

FEASIBILITY STUDY OF THE VISUALIZATION OF SANDY BOTTOM TRANSPORT BY USING PARTICLE TRACKING VELOCIMETRY

Jens Michiels

Student number: 01102798

Promotor: Prof. dr. Frank Mostaert

Copromotor: Styn Claeys

Jury: Prof. dr. Sébastien Bertrand, dr. Ir. Dieter Meire

Master's dissertation submitted in partial fulfilment of the requirements for
the degree of Master of Science in Geology

Academic year: 2016 – 2017

PREFACE

During the Spring semester of 2015, I did my Bachelor Project at Flanders Hydraulics Research, with prof. dr. Frank Mostaert and Styn Claeys as my mentors. A year later, Styn got back in contact with me and offered me the possibility to write my Master's dissertation here as well. I accepted, a subject was found and this report is the product of my research.

First, I want to thank Styn Claeys for his offer, the subject that he has found and his insights when I presented intermediate and final results; prof. dr. Frank Mostaert for being promotor of my dissertation; and dr. Ir. Dieter Meire, also for his insights when I presented intermediate and final results and to point me in the right direction when I wasn't sure what to do next.

I want to thank everyone at Flanders Hydraulics Research who helped me. That list includes Stefan Geerts (camera equipment and software), Bart Deckers (operation current flume and mounting the top-down camera), Glen Heyvaert (introduction to the Streams software), Ellen and Lia of the sediment laboratory (grain size analysis) and many others.

I want to thank prof. dr. Roger Nokes of the University of Canterbury, for giving me permission to use Streams, for providing me with the latest version and for answering my questions.

I want to thank Joachim Beckers, for sending me the results of his research, for operating the flume and for helping with my experiments.

Finally, I want to thank my parents and my sister for their support while writing this dissertation.

Without any of them, this dissertation would not be what it is right now

Jens Michiels, Summer 2017

TABLE OF CONTENTS

Introduction.....	6
Literature review and Prior research	8
Direct methods to measure bedload transport	8
Direct methods to measure bedload transport	9
Particle Tracking Velocimetry	12
Material and methods.....	14
The current flume	14
Camera equipment	15
Camera software	19
PTV analysis software	21
Experiments.....	25
Revised recording software and settings	35
PTV analysis.....	36
Results.....	42
Particle velocity.....	42
Sediment transport	45
Discussion	49
Conclusion.....	52
References	53
Appendix.....	55

TABLE OF FIGURES

Figure 1: Particle movement in a flow..	6
Figure 2: Sketch of the Helley-Smith bedload sampler.	8
Figure 3: Sketch of the BTMA. (van Rijn 1993).	9
Figure 4: The necessity of a long transmitter pulse in bottom tracking.	11
Figure 5: Current flume.	15
Figure 6: Grain size distribution of the sand substitute, measured by a Malvern Mastersizer 2000.	15
Figure 7: UI – 308xCP – C camera.	16
Figure 8: Ricoh FL-BC7528-9M lens.	17
Figure 9: Opto 100-BTC-005 lens.	18
Figure 10: Camera and lens in dry cell, behind the window.	18
Figure 11: Camera in aquarium above the flume.	19
Figure 12: Illustrating the effect of extracting or removing colours.	24
Figure 13: Cycle of experimenting, evaluating results and changing setup and settings.	25
Figure 14: Close-up of frame 130 of the 49 ms exposure time image sequence.	29
Figure 15: Particle identification of the linear features of the 49 ms exposure time image sequence.	29
Figure 16: Image sequence of the 40 Hz recording of the sideways camera, filter and PID applied.	31
Figure 17: Particle record of the single threshold PID of the 40 Hz recording of the side-looking camera.	32
Figure 18: Eulerian field of the single threshold PID of the 40 Hz recording of the camera sideways.	33
Figure 19: Images of the recordings by the top-down camera.	34
Figure 20: Close-up of particle paths from state-based costings.	38
Figure 21: Effect of repeating processes within a pipeline.	41
Figure 22: Part of the image in which the particle velocity, particle density and transport are calculated.	42
Figure 23: Comparison of the paths of frames 44 (left) and 4 (right) of the image sequence.	43
Figure 24: Eulerian velocity field, averaged in time.	43
Figure 25: Average particle velocity near the sediment bed	45
Figure 26: Close-ups of the particle identification applied to frame 11 of the filtered image sequence.	46
Figure 27: Average particle density near the sediment bed.	47
Figure 28: Sediment transport near the sediment bed.	48

STORY OF MY RESEARCH

In my third year of my Bachelor of Science in Geology, I chose to do my Bachelor Project at Flanders Hydraulics Research. I was curious about how things went outside of the university. It was a good experience, and a year later, when my mentor said that I could do my Master's dissertation here if I wanted as well, I accepted. A subject was found: the application of Particle Tracking Velocimetry to visualize bottom transport in a laboratory flume. The subject has nothing to do with my Bachelor Project (Acoustic Doppler Current Profiler to measure concentration of suspended sediment), but the methodology is similar. The research involves an existing technique and combining the available equipment and software to answer the following questions: What can we measure? What can be improved? What can we use for future research?

In this research, we started with individual components: two cameras with lenses, a current flume, a way of mounting the cameras, software to record and software to process images, instruments to provide comparable data. Through the research, I learned to work with people who could help me with one or two individual parts of the process each, but I had to combine all of these elements into one project. The research was a cycle of choosing settings, carrying out experiments, processing videos based on what I learned from one person, discussing the results with a second person, changing the setup with the help of a third person and repeat. Through trial and error, it becomes clear what is possible to measure and what not, what can or should be changed and what not, and even which research questions can be answered within the scope of my work and which ones cannot. My final result is the sediment transport visualized in the lowest 10 cm of the water column at one specific water velocity by one camera in specific settings. If my progress had been going faster, I might have been able to collect data at different water flow rates or with the second camera or vary the settings. If it had been going slower, I might have not been able to collect or process even these images. It was an informative experience having an own scientific project, starting from an initial situation and expectations up to the final results and what is left for the next researcher.

1. INTRODUCTION

Of all sediment transport mechanisms, sediment transport in water is the most important (Nichols 2009). Traditionally, it has been divided in two categories: suspended transport and bedload transport (Figure 1). Their definitions can vary, as illustrated by Bagnold (1956) and Einstein (1950). Bagnold (1956) defines the bedload transport as that in which the particles have successive contacts with the bed by the effect of gravity, while the suspended-load transport is defined as that in which the excess weight of the particles is supported by upward impulses from turbulence in the flow. Bedload is transported by rolling, sliding and jumping (also known as saltation). Einstein (1950) defines the bedload transport as the transport of sediment particles within a thin layer (thickness: 2 particle diameters) just above the bed by sliding, rolling and sometimes by making jumps with a longitudinal distance of a few particle diameters. This view does not include saltating particles, whose jump lengths are considerably larger than a few grain sizes. What these definitions have in common is that bedload transport takes place in a region near the bottom of the water column and near the top of the sediment bed. Particles in bedload transport move at a speed that is less than the velocity of the transporting current (Gomez 1991), whereas the velocity in longitudinal direction of particles in suspension is almost equal to the fluid velocity.




Figure 1: Particle movement in a flow. Reproduced from Nichols (2009).

There are many devices to collect sediment from bottom transport directly. But bedload samplers have certain disadvantages (Nevelsteen & Van Hoestenberghe 2009; Yager et al. 2015; Ryan & Porth 1999; Rachlewicz et al. 2017). Indirect methods to determine sediment transport do exist, but before using them in the field, laboratory tests and calibration are required (Nevelsteen & Van Hoestenberghe 2009; Yager et al. 2015; Beylich & Laute 2014). One of these methods is the Acoustic Doppler Current Profiler Bottom-Track (ADCP-BT). Experiments with acoustic instruments will be carried out at Flanders Hydraulics Research. A current flume that was already built was modified to fit the research terms. The most important modifications include a pit sampler at the end of the sediment bed to collect all sediment transported by bedload transport, an acrylic glass window and a dry cell, where a camera can be placed to monitor the sediment transport, and mounting structures above the flume, to which ADCP's, a camera and other instruments can be attached. The results obtained by the acoustic instruments should then be compared with data from other devices and methods to be verified. We are searching for a method that can visualise bottom transport and can measure the velocity of the sand that is being transported. A good candidate for this is the use of cameras and the use of Particle Tracking Velocimetry (PTV) to process the images, in the hopes of obtaining a good image of the particle movement and the particle velocities of sand-sized particles near the bottom.

This study will investigate if, using the equipment and software currently available, a methodology can be found to record camera images, process these to particle velocities and sediment transport and visualize it. If this has been obtained, we will try to measure particle velocities and transport rates under conditions where this is possible. If good data has been acquired, it will be compared to ADCP and bottom-track measurements.

2. LITERATURE REVIEW AND PRIOR RESEARCH

2.1 Direct methods to measure bedload transport

There is a large diversity of methods and devices to measure bottom or bedload transport directly. Reasons for the poor comparability and representativeness of bedload research include difficulties regarding measurement techniques, limited budget, available material, experience of the users with specific methods or availability of operators (Rachlewicz et al. 2017). Bedload samplers are instruments that capture and collect bedload sediment. Examples include the Helley-Smith bedload sampler and the Bedload Transport Meter Arnhem (BTMA). These devices are placed in the flow, near the bottom. Their presence may affect local patterns of the flow and transport (Ryan & Porth 1999). Construction (type of material in particular) can result in significant variability in the amount of captured sediment (Rachlewicz et al. 2017). Issues have been reported related to the limited duration of bedload sampling, size of samplers, and disturbances on the bed produced by the same samplers (Mao et al. 2016). Descriptions of two bedload samplers (the Helley-Smith sampler and the Bedload Transport Meter Arnhem) are given by van Rijn (1993). The Helley-Smith bedload sampler consists of a nozzle, sample bag and frame (Figure 2). The Helley-Smith is one of the most widely-used instruments to measure bottom transport, and it has been calibrated extensively (based on about 10 000 samples). The lower and upper threshold of the grain size of the captured sediment depend on the mesh size of the basket and the opening of the sampler mouth respectively.




Figure 2: Sketch of the Helley-Smith bedload sampler, including sample bag (captures sediment), frame (holds bag) and nozzle (opening through which bedload enters). Reproduced from van Rijn (1993)

The Bedload Transport Meter Arnhem (BTMA) collects sediment particles in a basket type sampler. This basket is pressed on the channel bed when the frame has been lowered on the sediment bed by means of a spring (Figure 3). This results in a pressure reduction behind the instrument, which causes water and bedload to enter the sampler mouth at the same velocity as that of the local flow, under the condition that the sediment content already in the basket is not too high. The lower and upper threshold of the grain size of the

captured sediment again depend on the mesh size of the basket and the opening of the sampler mouth respectively.




Figure 3: Sketch of the BTMA. 1) wire mesh basket (capture sediment) 2) sampler mouth (bedload enters) 3) tail 4) frame 5) leaf spring 6) stabilizer. Reproduced from van Rijn (1993)

Ryan & Porth (1999) showed that three comparable bedload samplers (out of which two Helley-Smith samplers) can produce significantly different results. Flow remains the most important element predicting the weight of the sediment samples taken, but the type of sampler cannot be ignored. Bedload data from samplers needs to be calibrated therefore. Rachlewicz et al. (2017) also compared three bedload samplers (including one Halley-Smith), and concluded that placement of samplers on the river bottom changes the dynamics and character of the movement of bedload particles locally. The construction of the sampler, its way of placing and stabilization, the operator all have a significant influence on discrepancies in estimating the amount of transported bedload.

2.2 Indirect methods to measure bedload transport

Indirect methods to determine sediment bottom transport include geophones, hydrophones, vibrational sensors, pipe-microphone acoustic systems, acoustic sensors fixed beneath a steel plate and impact sensors fixed to steel plates (Beylich & Laute 2014). These methods are less invasive than trapping or collecting sediment or adding tracers, they minimize local and temporal changes to the flow field, and they could provide continuous records in space and time. Laboratory experiments in a controlled environment are important for a better understanding of the factors that can influence the calibration of these devices (Beylich & Laute 2014). Experiments such as these are needed because of the following reasons (Yager et al. 2015):

- 1) field measurements may include too many spatially and temporally variable parameters that could influence bedload transport, on the other hand flume experiments allow for the isolation of the effects of individual parameters;
- 2) bedload transport in the field is notoriously difficult to quantify and has large uncertainties, as discussed earlier;
- 3) changes in bed surfaces (e.g. grain size distribution, elevation, bed roughness) during flow events can also be difficult or impossible to measure;
- 4) grain scale mechanics of sediment motion and transport are much easier to quantify in controlled laboratory conditions rather than in the field.

In Flanders, sediment transport measurements in waterways only really started in 2000 (Nevelsteen & Van Hoestenbergh 2009). Since then, the Flemish Environment Agency has been able to measure suspended sediment transport precisely and continuously using optical turbidity sensors and analysis of water samples in the upper catchment areas of the Scheldt and Demer rivers, where suspended transport is dominant. However, in waterways in regions of Flanders that are dominated by sandy soils, bedload transport is a much more important part of the total sediment transport. To find a method to measure bedload transport continuously in time and space, Nevelsteen & Van Hoestenbergh (2009) reviewed a list of potential bedload measurement devices and techniques:

- 1) Optical instruments: these include turbidity sensors, ASM = Argus surface meter (by Argus environmental instruments) and LISST Streamside (Sequoia). The downside for these instruments is that they take point measurements instead of measuring a large part of the water column. They are more fitted to measure suspended sediment transport than bedload transport.
- 2) Sampling: sampling can be difficult, even dangerous, is intrusive and discontinuous, and sampling efficiency decreases as the samplers are filled. Nevelsteen & Van Hoestenbergh (2009) reviewed the pit sampler, the Helley-Smith and the BTMA.
- 3) Image capture and processing: recording underwater videos requires a lot of material to be brought into the field, and the processing itself is labour-intensive. There is potential to use cameras in the laboratory experiments however.
- 4) Bedload formulae: Transport rates can be calculated from hydraulic parameters, but these formulae need verification.
- 5) Tracers: tracers can be added to the flow to visualize transport, but this is not fit for continuous measurements, only as a first indication of transport rates.
- 6) Acoustic instruments: devices such as the AZTM = akoestische zandtransport meter (acoustic sand transport meter, by Delft Hydraulics) suffer from the same flaws as optical instruments, but ADCP devices can measure nearly all of the water column and the river bottom (if they have a bottom-track function) and do this continuously.

Out of these options, it was concluded to further investigate the ADCP-BT capabilities to measure bedload transport. More specifically, the Workhorse 1200 kHz and the Streampro 2000 kHz by Teledyne RD Instruments were chosen because of their sufficiently small cell size, inclusion of bottom tracking and their working ranges of water depth and flow velocity. For more information on ADCP-BT, we refer to the RDI Instruction Manual and the thesis by Joachim Beckers (Beckers 2017). In short, an Acoustic Current Doppler Profiler (ADCP) emits beams of sound pulses that are reflected by scatterers such as suspended sediment particles and the river bottom. When these objects move away or towards the transducer, the frequency of the incoming soundwave changes: this phenomenon is called the Doppler effect, and the increase or decrease in frequency is called the Doppler shift. An ADCP reads the echo wave that was reflected on suspended particles and calculates their velocity relative to the instrument. Hence the ADCP can measure the suspended sediment transport. To measure bedload transport, the ADCP must have a bottom-track feature. Bottom tracking requires long acoustic pulses, for the beam to completely illuminate the bottom at one time, as shown in Figure 4. The downside of long transmit pulses is that a considerable part of the echo can come not only from the bottom, but also from the sediment in the water column just above the bed. In conditions of high sediment concentrations near the bottom, this could cause an underestimation of the water depth and bias the bottom-track velocity towards the water velocity: this is called the water bias. What is being measured by the bottom-track is therefore not always clear and this demands further laboratory research before using BT-ADCP in the field.




