PAPER

Spatio-temporal image-based parametric water surface reconstruction: a novel methodology based on refraction

View the article online for updates and enhancements.

Related content

IOP ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Spatio-temporal image-based parametric water surface reconstruction: a novel methodology based on refraction

L Engelen[1], S Creëlle[2], L Schindfessel[2] and T De Mulder[3]

[1] Ph.D. Student, Hydraulics Laboratory, Dept. of Civil Engineering, Ghent University, Sint-Pietersnieuwstraat 41 B5, B-9000 Ghent, Belgium
[2] Postdoctoral Fellow, Hydraulics Laboratory, Dept. of Civil Engineering, Ghent University, Sint-Pietersnieuwstraat 41 B5, B-9000 Ghent, Belgium
[3] Associate Professor, Hydraulics Laboratory, Dept. of Civil Engineering, Ghent University, Sint-Pietersnieuwstraat 41 B5, B-9000 Ghent, Belgium

E-mail: Lukas.Engelen@UGent.be, Stephan.Creelle@UGent.be, Laurent.Schindfessel@UGent.be and TomFO.DeMulder@UGent.be

CrossMark

Abstract

This paper presents a low-cost and easy-to-implement image-based reconstruction technique for laboratory experiments, which results in a temporal description of the water surface topography. The distortion due to refraction of a known pattern, located below the water surface, is used to fit a low parameter surface model that describes the time-dependent and three-dimensional surface variation. Instead of finding the optimal water depth for characteristic points on the surface, the deformation of the entire pattern is compared to its original shape. This avoids the need for feature tracking adopted in similar techniques, which improves the robustness to suboptimal optical conditions and small-scale, high-frequency surface perturbations. Experimental validation, by comparison with water depth measurements using a level gauge and pressure sensor, proves sub-millimetre accuracy for smooth and steady surface shapes. Although such accuracy cannot be achieved in case of highly dynamic surface phenomena, the low-frequency and large-scale free surface oscillations can still be measured with a temporal and spatial resolution mostly limited by the available optical set-up. The technique is initially intended for periodic surface phenomena, but the results presented in this paper indicate that also irregular surface shapes can robustly be reconstructed. Therefore, the presented technique is a promising tool for other research applications that require non-intrusive, low-cost surface measurements while maintaining visual access to the water below the surface. The latter ensures that the suggested surface reconstruction is compatible with simultaneous image-based velocity measurements, enabling a detailed study of the flow.

Keywords: parametric shape fitting, shape from refractive distortion, optical water surface measurement

(Some figures may appear in colour only in the online journal)

1. Introduction

Accurate registration of a liquid's free surface is important for numerous hydraulic research fields that study flow characteristics using laboratory experiments. Non-intrusive, optical techniques are preferred over traditional point measurements because they offer instantaneous information of the entire surface area. Although the extensive overview given by [15] showed that a large number of methodologies are possible to reconstruct refractive and/or transparent objects, only a

limited set of optical techniques seem applicable to reconstruct a specular and highly dynamic water surface. The three most commonly used types of optical surface measurement techniques are listed hereafter, while focusing on their ease of implementation and robustness to record shallow-water surface oscillations.

A first stereo-based approach is to apply triangulation to characteristic points on the water surface, of which [5] described different implementations mentioned in literature. Due to the lack of distinguishable points in steady surface conditions, surface seeding (e.g. [9]) or neutrally buoyant seeding particles in combination with interfacial particle image velocimetry [26] has been employed to improve the resolution in homogeneous regions. Nonetheless, this approach remained difficult in combination with water due to the latter's specular nature. Although these difficulties could be mitigated by additives that change the fluid's optical properties, clear water might be needed to acquire subsurface (velocity) information and maintain its rheological properties. Akutina quantified the water depth using 3D-PTV, in which the position of the detected particles located the furthest from the bottom was considered as a reliable estimate of the local water depth [1]. This provided detailed, simultaneous information about the 3D-flow field and surface shape, but neglected refraction of the viewing rays which is only valid in case the water depth and surface inclination remain small. Moreover, the recent 3D-PTV technique seemed to require a lot of experience and fine-tuning to achieve reliable results.

As a second alternative, reflection-based techniques have been developed that use the polarisation of reflected light (e.g. [30]) or the deformation of a reflected pattern (e.g. [24]). However, complications often arose when applied to water, such as specular reflections, occlusions, and the fact that only a small fraction of the light hitting the water surface is reflected. To overcome some of these problems, several authors used dye or other additives to make the water opaque and relied on the deformation of a regular grid [12], irregular grid [25], or fringe pattern (e.g. [10]) projected on the surface. Despite their improved performance, such solutions impeded visual access for subsurface velocity measurements and usually required multiple views of different cameras for consistency and accuracy. Rupnik *et al* presented two approaches, in which they treated the water as a diffuse or as a specular surface and fitted a sinusoidal surface function to either a set of triangulated points or to travelling specularities, respectively [23]. Although limited by the simplifications of the adopted model, a sinusoidal surface could be reconstructed with mm accuracy.

Finally, highly accurate refraction-based techniques have gained more and more attention because their performance has been reported superior in case of strong surface gradients. Kutulakos and Steger summarised the numerous solutions to resolve the ambiguity that arises when a refractive approach is applied, of which some of the most applicable techniques are discussed hereafter [17].

Several authors related the local surface slope to the apparent colour of the water surface, originating from a refracted (e.g. [31]) or reflected (e.g. [11]) two-dimensional colour pattern. While limited to small capillary or wind-driven waves, [16] determined the surface slope and surface height both simultaneously and independently by measuring the refraction of colour-coded light and the difference in absorption rate of light rays with different wave length. Recently, [3] have presented an adapted version of this methodology in which the need for a telecentric system was avoided by simultaneously processing visible and near-infrared digital images. This allowed highly accurate water depth and surface slope measurements, although the accuracy decreased with increasing water depth. Seeding the flow with neutrally buoyant particles has often been combined with illumination of particles present at a certain depth using a horizontal, subsurface laser sheet. Ng *et al* subsequently searched the surface gradient iterative by comparing the views from below (without refraction) and from above the water surface (with refraction) [22]. Others applied the same methodology on an irregular, fixed pattern [13, 19]. Morris avoided the need for lasers and proposed a simple but accurate multi-view stereo approach that combined 'shape from refractive distortion' with a traditional stereo set-up [21]. Using a 'verification camera' to solve the depth-normal ambiguity related to refraction, submillimetre accuracy was achieved in combination with a fixed pattern [20] or flow seeding [14].

