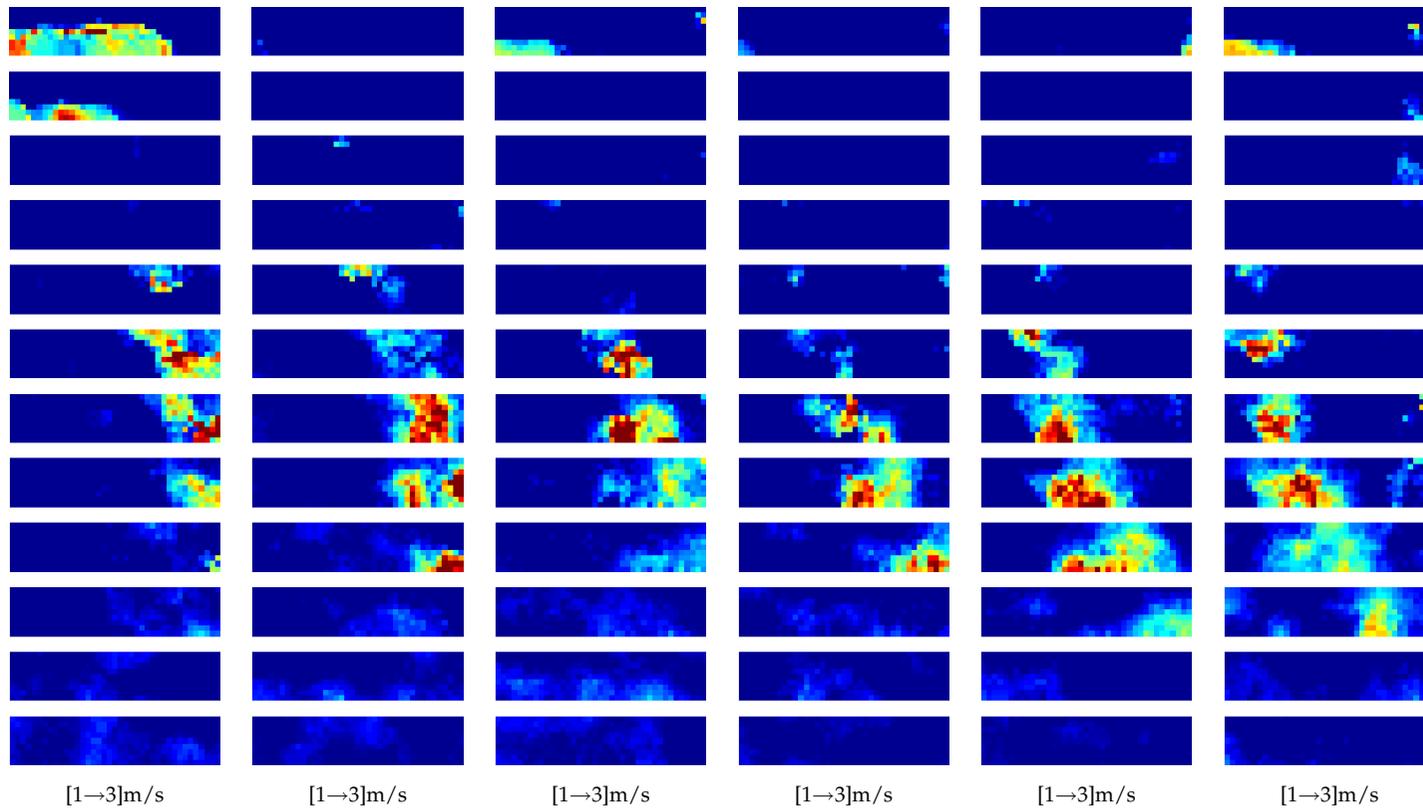


Figure 3.38: *I-Wake flight test: the detection field for scans 09-09 to 09-14.*

Conclusion

We have presented three photogrammetry techniques that successfully capture different short-lived events of increasing complexity. First we have shown that the first simple unidimensional motion of vehicles on a highway can be analysed with accuracy by a set of two line-scan cameras. This new technique also provides a large number of other information about the traffic, making it a very good choice for both traffic management and law enforcement. The high commercial potential of this approach will lead to a patent application in the near future. The second theme developed shows how a close relationship between the fluid mechanics and the vision techniques can lead to a better confidence in the measurements compared to what would be obtained by each method individually. The recent advances in the theoretical formulation of the antidunes behaviour will have to be compared with surfaces like the ones presented here, although the unstable and evolutive character of this phenomenon will make this task far from trivial. At last, a new model for the spectrum returned by an on-board lidar is compared with several other spectral estimation approaches. The higher performance of the model enables the detection of the aircraft wake vortices that were sought. Further developments of the technique will be made during two following projects and hopefully lead to the commercial exploitation of these results.

Appendix A

3D matching with perspective distortions

A.1 Introduction

One of the most difficult problem in 3D image reconstruction from several images is the matching of physical features between two images in order to use triangulation for the recovery of 3D coordinates. Several techniques have been developed to solve this ill-posed problem of correspondence. These are usually based on image features (e.g. points, lines) and/or image regions [87]. While most feature-based approaches yield sparse three-dimensional data, area matching has the significant advantage to provide dense results but also supposes to overcome the specific challenge of perspective distortions. Our method, based on area matching, proposes to solve for these perspective effects not by eliminating them but by using them as primary source of information.