Figure 4: The necessity of a long transmitter pulse in bottom tracking. A) Pulse too short B) Pulse long enough to completely illuminate the bottom at one time. Reproduced from the Teledyne RDI Instruction Manual.

2.3 Particle Tracking Velocimetry

Particle Tracking Velocimetry (PTV) and Particle Image Velocimetry (PIV) are two image analysis methods commonly used in fluid dynamics research (Smith & Friedrichs 2015). Both techniques are based on the principle of capturing video images of a specially illuminated particle seeded fluid flow, to then extract quantitative information about the flow field from that video record (Nokes 2009a). PIV involves correlating the movement of particle groups (Smith & Friedrichs 2015). The technique relies on the cross-correlation of the intensity fields in two consecutive frames. Fluid velocities are obtained from matches based on peaks in the cross-correlation field. For this cross-correlation to be meaningful and PIV to work efficiently, particle density needs to be sufficiently high (Nokes 2009a). PTV is based on tracking individual particles from frame to frame. This method produces particle-centred velocity estimates, which can be interpolated to a regular grid if necessary. PTV works better in sparsely seeded flows, but in general particle density is less of an issue with PTV than PIV (Nokes 2009a). O'Brien & McKenna Neuman (2016) reviewed PIV, PTV and Laser Doppler Anemometry (LDA) as techniques for the measurement of saltation of sand, including a short historical overview of the use of PTV. LDA can obtain particle population statistics, but this method does not track particles individually and photo sensors can be overloaded at the sand bed. PIV can track individual sand particles, but according to O'Brien & McKenna Neuman (2016) this technique can only measure an insignificant fraction of the saltation trajectory of particles and does not perform reliably when there is a high density of sand particles of various size. PTV is widely regarded as the most desirable method for obtaining information about saltation dynamics as this technology can track a particle's displacement and velocity throughout a portion of its trajectory, in rare instances from ejection through to impact (O'Brien & McKenna Neuman 2016). Like LDA and PIV technologies, however, existing PTV systems generally perform poorly in flows with a high density of particles. PTV experiments started using cine film photography and relied on manual particle identification and tracking only, which is subjective and labour intensive (unfit for large datasets). Over the last decade, scientific-grade, digital cameras have replaced film in the application of PTV. Particles can now be identified and assigned spatial coordinates using computer programs. The most recent innovations are that cameras delivering both high resolution and a high frame rate have gotten cheaper and more accessible, and the ongoing development of software that can detect both the sand particles and their trajectories (O'Brien & McKenna Neuman 2016).

Particle Tracking Velocimetry (PTV) involves the matching of particles between subsequent video frames. When the time between frames is known (most often by having a fixed framerate) and the distance travelled between frames is measured, the individual particle velocities can be calculated. In order to do this successfully, the particles first need to be recognised in the frames correctly. To make this process easier or even possible, PTV is most often used on low concentrations of relatively large particles (Smith & Friedrichs (2015) use diameter > 30 μm , so this includes coarse silt, sand and gravel, not fine silt and clay) that differentiate strongly from the image background. Preferably the particles are bright objects in a dark

background, clearly separated from each other and easy to follow even with the naked eye between frames. Some potential limitations for video analysis of bedload transport include the following (Yager et al. 2015):

- 1) very high transport rates often cannot be measured because of problems when trying to identify individual particles;
- 2) top-down view videos might only capture sediment transport immediately at the sediment bed surface; and
- 3) sideview videos might only capture limited spatial variation in bedload transport.

Based on the described advantages of PTV (good for sand-size particles, saltation-type motion, can track individual particles, innovations in camera and software) and to study its potential limitations, PTV analysis was chosen to try to visualize bottom transport and hopefully compare data with ADCP measurements.

3. MATERIAL AND METHODS

3.1 The current flume

The experiments are performed in a current flume at Flanders Hydraulics Research. The installation consists of an elevated water basin, fed by a larger underground basin by two large pumps, two valves, a control room, the flume itself, a dry compartment where a camera can be placed behind a transparent window, a filter (150 μm mesh size) and a tilting sluice gate, after which the water flows back to the underground basin. Figure 5 shows the most important parts of the flume. Two pumps replenish a basin that is elevated compared to the flume, so water can flow gravitationally through two pipes into the flume. Two valves are installed on these pipes, to control the flow rate. Each pipe has a maximum flow rate of 200 litres per second. The valves can be operated manually or from a computer in the control room. A computer program, after loading a calibration file, controls and displays the current flow rate for each valve and allows the user to set the desired flow rate. Most experiments were performed at a flow rate of 200 l/s: this flow rate is high enough to initiate a continuous bottom transport regime. Higher flow rates bring more material in motion, increasing the amount of particles and their velocity, which makes it more difficult to track them. The particles themselves are not regular quartz sand grains: prior research concluded that it was too difficult to get quartz sand into motion in the current flume. A replacement was chosen with a lower density, but with a coarse to medium sand-like grain size distribution (Wentworth scale). The density of quartz is 2650 kg/m^3 . The sand-like substitute was manufactured from old sewing buttons (most likely plastic, it can even be an old type of plastic such as Bakelite). The density of this material is 1220 kg/m^3 and the grain size distribution is given in Figure 6. The grain size distribution was determined with a Malvern Mastersizer 2000 (Figure 6). Full results can be found in the Appendix (A.4). This material is lighter and can be transported more easily through the current flume, which makes it ideal for the first test with PTV. The particles are sometimes coloured (red or black mostly), but most of the material isn't: the pale particles can be seen clearly, even under water with appropriate lighting. Lighting is provided by two options: a light screen opposite of the transparent window of the dry cell and a lamp above the water, shining down (Lupolux Dayled 1000 90 Watt). The 'sand' forms a bed in the narrow part of the flume. At the end of the sand bed, there is a pit sampler with a load cell that should trap the bedload material. The amount and weight of material captured in the pit sampler could then be used to estimate the amount of bedload transport over a longer time span. Behind the pit sampler, the flume broadens again and a large sieve filters all the water before it flows to the underground basin again. A sluice gate can be tilted up and down to empty or fill the flume with water. Most of the time, the tilting gate is placed up, but not higher than the top of the filter. However, if the water rises above the sluice gate level but is still below the top of the filter, the water level can still rise. The filter does not stretch to the top of the flume, so the water level (even at maximum flow rate) is limited to up to ~ 10 cm above the filter.




Figure 5: Current flume. A: tilting sluice gate. B: filter. C: laptop (upper left corner), camera and lens in dry cell (bottom), lamp and light screen (right side). D: dry cell (bottom left corner), lamp, aquarium, camera and lens above the water (centre). E: view of the flume from the control room. Length of the narrow part of the flume: 23,7 m. 1): outer wall of the current flume. 2): dry cell window. 3) & 4): inner wall of the flume. 5): end of the narrow part of the flume. 6): scale. Water flow direction from right to left

d(0.1): 365.692 μm

d(0.5): 501.434 μm

d(0.9): 686.954 μm




Figure 6: Grain size distribution of the sand substitute, measured by a Malvern Mastersizer 2000.

3.2 Camera equipment

To capture images during experiments in the current flume, two cameras will be used: one will be positioned in a dry compartment behind a transparent window next to the flow, positioned to record the bottom and the region just above the bed, where the majority of sand transport (bedload transport) takes place (Figure 10). A second camera provides an image from above the water surface, looking down to the bottom (Figure 11). These cameras and their software are products of IDS Imaging Development Systems GmbH. A manual was provided with the software. Both cameras are UI – 308xCP – C models (see Figure 7 and Appendix A.3 for more details). The software packages used to operate the cameras and to record are IDS Camera Manager

and uEye Cockpit, also by IDS Imaging Development Systems GmbH. Everything in this work that relates to IDS Camera Manager and uEye Cockpit comes from personal experience and the IDS uEye Camera Manual, version 4.82. Before positioning and using the cameras in the current flume, sufficient time was spent getting familiar with how to operate them, in preparation to find the optimal recording setup. Most preparation time was spent with the lens and camera for the side-looking view, because the other camera got humid and had to be dried for a week. The first tests after drying showed no indications of damage or failing to respond to software commands.




Figure 7: UI-308xCP-C camera. Top left: front view, where the lens is attached. Top right: rear view, with a Micro USB 3.0 port. Bottom left: schematic frontal view. Bottom right: schematic top view. Dimensions in mm. Reproduced from the camera specifications by IDS Imaging Development Systems GmbH.

The cameras are connected to a laptop through a USB 3.0 cable. A green light on the camera indicates when a good connection has been established. In IDS Camera Manager, a list of all cameras connected is shown, and whether they are free and available. The right camera is then opened in uEye Cockpit in expert mode: videos can only be recorded in expert mode. A pop-up window allows for the selection of the initial parameters: optimized for optimal colour imaging, monochrome imaging, live video imaging, loading a custom profile of parameters, or default parameters (no profile). This choice is not very important, the parameters can easily be changed and will be changed depending on the circumstances. uEye Cockpit displays the image that the camera is capturing in real-time. Various toolbars allow for the user to change a multitude of camera parameters. However, these do not include the iris aperture and the working distance of the lens.

The lens for the camera placed above the flume, is a FL-BC7528-9M lens made by Ricoh Company Ltd. (see Figure 8 and Appendix A.1). It allows for manual adjustment of the focus distance and the iris diaphragm

aperture. The advantage is that the focus can be adjusted so that the bottom is captured as sharply as possible, and the amount of incoming light can be altered with the diaphragm. The disadvantage is that for every new water depth and distance to the bottom, new adjustments are obligatory. The original setup was a closed transparent polycarbonate container surrounding the camera and lens, with its lower surface just below the water surface (to avoid reflections of a rough water surface). This closed container protects the camera from the water, but also prevents the described manual adjustments to the lens when the camera was mounted and water was flowing in the current flume. An effort was made to try to get the focus and iris aperture in the optimal setting before mounting. Later a larger glass 'aquarium' replaced the container: open at the top so that the lens can be manually adjusted while the camera is mounted and the flume is filled with water (Figure 11). The other lens (for the camera sideways) is a 100-BTC-005 lens by Opto GmbH (see Figure 9 and Appendix A.2). It is bitemcentric: this means that the size of the image of an object does not depend on the distance between the object and the lens, and the distance between the lens and the camera sensor. This has important advantages: particles moving away or towards the camera remain the same size on the image, which should help with their identification and to match them correctly. The image is not distorted, corrections are not necessary. Different particles of the same size moving at different distances near the working distance are displayed with the same image size, reducing the window of possible particle sizes in the process of particle identification. The focus and diaphragm aperture of this lens are fixed: no manual adjustments are needed or allowed. Working distance and depth of focus are fixed: 416 mm and 254 mm respectively, independent of the medium (air or water) or the presence of a window. Pixel size is also fixed and independent of the medium: 15 mm/pixel or 0.066.. mm/pixel. For the other lens and camera, this has to be measured every time when the settings are changed.




Figure 8: Ricoh FL-BC7528-9M lens. Dimensions are in mm. The first and third cylinder from the top can be rotated to adjust focus distance (from 0.25 m to infinity) and aperture ratio of the iris diaphragm (1:2.8 to 1:32). Reproduced from the product specifications by Ricoh Company Ltd.



Figure 9: Opto 100-BTC-005 lens. Dimensions are in mm. Reproduced from the product specifications by Opto.GmbH.




Figure 10: Camera and lens in dry cell, behind the window. A: USB 3 cable to laptop. B: camera. C: camera mount. D: lens. Water flow indicated by blue arrow



Figure 11: Camera in aquarium above the flume. A: aquarium. B: USB 3 cable to laptop. C: camera. D: lens. Water flow direction indicated by blue arrow

3.3 Camera software

After the manual settings are applied and have been verified visually in uEye Cockpit, the software allows for more parameters to be altered. A shorter settings overview can be found in Table 1. Important submenus of the Camera Properties include Camera and Image, and to a lesser degree Format and Size. 'Camera' lets the user change three very important settings: pixel clock, framerate and exposure time. Pixel clock is a measure of how often the sensor is being read, and gives an upper limit of the framerate. There are only 3 settings, and it should be chosen to allow the preferred framerate, without being unnecessarily high. The lowest setting allows for a framerate up to 22 frames per second or 22 Hertz. Framerate should be set on 'Hold' at the preferred setting. Otherwise the framerate might vary during the recording. An unknown timespan between frames prevents calculation of particle velocities. Exposure time is important for illumination and particle capture: a long exposure time gives a bright image with particles as short tracks smeared out, a short exposure time gives a dark image with particles as dots. An exposure time of 2 to 5 milliseconds results in the display of particles as dots, longer than 10 ms leads to tracks (in general; it depends on the particle velocities themselves of course). If the exposure time is low and the image is too dark, the software provides some tools to brighten the frames. In 'Image', the master gain is a setting to get a sufficiently bright image manually or automatically. If this is still insufficient, a gain boost can be toggled on. Automatic gain can be useful if the lighting varies during capturing, but a stable light source is highly

recommended. When the lighting is constant, a fixed gain is better, to have the same average pixel intensity in all frames. The tab 'Size' lets the user choose the size of the image. Maximum size is 2456 x 2054 pixels, which has been used in all tests. Choosing an area of interest instead did not seem to be necessary yet. 'Format' contains a list of all possible data formats: RAW, RGB 32 bit, Grey 8 bit etc. Especially the latter two are preferred: video capture is limited to only a few of these data types, and not all of them can then be processed in the PTV software. Grey 8 bit and RGB 32 bit are the best options for respectively monochrome and colour videos.

Useful tools are the 'Draw Measure' tool, 'Save Image' and 'Sharpness measure AOI'. 'Draw measure' allows the user to draw a line on the image, displaying it and showing how many pixels are in this line. It is a valuable tool to calculate the pixel size. 'Save Image' saves the current frame displayed as a PNG-image. Such a single image can be useful for example as a background image. 'Sharpness measure AOI' lets the user draw a square and gives the current and maximum sharpness of the image in return. It is very useful when focussing the lens or determining the working distance. The toolbars furthermore contain buttons to rescale the display and to capture single frames or display a continuous live video.

uEye Cockpit can record videos in AVI format. There is no way to save frames individually automatically with this software. Before recording, the maximum file size can be set (in Megabytes or number of frames) when necessary. It is important to set the framerate of the video recording to the same and fixed amount as the framerate of the camera (which was set under Camera Properties). Otherwise the framerate would not remain constant or frames are not fully recorded. There is a slider for the frame compression, called JPEG quality. By default, this is 75 out of 100. A higher quality is often wanted for capturing small particles in large quantities. A lower quality however lowers the file size and lowers the risk of frame loss (see next paragraph). Then the recording can be started.

Table 1: Settings overview

Options	Range settings	Settings final recording
Pixel Clock	3 settings	Middle setting
Framerate	Depending on pixel clock	40 Hz
Exposure time	Depending on framerate	2 ms
Format	Plenty, but only Grey 8 bit and RGB 32 bit are relevant	RGB 32 bit
Size	Plenty, but no need for smaller size	2456 x 2054 pixels
Master gain & gain boost	0 to 100 & on/off	No information, depending on other settings and lighting
JPEG quality	0 to 100	Not relevant: bitmaps instead of video recording

When recording, the number of frames received, saved and dropped are shown. Losing frames is bad: the framerate needs to be constant to calculate the timestep between frames and the velocities of particles during these timesteps. There is no way of knowing when this loss of frames occurred exactly from the video file after recording, so losing frames should be avoided completely. When it occurs, the recording must be stopped: as soon as the loss of frames has started, lost frames will start to pile up. Sometimes the first few frames of a recording are lost before stabilizing. Stopping and restarting the recording solves this problem. When the user has recorded enough frames but lost some at the end, the last frames can be ignored during frame extraction or when loading frames in the PTV software. When the recording had to be stopped prematurely, a new recording might be required. In the worst case, this has to be tried three or four times, or else the settings have to be lowered (framerate or JPEG quality mostly). 10 Hertz at 100% of the quality should not lose any frames during recording. 20 Hz or even 25 Hz at 100% quality has to be kept in check by the user: frames can be lost, most often at around 200 frames recorded. For higher framerates, the quality has to be lowered: 70% or 80% quality allows for the camera to record at 40 Hz. When enough frames are recorded, the recording can be stopped and saved.