The methodology of [21] seems preferential in case a simple but accurate technique is preferred to reconstruct 3D surface variations, in which problems related to specular reflections are avoided. Moreover, the preserved transparency of the water gives visual access to the water mass. This enables simultaneous (subsurface) velocity measurements that rely on the registration of particles dispersed in the flow. However, [21] tracked individual feature points, i.e. the corners of a chequerboard pattern, to determine the optimal water depth for each of these characteristic points on the surface. When such methodology would be applied in presence of high-frequency surface effects or seeding particles that partly mask the pattern located below, the risk of loosing a large part of these points during tracking could increase significantly. In the region where features would be lost, further reconstruction becomes less accurate due to a lack of surface information.

To solve these limitations, this paper proposes a modified methodology in which feature tracking is avoided to improve its robustness. A single-view and purely image-based technique is envisaged, without the need for lasers or an extensive optical set-up, to ensure ease of implementation and limit the cost of the equipment. Avoiding the need for feature tracking also ensures compatibility with simultaneous image-based velocity measurements, in view of future research for which also the flow velocity is of interest.

The proposed methodology is initially intended for hydraulic research in which the surface remains relatively smooth and is dominated by long surface gravity waves. In that case, a low parameter surface model can be derived, based on the expected surface shape, that describes the reconstructed surface area. By combining this surface model with the refractive approach of [21], the parametrised model can be fitted to describe the instantaneous surface deformation. Instead of determining the optimal water depth for a sparse set of surface points, the deformation of the entire pattern is used to fit the

surface model describing the global surface shape. Although the accuracy that can be reached depends on the choice of surface model, this modified approach avoids problems related to feature tracking and makes reconstruction of a single image-frame independent of previous time-steps. It also avoids the need for overlapping views of multiple cameras, resulting in a more flexible single-view experimental set-up that can cover a larger surface area.

If, on the other hand, the total surface area is too irregular to describe sufficiently accurate with a low-parameter surface model, subdivision of the surface area into a mesh of smaller surfaces is proposed. Subsequently, the same methodology can be applied locally for each of these smaller areas, in which the surface can be approximated by a more simplified surface model (e.g. a straight plane). The results in each of these local areas can then be combined to describe the entire surface topography. Although this increases the computational cost of the surface reconstruction, irregularities of the surface shape can as such be incorporated. Moreover, this local methodology allows to adjust the spatial resolution according to the dynamics of the phenomenon that needs to be measured. For the dynamic validation tests presented in this paper, application of such a small reconstruction domain proved successful to obtain robust and reliable results.

In section 2, a theoretical description of the reconstruction methodology is presented. The methodology is then validated experimentally by comparing the results with alternative measurement techniques, both on still (section 4.1) and non-stagnant water (section 4.2). The results on still water prove that in case the optimisation settings and surface model are adjusted to the specific test conditions, the image-based methodology can achieve sub-millimetre accuracy for regular and simple surface shapes over a relatively large range of water depths. More importantly, application of the technique to oscillating surface shapes proves that also dynamic surface phenomena can robustly be reconstructed.

2. Presentation of the image-based methodology

2.1. Theoretical description of the refractive surface reconstruction

The presented methodology employs the refraction of light rays at the air–water interface. A fixed colour pattern is positioned on a plane surface F, i.e. the bottom of the flume, and its projection on the surface is viewed by a camera from above (figure 1). Refraction causes each light ray to change direction at the corresponding surface point **p**. For the same point **f** on plane F, the projection in the image plane therefore shifts from image point **q** (in absence of water) to **q'** (in presence of water). This apparent shift in the image plane is further denoted as the 'image disparity shift'. These image disparity shifts result in a deformation of the projected pattern on the surface (surface points $p_1 - p_N$) and corresponding $q'_1 - q'_N$ compared to the physical pattern $f_1 - f_N$ on F. The shape of the distorted pattern, which depends on the local surface position and orientation around every point **p**, is then used to quantify the local water depth.
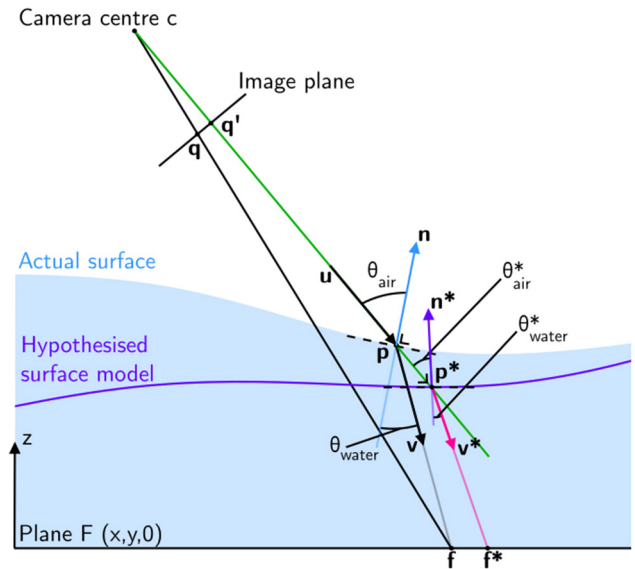
Figure 1. Image-based surface reconstruction using the refraction of the viewing rays at the air–water interface: the discrepancy between the actual surface (**p**) and hypothesised surface (p^*) results in a different intersection with bottom plane F (f^* instead of **f**) (after [21]).