Perspective distortions are related to the local surface orientation of the 3D object relative to the position of the cameras. These position and orientation differences have an effect on the image blocks used for correlation, as shown on fig. A.1: a rectangular image block in the first image I_1 will not generally correspond to a rectangular block in the second image I_2 . Not only can the block be rotated or rescaled, but its shape will be transformed into something that depends on the camera (orientation, position and model), and the relative positions of the cameras with respect to the object. Not taking these distortions into account will inevitably lead to matching inaccuracies. Significant perspective effects can even make a block matching process close to impossible, especially in fine structured

areas of the image where little changes in block shapes will prevent features to correlate.

Methods that solve the 3D reconstruction problem while taking those perspective distortions into account have started to arise in the 90s. Among others, we would like to stress the work of Devernay [7] [88] who retrieves the local tangent plane with an analysis of disparity differential, followed by an optimisation based on the rectified images content. His approach elegantly fine-tunes the results from a classic correlation but does not allow large perspective distortions since the first correlation step does not take the local orientation into account. The present work [89] can be seen as a generalisation of this approach where we operate directly on the original image content, avoiding rectification or an initial classic block matching step. While we can expect the generalised approach to be more time consuming, it is also able to deal with strong perspective differences, something quite difficult for methods that use a fronto-parallel scene assumption in a first step.

A.2 Local planarity constrain

The underlying idea of our approach is to cast all the geometric aspects in 3D, therefore avoiding any 2D hypothesis in the image plane. Let us associate an image plane I_i to the i^{th} camera of the system. Fig. A.1 shows the case of a two camera system ($i = 1, 2$). For every 3D point $P = (x, y, z)$, we can define a projection line (or *ray*) R_i^P , for each camera, as the line passing through P and the focal point of that camera. This projection ray forms the image point $p_i^P = (m_i^P, n_i^P)$ of P at its intersection with the image plane I_i . It is important to note that since we are working in a projective space, the ray R_i^P is equivalent to its associated projected image point p_i^P [8]. In other words, each ray has only one image point, and vice-versa, so that we can write with little abuse that

$$I_i (R_i^P) \equiv I_i (p_i^P) . \quad (\text{A.1})$$

This relation has the interesting consequence that most of our further reasoning can be done in 3D using rays, instead of working on image coordinates.

Suppose now that the two cameras of our stereo setup are calibrated so that both extrinsic and intrinsic parameters are known. Furthermore, we assume that the equation of the plane L that is locally tangent to the object around P is known. The intersection of the ray R_1^P with this plane

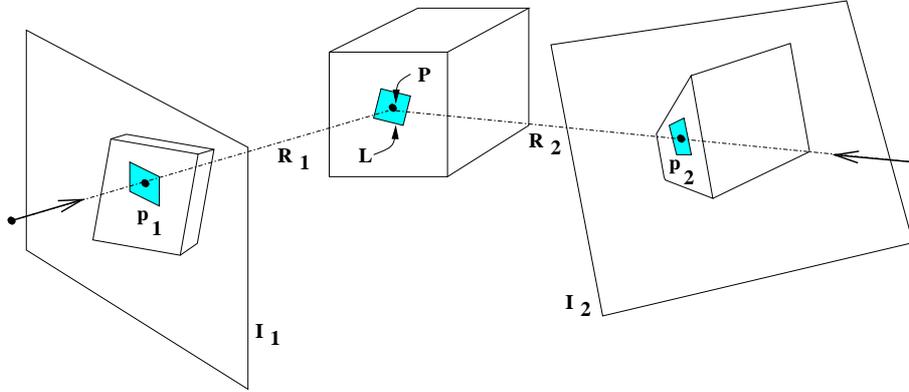


Figure A.1: *Geometric aspects*

will give the 3D point P . Once P is known we can exactly determine R_2^P from the *extrinsic* parameters of the camera. One can thus associate a ray $R_{2,L}^P$ for I_2 to every pair of ray R_1^P and plane L .

Our method for recovering L is based on a modified optical flow, block-matching algorithm (BMA). We wish to compare a region around p_1^P with image points in I_2 , but without imposing any region shape or scale in I_2 . Let M be an estimate for plane L . We can project on plane M a neighbourhood of n image points of I_1

$$p_{1,j} \equiv R_{1,j}, j = 1..n, \quad (\text{A.2})$$

as shown on fig. A.2, (1). The intersection of this bundle of rays with the plane M yields a set of 3D points $P_{1,j,M}$. Using the second camera extrinsic calibration data we can obtain the rays $R_{2,j,M}^P$ (fig. A.2, (2)). Comparing the pixel intensities of these two neighbourhoods $R_{1,j}$ and $R_{2,j,M}^P$ using a sum of absolute differences estimator (SAD) provides a correlation information associated to plane M for an initial image point p_1 :

$$SAD(p_1, M) = \sum_{j=0}^n |I_1(R_{1,j}) - I_2(R_{2,j,M}^P)|. \quad (\text{A.3})$$

Finding the minimum of the SAD for M by varying the plane parameters yields an estimation of the locally tangent plane L (fig. A.2, (3)). Note that this problem is significantly more complex than the classical stereo: M has three degrees of freedom while there is only one in classical stereo.