Most PTV software, including Streams, cannot work with video files. Instead they use image sequences consisting of the individual frames as PNG, BMP, JPG files or similar. For the purpose of extracting the frames from the AVI video file, three software packages were tried. VLC and ImageGrab are not suited to extract every single frame from a 10 Hz video file or higher. VirtualDub can do this and at a decent rate, so it is recommended. If needed, only a part of the video can be selected to export. Otherwise the whole video is exported to individual frames. The user can choose between different formats: PNG, BMP, JPG, TIFF etc. The PNG format was chosen: no compression and loss of information (compared to JPG) without being unnecessarily large (compared to BMP). The frames are then extracted and saved one by one.

3.4 PTV analysis software

For the actual PTV analysis of the images, Streams was used. Everything written in this work relating to Streams is based on personal experience and the three manuals: System Theory and Design (Nokes 2009a), User Guide: Core Objects (Nokes 2009b) and User Guide: Field Objects (Nokes 2009c). Streams can be downloaded freely from its website, but a registration code is required to launch the software for the first time. Permission to use the software and a registration key can be asked from dr. Roger Nokes. Streams is standalone software, an advantage it has over PTVlab for example. PTVlab is a MATLAB add-on that requires (besides a MATLAB licence itself) the Image Processing Toolbox. Still, an effort was made to use PTVlab to compare results with Streams. However, the extensive set of filters in Streams and the PTV analysis algorithms that could be combined in pipelines made it obvious that in order to obtain the best results, Streams was clearly a better candidate. Three extensive manuals for using Streams can be found on the same website as the download for the software itself. Streams works with objects such as image



sequences, particle records, velocity fields etc. For every type of object, a set of views are defined, as well as a set of processes that can work on them.

Every PTV analysis starts with an image sequence that's newly created or loaded in. When creating a new image sequence, the user has to define the timestep between frames (in seconds) and the scale or pixel size (in millimetres per pixel). These were already obtained from the framerate during recording and the measured pixel size before recording. Then the image type is set (RGB 32 bit, Grey 8 bit etc.) and image files are added to the image sequence. The two most important views of an image sequence are the process view and the image view. The process view is a list of all processes created by the user to operate on the image sequence. In the image view, the frames can be viewed, and filters and particle identifiers created in the process view can be turned on and off to view their effect.

Before the actual PTV analysis, the frames need to be pre-processed. Streams has a few useful processes for this. First a form of background needs to be removed. This background can be a picture taken by the same camera as the image sequence. This should be taken at a moment when there is no sediment transport, so at low flow velocities. This leads to the problem then that it cannot be a picture taken at the moment of (or even right before or right after) the video recording. The resulting time difference can be enough for the bottom topography to alter, and the picture might be useless as a background proxy. A second option is to take an average of all captured frames, and use this as a background picture. This has produced better results near the bottom: the total timespan of one video recording is never more than 30 seconds. Within this period, the bottom topography does not change much. To get this average picture, Streams offers a process called 'Average Images'. Executing this process creates a new image sequence object that consists of a single frame: the averaged image. The picture itself is saved automatically under a user defined name, the image sequence is not saved initially. The newly created background image can then be subtracted from every single frame as part of a filter pipeline. Within a filter pipeline, the user can select and create a number of filters to operate in sequence. The 'Subtract image filter' is what is needed to subtract one image from all images in the sequence (in our case: background image from the captured frames). Do not confuse this with the 'Subtract Images' process! This is used to subtract subsequent frames from each other. Trying this as a form of pre-processing did not produce beneficial results.

When using RGB images instead of monochrome ones, it is possible to extract one of three colours: red, green or blue. Sometimes only using one or two of the colours results in a less noisy image (from particles that are out of focus mostly). To do this, use the Extract filter in the filter pipeline to remove two colours and retain one. Alternatively, the Remove filter removes one colour and retains the other two. An example is shown in Figure 12: most of the noise of particles out of focus is seen in the blue colour spectrum, and when the blue colour is removed, particles in focus can be seen and identified much more clearly.

A last filter that deserves to be mentioned is the Eliminate Objects filter. This filter was designed to remove large bright objects from the image that are not particles. In our case, a filter to eliminate particles that are out of focus would be a welcome addition to the filter pipeline. Unfortunately, the particles out of focus are less bright than the particles in focus, and more bright than the background. To remove them, a dual threshold of intensity or colour would be necessary, and the Eliminate objects filter is single threshold only. Choosing the right minimum and maximum pixel size of these particles in the filter removes them, but at the risk of also removing particles in focus. With the right colour(s) extracted, the particles out of focus are not as much of a disturbance as in full colour images or monochrome images, so the Eliminate objects filter is not necessary anymore.



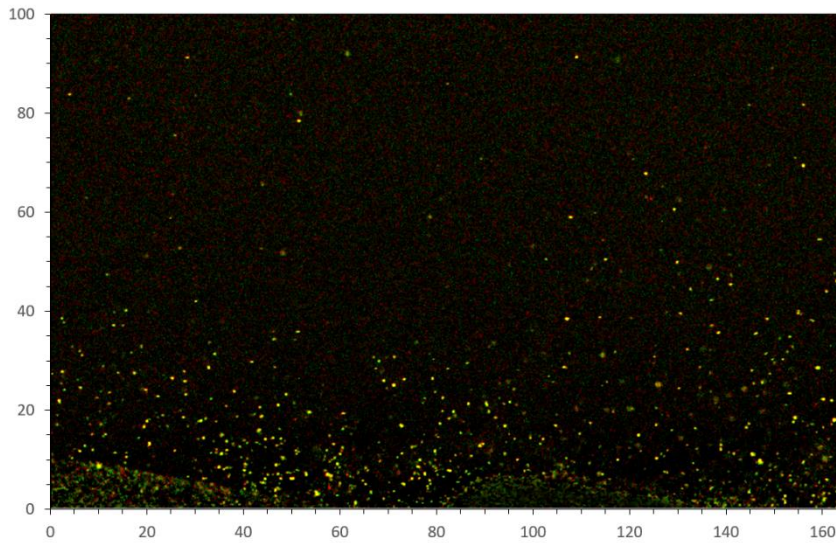


Figure 12: Illustrating the effect of extracting or removing colours. First filter: subtract background, amplify all colours 3x. Second filter: same as filter 1 + extract Blue. Third filter: same as filter 1 + remove Blue. Dimensions in mm.

After filtering, a new process can be created to identify particles (aptly named 'Identify particles' in the process view). This is a single process, not a pipeline of processes. There are four kinds of particle identifiers (PID's), each with their own set of criteria and algorithms, but they can't be combined in any way. These are the Single Threshold, Single Threshold (colour), Dual Threshold and Dual threshold Gaussian PID's. All of them use an intensity threshold to distinguish bright particles from a dark background. To avoid particles that are too large or too small, the user can define an upper and lower limit to the particle diameter.

Single threshold PID's have only a lower limit: every pixel with a larger intensity is potentially part of a particle. When running this algorithm, the PID looks for pixels with an intensity value higher than the user defined threshold. When such a pixel is found, all surrounding pixels that also match this criterium are added to this particle. Then the size criterium is evaluated: when the particle size is larger or smaller than the user defined threshold, the particle is dismissed: the algorithm does not produce any particle containing any of these pixels in the end result. An aspect ratio can be set by the user as well. The aspect ratio of an ellipse is the ratio between the largest and smallest axis. The colour variant of Single threshold PID allows the user to set the intensity of the three colours (red, green and blue) independently. The user can even decide if the algorithm should allow particles that fulfil only one of three colour criteria or must fulfil all three.

Dual threshold PID's let the user choose a peak intensity threshold (from 0 to 255) and an edge threshold intensity (0 to 1). The algorithm searches for pixels with an intensity value higher than the peak intensity threshold. When a pixel is found, it searches in the immediate vicinity for the pixel with the highest intensity value. Starting from this pixel, the PID creates a particle containing all connected pixels with an intensity value between the highest intensity value found (upper limit) and the highest intensity value multiplied by the

edge threshold intensity (lower limit). Once again the user can define a minimum and maximum diameter and a maximum aspect ratio to retain or dismiss a particle.

The Gaussian PID is unique in the way that it constructs elliptical particles around a maximum intensity in a region, based on the intensity thresholds. After using all PID's, the Single threshold (colour) and Dual threshold algorithms seem to give the best and most similar results (example in Figure 16). There are a rare few particles that one PID misses and the other doesn't, but there is no method to combine them. The user can decide to store info regarding the shape and size of particles. This is useful information for a few PTV analysis algorithms, so it is strongly advised to choose this option. The particles have a diameter between 0.3 and 0.7 mm, which in theory should then be used for minimum and maximum diameter in the algorithm. In practice, it is confirmed that these are good values to use: a smaller lower limit will add too many particles to the record that are actually noise. A larger upper limit would result in particles too close to each other being counted as one particle. The intensity (a value from 0 to 255) can be set relatively low (as low as 20): when the background is subtracted, it should be completely black (intensity value of zero for all colours).

The result of a PID is a particle record: a file that contains all frames with each of them the particles that were identified. PTV analysis of a particle record will be described later.

3.5 Experiments

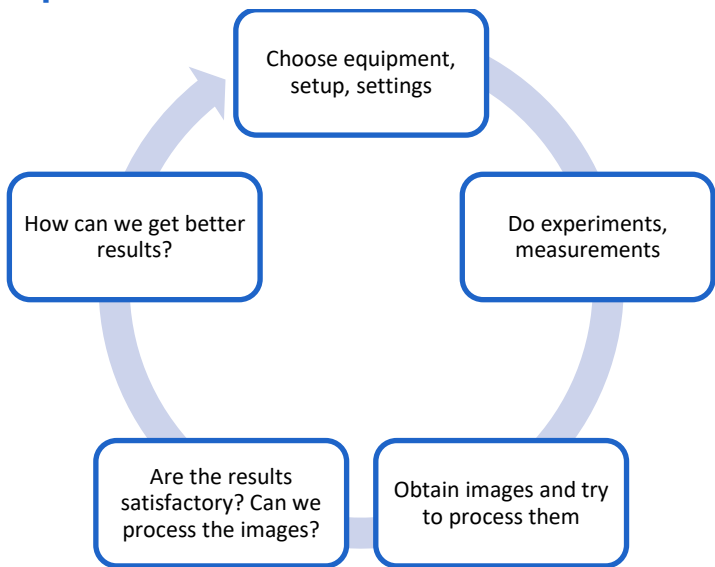


Figure 13: Cycle of experimenting, evaluating results and changing setup and settings.

Figure 13 shows how, systematically, different parts of the total process from recording images until visualizing sediment transport in function of depth were improved. The following list shows an overview of the series of tests that have been performed (a shorter, more schematic overview can be found in the Appendix A.5):

- 1) First tests with the bitelecentric lens, the camera and the uEye camera software. These tests were performed at a desk, with the camera and lens next to the laptop that was going to be used. These tests had the goal to familiarise the user with the different settings and their effect on the captured images. Settings that were investigated include the expert modus, pixel clock, the framerate, the exposure time, the image quality, size, format and gain. Expert modus is necessary to capture videos. The pixel rate determines the maximum framerate, and is best set as low as possible (lowest setting for 20 Hz framerate, medium setting for 40 Hz). The framerate and the image quality were considered to be limiting factors to not drop frames, as indicated by the software during recording. Format influences file size: a grey-scale image is smaller than a RGB image, and should therefore decrease the dropping chance. A short exposure time will freeze a frame better so that the particles are displayed more precisely, but it will also decrease the amount of light in the image. The Gain settings can help to illuminate the image when using the preferred short exposure times of ~2 ms. The working distance and pixel size were measured (later verified to be the same under water behind a window). The video file had to be converted to individual frames: 3 software packages were tested, out of which VirtualDub was considered to be the best (the only one to obtain all frames from a video file).
- 2) The first test of the bitelecentric lens and camera mounted in the dry cell of the current flume. First the window alone, then the window and water: these elements did not seem to distort the image. The following settings were used in these first tests: Greyscale colour format (smaller image files), 20-25 Hz framerate (a good balance of particle displacement distance and risk of frames being dropped), the use of a light screen (already in place), a water flow rate of 160 – 200 l/s (just enough to have sediment transport in the lowest part = 10 cm of the water column), exposure rate 1-2 ms (clear view of the particles). The first efforts were made to pre-process the images in Streams. Backgrounds from single images captured or average images were subtracted. Average images have the advantage that they are more representative for measurements that last a long time, especially at high water flow rates. The images captured for background specifically need to be taken when there is no flow: so before or after the measurement, when the water flow has been stopped. Efforts were made to further decrease the noise in the frames: the Eliminate objects filter was able to remove noise from particles out of focus, but at the cost of some particles in focus as well. Then the first particle records were obtained by experimenting with the different PID's. First the simplest PID was used: single threshold, monochrome (we were working with greyscale images). This provided good results when using the estimated grainsize in conjunction: 0.3 to 0.7 mm.
- 3) PTV analysis was practised, on particle records created by models in the Streams software. Then the particle records from image sequences captured were analysed (PTV analysis is discussed in a later chapter in more detail). The strategy of a first global optimization using matching-based costings such as Pseudo-correlation, followed by residual optimizations and clean-up with matching-based costings (Recent acceleration, Recent velocity, Time-averaged velocity, Magnitude velocity) was

developed. Forwards and backwards iterations resulted in better matches, which were then interpolated to create the first velocity fields.

- 4) The camera for the top-down view with the adjustable lens was tested. First measurements at the desk included different lens settings to evaluate their influence and to practice focussing at the same distance as the bottom when mounted above the current flume. The old method of setup was to get the lens settings correctly, then lock the camera and lens up in a transparent polycarbonate container so that it is protected from the water. The bottom of the container is below the water surface: this is to avoid reflections and distortions by the moving water surface when the camera looks through the air and water to the sediment bed. A detrimental flaw was discovered: the focus is not only dependant on the distance to the bottom, but also on the relative amounts of air, water and the material of the container between the lens and the sediment bed. These mediums have different optical properties (e.g. the refractive index), so trying to get the settings optimised before mounting seemed too impractical and complicated. It is practically impossible to get the lens settings right before filling the current flume, and with the closed container the operator can't adjust the lens anymore: this can only be done manually. To fix this issue, the container was replaced by a large aquarium made of glass, mounted between the camera lens and with an adjustable height. The aquarium is open at the top, but the walls are high enough so that there is no chance that water could reach the camera (in normal conditions, when there are no moving objects in the water, the water surface is smooth even at the highest water flow rates). The size and the open construction of the aquarium allow the operator to manipulate the lens from above while the current flume is active (water is flowing). The height of the camera and the aquarium can be changed, but this can be done more easily and safely when the current flume is empty. The water level rises slightly (in the order of 10 cm from 200 to 400 l/s flow rate) when increasing the flow rate, and it is best that the aquarium is only 1 to 5 cm deep in the water, to limit the force of the water flow pushing on this large surface.
- 5) A series of test with ADCP instruments was carried out in the current flume by Joachim Beckers, at water flow rates increasing from 200 to 400 l/s step by step. During these experiments, videos were recorded by the camera mounted in the dry compartment. Settings: monochrome, 20 Hz framerate, 2 ms exposure time. For every flow rate, three videos were recorded: a Lupolux Dayled 1000 90-Watt lamp as a light source, the earlier mentioned light screen as a light source, and both. The light screen results in an image with dark particles against an illuminated background, for the lamp it's the opposite, and when using both there are dark and light particles, where the relative amount varies in different parts of the image. Out of these 3 options, the lamp was concluded to be the best light source. In the images with the lamp as light source, there is less noise from particles out of focus (see Figure 5: the lamp only illuminates the part of the water column that is in focus, while the light screen illuminates the whole water column). The lamp is also more consistent in all directions and in time when it comes to brightness of the images, the light screen shows a noticeable variation in brightness across individual frames and across time. Combining both of them is needlessly

complicated: the part of the frame that is illuminated by the screen shows a white background with black particles, the part where the lamp is stronger displays white particles to a darker background. This makes pre-processing needlessly difficult. The reason for these recordings was to have a series of image sequences available for different flow rates. However, the recording settings have been improved after these experiments (see 6)): the framerate was increased (to have shorter displacements between frames), and the format was changed to RGB. The use of a lamp resulted in bright particles instead of black ones when using the screen, so it became interesting to retrieve the colour information of particles from frames and use it in the different processes and algorithms. The illuminated particles contain more useful information than dark ones in single colour images: red, green and blue, which can be extracted in Streams (see also Figure 12).