Figure 1 depicts the change in direction of the incident light ray $\vec{u} = \vec{cq'}$, originating at the camera centre **c** to the refracted ray $\vec{v} = \vec{pf}$. The change in angle relatively to the local surface normal $\vec{n}$ is different for every viewing ray and can for every pixel in the camera's view be computed according to Snell's law. Because the surface normal, incident and refracted light rays are located in the same plane, two straight lines describe the total trajectory of the viewing ray within this plane of refraction. This simplifies the geometric problem because the two remaining unknowns, i.e. the location of surface point **p** and the direction of the surface normal $\vec{n}$, must be located in the plane determined by the known position of **f q'** and **c**. Nonetheless, a closed-form solution is impossible to derive because only a single, underconstrained expression (Snell's law) in two unknowns ($p, \vec{n}$) is available.

To relieve this indeterminacy, [21] assumed the surface location **p** to be known for a discrete set of feature points and verified these hypothesised locations by a complementary stereo-view of a 'verification camera'. In this paper, a single-view reconstruction is achieved by combining the observed image disparity shifts with a parametrised surface model that represents the surface topography within the reconstructed area. This 'surface model' is composed of multiple terms, each scaled with individual coefficients that can be fitted to describe the instantaneous surface shape. Although the accuracy depends on how well the surface can be approximated with only a limited amount of terms, the resulting parametrised surface model $\eta(x, y)$ quantifies the surface height (z-coordinate) above plane F ($z = 0$) for a given combination of coefficients at every position (x, y) within the reconstructed area. For a chosen surface model, a hypothesised set of the model's coefficients determines the surface location of the imaged area. Parameters that are based on a single set of such hypothesised coefficients will further be indicated by an

asterisk (*). The junction between each incident ray $\overrightarrow{\boldsymbol{u}} = \overrightarrow{\boldsymbol{cq'}}$ and this hypothesised surface $\eta^*(x, y)$ is used to determine the location of the corresponding surface point $\boldsymbol{p}^*$:

$$\begin{cases} \boldsymbol{p}^* \in \eta^*(x, y) \\ \boldsymbol{p}^* \in \overrightarrow{\boldsymbol{u}} = \overrightarrow{\boldsymbol{cq'}} \end{cases}. \quad (1)$$

Additionally, $\eta^*(x, y)$ enables the derivation of the local surface normal $\overrightarrow{\boldsymbol{n}}^*$ at the same point $\boldsymbol{p}^*$:

$$\overrightarrow{\boldsymbol{n}}^* = \left(-\frac{\partial \eta^*(x, y)}{\partial x}, -\frac{\partial \eta^*(x, y)}{\partial y}, 1\right)$$

$$\text{because } n_z^* = \frac{\partial \eta^*(x, y)}{\partial z} = 1. \quad (2)$$

Figure 1 illustrates that the direction of the refracted ray $\overrightarrow{\boldsymbol{v}}^*$ is subsequently computed by applying Snell's law (equation (3)), a fixed refraction index of water r_w (at the ambient water temperature) and basic trigonometric identities:

$$\sin(\theta_{\text{air}}) = r_w \sin(\theta_{\text{water}}) \quad (\text{Snell's law}) \quad (3)$$

and when applied to the hypothesised surface point $\boldsymbol{p}^*$:

$$\sin(\theta_{\text{air}}^*) = r_w \sin(\theta_{\text{water}}^*) \quad (4a)$$

$$\theta_{\text{water}}^* = \arcsin(\sin(\theta_{\text{air}}^*)/r_w) \quad (4b)$$

$$\theta_{\text{water}}^* = \arcsin\{\sin[\arccos(-\overrightarrow{\boldsymbol{u}} \cdot \overrightarrow{\boldsymbol{n}}^*)]/r_w\}. \quad (4c)$$

The refracted ray $\overrightarrow{\boldsymbol{v}}^*$, starting at position $\boldsymbol{p}^*$ and with a direction determined by θ_{water}^*, is then traced back to point $\boldsymbol{f}^*$ on the bottom plane at $z = 0$. By minimising the difference between $\boldsymbol{f}$ and $\boldsymbol{f}^*$, the optimal surface parameters can be found that best describe the instantaneous surface location and orientation around $\boldsymbol{p}$.

2.2. Optimisation of the surface model

For each point $\boldsymbol{f}^*$, its $(x,y,0)$ coordinates correspond to a certain colour of the fixed pattern. In this paper, F is coloured using a black-white chequerboard for which the reciprocal black and white squares can robustly be distinguished from each other. The expected colour, in monochrome images represented by a single pixel intensity, is then compared to the corresponding pixel intensity of the image ($I_{q'}$) to verify the hypothesised surface coefficients. A discrepancy between both indicates an inappropriate estimation of the instantaneous surface location and/or orientation, i.e. the computed surface point $\boldsymbol{p}^*$ (equation (1)) and/or surface normal $\overrightarrow{\boldsymbol{n}}^*$ (equation (2)). This comparison therefore transforms the pixel intensity of every pixel in the image plane to an error value $E_{q'}^*$:

$$E_{q'}^* = \begin{cases} 0, & \text{if } I_{q'} > I_{\text{max}} \text{ and } \boldsymbol{f}^* \text{ is white} \\ 0, & \text{if } I_{q'} < I_{\text{min}} \text{ and } \boldsymbol{f}^* \text{ is black} \\ 0, & \text{if } I_{\text{min}} < I_{q'} < I_{\text{max}} \\ & \hookrightarrow \quad \text{ambiguous pixel intensity} \\ 1, & \text{otherwise} \end{cases} \quad (5)$$

where I_{max} and I_{min} are user-defined intensity-thresholds that characterise the white and black parts of the image, respectively. By discarding the grey zones in the image ($I_{\text{min}} < I_{q'} < I_{\text{max}}$), local areas where the distinction between black and white becomes ambiguous are excluded from the optimisation. As such, errors related to suboptimal optical conditions (section 3) are avoided as much as possible.