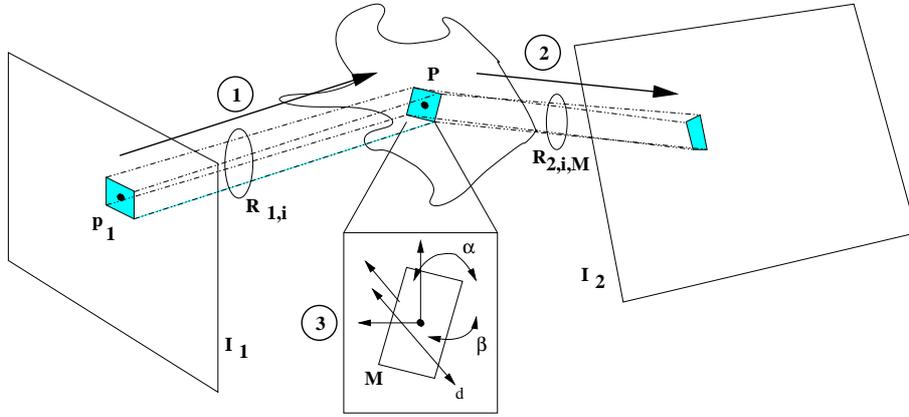


Figure A.2: *Projection of the neighbourhood and optimisation of the tangent plane*

A.3 Implementation

The locally tangent plane L has a 3D equation that can be written as:

$$L \equiv ax + by + cz + d = 0. \quad (\text{A.4})$$

This equation is homogeneous and has four degrees of freedom that can thus be reduced to three without loss of generality by adding the following constraint:

$$\| (a, b, c) \| = 1. \quad (\text{A.5})$$

The triplet (a, b, c) being normalised, we can express it in polar coordinates as $(\alpha, \beta, 1)$, yielding the parameters triplet of the tangent plane L : (α, β, d) . This formulation separates the plane parameters in two classes: orientation and offset. As comparison, a classic block matching process for stereo would only optimising for d , leaving the orientation fixed and equal to that of the image plane.

This is the basis of a two-step optimisation: first find P with a classic single degree of freedom optimisation along the ray R_1 using d , then find the final plane equation by optimising on the three parameters (α, β, d) . Because L is at first unknown, finding P with a classic BMA process (like Devernay) will not work in regions of intense perspective distortions. This is precisely where we aim to provide better results so that we must somehow integrate M in the first look for P .

This chicken-and-egg problem can be partially solved with an hierarchical approach [90]: plane orientations obtained at lower resolutions are

used for a better search along the ray R_1 . We thus shift the plane estimate along the ray R_1 and look for the best correlation (α and β fixed estimates - d varying). Once this first estimation of P is found we start a full optimisation on the three plane parameters.

Although this approach is simple and tempting, it will work well only if the scene has a smooth depth map. If large discontinuities are present, the hierarchical mechanism will fail in those regions.

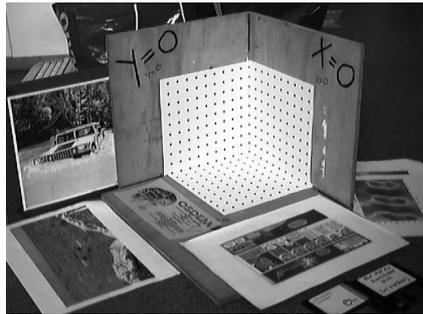
We found that a better stability can be achieved by testing several orientations during this first search. The processing time becoming quite large in that case, we process the image pair at each resolution with a region-growing BMA in order to propagate the first search result to neighbours that will likely have the same local orientation. The criterion used for region growing is that L should remain approximately constant across a region so that we can directly optimise on the three plane parameters with a Nelder-Mead technique [91], [92].

A.4 Results and discussion

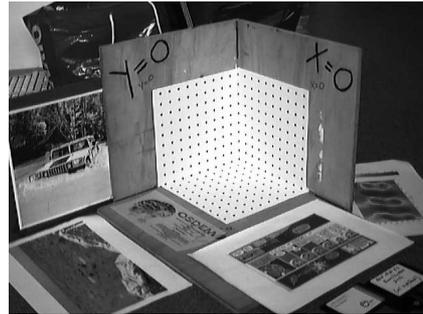
The system has been tested on image pairs like the one shown on fig. A.3a-b. This image pair contains the calibration target consisting in a trihedral with calibration dots on each of the principal planes $X = 0$, $Y = 0$ and $Z = 0$. Two dimensional features have been added on (or close to) these principal planes. Let us represent the results as four images for the four coefficients a, b, c and d of (A.4). In the area of the image corresponding to $X = 0$, which should theoretically be as $(a, b, c, d) = (1, 0, 0, 0)$, we expect to see high values (white) for the a coefficient image, while other images should remain relatively dark in that zone (intensity values near zero). $Y = 0$ and $Z = 0$ are described in the same way.

Fig. A.3c through A.3e show the three components (a, b, c) of the recovered planes, at a reduced resolution. As expected these images show almost only black or white pixels since the scene consists almost only in principal planes. These have been correctly extracted, and their edges are in general well defined. Sometimes, however, the region growing leads to incorrect results, like plane $X = 0$ which extends in a featureless zone.

Regions defined by the region-growing BMA are shown on fig. A.3f. In general, the algorithm spreads correctly on plane surfaces, gaining much processing time and stability.



(a)



(b)



(c)



(d)



(e)



(f)

Figure A.3: Results from a test image pair

A.5 Conclusions

We have presented a method for recovering locally tangent planes of a scene from a multi-camera setup. The first implementation produced some encouraging results, although it should still be seen a basic platform for further experiments. It is clear at this point that two major drawbacks must be overcome: the computational power required is too large and the overall quality of the results must be enhanced. In that perspective, we are now working on using image properties that are not subject to perspective distortions in order to guide the matching process.