- 6) A series of experiments with both cameras where the exposure time was varied. Long-exposure videos were recorded, with the exposure time increasing from 16 to 100 ms. The flow rate was set at 200 l/s, both formats were used (RGB + lamp and monochrome + screen) and the framerate varied between 10 and 40 Hz, limited by the exposure time. The long-exposure images provide a good visual representation of the paths of particles: they are smeared out as lines which represent the path that a particle took within the exposure time. These lines showed that the particles mostly follow near-horizontal linear paths, especially higher up in the water column. Closer to the bottom, particles start to take off or come down after what looks like long jumps at 45 degrees or lower. These jumps span horizontal distances much larger than the width of the frames. Much less often, you can see clusters of particles do much smaller, more vertical hops. And finally, immediately after a dune crest, particles will fall or make turbulent motions. These types of movement are shown in Figure 14. As mechanics, they are very interesting, but to the total sediment transport, they contribute only a very small part. They are much more limited in time (short jumps do not occur continuously), space (they only take place near dune crests) and distance (distance covered is only a few centimetres). This type of images is very useful to visualize particle movement for the naked eye, but they could not be processed by Streams (or any other software available at that moment) to obtain particle records, velocity fields, density fields and sediment transport values. There was no software available to draw and identify particle tracks based on these images. The human eye can relatively easily interpret the lines correctly, but for computer algorithms, including those of Streams, it is considerably more difficult. Streams can draw ellipsoid particles, so PID algorithms were used to try to identify the linear features as elongated particles, and use their long axis as a measure for particle displacement during exposure. The Streams PIDs all use intensity or colour as thresholds, and on pixel level, this varies substantially within the particle paths of long-exposure images. With the Dual threshold (Gaussian) PID, instead of one large elongated ellipse, multiple small and narrow ellipses were drawn inside the linear features, perpendicular on the particle path. When using the other PID's, it is somewhat possible to draw elongated particles based on the linear features. Filtering will severely lower the depth of view: only the most intense lines (those that are most in focus) are retained (see Figure 15).

Despite removing a substantial number of particles, linear features that cross each other are still present in the filtered image. In this condition, the PID's have trouble distinguishing these particles individually, and will instead draw a large particle that comprises of two or more particle paths. The start and end points of the linear features have a lower intensity, so these intensity-based algorithms will often draw particles that are shorter than the linear features.




Figure 14: Close-up of frame 130 of the 49 ms exposure time image sequence. Top: suspended and long-range jumping particles. Right centre: shorter jump of cluster of particles. Bottom left: particles falling and moving turbulently after dune crest.




Figure 15: Particle identification of the linear features of the 49 ms exposure time image sequence. Left: close-up of frame 130 of the unfiltered image sequence. Centre: filtered image sequence. Right: Single threshold (colour) PID, 'particles' identified in red.

By taking the coordinates of the start and end of the linear features in the image view of Streams and knowing the exposure time, it was possible to calculate a limited amount of particle velocities. This can only be done one point at any given time by the user, not automatically or for multiple paths at once. It was concluded that the long-exposure images were useful to get an early overview and visualization of the particle displacement, but to get hard data, short-exposure images will be used. The particle velocities based on the tracks seemed to be consistent through their trajectory and close to the estimated water flow velocity, but never higher.

With all this in mind, the penultimate tests were executed on March 11, 2017. A flow rate of 195-190 l/s was chosen, and video recordings were made with both cameras (from above and from the side), one at a time. The Lupolux 90 W lamp was used as light source. Water depth was measured to be 73 cm. The adjustable lens of the camera above the water was set to focus on the bottom. Settings for the camera in the dry cell, first recording: framerate 40 Hz, exposure time 2 ms, RGB format, JPEG quality 75%. Settings for the camera in the dry cell, second recording: framerate 30 Hz, exposure time 2 ms, RGB format, JPEG quality 80%. In the first recording, a video file of 980 frames was recorded without dropping a single frame according to the software. In the second recording, 997 frames were recorded without any warning of dropped frames. Settings for the camera above the current flume, first recording: framerate 40 Hz, exposure time 2 ms, RGB format, JPEG quality 75%. 613 frames were recorded when the software warned that frames were being dropped: the last 40 frames of the video file cannot be used. Settings for the second recording: framerate 30 Hz, exposure time 2 ms, RGB format, JPEG quality 75%. Once again the software warned about dropped frames during the recording: the last 40 frames were removed, resulting in 484 frames with a supposedly constant framerate. The frames were extracted by using VirtualDub.

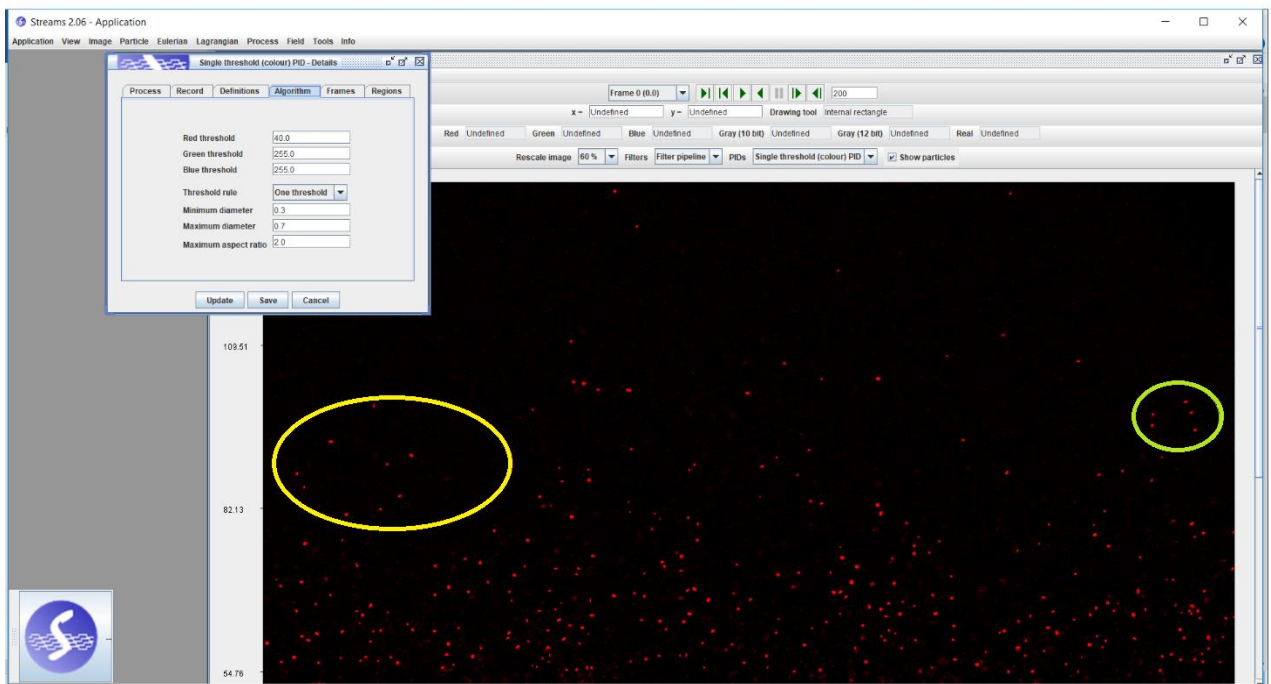



Figure 16: Image sequence of the 40 Hz recording of the sideways camera, filter and PID applied. Top: Dual threshold PID. Bottom: Single threshold (colour). The green and yellow circles show regions representative of the difference in which particles the PID's can find: 1 of the 5 particles in the green region of the single threshold is not found in the dual threshold. Out of the 11 particles in the yellow region, 3 are not identified by both algorithms.

The frames could then be loaded into Streams as image sequences. Filters and particle identifiers were created. The 40 Hz recording of the sideways camera was concluded to be the most promising. For pre-processing, an average image was created and then subtracted from all frames in the image sequence. The red colour was evaluated to provide the best particle signal compared to green and blue, so the red colour was extracted from the RGB images. The two best PID's were chosen and optimised for this recording specifically (see Figure 16). Both produced similar results, but in every frame there are particles that one PID detects and the other does not. However, this has barely no influence on the result after PTV analysis. Each PID produced a particle record: a sequence of frames that contains all particles per frame, with their shape and size information stored. In the particle record, PTV analysis pipelines were constructed, executed, evaluated and modified (more information on PTV analysis later). Search windows were set sufficiently large: the search window was shown to be one of the most determining factors in the spread of the calculated particle velocities (increasing the search window would increase the spread of calculated velocities significantly). The resulting particle paths could then be exported as Lagrangian paths or Eulerian fields (grids constructed by interpolation). The Lagrangian paths were overlaid on the original image sequence: this allowed the user to follow a particle visually along its proposed path, to then compare it with the path that can be followed with the naked eye. The most useful algorithms were deemed to be the Shape state-based algorithm, and for the matching-based ones the Time-averaged velocity, Recent velocity and Recent acceleration. The Eulerian field of one of these algorithms is shown in Figure 18.

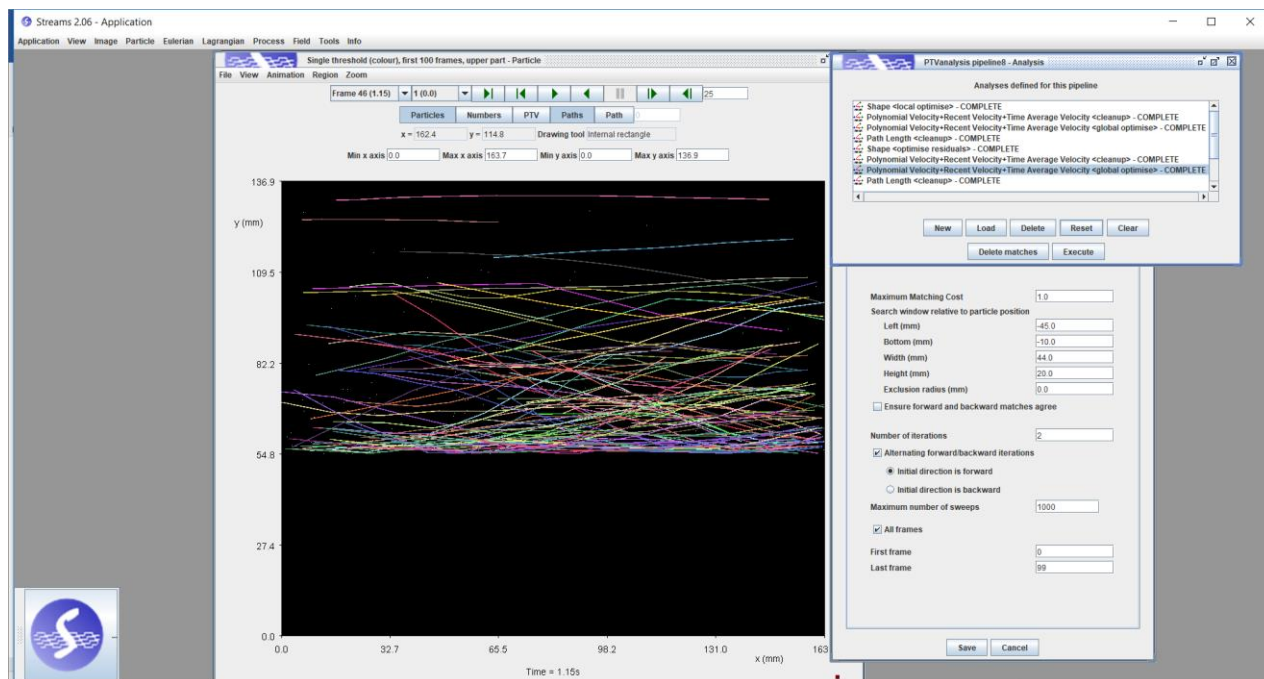


Figure 17: Particle record of the single threshold PID of the 40 Hz recording of the side-looking camera. A PTV analysis pipeline was executed. Left centre: particle view of the particle record showing the paths found by the PTV analysis. Upper right: PTV analysis pipeline. Lower right: optimization window of one of the algorithms in the PTV analysis pipeline, showing the maximum matching cost, search window and number of iterations.

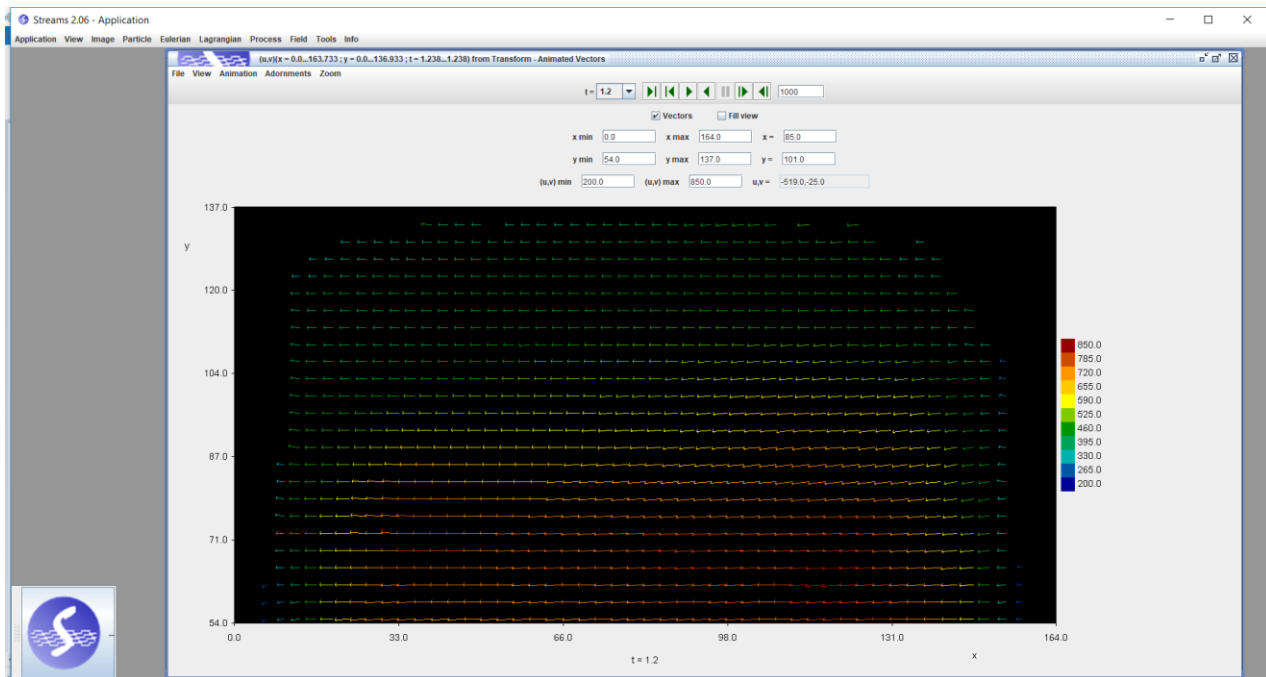


Figure 18: Eulerian field of the single threshold PID of the 40 Hz recording of the camera sideways. It shows a grid of interpolated velocity vectors, averaged over the first 100 frames.