These computed error measures are subsequently combined in a global error function E_{tot}^* that quantifies how well the surface model $\eta^*(x, y)$ matches with the processed image:

$$E_{\text{tot}}^* = \sum_{\text{image}} E_{q'}^*. \quad (6)$$

The surface coefficients that minimise this single error function are finally assumed to provide the best representation of the instantaneous surface topography.

Compared to [21], who only computes an error measure for a discrete set of feature points, E_{tot}^* incorporates the information of every pixel in the image plane. This alternative approach avoids problems related to feature detection/tracking (loss of feature points) and increases the accuracy, robustness, and spatial resolution of the estimated surface shape. Hence, adopting a surface model reduces the temporal reconstruction of the entire surface area to a sequence of multivariate optimisations of the model's time-dependent variables by processing each image frame separately. Additionally, the robustness and/or accuracy of the surface reconstruction can further be improved by adopting multiple (simultaneous) camera views. Although a single view is sufficient to fit the surface parameters, multiple views of the same surface area provide even more information about the imaged surface area. This might improve the algorithm performance and minimise possible error sources (section 3) in case of suboptimal reconstruction conditions.

2.3. Implementation details

The reconstruction algorithm explained in the previous paragraph is implemented in a modular C++ program, in which the different steps in the reconstruction procedure are separated and called upon progressively. Image processing is performed using classes and functions provided by the OpenCV library [6].

Firstly, calibration of the camera is needed to obtain the intrinsic camera parameters and distortion coefficients, the latter being required to correct for optical aberration. Before the camera is set-up, a set of calibration images is taken which allows to estimate the (radial and tangential) distortion coefficients according to the Brown–Conrady model [7], as well as the intrinsic camera parameters. A reliable estimate for both types of parameters is found using a Levenberg–Marquardt-based algorithm. Secondly, the transformation from pixels to 3D coordinates requires knowledge about the position of the cameras with regards to a chosen reference system, which is characterised by the camera's extrinsic parameters. To that end, an image of a reference pattern at a known position is adopted to compute the camera position and orientation during the measurements in the chosen frame of reference.
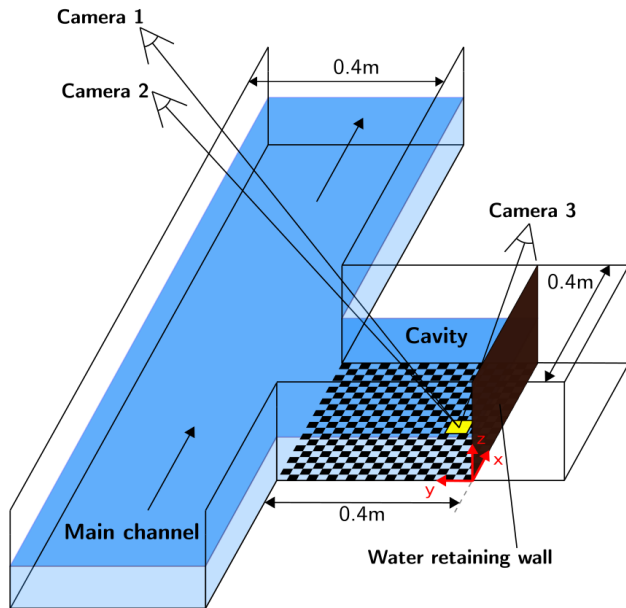
Figure 2. Experimental set-up in the confluence flume of the Hydraulics Laboratory of Ghent University. A water retaining wall is used to create a lateral cavity, where a white–black chequerboard pattern is fixed on the bottom. The chosen reference system, as adopted in section 4, has an origin at the upstream corner of the cavity and an upward directed z-axis with $z = 0$ at the bottom of the flume. The reconstructed surface area (yellow) and the position of the cameras that are used in this paper are indicated.

The location of the surface points $\boldsymbol{p}^*$, i.e. the intersections of the viewing rays $\overrightarrow{\boldsymbol{u}}$ with the hypothesised surface, are the solution of a set of non-linear equations (see equation (1)). Based on the estimated camera parameters and hypothesised surface coefficients, this 3D geometric problem is solved using a nonlinear solver from the open-source library MINPACK [8].

The multivariate optimisation of the surface coefficients of $\eta(x,y)$ is performed by application of the Levenberg–Marquardt optimisation algorithm, implemented in the open-source library ALGLIB [2]. Assuming temporal smoothness of the surface oscillations (and corresponding surface coefficients), the optimisation result from the previous time step is used as start for the iterative process to facilitate convergence.

Finally, quantification of the surface height in the domain of interest is done in a post-processing stage by combining the fitted, time-dependent coefficients with the mathematical formulation of $\eta(x,y)$. For a more detailed explanation about the use and implementation details of the C++ program, we refer to the user guide of the open source software[4].

2.4. Experimental set-up

The experiments are performed in the confluence flume of the Hydraulics Laboratory of Ghent University (figure 2), in which a water retaining wall is installed in the tributary channel to create a lateral cavity with a length and width of 0.4 m. Both the sides and the bottom of the flume are made
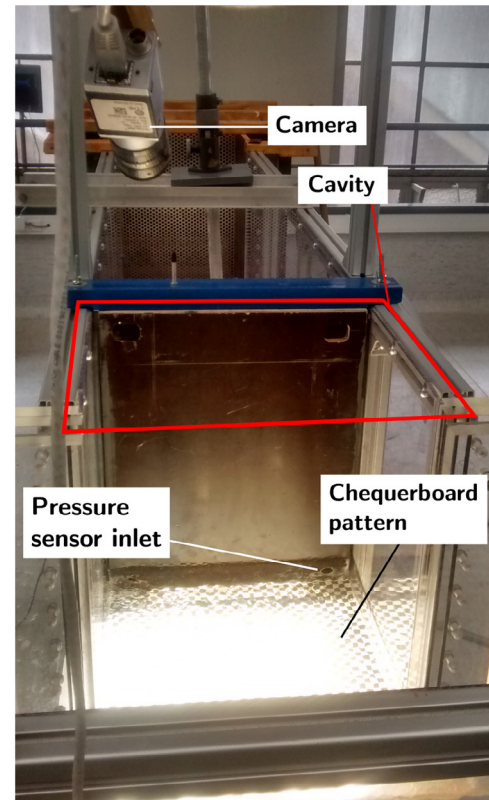
Figure 3. Lateral cavity with water retaining wall. The inlet of the pressure sensor used for the dynamic validation (section 4.2) is located at the upstream end of the cavity.

of translucent polycarbonate (figure 3), providing easy visual access for the application of optical measurement techniques.