Appendix B

Flight Logs

The following flight logs have been created by Stéphane Vaux of Airbus during the two IWAKE flight tests in Toulouse on June 22 and 25 2004. They are reporting on the events of the flight of the A340-600. Following the logs are the flight data recorded in the a340-600

AI/EV	Flight Engineer Log	A3456	S/N 360
CAPT : RM		F/O : ES	
Date : 22/06/2011		Fit Nr : F460	
FTE1 : SV		TFE : —	
FTE2 : —		Crew Obs : —	
TECH		Orig	
OBS		Orig	
DATA ANALYSIS : YES / NO		PAX Nr See PAX manifest (6 pers)	
Loading Plan (kg) Lb		Fuel before Flight	
N° : 62	Outer L 4880	TK 1 15420	TK 2 14560
ZFW CG : 24.0 %	CTR 300	TK 3 14580	TK 4 15420
ZFW : 158225	TOT FUEL 70100	Trim 40	R.C.T. —
FUEL : 70100	Fuel after Flight		
RAMP WT : 258375	Outer L 4820	TK 1 14200	TK 2 14000
RAMP CG : 24.5 %	CTR 0	TK 3 10340	TK 4 14620
Outer R 4900	Outer R 4820		
Water Ballast	TOT FUEL 54500	FUEL USED 15780	Trim 40
Fwd 1 2 3 4 5 6 7 8	R.C.T. —		
Aft 1 2 3 4 5 6 7 8	ACT —		
Take Off 06:06 G	(Dry) Wet - Soak - Snow - Icy	Landing 08:21	(Dry) Wet - Soak - Snow - Icy
ARPT : LFBO	RWY : 14R	ARPT : LFBO	RWY : 14R
Wind 150°/11 kt	Temp/DP 19°/16°	Wind 130°/11 kt	Temp/DP 21°/16°
QNH 1017	QFE 999	QNH 1017	QFE 999
Visi 10000	Clouds F. 5000 ft	Visi 10000	Clouds 5000 ft
A/C times	MAX FL 030	TOTAL	Location of mvts outside destination Nb
OFF Block 07:37	T/O 07:46	1	Ldg —
ON Block 09:15	LDG 09:09	0	GA —
Block Time 1:35	Fit Time 1:25		
Night time	Engine times		Total time/Starts
Eng 1 07:34 09:15			1:40 1
Eng 2 07:34 09:15			1:40 1
Eng 3 07:36 09:15			1:40 1
Eng 4 07:36 09:15			1:40 1
APU 06:54 07:37 09:10 09:16			0:45 2
Prog 0540	Test ref 09:05 Run	Time sharing 1:35	Purpose / Remark
105	00 FOD 1	1:35	In water research project



EV	A3 60	S/N	Flt nr	Engineer	Page
		360	F460	SV	1
					Snag pointer
06:54	FTI on (FM prog. 4 ⁵)				
06:54	APU start				
07:26	Pushback				
07:34	ENG 1+2 start				
07:36	ENG 3+4 start				
07:37	APU s/D				
07:37	Brakes off				
07:40	F/CIL check FRIM OFF				
07:41	F/cil check normal with FRIM ON				
07:46	T/O MR CONF3 FLEX 50° THIS = -3.7° UP 140/147/148 257.9°/24.39° Wind = 130°/9 kt				
07:55	CONF Full/UP, 154 kt, 9041 ft, 254.7°/24.79° Close formation with the STATION				
07:58	End of calibration.				
08:04:25	Interception of AFRIC 3 /ASPE1 route + CONF Full/BN 154 kt				
08:10	start of measurement. HDG = 240° Wind = 249°/17 kt				
08:22	Wind = 255°/17 kt				
08:24	Turn on Rb side 180°, and back on the track to AFRIC 3 .				
08:27	On Track to AFRIC (70 km) - HDG = 060° Wind = 245°/19 kt				
08:37	Wind = 238°/19 kt				
08:48	End of measurement. Turning to HDG = 270° and accelerating to CONF CLEAR/UP, 250 kt -				