Efforts were made to process the images of the top-down camera in Streams, but the images proved to be too complex. The background in these images is the sediment bed itself, so the whole frame consists of particles at all times, some moving, some resting. Subtracting an average image or subtracting subsequent frames from each other were used to only retain moving particles (see Figure 19). Other filter subprocesses that were tried are the Discrimination filters 1 and 2 (filters to distinguish particles from each other and from the background) and the colour Extract filter. Even when providing a (debateable) decent result, particle movement seemed too erratic and complex to match particles across frames correctly. The camera lens is focussed at the bottom, and this close to the sediment bed, particles seem to move by rolling or turbulently between the dunes, also perpendicular to the water flow. Particles will (partially) cover and (partially) expose other particles, so a substantial number of particles will appear, disappear, or change across subsequent frames without actually moving, so having these particles in the particle record causes a lot of errors. Rolling particles themselves change shape or intensity across frames, more so than jumping or suspended particles, so they are not easy to track either. But mostly the amount of particles of a moving dune rolling across the frame makes it impossible to use PTV analysis. This study will focus on the images captured by the lateral bitemporal camera: the chances of using PTV analysis on this part of the water column (0 to 10 cm above the sediment bed, instead of the sediment surface itself) are much better: lower sediment concentration and more predictable particle movement. Future research can revisit the top-down view of the sediment view: these first videos already indicate that a larger part of the sediment is moving by rolling and sliding than by

jumping or in suspension. But to quantify and visualise this, it is best to first to research the feasibility of PTV analysis on the lowest part of the water column.




Figure 19: Images of the recordings by the top-down camera. Upper row, from left to right: unfiltered colour image, image after subtracting average image (background), image after subtracting previous frame (background). Lower row, from left to right: monochrome image, image after subtracting average image (background) and Discrimination filter 1, image after subtracting average image (background) and Discrimination filter 2

Figure 18 shows that the particle velocity varies between 200 and 850 mm/s according to the PTV analysis on the captured frames. However, acoustic measurements with two ADCP's, performed by Joachim Beckers, result in a water flow velocity between 228 and 268 mm/s. This also fits the theoretical value of 228 mm/s much better (calculated based on the water flow, water depth and width of the flume) and those of the long-exposure images for the same water flow (~70 to 200 mm/s). Particle velocities that are larger than the water flow velocity are highly unlikely, a strong indication that something went wrong. More careful visual observation of the image sequences shows that the framerate of the video recording was not kept constant, despite the fact that the software did not report any dropped frames. The resulting video file is indeed of the user specified framerate, but the real timestep between the frames used does not correspond with the timestep that is associated with this framerate. At low framerates (15 to 25 Hz), sometimes subsequent frames are partially or completely identical. This can be seen when scrolling through the frames in Streams, or by using the Subtract Images process: subsequent frames are subtracted from each other, and when the frames are (partially) identical, the resulting image is (partially) black. When increasing the framerate (up to 30 and 40 Hz), it becomes increasingly more common that the real timesteps between the frames used is larger than the timestep of the video framerate. It is clear that the laptop that has been used up to now is

unable to handle the incoming data sufficiently and/or that the currently used method (recording an AVI video file) is not suited for this research. Because of this, a new computer and a new method of capturing images were tested.

Computer specifications: Windows 10 Pro N operating system, Intel® Xeon® CPU E3-1240 v3 @ 3.40 GHz processor, 32 GB RAM, Intel® SSD hard drive.

Image capturing tool: instead of uEye Cockpit, a software tool developed at Flanders Hydraulics Research was used, under the name 'WL Cam Triggered', supplied by Stefan Geerts. The operator only uses uEye Cockpit to create a parameter file. This file contains all camera settings: framerate, exposure time, format etc. The new camera tool reads this parameter file and sends it to the camera. The camera creates bitmaps with these settings that the software tool can save directly.

3.6 Revised recording software and settings

With the new computer and software tools, the new limits of this setup were tested again. The software tool captures images directly, instead of videos. These images are given a timestamp. Using this, the framerate can easily be checked for irregularities. At 40 Hz, frames will still be lost after some time. The images captured are bitmaps (.BMP files). Data transmission is dependent on the framerate and the image quality: multiplying the image size of one frame (in Mbytes/frame) and the framerate (in frames/second) results in the data transmission (Mbytes/second). At 20 Mbytes per frame, bitmaps are relatively large compared to video files and other image formats. A 20 Mbyte file every 0.025 seconds results in an 800 Mbytes/second data flow. This is a lot, even for a powerful computer such as described earlier. The uEye manual proposes a bandwidth of 400 Mbytes/s for USB3 connections. The computer performance was monitored during following tests: the CPU, the RAM memory usage and the hard drive. Tests were performed at 40, 20, 16 and 10 Hz framerate (which corresponds with 800, 400, 320 and 200 Mbytes/s respectively). The CPU was able to handle all of these data inflows, and there was plenty of RAM memory to be allocated, even at 40 Hz. The limiting factor is the writing and reading speed of the hard drive: at 40, 20 and 16 Hz, the maximum capabilities of the hard drive are reached constantly, and at some point, frames are dropped. Only at 10 Hz (or 200 Mbytes/s) is there no loss of frames and is the hard drive working below its maximum speed. 10 Hz is a low framerate to be used for particle tracking in our current setup, even at the lowest water flow rates (200 l/s). To get a data flow of 200 Mbytes/s together with a high framerate, the file size needs to be reduced. To do this, pixels can be binned: a row or column of pixels can be grouped together to form one pixel. This however reduces the amount of detail in a frame. The smallest amount of binning (and smallest loss of detail) while still having equidimensional pixels is a double horizontal and vertical binning: a square of 4 pixels (2 by 2) is reduced to 1 pixel. This reduces the image file size by a factor of 4: 20 Mb bitmaps become 5 Mb bitmaps. To stay within the limit of 200 Mbytes/s, a framerate of 40 Hz was used for the final measurements. The loss of information

could lead to problems with identifying particles and matching algorithms, but no problems were encountered when processing the binned images.

The settings for the final measurements are: flow rate 200 l/s (lowest flow rate to get consistent bottom transport, so small displacements of a low density of particles, which are easy to track and match); framerate 40 Hz and 2x horizontal and vertical binning(to stay within the limit of 200 Mbytes/s data flow, with a minimal loss of information and a sufficiently small displacement of particles between frames); RGB images and the use of a lamp(so that specific colour information of particles can be used or filtered), direct capture of individual frames as images with timestamps with the software tool (to easily verify if the framerate remains constant). 1007 frames were captured. The total time between the first and last image is only 0.256 seconds or 1% longer than expected from a 40 Hz framerate. No skips in the framerate were found. A schematic workflow of the final recordings and PTV analysis can be found in the Appendix (A.6).

3.7 PTV analysis

The images are loaded into Streams as an image sequence. An averaged image was taken to subtract the background. Then the three individual colours were evaluated (red, green and blue). The Blue-coloured image was considered to be too noisy, but the Green and Red images provided a clear distinction between particles in focus and the dark background. The blue colour gun was eliminated, red and green were retained together (see Figure 12: this is frame 50 of this image sequence). Regions to exclude from the particle identification algorithm were drawn: the sediment bed itself, including the ripples. When working with a small subset of frames (e.g. 100 frames or 2,5 seconds), the ripples could be considered stationary, but for the whole set of frames, it is important too enlarge the regions so that the dunes are covered in every frame. Regions can be set as 'internal' or 'external': to include the area of the frame within the boundaries drawn or exclude. In all processes however, the user is asked to add the regions that should be excluded from the process. In order to exclude the sediment bed and the dunes, internal regions should be drawn and added to the PID process. To reduce calculation time when optimizing algorithms, the first 101 frames were used instead of all 1007 frames.

For particle identification, the Dual threshold PID was chosen. The Single threshold algorithms struggle more when particles are clustered. A Dual threshold PID can at least try to distinguish between two particles that touch each other, by having a lower limit in intensity based on the maximum intensity found. Single threshold PID's don't have such a lower boundary, and will always see two touching particles as one particle (that is then often considered to be too large, and ignored completely). The Dual threshold Gaussian algorithm is complex and creates elliptical particles. If we want to use shape-related information in our PTV analysis, the normal Dual threshold could provide more detailed information.

One of the most important steps in Particle Tracking Velocimetry (PTV) is matching particles to the same particles in previous or following frames correctly. The result of this are particle paths from one frame to the next frame. With this temporal and spatial information, particle velocities can be calculated. To create paths, Streams allows the user to construct a pipeline of processes. The pipeline process is called 'PTV analysis pipeline', and can be found in the list of processes of particle records. For every subprocess in the PTV analysis pipeline, three main components have to be defined: the analysis type, the costing and further optimisation.

Every pipeline starts with a global optimisation type or local optimisation type of process. The first process will result in the initial matching of particles, from which the second process starts, which adds, removes, or creates different matches. The third process starts with the result of process 2 etc. A global optimisation results in particle matches for every frame pair where the total cost of these matches is minimal. Cost is a parameter defined differently for every costing, but the principle is the same: the best matches according to the costing algorithm have the lowest cost. For example, in the Intensity costing the lowest difference in intensity of two particles results in the lowest cost to match these particles. A local optimisation pairs particles p1 and p2 when the best match (= lowest cost) for p1 is p2 and vice versa. This does not take the total cost of a frame pair into account and generally results in less, but higher quality matches. After the first local or global optimisation, the result is often not yet good enough. For the next process, four types of analysis are available: another local or global optimisation, a residual optimisation, or a clean-up. A residual optimisation takes the unmatched particles and tries to match them, it does not change the existing paths. A clean-up evaluates all existing matches and removes the matches with costs that are too high. So no new matches are created, only existing ones that are retained.

There are 17 different costings. 8 of them are state-based costings, 9 are matching-based. Before matching, the only characteristics known of particles are their x- and y-coordinates and intensity (and shape and/or colour, when specified during particle identification). State-based costings only use this information in their algorithms. The Adjacency, Centre of mass, Correlation and Pseudo-correlation costings are based on the grouping of particles, patterns in the particle's locations and the cross correlation of intensities in user defined regions. All of them can be used to a certain extent, but they tend to make mistakes when particle densities are low (not enough particles present to correlate with or to determine centre of mass or adjacency). Such low particle densities can occur further from the bottom, in the upper part of the images. This type of error is shown in Figure 20. The left figure shows the correct path of the particles, the other state-based costings in the right figure match the lowermost particle with the uppermost particle in the next frame. This type of error tends to occur more when using Adjacency, (Pseudo)Correlation, Centre of Mass costings, and especially Correlation. The Distance costing uses a prediction of the distance travelled of a particle between frames. It is not very useful in our case: it already assumes a partially known particle velocity, and particle velocity can vary across the water column. What remains are the Shape, Intensity, and Colour costings. They

use properties of individual particles that can be obtained from our images. Colour can be considered a more sophisticated form of the Intensity costing, so when using colour images (in the final version, Red and Green were retained) there is no reason not to use Colour instead of Intensity. Shape gives similar results compared to Colour, so they can be used separately or combined (the software allows for one process to hold multiple costing algorithms. Even their relative contribution to the total cost can be defined). The number of matches, number of paths and average path length when using a state-based costing are grouped together in 'Single state-based costing' in Table 2. The statistics of Colour + Shape is shown separately. Note that this information does not represent the quality of the matches made by these pipelines in any way, it only quantifies how many particles are matched in how many paths and the average path length!




Figure 20: Close-up of particle paths from state-based costings. Left figure shows particle paths (green) of two particles (red) based on the Colour and Shape costings. Right figure shows the same region and particles, paths are the result of Adjacency, (Pseudo)Correlation and Centre of Mass costings

Matching-based costings use particle velocities and accelerations in their algorithms. This means that they can only work for particles that are already part of a path. Because of this limitation, the first process in any PTV analysis pipeline must use one or more state-based costings, to create the initial paths. Matching-based costings can then be used to refine these paths. Least-squares velocity, Local velocity, and Space-averaged velocity use the velocities of neighbouring particles to determine the best match for a central particle in their algorithms. These options are not useful for sediment particles: although all particles follow the current flow more or less, individual particle behaviour might vary enough to favour alternatives. Magnitude velocity, Recent velocity and Recent acceleration will match particle p1 in the frame f1 and p2 in the following frame f2 based on the velocity or acceleration of p1 from frame f0 to f1 and of p2 from f2 to f3. These are better options than those earlier described: visually, the acceleration and velocity of particles barely changes along their individual paths. However, there are even better candidates: Polynomial velocity and Time-averaged velocity costings. These algorithms take into account the whole path of a single particle and its velocity along this path, not just the velocity in the previous or next frame pair. There is one costing left: Path length. This allows the user to remove short paths (path length is based on the number of frames that they span, not the physical distance that a particle travels). A good path should ideally span a lot of frames: a particle was identified and successfully tracked along many subsequent frames. Short paths are more likely to be wrong and the user might prefer to remove them.

For every subprocess in a pipeline, further optimisation options include the Maximum Matching Cost (MMC), Number of iterations and Search window. The Maximum Matching Cost defines an upper limit to the

cost of potential matchings to take into consideration. All particle pairs with a cost higher than the Maximum Matching Cost will not be considered during a local optimisation, reducing the computing time and effort. In a clean-up, all pairs with a higher cost than the MMC are removed. For example, when using the Path length costing in a clean-up, all paths shorter than the inverse of the MMC are removed. In general, increasing the MMC retains or creates more matches. Number of iterations will determine how often a single process goes through all frames. The standard setting is one iteration, from the first frame until the last one. Two iterations, the first going forward and then backwards, has proven to be beneficial. More iterations did not change the resulting matches. The search window defines an area of the frame in which the algorithm searches for potential matching candidates. This leads to a significant decrease in processing time, not having to take the whole frame into account. The search window is rectangular and is defined in relation to the particle for which a match is sought. The location and borders of this rectangle are determined by an estimate of the minimum and maximum distance travelled possible by any given particle. So, the search window in frame f2 for particle p1 in the previous frame f1 will be located against the water flow direction. With a water flow estimate of 230 mm/s, a framerate of 40 Hz and the assumption that particles will never go faster than the water velocity, all potential candidates to match with p1 are found within 0 and 6 mm of p1, against the direction of the current. Assuming all particles are in motion when not resting at the bottom, a window of 0.5 mm to 6 mm has proven to be adequate. The upper and lower boundary of the searching window were then set at 6 mm above and below p1.

After careful optimisation and trial-and-error, the following algorithm was constructed:

- 1) Local optimisation, Colour and/or Shape, MMC = 2
- 2) Clean-up, Polynomial velocity + Time-averaged velocity, MMC = 0.5
- 3) Global optimisation, Polynomial velocity + Time-averaged velocity, MMC = 1
- 4) Clean-up, Path length, MMC = 0.3
- 5) Residual optimisation, Colour and/or Shape, MMC = 2
- 6) Clean-up, Polynomial velocity + Time-averaged velocity, MMC = 0.5
- 7) Global Optimisation, Polynomial velocity + Time-averaged velocity, MMC = 1
- 8) Clean-up, Path length, MMC = 0.3

The first process is to create an initial set of matches, using only state-based costings (first row in Table 2). As described earlier, the Colour and Shape costings were preferred (second row in Table 2). Using either of them or both does not influence the final result by much however. The Maximum Matching Cost was set relatively high, to get a large number of paths, of which the quality could then be further improved. The first step is to remove outliers: matches within paths that are incorrect and are expressed as sudden changes within a larger path. Often, when a particle disappears in the next frame, another particle is chosen: the second-best option, yet still an incorrect match. This results in a sudden change of direction of the normal path, which can be detected by the Polynomial and Time-averaged velocity algorithms. Combining them (alternatively, Polynomial velocity + Recent velocity give a similar result), and with a low MMC of 0.5, wrong

matches are deleted effectively. The third subprocess is a global optimisation, going through all particles, with the Polynomial Velocity and Time-averaged costings. Particles that have been matched previously now have velocity information that can be used to create new and improved paths. A MMC of 1 is not too strict or too forgiving. A clean-up then removes all paths shorter than 4 frames: these paths are more likely to be incorrect. Additionally, these particles are then made available for residual optimisation. A residual optimisation based on Colour and/or Shape tries to create matches using only unmatched particles. This means that the matches that have been optimised by the velocity-based algorithms remain intact. A second clean-up and global optimisation with Polynomial and Time-averaged velocity costings are added to the pipeline: this method of repeating results in a large number of high-quality matches and long paths. See Table 2, last two rows: the increase in the number of matches and paths is apparent, but to verify if the correctness of paths has been improved, the user has to take a good look at individual paths. Figure 21 shows how adding this second set of processes improves the matching quality. Comparing the result after step 4 of the pipeline and the result of the whole pipeline, it is clear that correct matches are retained (in the upper part of the images, all paths were retained, and these are all correct). Meanwhile, some incorrect paths, seen here mostly in the lower centre and right corner of the left figure, have been removed. More paths have been found close to the bottom (right corner of both figures). The final process is a second clean-up, removing all paths shorter than 4 frames. All global, local and residual optimisations have 2 iterations each, one forward and one backwards through all frames.