The adopted colour pattern consists of a chequerboard pattern with a square-size of 10×10 mm. This pattern is printed on a transparent foil, which is glued with translucent adhesive to the inner side of the bottom plate. During the experiments, the bottom of the flume is lit from below. This avoids specular reflections on the surface and maximises the amount of light received by the camera, which views the pattern from above. As such, good contrast in the images is maintained while the aperture and opening time of the camera can be reduced to avoid image blur. The backlighting (figure 4) consists of an array of white light-emitting LED's, made uniform by positioning two diffusive, acrylic plates between the LED array and the bottom of the flume. Images are acquired using one or multiple synchronised Basler ace GigE cameras (model acA1920-25gc) at a frame rate of 40 Hz and a shutter speed of 1/1000 s. The camera contains a CMOS digital image sensor (MT9P031), which is equipped with a red-green-blue (RGB) Bayer pattern colour filter. This Bayer pattern consists of twice as many green filters compared to red and blue, making the green channel much more responsive. The corresponding higher contrast and less noise supports the choice of processing only the green channel of the RGB images, as was done for the results presented in this paper. Given that the sensor's peak spectral response in the green channel equals 520 nm, a corresponding refractive index r_w of 1.336 is used based on an average water temperature of 20 °C during the experiments (4).
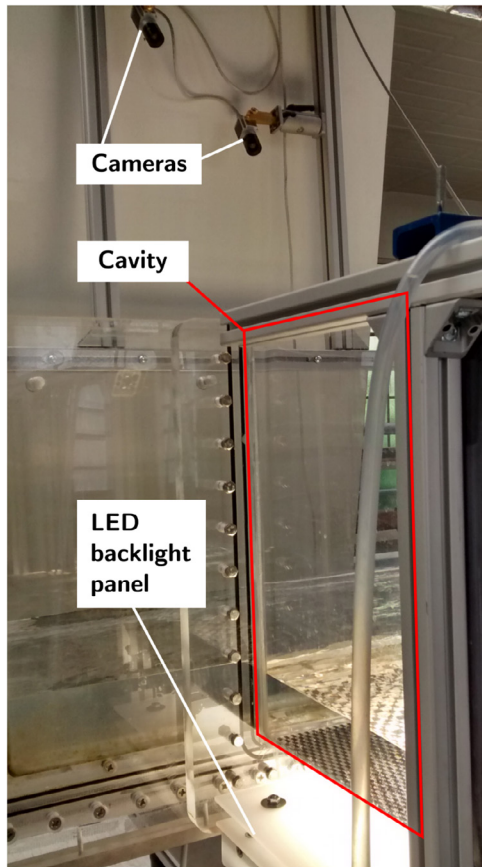
Figure 4. Optical set-up, with a backlight panel consisting of a LED array below two diffusive plates lighting the flume from below.

## 3. Possible error sources and limitations

Different aspects in the reconstruction methodology are affected by possible inaccuracies and/or wrong estimates of important parameters. These error sources may influence the final reconstruction accuracy and can roughly be divided into five main categories:

(a) Camera calibration

The presented image-based technique requires a reliable camera calibration to correct for lens distortion. Typically, the camera is calibrated before its installation, so that after installation, the calibration result can directly be employed to undistort the images before they are used as input for the reconstruction algorithm. Additionally, the camera's 3D position in a chosen frame of reference needs to be estimated by comparing a physical reference grid and its projection in the image plane. Combined with the results of the previous calibration, it defines the geometry of the refractive problem that needs to be solved. Hence, errors related to their estimated values have a non-linear, detrimental influence on the reconstruction result.

(b) Suboptimal contrast in the images

A first prerequisite for the input images is that the black and white pattern can reliably be detected. Good image contrast between these two colours can be affected by
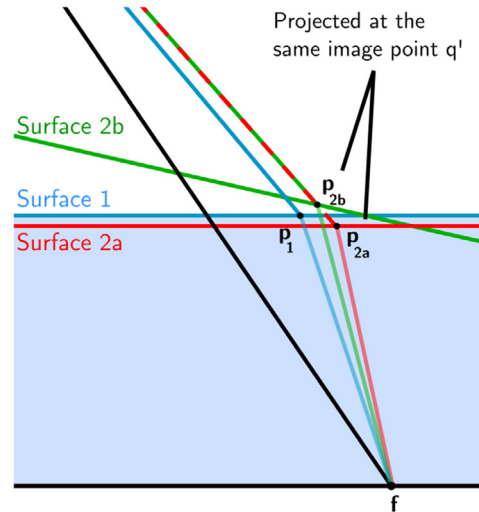
Figure 5. For the same point **f** on the bottom plane, the image disparity shift corresponding to the change in position from surface point p_1 towards p_{2a} or p_{2b} is identical for a varying water depth (Surface 2a) or varying inclination in the direction of the viewing ray (Surface 2b), respectively. The movement from p_1 towards p_{2a} or p_{2b} is exaggerated for clarity and ease of understanding.

suboptimal lighting conditions, image noise, and blurry images due to a limited depth of field or shutter speed of the camera. Then, the final result becomes more sensitive to the adopted intensity-thresholds, potentially reducing the robustness and reliability of the methodology.