AIRBUS																	
A/EV		Flight Engineer Log						A3456		S/N 360							
CAPT : <u>VSS</u>		F/O : <u>MAG</u>						Date : <u>25/01/04</u>									
FTE1 : <u>SV</u>		TFE : <u>/</u>						Flt Nr : <u>FL452</u>									
FTE2 : <u>/</u>		Crew Obs : <u>/</u>						F-Flight / R-eng Run / T-Taxi / S-Static									
TECH <u>/</u>		Orig		OBS		Orig		DATA ANALYSIS : <u>YES</u> / NO									
				<u>E. DEKKER</u>		<u>(NLR)</u>		PAX Nr <u>/</u>									
				<u>n. VERBEEK</u>		<u>(NLR)</u>											
Loading Plan			(Kg) Lb		Fuel before Flight												
N° <u>42</u>				Outer L <u>4840</u>		TK 1 <u>14300</u>		TK 2 <u>15140</u>		CTR <u>0</u>		TK 3 <u>15100</u>		TK 4 <u>16200</u>		Outer R <u>4320</u>	
ZFW CG <u>24.5</u>		%		TOT FUEL <u>69900</u>		Trim <u>40</u>		R.CT <u>/</u>		ACT <u>/</u>							
ZFW <u>187870</u>				Fuel after Flight													
FUEL <u>69900</u>				Outer L <u>4800</u>		TK 1 <u>14060</u>		TK 2 <u>14780</u>		CTR <u>0</u>		TK 3 <u>14100</u>		TK 4 <u>15200</u>		Outer R <u>4880</u>	
RAMP WT <u>257730</u>				TOT FUEL <u>54800</u>		FUEL USED <u>15130</u>		Trim <u>20</u>		R.CT <u>/</u>		ACT <u>/</u>					
RAMP CG <u>24.7</u>		%		Water		Fwd		1 2 3 4 5 6 7 8		Aft		1 2 3 4 5 6 7 8					
				Ballast													
Take Off <u>LFBO</u>		Dry - Wet - Soak - Snow - Icy		Landing <u>LFBO</u>		Dry - Wet - Soak - Snow - Icy											
ARPT <u>LFBO</u>		RWY <u>32L</u>		ARPT <u>LFBO</u>		RWY <u>32L</u>											
Wind <u>320/6 kt</u>		Temp/DP <u>18/16</u>		Wind <u>320/6 kt</u>		Temp/DP <u>18/13</u>											
QNH <u>1020</u>		QFE <u>1001</u>		QNH <u>1020</u>		QFE <u>1002</u>											
Visi <u>10000m</u>		Clouds <u>B3000ft</u>		Visi <u>10000m</u>		Clouds <u>B3300ft</u>											
A/C times		MAX FL <u>080</u>		TOTAL		Location of mvts outside destination		Nb									
OFF Block <u>07:27</u>		T/O <u>07:35</u>		1 Ldg													
ON Block <u>08:04</u>		LDG <u>08:56</u>		0 GA													
Block Time <u>1:35</u>		Flt Time <u>1:20</u>															
Night time <u>/</u>		Engine times							Total time/Starts								
Eng 1 <u>07:22 08:04</u>										1:45		1					
Eng 2 <u>07:22 08:04</u>										1:45		1					
Eng 3 <u>07:23 08:04</u>										1:45		1					
Eng 4 <u>07:23 08:04</u>										1:45		1					
APU <u>08:44 07:25</u>		<u>08:57 08:06</u>								0:30		2					
Prog <u>0140</u>		Test ref <u>0210</u>		Run		Time sharing		Purpose / Remark									
<u>105</u>		<u>00FOOD</u>				<u>1:35</u>		<u>I-WAVE research project</u>									



EV	A340	S/N	360	Fit nr	F462	Engineer	8V	Page	1
									Snag pointer
06:38	FTI ON (Fit prog 4 ^e 5)								
	CWS limit = 257790 kg / 24.7%								
07:16	APU start								
07:18	Pushback								
07:22	ENG 1+2 start								
07:23	ENG 3+4 start								
07:25	APU stop								
07:25	F/CIT check with PRIMA OFF								
07:25	idle with PRIMA ON								
07:26	Blocks OFF								
07:26	PACK1 = OFF → PAIR AIR outlet from PACK1 closed.								
07:27	Blocks off								
07:28	A/SKID reset → OK.								
07:35	T/O 32L CONF3 FLEX 50° TAS = 3.2° UP PACK1 OFF 140/142/148 257.4°/24.6% Wind = 310°/6 kt								
07:56	FLOED / CONF Full / UP / 154 kt 253°/25.5% Start of measurement above T80, HDG = 096° Wind = 255°/13 kt								
08:02	Start of measurement → Wind = 255°/23 kt, HDG = 098°								
08:06	Wind = 266°/13 kt, HDG = 096°								
08:08	PACK1 ON								
08:12	HEAT 2 C/B engaged -								
08:15	Climb to FLO80, same HDG = 097°								
08:21	Wind = 281°/25 kt -								
08:24	Turn RH side chair CS (Grossesime).								

Bibliography

- [1] P. Piscaglia, A. Cavallaro, M. Bonnet, and D. Douchamps, "High level description of video surveillance sequences," in *Proc. European Conf. on Multimedia Applications, Services and Techniques (ECMAST'99)*, pp. 316–331, (Madrid, Spain), May 1999.
- [2] P. Remagnino, A. Baumberg, T. Grove, D. Hogg, T. Tan, A. Worral, and K. Baker, "An integrated traffic and pedestrian model-based vision system," in *Proc. British Machine Vision Conf. (BMVC'97)*, **2**, pp. 380–389, (Colchester, UK), Sept. 1997.
- [3] B. Abreu, L. Botelho, A. Cavallaro, D. Douchamps, P. Figueiredo, B. Macq, *et al.*, "Video-based multi-agent traffic surveillance system," in *Proc. Intelligent Vehicles Conf. (IV'00)*, (Dearborn, MI), Oct. 2000.
- [4] A. Cavallaro, D. Douchamps, T. Ebrahimi, and B. Macq, "Segmenting moving objects: the MODEST video object kernel," in *Proc. Workshop on Image Analysis For Multimedia Interactive Services (WIAMIS'01)*, (Tampere, Finland), May 2001.
- [5] H. P. Moravec, "Towards automatic visual obstacle avoidance," in *Proc. Intl. Joint Conf. on Artificial Intelligence*, p. 584, (Cambridge, MA), Jan. 1977.
- [6] H. P. Moravec, "Visual mapping by a robot rover," in *Proc. of the 6th Intl. Joint Conf. on Artificial Intelligence*, pp. 598–601, (Tokyo, Japan), Aug. 1979.
- [7] F. Devernay, "Computing differential properties of 3-d shapes from stereoscopic images without 3-d models," Tech. Rep. RR-2304, INRIA, Sophia Antipolis, 1994.
- [8] O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, Cambridge, Massachusetts, 1993.