Table 2: Improvements to the PTV analysis pipeline, and how this affect the number of matches, matched particles and path length. Total number of particles: 68530. Note that these numbers give no indication of matching quality.

Algorithm improvement	Number of matches	Unmatched particles	Matched particles	Total paths	Average path length (frames)
Single state-based costing	34290 - 37602	14790 - 17988	50542 - 53740	15179 - 17145	3 - 3,38
Colour + Shape	36097	18018	50512	15127	3,39
Matching-based clean-ups and optimisation	26934 - 60510	4257 - 40369	28161 - 64273	1939 - 4475	12,69 - 14,89
Polynomial + Time-averaged velocity, Path length clean-up	44415	21206	47324	3621	13,27
Two iterations	50749	13666	54864	4827	11,51




Figure 21: Effect of repeating processes within a pipeline. Left figure shows a Lagrangian path field overlay on frame 48 of the image sequence. PTV analysis pipeline: steps 1 to 4. Right figure shows the same region, but the PTV analysis consisted of steps 1 to 8.

4. RESULTS

4.1 Particle velocity




Figure 22: Part of the image in which the particle velocity, particle density and transport are calculated. X- and Y-axis in mm. Top: image before filtering and particle identification. Bottom: image after filtering and particle identification. Particles identified in red, particle (clusters) unidentified in yellow-green.




Figure 23: Comparison of the paths of frames 44 (left) and 4 (right) of the image sequence. Green lines indicate particle paths from the current frame to the next and (if present) the frame after (path length = 3). Region corresponding with Figure 22 in purple.

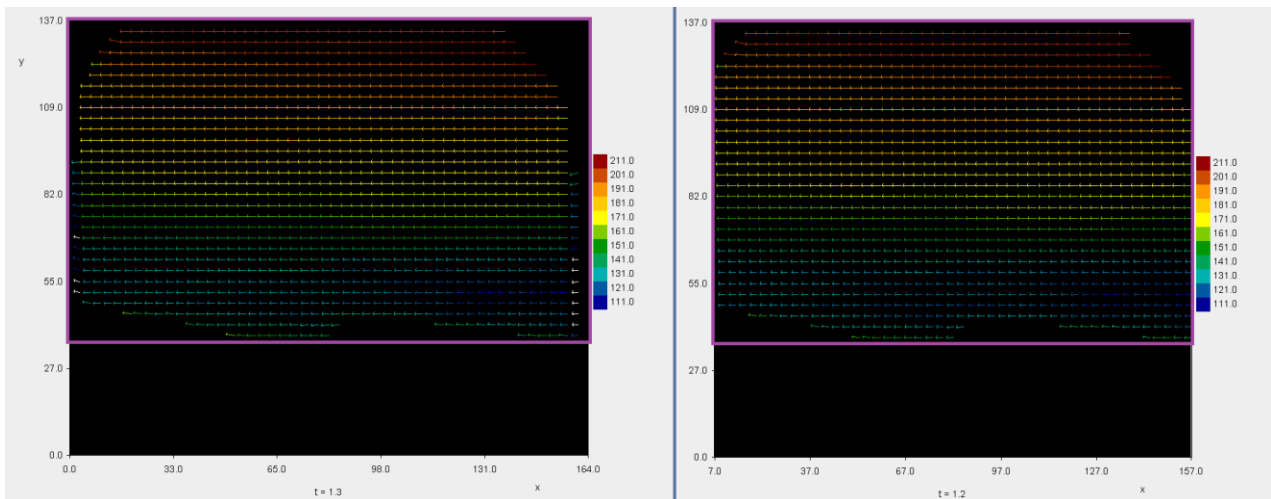


Figure 24: Eulerian velocity field, averaged in time. The right one has removed the first and last 10 frames of the record, as well as the left- and rightmost 6 mm of the frame. Region corresponding with Figure 22 in purple.

Figure 23 shows two frames of the filtered image sequence, with a Lagrangian path field overlay. A Lagrangian path field holds all paths from a PTV analysis of a particle record. It can be used as an overlay on the original or filtered image sequence, so the user can follow particle across frames to see if their real paths are the same as the paths from the PTV analysis. The user can set the path length of the paths shown: a path length of 3 means that the paths between this frame and the 2 following frames (when available) are shown. The left frame (frame 44 of 101) is representative of most frames: particle paths are in general linear, mostly horizontal and to the left. In the initial frames (and to a lesser degree the last frames) of a path field, there are more incorrect paths. Frame 4 (right part of Figure 23) for example shows a lot of steeper paths than found commonly. Tracking the particles with the naked eye shows that many were matched incorrectly.

The particles in the first 10 frames do not have enough previous frames to create long paths and to use the particle properties in previous frames for the different steps in PTV analysis. This causes incorrect matches, especially near the left edge, where particles leave the frame. For these particles, there are no correct matches in previous frames (no frames yet) and no correct matches in following frames (particles have left the frame). These particles are only a very small part of the total particle record, so the error caused is not noticeable when using 101 frames, but leaving the first and final 10 frames out of the final calculations will remove this small error easily. This is shown in Figure 24: the figure on the left is a Eulerian velocity field of the time-averaged velocities of all 101 frames, the grid on the right did not take the first and last 10 frames into account. The difference is barely noticeable. A related phenomenon in the left grid is the low velocities at the left and right boundary of the screen. As explained earlier, for particles very close to the edge of the frame, their correct match might not be present in the previous or next frame, because they have not entered the frame yet or they have left it. This leads to incorrect matches with other particles that are very close near the frame boundary too, hence the low velocities. In the right figure, 6 mm (which equals 1 search window) is taken from the left and right edge of the grid. For the top of the frame, this adjustment is not required: the particle density is so low that incorrect matches are much less likely to occur. The lowermost 15 mm of the purple region of the grid, between the top of the ripples and the sediment bed, shows an increase in particle velocity. This will be discussed later.

In order to evaluate the particle velocity in function of the depth, a Eulerian field was created. A Eulerian field interpolates the particle information (particle velocity, particle density, particle shape etc.) to a grid. A velocity field for example is a grid-based representation of the particle velocities. The standard (x, y) grid is 51 by 41 grid points, 3.275 mm by 3.42 mm spacing, which works well for the particle velocity calculation in this case. For every frame of the particle record containing particle velocities, there is a velocity field frame. 2 types of information can be extracted from a Eulerian field: vectors or scalars. Vectors include velocity vectors (with a u and v component, in the x and y directions respectively) and acceleration. Scalars include the velocity magnitude, kinetic energy and partial derivatives of u and v. A complete list can be found in the calculator view of a velocity field. The software allows for the user to create custom calculators. These calculations can include u, v, mathematical operators (+, -, *, /, square root etc.) and averages (average over x, y, x and y, t or x, y and t). In our case, it is interesting to know how the particle velocity evolves in function of the depth. A standard calculator can be used that calculates the average particle velocities in the (x, y)-grid. Then these values were averaged from the first to the last frame. A second, custom calculator was constructed to calculate the standard deviation of the individual horizontal velocity interpolations compared to the time-average horizontal velocity of each depth interval. Both calculations can be exported to a CSV-file, that could then be loaded into Microsoft Excel. Here, the particle velocities averaged for each depth interval were calculated, which are shown in Figure 25.

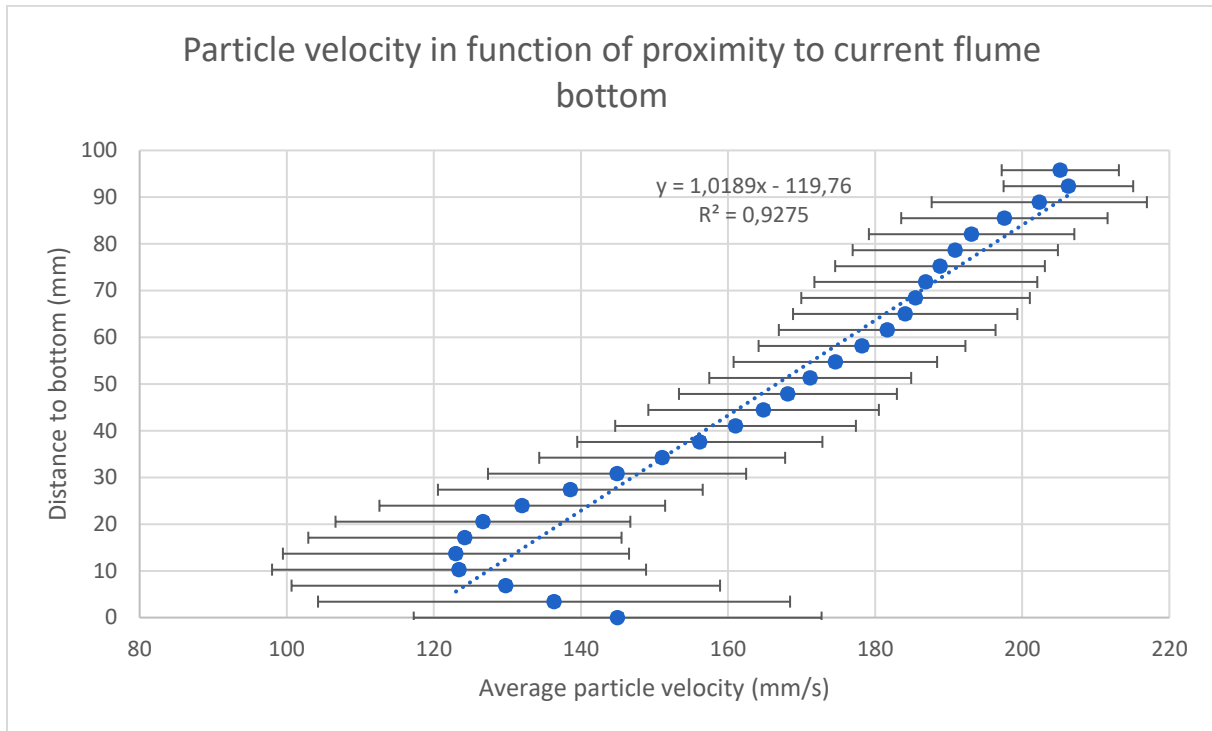


Figure 25: Average particle velocity near the sediment bed. Error bars represent one standard deviation.

4.2 Sediment transport

Having calculated the particle velocities, it should be possible to make an estimation of the sediment transport rate if also the particle density through the water column is known. Streams allows to create a density field out of a particle record. The process creates a grid, counts all particles inside each cell and calculates the particle density for each cell and for each frame. It is not possible to create custom calculators for density fields. However it is possible to use the standard calculators to obtain values for the time-average density for depth intervals and the error/variation for each depth interval over the total timespan. For the average density, the same grid can be used as in the creation of the particle velocity field. The time-average can be calculated in Streams, and in Excel the average for each depth interval can be calculated (Figure 27). The standard calculator that approaches a standard deviation most closely is the intensity of the fluctuating density (named rho'rho' in the calculator view). This is the difference between the density (rho) of a grid point for a frame and the time-averaged density of this grid point, squared. By creating grid elements that stretch horizontally across the whole frame (so the grid is composed of a single column of vertically stacked cells), and calculating the time-average of rho'rho' across all frames in Streams, a variance is calculated. This can be exported as a CSV-file. In Excel, the square root is calculated, which equals the standard deviation of the density for this grid. Taking smaller grid elements will result in a much larger standard deviation, often larger

than the average density. At such a small scale, particle density can vary severely and rapidly: there are frames or parts of where there are almost no particles sometimes, followed by an influx of a large particle cluster. By taking the average of a depth interval, the variation is reduced to a more acceptable degree. Added to the standard deviation, a positive systematic error of +10% is added. After studying the frames and the particle identification results visually, it was concluded that there will always be particles that can't be identified: in a few rare cases, single particles are too large, but more often it's clusters of particles that are tough to interpret for the software. When 2 or more particles are too close to each other, the PID will consider this object to be too large to be a particle, and only identify one or none of the particles. Clustering seems to increase linearly with the particle density, so a relative error of +10% of the average density was concluded and added to the positive error.




Figure 26: Close-ups of the particle identification applied to frame 11 of the filtered image sequence. These illustrate the positive error of the particle density. Top left: two particles are too close to each other and are not identified by the PID. To the right, only one particle out of a cluster of four is identified. Top right: one abnormally large particle at the top and a pair of two large particles at the right are not identified by the algorithm. Bottom left: a gust of sediment can make it impossible to identify individual particles, which makes it difficult to estimate the error at the lowest part of the frame (10 – 15 mm).

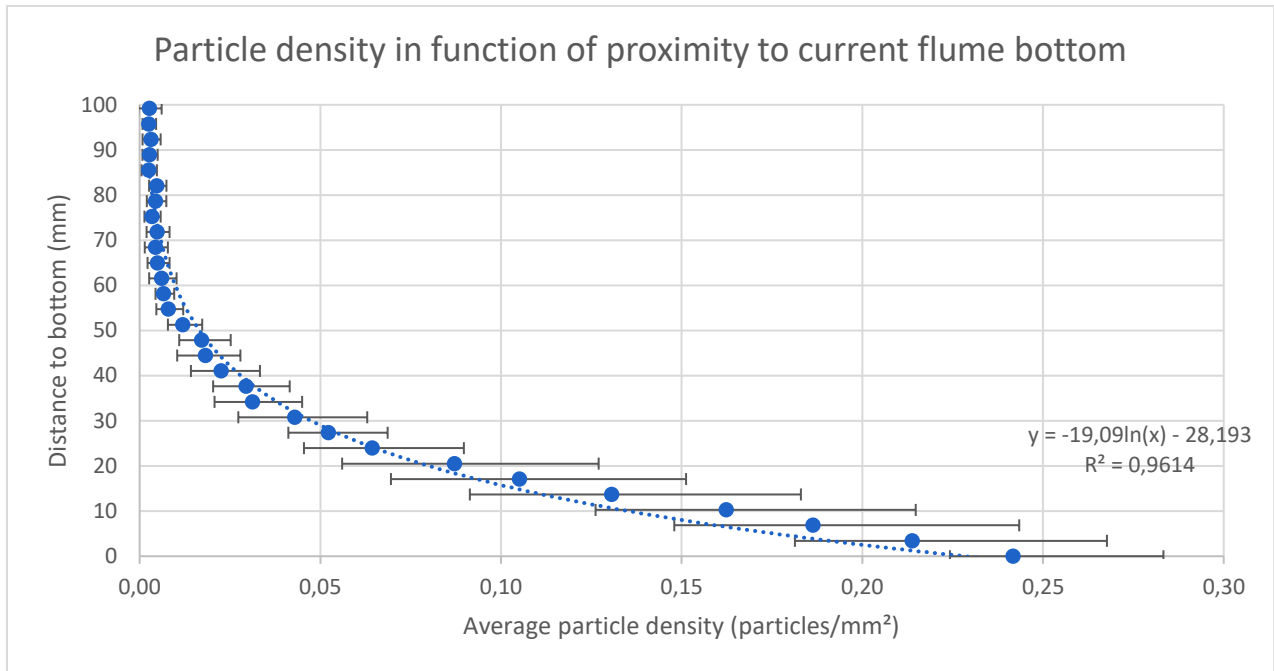


Figure 27: Average particle density near the sediment bed. Error bars represent one standard deviation + 10% positive error.