(c) Multivariate optimisation

The final result of a non-linear optimisation solver is always an approximation of the perfect solution. In this paper, the Levenberg–Marquardt optimisation algorithm is adopted to estimate the time-dependent surface parameters for a single time step. Due to the inherent uncertainty about the perfect outcome of non-linear optimisation problems, local minima of the error function might result in surface coefficients that are a sub-optimal representation of the instantaneous surface shape. Especially in case only a small surface area is reconstructed, multiple solutions with a different average water level and inclination might give a similar error value (see next paragraph).

(d) Single-view approach

When only a single image and a small part of the distorted pattern are available to fit the surface coefficients, distinguishing local water depth changes from variations in inclination of the surface in the direction of the viewing rays becomes difficult. Both result in an image disparity shift in the direction of the viewing rays, as depicted in figure 5, and changes in the water depth may be erroneously interpreted as changes in surface inclination. To avoid this ambiguity, it can be useful to not incorporate a surface inclination in the viewing rays' direction during optimisation, as was done in section 4.2.

On the other hand, an inclination of the surface in the direction perpendicular to the viewing rays is easily distinguished from water depth changes. It moves the projection of the pattern perpendicular to the direction of the viewing rays, which cannot be accounted for by a vari-
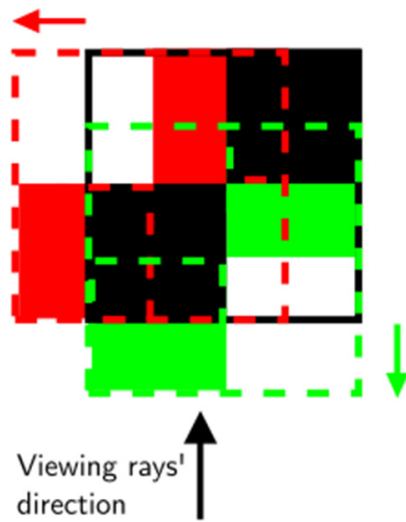
Figure 6. A change in water depth or change in inclination in the viewing rays' direction both shift the pattern in the direction of the viewing rays (green). A change in inclination in the perpendicular direction shifts the pattern also perpendicular (red).

ation of the average water depth. To find a better match with the distorted pattern, it is advised to incorporate this inclination during optimisation. For clarification, figure 6 depicts a schematic representation of the two possible, perpendicular directions of refractive displacement.

(e) Simplification of the surface.

The presented methodology uses a mathematical description of the surface, i.e. a parametrised surface model, to approximate the instantaneous surface shape. Depending on the irregularity of the actual surface, this simplification might not always be sufficiently accurate. As mentioned earlier, this inaccuracy can be minimised by subdividing the entire surface into smaller areas, where the surface shape can be approximated by a straight plane. In section 4.2, the surface height within the reconstructed area is quantified using a single parameter to obtain a straightforward comparison with the single-point pressure measurements. However, the model is extended with an extra varying inclination to more accurately describe the recorded refraction, as will be explained in section 4.2.

Anther solution might be to include a more extensive surface model, which should be able to describe a complex surface shape within the reconstructed zone more accurately. However, adding more parameters that need to be fitted complicates the optimisation problem, and could require multiple views or a different error function to obtain a reliable solution of the best-fitted surface shape.

Next to these error sources, the accuracy and applicability of the image-based technique suffers from three limitations that are the result of the optical geometry. Figure 7 illustrates two limitations, inherent to refraction-based techniques.

A first limitation is the maximum amplitude of the surface oscillations, which needs to remain limited with respect to the pattern-to-surface distance to avoid crossings of light
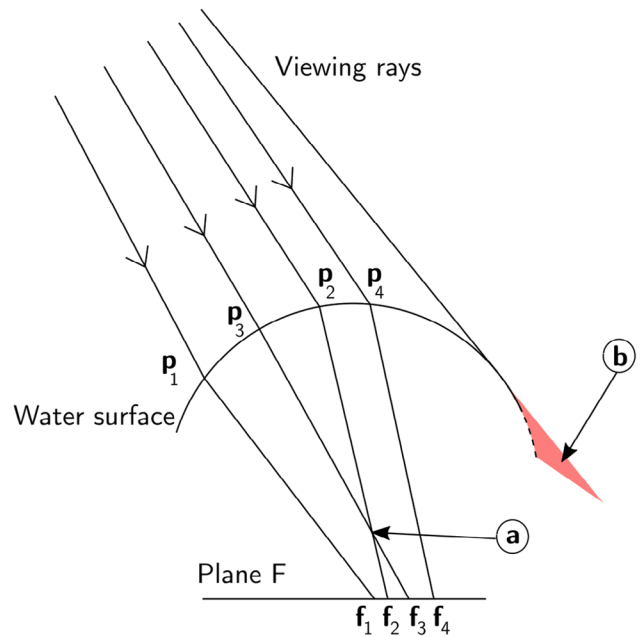
Figure 7. Limitations of refractive-based techniques. (a) The pattern-to-surface distance must remain smaller than a critical distance to avoid caustics. (b) Shaded area becomes invisible for the camera due to the strong inclination of the surface.

rays (caustics). In case of caustics, two neighbouring points on F (e.g. f_2 and f_3) are viewed on the surface in the opposite order (p_3 and p_2). For parallel incident rays and a sinusoidal surface shape, [19] proposed a maximum critical distance between the surface and the pattern on the bottom in function of the surface waves' amplitude and wavelength. In a more realistic situation, for which the incident light rays converge towards the camera centre, [14] mention that this critical distance also depends on the viewing angle and the curvature of the surface. Therefore, a direct guideline to avoid such caustics is difficult to derive.

Secondly, the inclination of the surface must remain smaller than the inclination of the viewing rays in order to remain visible for the camera. Because the surface area in the invisible, shaded area in figure 7 does not contain any surface points seen by the camera, it cannot be included in the optimisation.