-
- [9] S. Smith, "A new class of corner finder," in *Proc. British Machine Vision Conf. (BMVC'92)*, **1**, pp. 139–148, (Leeds, UK), Sept. 1992.
- [10] S. Smith, "Flexible filter neighbourhood designation," in *Proc. Intl. Conf. Pattern Recognition (ICPR'96)*, **1**, pp. 206–212, (Vienna, Austria), Aug. 1996.
- [11] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes*, Cambridge University Press, 1989.
- [12] J. Hammack, N. Scheffner, and H. Segur, "Two-dimensional periodic waves in shallow water," *Journal of Fluid Mechanics* **209**, pp. 567–589, 1989.
- [13] B. Jähne, J. Klinke, and S. Waas, "Imaging of short ocean wind-waves - a critical theoretical review," *J. Opt. Soc. America A* **11**, pp. 2197–2209, 1994.
- [14] A. W. Lang and M. Gharib, "Experimental study of the wake behind a surface-piercing cylinder for a clean and contaminated free surface," *Journal of Fluid Mechanics* **402**, pp. 109–136, 2000.
- [15] D. Stillwell and O. R. Pilon, "Directional spectra of surface waves from photographs," *J. Geophys. Res.* **79**, pp. 1277–1284, 1974.
- [16] K. R. Nicolas, "Radar imaging of water surface flow fields," *Experiments in Fluids* **23**, pp. 14–19, 1997.
- [17] C. Craeye, P. W. Sobieski, L. F. Bliven, and A. Guissard, "Ring-waves generated by water drops impacting on water surfaces at rest," *IEEE J. Oceanic Eng.* **24**, pp. 323–332, 1999.
- [18] G. K. Gilbert and E. Keller, "The transportation of debris by running water," Tech. Rep. 86, U.S. Geological Survey, 1914.
- [19] J. F. Kennedy, "The mechanics of dunes and antidunes in erodible-bed channels," *Journal of Fluid Mechanics* **16**, pp. 521–544, 1963.
- [20] A. J. Reynolds, "Waves on the erodible bed of an open channel," *Journal of Fluid Mechanics* **22**, pp. 113–133, 1965.
- [21] J. R. Allen, *Current Ripples: Their Relation to Patterns of Water and Sediment Motion*, North-Holland, Amsterdam, Netherlands, 1968.

- [22] J. Alexander and C. Fielding, "Gravel antidunes in the tropical burdekin river, queensland, australia," *J. Sedimentology* **44**, pp. 327–337, 1997.
- [23] T. C. Blair, "Sedimentology and progressive tectonic unconformities of the sheetflood-dominated hell's gate alluvial fan, death valley, california," *Sediment Geology* **132**, pp. 233–262, 2000.
- [24] G. V. Middleton, "Antidune cross-bedding in a large flume," *J. Sediment. Petrology* **35**, pp. 922–927, 1965.
- [25] J. Shaw and R. Kellerhals, "Paleohydraulic interpretation of antidune bedforms with applications to antidunes in gravel," *J. of Sediment Petrology* **47**, pp. 257–266, 1977.
- [26] A. C. E. Wessels, A. E. Mynett, J. K. Kostense, and J. E. Prins, "Modern laboratory techniques," Tech. Rep. 418, Delft Hydraulics Publication, Delft, Netherlands, 1989.
- [27] A. Weigand, "Simultaneous mapping of the velocity and deformation field at a free surface," *Experiments in Fluids* **20**, pp. 358–364, 1996.
- [28] D. Dabiri and M. Gharib, "Simultaneous free-surface deformation and near-surface velocity measurements," *Experiments in Fluids* **30**, pp. 381–390, 01.
- [29] D. Devriendt, *Reconstruction tridimensionnelle de la surface d'un écoulement sur antidunes à partir de ses lignes de courant acquises par imagerie digitale*. Master thesis, Université catholique de Louvain, Louvain-la-Neuve, Belgium, June 1998.
- [30] D. Douxchamps, "Stereometric reconstruction of the free surface associated to a flow over three-dimensional antidunes," master thesis, Université catholique de Louvain, Louvain-la-Neuve, Belgium, June 1998.
- [31] D. Devriendt, D. Douxchamps, H. Capart, C. Craeye, B. Macq, and Y. Zech, "Three-dimensional reconstruction of a periodic free surface from digital imaging measurements," in *Proc. Second Intl. Conf. on Advances in Fluid Mechanics*, M. Rahman, M. Comini, and C. Brebbia, eds., *Advances in Fluid Mechanics*, pp. 203–212, Computational Mechanics Publications, Boston, MA, 1998.

- [32] H. Capart, D. Devriendt, D. Douxchamps, and C. Craeye, "Three-dimensional reconstruction of the oscillatory free-surfaces of a flow over antidunes: stereoscopic and velocimetric techniques," in *Proc. OCEANS Conf. (OCEANS'00)*, (Providence, RI), Sept. 2000.
- [33] B. Jähne, *Digital Image Processing*, Springer, Berlin, Germany, 1995.
- [34] H. Capart, D. L. Young, and Y. Zech, "Voronoi imaging methods for the measurements of granular flows," *Experiments in Fluids* **32**, pp. 121–135, 2002.
- [35] P. Veber, J. Dahl, and R. Hermansson, "Study of the phenomena affecting the accuracy of a video-based particle tracking velocimetry technique," *Experiments in Fluids* **22**, pp. 482–488, 1997.
- [36] A. Okabe, B. Boots, and K. Sugihara, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, July 1992.
- [37] R. Y. Tsai, "A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses," *IEEE Trans. Robot. Automat.* **3**, pp. 323–344, Aug. 1987.
- [38] R. Jain, R. Kasturi, and B. G. Schunck, *Machine Vision*, McGraw-Hill, New York, USA, 1995.
- [39] B. Spinewine, H. Capart, M. Larcher, and Y. Zech, "Three-dimensional voronoi imaging methods for the measurement of near-wall particulate flows," *Experiments in Fluids* **34**, pp. 227–241, 2003.
- [40] P. R. Wolf and B. A. Dewitt, *Elements of Photogrammetry*, McGraw-Hill, 2000.
- [41] H. L. Cameron, "The measurement of water current velocities by parallax methods," *Photogrammetric Eng.* **18**, pp. 99–104, 1952.
- [42] R. O. Duda and P. E. Hart, *Pattern classification and scene analysis*, Wiley-Interscience, New York, USA, 1973.
- [43] R. Frühwirth, M. Regler, R. K. Bock, H. Grote, and D. Notz, *Data analysis techniques for high-energy physics*, Cambridge University Press, 2000.
- [44] L. D. Vinci, *The Notebooks of Leonardo Da Vinci*, Dover, Mineola, NY, ca. 1500 / 1975.