Multiplying the average densities (in particles/mm²) for 101 frames for each depth interval with the average interpolated particle velocities described earlier (also for 101 frames and for the same depth intervals, in mm/s) results in an estimate for the particle transport rate (in part/mm/s). The average particle diameter was determined from the grain size distribution from the results of the Malvern Mastersizer: $d_{50} = 501 \mu\text{m}$. Using this value, the average volume of a particle is 0.066 mm^3 . Particle material density (not concentration, but the mass/volume ratio of a single particle) is also known: $1220 \mu\text{g/mm}^3$. Adding these to the calculation, the transport rate is known as the mass of particles within working distance per unit of depth per unit of time ($\mu\text{g/mm/s}$) (Figure 28). The error bars are the combined error of the particle velocity (standard deviation) and density (standard deviation, +10% positive error). Integration over the complete vertical leads to a transport rate in view of $56,2 \text{ mg/s}$ (Table 4).

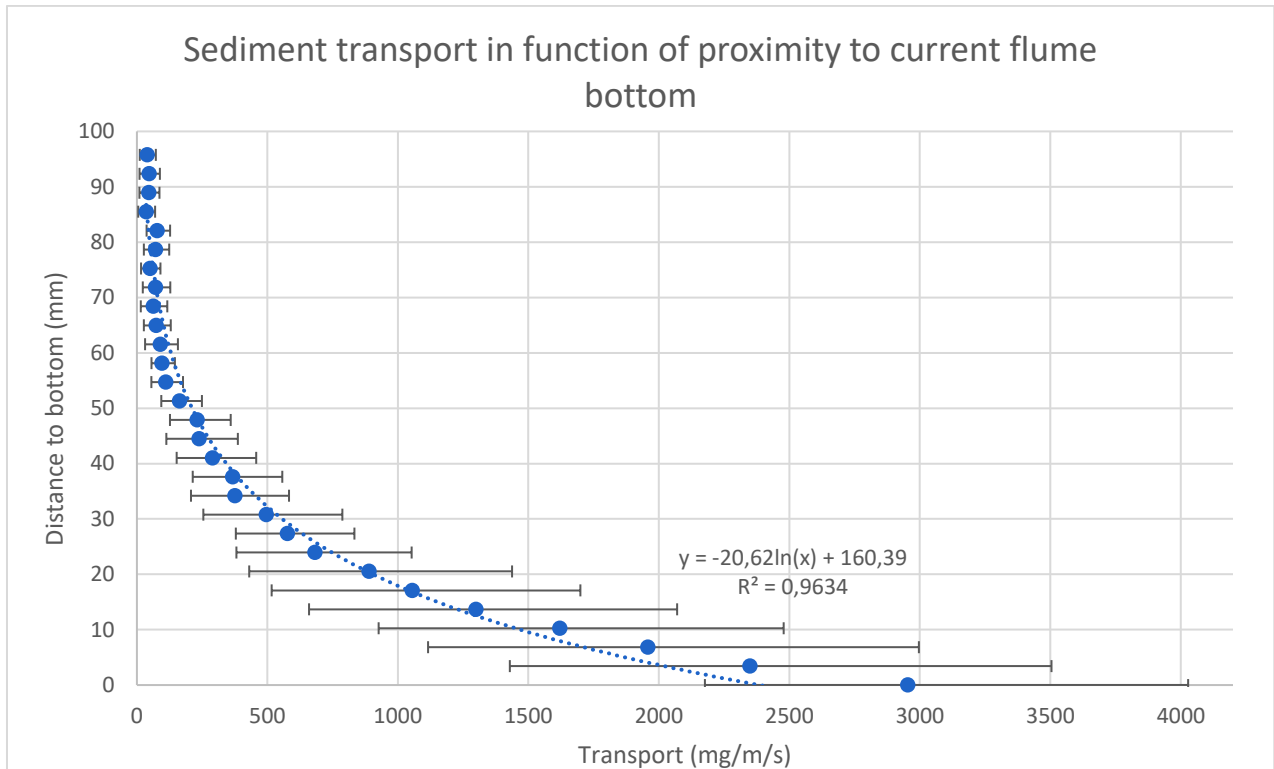


Figure 28: Sediment transport near the sediment bed.

The top of the ripple seen in the image sequence moved about 6 mm across 1000 frames = 25 seconds. The calculated bedform velocity is then 0.24 mm/s. The maximum height of the ripple is 9 mm. In the foreground, two ripple crests are visible, with a distance of 13 cm between them. Using this information, assuming a triangular cross-section of the ripple and the width of the flume (120 cm), an approximation of the volume of the ripple of 702 cm³ can be calculated. The material density (1220 kg/m³) and the porosity in dry atmospheric conditions of the sediment (36%) are known. A rough estimate for the mass of the sediment in the ripple is 0.548 kg. Using the ripple displacement velocity and the ripple wavelength, a transport rate of 1.01 g/s can be estimated (Table 3).

Table 3: Sediment transport based on ripple displacement

length ripple (m)	height ripple (m)	width flume (m)	volume ripple (m ³)	material density (kg/m ³)	porosity (m ³ /m ³)	mass sediment in ripple (kg)	velocity ripple (m/s)	transport rate (kg/s)
0,13	0,009	1,2	0,000702	1220	0,36	0,548	0,00024	0,00101

5. DISCUSSION

Figure 25 shows the decrease of the particle velocity towards the bottom. At the top of the frames, furthest from the bottom, the average particle velocity of ~205 mm/s is close to the water velocity, which is approximately 230 mm/s theoretically. The deviation is also small, around 10 mm/s: all particles move at the same speed. This is to be expected: higher up in the water column, particles are being transported as suspended load, they move quasi-horizontal near the water flow velocity. The average particle velocity decreases linearly toward the bottom whereas the standard deviation increases. This increased standard deviation could be caused by an increase in particle behaviours (suspended, jumping, falling down) towards the bottom and/or an increase in the amount of incorrect matches towards the bottom (more unpredictable behaviour, larger particle density). The amount of each of these factors contributing to the total error is not evaluated yet. When approaching the bottom topography (top of the underwater ripples) and until the deepest point (valleys between ripples), the average particle velocity shows a slight increase (between 0 and ~15 mm in Figure 22 and Figure 25). This increase is unexpected; however it still falls within one standard deviation, which is at its largest here. Velocity calculations could be considered to be the least reliable in this part of the water column: examination of images of the top-down camera and long-exposure images have shown that there is a lot of irregular movement of particles here (see also Figure 14). This might lead to more incorrect matches, thus giving a false (too high) representation of the particle velocity here.

To know the exact transport rate (in $\text{mg}/\text{m}^2/\text{s}$), the depth of view of the camera needs to be known exactly: what is the minimum and maximum proximity of a particle to the lens to be identified? This depends on the camera lens, the particles, the lighting but ultimately to the particle identification algorithm: particles that are too close or too far away are not intense enough to be identified. To get an approximation for this is very difficult, but a good estimate should not be larger than 25 cm (depth of focus of the lens) and not smaller than 1 cm. However, the evolution of the transport rate with proximity to the bottom is clear from the graph. In Figure 28, the transport rate is nearly constant first, then increases exponentially towards the bottom. Particle velocities (Figure 25) decline towards the bottom, but the particle density (Figure 27) increases much faster, which explains the overall increase of the transport rate. The transport rates closest to the bottom are even higher than the best fitted logarithmic curve. This could be caused by the increase in particle velocity discussed earlier.

As mentioned earlier, acoustic measurements with ADCP instruments were performed in the current flume as well. For more details, see Joachim Beckers' "Application of ADCP-BT to measure bed load transport" (Beckers 2017). From these measurements, water flow velocities, bed load velocities and bottom transport rates (G_b) were calculated for different flow rates (including 200 l/s, the flow rate used in the most camera experiments, including the last). The water flow velocity estimated by taking the water flow rate and the cross section of the current flume and the velocities calculated by the two ADCP instruments at 200 l/s correspond

very well. This is a good confirmation of the maximum water velocity that we used to define our search window in the PTV analysis. The bed load velocities measured by the two bottom-track instruments is significantly lower than all velocities from the PTV analysis: 34 mm/s and 45 mm/s at 200 l/s. We established a decrease in velocity towards the bottom earlier: from 205 to 120 mm/s. The top of the ripple seen in the image sequence moves at a velocity of about 0.24 mm/s. This was calculated based on only one ripple, but it is still two orders of magnitude lower than the bottom-track velocity. We could assume that the bedload velocity is the speed at which the top of the sediment bed itself moves, by rolling and sliding of particles over ripples and other bedforms. This process should be faster than the movement of the whole ripple. The decrease of velocities fits well within the decreasing trend from PTV analysis. At 200 l/s, the bottom-track might measure the velocity of sediment transport at the top of the sediment bed, a region that is just below what we were able to visualise with a camera (the lowermost part of the water column). So at first sight, there is no contradiction between the results of both methods, but there is a difference because they measure in different areas. The top-down camera seemed to visualise the transport at the top of the sediment bed better, but to process these images a lot of work still has to be done. Meanwhile the acoustic (and other) instruments struggle to get good data from the lowermost part of the water column, because the sediment bed acts as such a strong reflector that it disrupts the measurements just above it. This makes PTV and the sideways camera interesting to further visualise this region.

The sediment transport rates measured by the load cell and by different ADCP models differ by 2 orders of magnitude (see Table 4). For the transport rate approximation by PTV analysis to match even the smallest transport rate from the weight cell or ADCP, the depth of view (the last unknown value) has to be extremely small: at 0.273 mm, it is smaller than the particle diameter. A good estimate based on camera use between 25 and 1 cm doesn't come close: 2 to 3 orders of magnitude difference with the weight cell transport rate and van Rijn, 4 to 5 orders compared to the other 3 ratios. An explanation for this can be the fact that, despite the fact that the bed load velocity is smaller than the particle velocities from PTV, the number of particles in the water column is much smaller than the amount of sediment that is transported by sliding and rolling as the top of the sediment bed. This is similar to the exponential increase in particle transport towards the bottom discussed earlier: particle velocity decreases linearly, but particle density increases exponentially. This phenomenon could be expected from the recordings from the top-down camera, showing the mass-movement of sediment ripples. Again, the results of the PTV analysis and the ADCP are potentially not contradictory, but instead show that the difference of sediment transport rates between suspended particles in the large upper part of the water column, jumping particles in the smaller lower part of the water column and the top part of the sediment bed of rolling and sliding particles differs with several orders of magnitude. The dune displacement gives an estimate of 1,01 g/s sediment transport. This is much closer to the sediment transport measured by PTV. In fact, it falls within the interval of total transport based on an estimated depth of view of 25 cm and 1 cm (Table 4). Still, based on only one ripple and a simplified calculation of sediment transport by this ripple, it is not possible to draw any further conclusions.

Table 4: Sediment transported in focus of the camera and after PTV analysis, compared to different transport rates based on the pit sampler, the ADCP measurements at 200 l/s flow rate by Beckers (2017) and the ripple displacement. To equal these transport rates, the corresponding depth of view of the PTV analysis is given. Finally, two estimated transport rates are given, based on the sediment transport in focus and estimated depth of views.

Transport in view (kg/s)	Gb Pit sampler [kg/s]	Gb ADCP Rennie [kg/s]	Gb ADCP Einstein [kg/s]	Gb ADCP van Rijn [kg/s]	Gb van Rijn [kg/s]	Gb ripple [kg/s]
5,62E-05	0,247	38,9	25,9	19,6	0,331	1,01E-3
Corresponding depth of view (m)						
	2,73E-4	1,73E-06	2,60E-06	3,44E-06	2,03E-4	0,0667
Total transport based on estimated depth of view 25 cm and 1 cm (kg/s)						
2,70E-4	6,74E-3					

6. CONCLUSION

We have tested the application of camera's and Particle Tracking Velocimetry to visualize bottom transport in a current flume. We tested different setups, equipment and software that was available, configurations and settings; and improved them when necessary and possible. Main components are the current flume, two cameras (one side-looking with a bitemcentric lens, one top-down with a manually adjustable lens) and the Streams software. Data obtained can be divided roughly in long-exposure images and short-exposure images, or top-down view and lateral view. Full PTV analysis was applied to the final image sequence. Particle velocity and sediment transport data were compared to ADCP data from Beckers (2017).

Long-exposure versus short exposure: long-exposure images provide a good visualisation of particle trajectories for the human eye. However, no software was found so far to extract particle velocity and sediment transport information from these images. Short-exposure images are more challenging for the human eye to track particles, but the Streams software can get valuable data from these frames and visualize the evolution of particle velocity, particle density and transport rates with depth. In an ideal world, having both types of images at the same time and at the same location would be extremely powerful when they can be compared (particle paths from Streams with long-exposure particle tracks), but this is not possible yet.

Top-down versus side-looking: processing and analysing frames from the top-down camera has proven to be difficult so far. Instead of a background with particles moving in front, the user looks at an image completely filled with particles. However, with the new computer, a renewed effort is encouraged. The top-down camera looks at a different level of transport than the side-looking camera, one that might be more comparable with the bottom-track measurements (at least at 200 l/s water flow rate). The side-looking camera resulted in good data of the transition zone of suspended transport and bedload transport. Identifying and matching particles in suspension and in long-distance jumps went well, but the algorithms struggle with high particle densities and discontinuous short jumps and turbulently falling particles after dune crests. So far, camera usage and PTV analysis provided good results in a region that is hard to monitor for other instruments (e.g. water bias of ADCP-BT). But the results of this work only contain only PTV data for one water flow rate. How PTV analysis and the current setup and methods will hold up for larger flow rates (and higher particle velocities and densities) or different sediments (such as quartz sand, when it can be brought in motion) remains an important research question. This study provides a good basis for continuing the work in these directions.

The final results show an exponential increase in sediment transport towards the bed, that is composed of sand-sized particles. This illustrates again the importance of a reliable, continuous method to measure bottom transport, especially in sandy conditions.