Finally, the maximum accuracy of the technique presented in this paper is linearly dependent on the camera resolution. To register a change in surface shape, the surface projection of the adopted chequerboard pattern has to move in the image plane of the camera. This image projection of the refractive pattern distortion, previously denoted as the image disparity shift, is determined by multiple variables:

- the resolution of the camera
- the distance between the camera and the imaged surface area
- the viewing angle (θ_{air})
- the (average) water depth
- the relative change in surface shape (water depth and/or inclination)

Appendix presents a theoretical estimate of the accuracy that can be reached, given a certain optical geometry.

Table 1. Overview of validation tests on still water between an initial (h_i) and final water depth (h_f), in which E_m is defined as the mean difference in water depth variation measured by the level gauge and the image-based methodology.

| | h_i (mm) | h_f (mm) | Camera 1 | | Camera 2 | | 2 Cameras | |
			E_m (mm)	E_m/h_m (%)	E_m (mm)	E_m/h_m (%)	E_m (mm)	E_m/h_m (%)
Test 1	28.2	47.7	0.22	0.30	0.29	0.41	0.48	0.67
Test 2	55.5	74.4	0.24	0.18	0.24	0.18	0.27	0.21
Test 3	87.8	117.0	0.13	0.06	0.15	0.07	0.24	0.12

Herewith, it is assumed that the image projection of the white–black border between two consecutive chequerboard squares should move at least 1 pixel in the image plane for a change in surface position to be detected. In section 4.1, it will be shown that in practice, this requirement is too stringent. By combining the refractive distortion of the entire pattern during optimisation of (a limited set of) the surface parameters, a still water surface could in this paper be reconstructed with a much higher accuracy than theoretically predicted.

4. Validation

4.1. Validation on still water

4.1.1. Validation methodology.
Because the presented methodology resolves the temporal reconstruction for every image-frame separately, a validation with a still and horizontal water surface is only a small simplification compared to a dynamic case. It allows on the other hand to compare the obtained water level differences with accurate and reliable measurements using a level gauge (measurement accuracy of the laboratory's level gauge = 0.1 mm).

To enable a straightforward comparison with the single-point measurements of the level gauge, only a small part of the cavity is reconstructed. The selected surface area, approximately 30 mm × 30 mm, is located around the measurement location of the gauging needle. For both the validation on still and dynamic water, a coordinate system with origin at the bottom of the upstream corner will be used, as depicted in figure 2. The adopted parameter model to quantify the water depth within the selected area represents a straight plane, of which the (x,y,z)-coordinates are determined by the following formula:

$$z \equiv \eta(x,y) = a_0 + a_1 \frac{(x - x_c)}{L_x} + a_2 \frac{(y - y_c)}{L_y} \tag{7}$$

where a_0, a_1, a_2 are the optimised surface parameters; x_c and y_c correspond to the approximate centre of the reconstructed surface area; and L_x and L_y are length scales chosen based on the dimensions of the reconstructed zone. The validation is done by comparing the water depth change between two level gauge measurements and between two optimised surfaces $\eta(x,y)$, corresponding to two different still water levels. Because each test is prone to the inaccuracy related to two individual level gauge readings, the water depth changes measured by the level gauge are reliable within an uncertainty range of 0.2 mm.

Table 1 gives an overview of the three uncorrelated validation tests, each characterised by an initial (h_i) and final water

Table 2. Position of cameras used in this paper. Camera 1 and 2 are employed during the validation on still water (test 1, test 2, test 3). For dynamic test 4, images of camera 2 are also used as input. Camera 3, positioned closer to the surface, has a higher pixel cm^{-1} ratio and is adopted for test 5.

	x (mm)	y (mm)	z (mm)
Camera 1	35	771	1048
Camera 2	297	764	858
Camera 3	413	22	531

Table 3. Summary of the theoretical maximum expected accuracy of the three tests on still water.

	h_m (mm)	$\theta_{\text{air},m}(°)$	$\Delta h_{\min}$ (mm)
	Still water test 1		
Camera 1	38.0	35	0.64
Camera 2	38.0	42	0.59
	Still water test 2		
Camera 1	65.0	35	0.50
Camera 2	65.0	42	0.48
	Still water test 3		
Camera 1	102.4	36	0.31
Camera 2	102.4	43	0.31

depth (h_f). Two different cameras (camera 1 and 2) are used to acquire images of the still water surface, of which the position with respect to the reconstructed surface area is depicted in figure 2. For every camera position used in this paper, table 2 gives the location of the camera centre in the chosen frame of reference (figure 2).

For each change in water level, the water level difference measured by the level gauge is compared with the image-based results. In table 1, the single-view results are presented in columns 3–4 and columns 5–6 for camera 1 and 2, respectively. Additionally, two simultaneous camera views of the surface (two cameras) are combined to reconstruct the overlapping surface area, of which the results are given by columns 7–8. An error metric E_m for each test is calculated as the mean difference between the image-based and level gauge measurements, within a fixed surface area located around the measurement location of the level gauge:

$$E_m = \sum_{(x,y) \in A_{\text{gauge}}} |\Delta h_{\text{gauge}} - \Delta h_{\text{optical}}(x,y)| \tag{8}$$

where A_{gauge} represents the surface area over which the difference in water depth is calculated; Δh_{gauge} is the change in water depth measured by the level gauge; and $\Delta h_{\text{optical}}(x,y)$ gives the difference in water depth at location (x,y) based on the optimised surface models corresponding to h_i and h_f.

Additionally, a non-dimensional error metric is provided as a percentage of the mean water depth of a single test: $h_m = \frac{h_i + h_f}{2}$.