- [45] D. Douxchamps, D. Devriendt, H. Capart, C. Craeye, B. Macq, and Y. Zech, "Stereoscopic and velocimetric reconstructions of the free surface topography of antidune flows," *Experiments in Fluids*, 2004. Accepted for publication.
- [46] W. Jackson, "Wake vortex prediction: An overview," Tech. Rep. TP13629E, Transportation Development Center, Transport Canada, Mar. 2001.
- [47] J. N. Hallock, C. Tung, and S. Sampath, "Capacity and wake vortices," in *Proc. Intl. Congress of Aerospace Sciences (ICAS'02)*, (Toronto, Canada), Sept. 2002.
- [48] *Air Traffic Control*, ch. 3-9-7. No. 7110-65P, Federal Aviation Administration (FAA), 2004.
- [49] *Aeronautical Information Manual*, ch. 7-3. Federal Aviation Administration (FAA), 2004.
- [50] J. M. Vaughan, K. O. Steinvall, C. Werner, and P.-H. Flamant, "Coherent laser radar in Europe," *Proc. IEEE* **84**, pp. 205–226, Feb. 1996.
- [51] B. Lamiscarre, B. Christophe, C. Fournet, J. Lemorton, L. Poutier, and A. Oyzel, "Nouveaux capteurs pour l'amélioration de la sécurité et de l'efficacité des mouvements en zone aéroportuaire," in *Proc. 5ème Journée AAAF*, (Toulouse, France), Jan. 1999.
- [52] W. L. Holforty and J. D. Powell, "Flight deck display of airborne traffic wake vortices," in *Proc. Digital Avionics Systems Conf. (DASC'01)*, (Daytona Beach, FL), Oct. 2001.
- [53] D. Soreide, "Airborne coherent LIDAR for advanced in-flight measurements," in *Proc. Conf. on Aviation, Range, and Aerospace Meteorology*, (Orlando, FL), Sept. 2000.
- [54] R. Bogue, R. McGann, *et al.*, "Comparative optical measurements of airspeed and aerosols on a DC8 aircraft," Tech. Rep. 113083, NASA, Edwards, CA, July 1997.
- [55] C.R. Tatnall, "An investigation of candidate sensor-observable wake vortex strength parameters for the NASA Aircraft Vortex Spacing System (AVOSS)," Tech. Rep. 206933, NASA, Hampton, VA, Mar. 1998.

- [56] D. A. Hinton, "An aircraft vortex spacing system (AVOSS) for dynamical wake vortex spacing criteria," in *Proc. 78th Fluid Dynamics Panel Symposium*, (Trondheim, Norway), May 1996.
- [57] D. Vicroy, J. Brandon, G. Greene, R. Rivers, G. Shah, E. Stewart, and R. Stuever, "Characterizing the hazard of a wake vortex encounter," in *Proc. Aerospace Sciences Meeting and Exhibit (AIAA'97)*, (Reno, NV), Jan. 1997.
- [58] D. Vicroy, P. Vijgen, H.M.Reimer, J. L. Gallegos, and P. R. Spalart, "Recent NASA wake-vortex flight tests, flow-physics database and wake-development analysis," in *Proc. World Aviation Congress and Exposition*, (Anaheim, CA), Sept. 1998.
- [59] M. Keane, D. Buckton, M. Redfern, C. Bollig, C. Wedekind, and F. Kopp, "Axial detection of aircraft wake vortices using doppler lidar," *Journal of Aircraft* **39**, pp. 850–861, Sept. 2002.
- [60] M. J. Levine, "Power spectrum parameter estimation," *IEEE Trans. Inform. Theory* **11**, pp. 100–107, 1965.
- [61] G. D. Nastrom, "Doppler radar spectral width broadening due to beamwidth and wind shear," *Annales Geophysicae* **15**, pp. 786–796, 1997.
- [62] R. J. Doviak, D. S. Zrnic, and D. S. Sirmans, "Doppler weather radar," *Proc. IEEE* **67**, pp. 1522–1553, Nov. 1979.
- [63] D. S. Zrnic and R. J. Doviak, "Matched filter criteria and range weighting for weather radar," *IEEE Trans. Aerosp. Electron. Syst.* **14**, pp. 925–930, Nov. 1978.
- [64] D. S. Zrnic, "Spectral moments estimates from correlated pulse pairs," *IEEE Trans. Aerosp. Electron. Syst.* **13**, pp. 344–354, 1977.
- [65] R. Frehlich and L. Cornman, "Estimating spatial velocity statistics with coherent doppler lidar," *Journal of Atmospheric and Oceanic Technology* **19**, pp. 355–366, Mar. 2002.
- [66] J. M. B. Dias and J. M. N. Leitaó, "Asymptotically efficient estimation of spectral moments," *IEEE Trans. Acoust., Speech, Signal Processing* **43**, pp. 2222–2225, Sept. 1995.