REFERENCES

- Bagnold, R.A., 1956. The Flow of Cohesionless Grains in Fluids. *Proc. Royal Soc. Philos. Trans., London*, 249.
- Beckers, J., 2017. Application of ADCP-BT to measure bed load transport. *Master's dissertation at University of Antwerp, Faculty of Applied Engineering, in cooperation with Flanders Hydraulics Research*, 69 pp.
- Beylich, A.A. & Laute, K., 2014. Combining impact sensor field and laboratory flume measurements with other techniques for studying fluvial bedload transport in steep mountain streams. *Geomorphology*, 218, pp.72–87.
- Einstein, H.A., 1950. The Bed-Load Function for Sediment Transportation in Open Channel Flow. *Technical Bulletin U.S. Dep. of Agriculture, Washington, D.C.*, 1026
- Gomez, B., 1991. Bedload transport. *Earth Science Reviews*, 31(2), pp.89–132.
- Mao, L., Carrillo, R., Escarriaza, C. & Iroume, A., 2016. Flume and field-based calibration of surrogate sensors for monitoring bedload transport. *Geomorphology*, 253, pp.10–21.
- Nevelsteen, K. & Van Hoestenbergh, T., 2009. Meettechniek voor het continue bemeten van bodemtransport. *Project of University of Leuven, research group Physical and Regional Geography, commissioned by the Flemish Environment Agency*, 74 pp.
- Nichols, G., 2009. *Sedimentology and Stratigraphy* Second Edition, Wiley-Blackwell, 419 pp.
- Nokes, R., 2009a. System Theory and Design. *Streams Version 2.06*, 164 pp.
- Nokes, R., 2009b. User's Guide: Core objects. *Streams Version 2.06*, 238 pp.
- Nokes, R., 2009c. User's Guide: Field objects. *Streams Version 2.06*, 120 pp.
- O'Brien, P. & McKenna Neuman, C., 2016. PTV measurement of the spanwise component of aeolian transport in steady state. *Aeolian Research*, 20, pp.126–138.
- Rachlewicz, G., Zwolinski, Z., Kociuba, W. & Stawska, M., 2017. Field testing of three bedload samplers' efficiency in a gravel-bed river, Spitsbergen. *Geomorphology*, 287, pp.90–100.
- van Rijn, L.C., 1993. *Principles of sediment transport in rivers, estuaries and coastal seas*, Aqua Publications.
- Ryan, S.E. & Porth, L.S., 1999. A field comparison of three pressure-difference bedload samplers. *Geomorphology*, 30(4), pp.307–322.
- Smith, S.J. & Friedrichs, C.T., 2015. Image processing methods for in situ estimation of cohesive sediment floc size, settling velocity, and density. *Limnology and Oceanography: Methods*, 13(5), pp.250–264.
- Yager, E.M., Kenworthy, M. & Monsalve, A., 2015. Taking the river inside: Fundamental advances from laboratory experiments in measuring and understanding bedload transport processes. *Geomorphology*, 244, pp.21–32.

Instrument manuals and product specifications

IDS Imaging Development Systems GmbH, 2017a. *Manual for uEye Cameras Version 4.82*

IDS Imaging Development Systems GmbH, 2017b. *Data sheet UI-3080CP-C-HQ Rev. 2.*

Opto GmbH, 2015. *Factsheet 100-BTC-005 Bi-telecentric lens 0.05x.*

Ricoh Company Ltd., 2017. *FL-BC7528-9M Product Specifications.*

Teledyne RD Instruments, 2011. *Acoustic Doppler Current Profiler Principles of Operation: a practical primer.*

APPENDIX

A.1 Product specifications of the lens of the top-down camera
[\(https://industry.ricoh.com/en/fa_camera_lens/lens/9m/\)](https://industry.ricoh.com/en/fa_camera_lens/lens/9m/)

FL-BC7528-9M Product Specifications

Resolution	9 Mega-pixel *1	
Format size	1" format	
Focal length	75mm	
Maximum aperture ratio	1:2.8	
Iris range	2.8~32	
Mount	C	
Horizontal angle of view	1/3" format	3.7°
	1/2" format	4.9°
	1/1.8" format	5.5°
	2/3" format	6.7°
	1" format	9.8°
Minimum object distance	0.25m	
Back focal length	21.3mm	
Filter size	40.5 P=0.5mm	
Dimensions	φ42mmx81mm	
Weight	189g	
Remarks	Focus & Iris lock screws	

*1 The lens is also suitable to 4 megapixel and 6 megapixel cameras.

A.2 Product specifications of the lens of the side-looking camera

100-BTC-005
Opto

FACTSHEET

Bi-telecentric lens 0.05x

Bi-telezentrisches Objektiv 0,05x
Objectif bi-télécentrique 0,05x




Specification

Article number 100-BTC-005		
Bi-telecentric		object and image side
Magnification	mm	0.05 x
Field of View (with 1/2" chip)	mm	125 x 94.1
Field of View (with 1/1.8" chip)	mm	139 x 104
Field of View (with 2/3" chip)	mm	167 x 125 (slightly vignetting possible)
Distortion	%	< 0.07
Telecentric failure	*	< 0.08
Resolution (MTF@70LP/mm)	%	50
Aperture	F/#	8
Depth of focus	mm	253
Working Distance	mm	415.7
Diameter x length	mm	205 x 502
Clamp diameter	mm	110
Clamp length	mm	120.7
Camera thread		C-Mount

02/2015 - Specifications are subject to change without prior notice.

Opto GmbH · Lochhamer Schlag 14 · 82166 Gräfelfing · Germany · Tel: +49 89 8980 55 0 · info@opto.de
www.opto.de

Technical Drawing / clamp dimensions



A.3 Camera product specifications (<https://en.ids-imaging.com/store/products/cameras/ui-3080cp-rev-2.html>)




UI-3080CP-C-HQ Rev.2



Specification

Sensor

Sensor type	CMOS Color
Shutter	Global Shutter
Sensor characteristic	Linear
Readout mode	Progressive scan
Pixel Class	5 MP
Resolution	5.04 Mpix
Resolution (h x v)	2456 x 2054 Pixel
Aspect ratio	5:4
ADC	12 bit
Color depth (camera)	12 bit
Optical sensor class	2/3"
Optical Size	8.473 mm x 7.086 mm
Optical sensor diagonal	11.05 mm (1/1.45")
Pixel size	3.45 µm
Manufacturer	Sony
Sensor Model	IMX250LQR-C
Gain (master/RGB)	24x/4x
AOI horizontal	same frame rate
AOI vertical	increased frame rate
AOI image width / step width	256 / 8
AOI image height / step width	4 / 2
AOI position grid (horizontal/vertical)	4 / 2
Binning horizontal	-
Binning vertical	-
Binning method	-
Binning factor	-
Subsampling horizontal	same frame rate
Subsampling vertical	same frame rate
Subsampling method	MVC automatic
Subsampling factor	2, 4, 6, 8, 16



Subject to technical modifications (2017-05-12)

Page 1 of 3

www.ids-imaging.com

IDS Imaging Development Systems GmbH

Dimbacher Str. 6 - 8 · 74182 Obersulm · Germany · Phone +49 7134 96196-0 · Fax +49 7134 96196-99 · E-mail info@ids-imaging.com

UI-3080CP-C-HQ Rev.2

Model

Pixel clock range	80 MHz - 474 MHz
Frame rate freerun mode	86.0 fps
Frame rate trigger (continuous)	86.0 fps
Frame rate trigger (maximum)	86.0 fps
Exposure time (minimum - maximum)	0.019 ms - 999 ms
Long exposure (maximum)	30000 ms
Power consumption	2.3 W - 3 W
Image memory	128 MB
Special features	Triggered linescan mode, Sensor source gain, Multi-AOI

Ambient conditions

The temperature values given below refer to the outer device temperature of the camera housing.

Device temperature during operation	0 °C - 55 °C / 32 °F - 131 °F
Device temperature during storage	-20 °C - 60 °C / -4 °F - 140 °F
Humidity (relative, non-condensing)	20 % - 80 %

Connectors

Interface connector	USB 3.0 micro-B, screwable
I/O connector	8-pin Hirose connector (HR25-7TR-8PA(73))
Power supply	USB cable

Pin assignment I/O connector

1	Ground (GND)
2	Flash output with optocoupler (-)
3	General Purpose I/O (GPIO) 1
4	Trigger input with optocoupler (-)
5	Flash output with optocoupler (+)
6	General Purpose I/O (GPIO) 2
7	Trigger input with optocoupler (+)
8	Output supply voltage, 5 V (100 mA)



Design

Lens Mount	C-Mount
IP code	IP30
Dimensions H/W/L	29.0 mm x 29.0 mm x 29.0 mm
Mass	52 g






Subject to technical modifications (2017-05-12)

Page 2 of 3

www.ids-imaging.com

IDS Imaging Development Systems GmbH

Dimbacher Str. 6 - 8 · 74182 Obersulm · Germany · Phone +49 7134 96196-0 · Fax +49 7134 96196-99 · E-mail info@ids-imaging.com



Subject to technical modifications (2017-05-12)

Page 3 of 3

www.ids-imaging.com

IDS Imaging Development Systems GmbH
 Dimbacher Str. 6 - 8 · 74182 Obersulm · Germany · Phone +49 7134 96196-0 · Fax +49 7134 96196-99 · E-mail info@ids-imaging.com

A.4 Grain size distribution of the sand-like substitute



Result Analysis Report

Sample Name:
OS20170317014 - Average

Sample Source & type:
Jens

Sample bulk lot ref:
20170317

SOP Name:

Measured by:
bastiaem

Result Source:
Averaged

Measured:
maandag 20 maart 2017 14:23:26


Analysed:
maandag 20 maart 2017 14:23:28

Particle Name: Default	Accessory Name: Hydro 2000MU (A)	Analysis model: General purpose	Sensitivity: Normal
Particle RI: 1.520	Absorption: 0.1	Size range: 0.020 to 2000.000 um	Obscuration: 8.82 %
Dispersant Name: Water	Dispersant RI: 1.330	Weighted Residual: 0.990 %	Result Emulation: Off

Concentration: 0.6377 %Vol	Span : 0.641	Uniformity: 0.204	Result units: Volume
--------------------------------------	------------------------	-----------------------------	--------------------------------

Specific Surface Area: 0.0123 m ² /g	Surface Weighted Mean D[3,2]: 486.535 um	Vol. Weighted Mean D[4,3]: 516.378 um
---	--	---

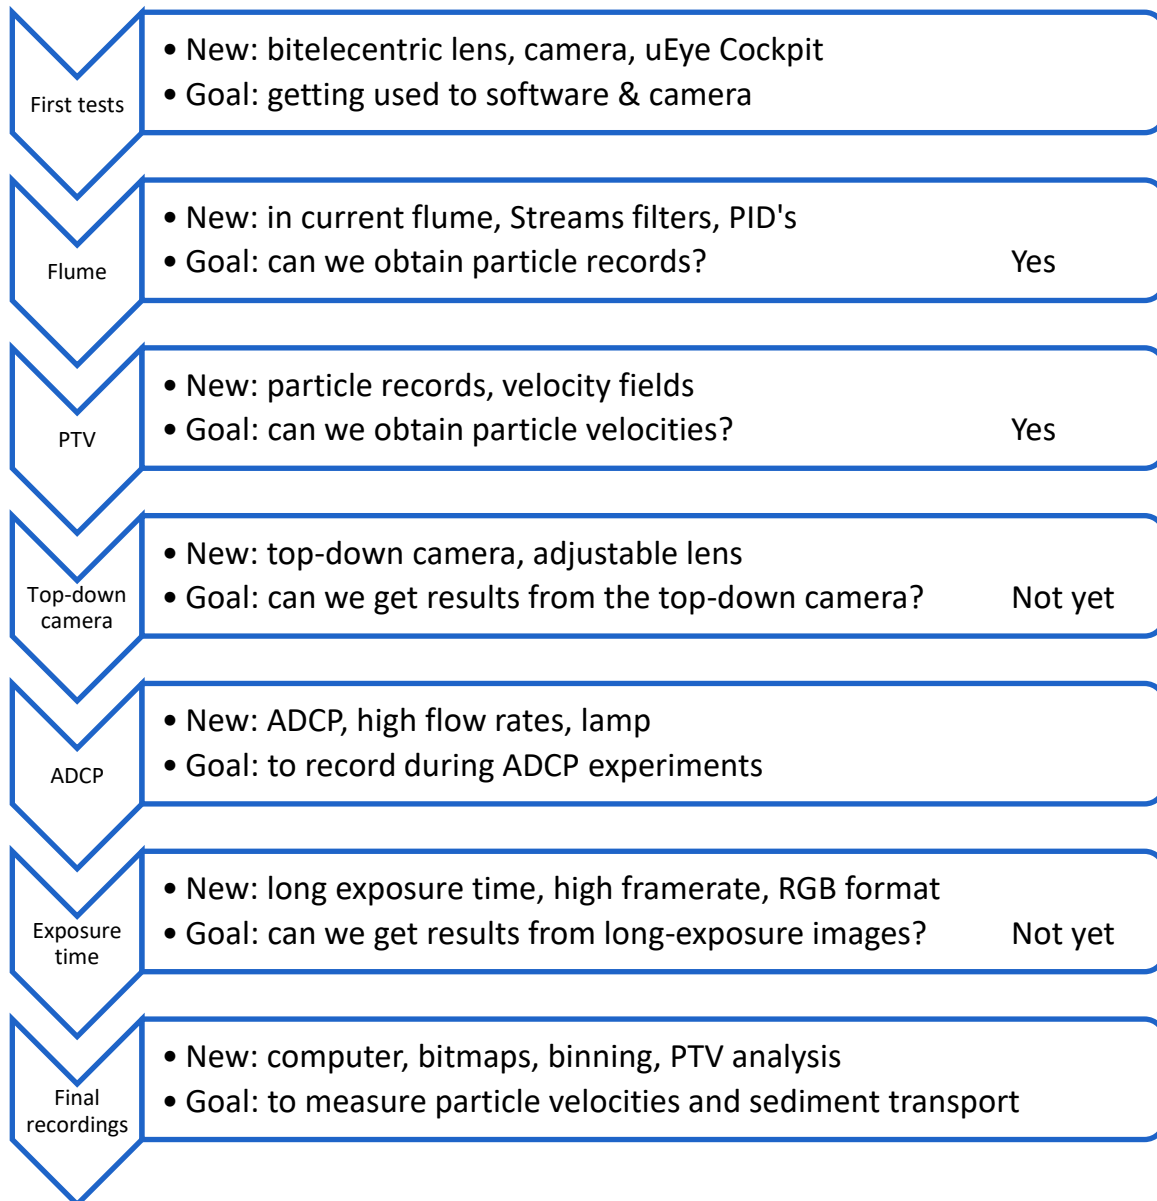
d(0.1): 365.692 um d(0.5): 501.434 um d(0.9): 686.954 um



Size (µm)	Volume In %	Size (µm)	Volume In %	Size (µm)	Volume In %	Size (µm)	Volume In %	Size (µm)	Volume In %	Size (µm)	Volume In %
0.010	0.00	0.105	0.00	1.096	0.00	11.482	0.00	120.226	0.00	1258.925	0.00
0.011	0.00	0.120	0.00	1.259	0.00	13.183	0.00	138.038	0.00	1445.440	0.00
0.013	0.00	0.138	0.00	1.445	0.00	15.136	0.00	158.489	0.00	1659.587	0.00
0.015	0.00	0.158	0.00	1.660	0.00	17.378	0.00	181.970	0.00	1905.461	0.00
0.017	0.00	0.182	0.00	1.905	0.00	19.953	0.00	208.930	0.01	2187.762	0.00
0.020	0.00	0.209	0.00	2.188	0.00	22.909	0.00	239.883	0.39	2511.886	0.00
0.023	0.00	0.240	0.00	2.512	0.00	26.303	0.00	275.423	0.39	2884.032	0.00
0.026	0.00	0.275	0.00	2.884	0.00	30.200	0.00	316.228	2.26	3311.311	0.00
0.030	0.00	0.316	0.00	3.311	0.00	34.674	0.00	363.078	6.82	3801.894	0.00
0.035	0.00	0.363	0.00	3.802	0.00	39.811	0.00	416.869	13.54	4365.158	0.00
0.040	0.00	0.417	0.00	4.365	0.00	45.709	0.00	478.630	19.66	5011.872	0.00
0.046	0.00	0.479	0.00	5.012	0.00	52.481	0.00	549.541	21.55	5754.399	0.00
0.052	0.00	0.550	0.00	5.754	0.00	60.256	0.00	630.957	18.01	6606.934	0.00
0.060	0.00	0.631	0.00	6.607	0.00	69.183	0.00	724.436	11.27	7585.776	0.00
0.069	0.00	0.724	0.00	7.586	0.00	79.433	0.00	831.764	5.05	8709.636	0.00
0.079	0.00	0.832	0.00	8.710	0.00	91.201	0.00	954.993	1.35	10000.000	0.00
0.091	0.00	0.955	0.00	10.000	0.00	104.713	0.00	1096.478	0.12		
0.105	0.00	1.096	0.00	11.482	0.00	120.226	0.00	1258.925	0.00		

Operator notes: met US

A.5 Short overview of experiments



A.6 Workflow final measurements