- [67] J. M. B. Dias and J. M. N. Leitaó, "Nonparametric estimation of mean doppler and spectral width," *IEEE Trans. Geosci. Remote Sensing* **38**, pp. 271–282, Jan. 2000.
- [68] J. M. B. Dias and J. M. N. Leitaó, "Nonparametric estimation of mean velocity and spectral width in weather radar," in *Proc. Intl. Geoscience and Remote Sensing Symposium (IGARSS'95)*, **2**, pp. 2121–2125, (Florence, Italy), Jan. 1995.
- [69] H. Linné, "Heterodyne DIAL at high repetition rate," in *Proc. Intl. Laser Radar Conf. (ILRC'98)*, (Annapolis, MD), July 1998.
- [70] R. M. Huffaker and R. M. Hardesty, "Remote sensing of atmospheric wind velocities using solid-state and CO₂ coherent laser systems," *Proc. IEEE* **84**, pp. 181–204, Feb. 1996.
- [71] P. R. Mahapatra and D. S. Zrnic, "Practical algorithms for mean velocity estimation in pulse doppler weather radars using a small number of samples," *IEEE Trans. Geosci. Remote Sensing* **21**, pp. 491–501, Oct. 1983.
- [72] D. S. Zrnic, "Estimation of spectral moments for weather echoes," *IEEE Trans. Geosci. Electron.* **17**, pp. 113–128, Oct. 1979.
- [73] E. S. Chornoboy, "Optimal mean velocity estimation for doppler weather radars," *IEEE Trans. Geosci. Remote Sensing* **31**, pp. 575–586, May 1993.
- [74] S. Kay, "A fast and accurate single frequency estimator," *IEEE Trans. Acoust., Speech, Signal Processing* **37**, pp. 1987–1990, Dec. 1989.
- [75] D. S. Zrnic, "Spectrum width estimates for weather echoes," *IEEE Trans. Aerosp. Electron. Syst.* **15**, pp. 613–619, Sept. 1979.
- [76] J. A. Cadzow, "High performance spectral estimation - a new ARMA method," *IEEE Trans. Acoust., Speech, Signal Processing* **28**, pp. 524–529, Oct. 1980.
- [77] S. M. Kay, "A new ARMA spectral estimator," *IEEE Trans. Acoust., Speech, Signal Processing* **28**, pp. 585–588, Oct. 1980.
- [78] K. S. Miller and M. M. Rochwarger, "A covariance approach to spectral moment estimation," *IEEE Trans. Inform. Theory* **18**, pp. 588–596, 1972.

- [79] B. Friedlander and B. Porat, "The modified Yule-Walker method of ARMA spectral estimation," *IEEE Trans. Aerosp. Electron. Syst.* **20**, pp. 158–173, Mar. 1984.
- [80] Y. T. Chan and R. P. Langford, "Spectral estimation via the high-order Yule-Walker equations," *IEEE Trans. Acoust., Speech, Signal Processing* **30**, pp. 689–698, Oct. 1982.
- [81] J. A. Cadzow, "Spectral estimation: an overdetermined rational model equation approach," *Proc. IEEE* **70**, pp. 907–939, Sept. 1982.
- [82] R. D. Palmer, J. R. Cruz, and D. Zrnic, "Enhanced autoregressive moving average spectral estimation applied to the measurement of doppler spectral width," *IEEE Trans. Geosci. Remote Sensing* **29**, pp. 358–367, May 1991.
- [83] B. M. Keel and E. G. Baxa, "An overdetermined approach to autocorrelation based spectral moment estimators for use in doppler weather radar," in *Proc. National Radar Conference*, (Syracuse, NY), May 1997.
- [84] S. Takeda, T. Nakamura, and T. Tsuda, "An improvement of wind velocity estimation from radar doppler spectra in the upper mesosphere," *Annales Geophysicae* **19**, pp. 837–843, 2001.
- [85] L. H. Jansen and G. V. der Spek, "The shape of doppler spectra from precipitation," *IEEE Trans. Aerosp. Electron. Syst.* **21**, pp. 208–219, Mar. 1985.
- [86] R. Frehlich, "Effects of wind turbulence on coherent doppler lidar performance," *J. of Atmospheric Science and Oceanic Technology* **14**, pp. 54–75, Jan. 1997.
- [87] "Special issue on 3-d and stereoscopic visual communications," *IEEE Signal Processing Mag.* **16**(3), May 1999.
- [88] F. Devernay, *Vision Stéréoscopique et propriétés différentielles des surfaces*. Phd thesis, INRIA, Sophia Antipolis, 1997.
- [89] D. Douxchamps and B. Macq, "Integrating perspective in stereo image matching," in *Proc. Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP'04)*, (Montreal, Canada), May 2004.

-
- [90] L. Cohen, L. Vinet, P. T. Sander, and A. Gagalowicz, "Hierarchical region based stereo matching," in *Proc. Intl. Conf. Computer Vision and Pattern Recognition (CVPR'89)*, pp. 416–421, (San Diego, CA), 1989.
- [91] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal* 7, pp. 308–313, 1963.
- [92] B. D. Bunday, *Basic Optimization Methods*, Edward Arnold, 1984.