4.1.2. Validation results. As was mentioned in section 3, the maximum accuracy of the presented methodology depends on the magnitude of the image-disparity shift corresponding to a certain water level difference. Appendix describes a theoretical approach to estimate the smallest change in water level ($\Delta h_{\min}$) that can be reconstructed for a specified optical set-up, in which it is assumed that every surface point has to move at least 1 pixel in the image plane. Given that 1 cm of the reconstructed surface corresponds with approximately 30 pixels in the image plane of both cameras, the theoretical maximum accuracy is calculated based on the average water depth (h_m) and average viewing angle ($\theta_{\mathrm{air},m}$) with respect to the reconstructed surface area. For more details on this theoretical estimation, see appendix. The results shown in table 3 suggest that the accuracy of both cameras should improve with increasing average water depth, as this increases the magnitude of the image disparity shifts for a certain water level difference. Moreover, the small difference in viewing direction between camera 2 and camera 3 is not expected to influence the measurement accuracy.

Table 1 shows that the difference between the image-based results and those measured by the level gauge are of the same order of magnitude as the uncertainty of the level gauge measurements (0.2 mm). As expected, the small difference in viewing direction between camera 1 and camera 2 does not seem to affect the measurement accuracy significantly. Moreover, the experimental results confirm that the accuracy and robustness of the methodology improves with increasing (average) water depth. Especially in case the dimensionless errors (E_m/h_m) are considered, the mean error with respect to the level gauge measurements reduces significantly.

However, the experimentally achieved accuracy seems significantly higher than theoretically predicted. This indicates that not every image point has to move in the image plane for the surface change to be detected, as was assumed during the theoretical estimation (appendix). This result proves the advantage of the integrated reconstruction approach, using the entire pattern deformation instead of a sparse set of characteristic points. Although the image disparity shift in figure 8 is smaller than 1 pixel, the white–black border indicated in green still moves from pixel q_i to pixel q_{i+1}. Even if such a change in pixel position might not occur for every point of the chequerboard pattern, a small amount of 'shifts' seems sufficient to fit the surface parameters and reconstruct the surface change.

Finally, table 1 indicates that combining the views of two cameras does not improve the reconstruction results. This proves that, in case sufficient care is taken to the correct positioning and calibration of the camera, a single view is sufficient to reconstruct the surface accurately. The superior performance of a single-view approach is nonetheless not expected, as error sources that are related to a single camera (e.g. a wrong estimation of the camera position, undistortion of the images, unfavourable camera position with regards to
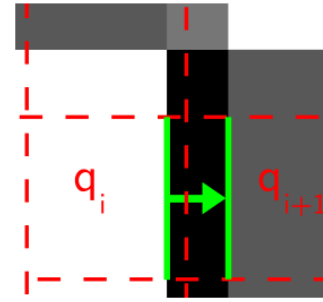
Figure 8. Although the disparity shift (indicated in green) of the pattern due to a change in surface level is smaller than 1 pixel, the image position of the border line between a black and white square still changes from pixel q_i to pixel q_{i+1}.

the surface...) should be mitigated by using more than one camera to reconstruct the same surface area.

The diminished performance of two cameras can be explained by the fact that it is usually difficult to select the exact same surface area for two different camera views. In theory, this does not impose a problem when a horizontal (flat) water surface is reconstructed. However, the flume used during the experiments does not have a perfectly horizontal bottom; the maximum local inclination equalling about 1 cm m^{-1}. Hence, a small water depth variation exists within the reconstructed surface area ($\approx$0.30 mm). When the two cameras are adopted separately, comparison of their reconstructed water level differences is still valid because a water level change is constant over the entire surface in case two horizontal surfaces are compared.

However, when two simultaneous views of the surface are combined, the reconstructed surface areas in the images of both cameras do not correspond with exactly the same water depth. Such a small difference is usually negligible when applied to practical applications that do not require sub-millimetre accuracy. Nonetheless, this discrepancy makes that the optimal solution for both cameras is not longer the same and a compromise for both cameras has to be made. Moreover, the variation in water depth is of the same order of magnitude as the measurement accuracy that is achieved in the present experiments, which explains the larger errors compared to the single-view results.

One should note that the excellent results presented in table 1 are achieved by fine-tuning the optimisation settings to achieve the best possible correspondence. These results were obtained with a well chosen differentiation step and an initial estimate of the water depth, although equal for every camera, close to the water level measured by the level gauge ($\pm$3 mm). Additionally, still water makes the fitted surface model (a straight plane) exactly correspond with the actual surface shape. This proves the inherent capability of the image-based methodology. However, such accuracy is usually not possible when applied to more challenging experiments where the surface is too irregular to describe exactly and the optimisation settings cannot be adjusted for every reconstructed surface. In the next section, the methodology will be tested in a more realistic experimental setting.

4.2. Validation on dynamic water surfaces

4.2.1. Validation methodology. In this section, the same cavity set-up is adopted yet with steady flow in the main channel next to the cavity. The resulting flow conditions are characterised by a non-steady surface shape and a mixing layer that develops at the interface between the main channel and the lateral cavity. This undulating mixing or shear layer is characterised by quasi-periodic, large-scale vortices that are shed from the upstream corner [27]. One or multiple recirculation cells or gyres generally form within the cavity, which interact with this vortex shedding. For certain combinations of geometry, water depth and flow velocity, standing gravity waves are excited within the cavity that oscillate at the cavity's natural frequency [18, 28]. Two dynamic surfaces are discussed in this paper (test 4 and test 5), for which the flow conditions are chosen to maximise these standing waves. An overview of the flow characteristics is given by table 4, in which the dimensionless Reynolds and Froude number are defined as

$$\text{Re} = \frac{U_m h_m}{\nu} \text{ and } \text{Fr} = \frac{U_m}{\sqrt{g h_m}} \tag{9}$$

where U_m and h_m are the average velocity and water depth in the main channel; g is the gravitational acceleration; and ν is the fluid's kinematic viscosity, equal to $1.0035 \cdot 10^{-6}$ m^2 s^{-1} for water at 20 °C.

A UNIK5000 pressure sensor of GE Measurement & Control is used to acquire simultaneous measurements for validation of the surface reconstruction technique. It registers the pressure fluctuations at the bottom of the cavity with a sampling frequency equal to the frame rate of the camera (40 Hz) and a measurement accuracy of ± 0.05 mm. The inlet of the pressure sensor is connected with a tube to a small hole at the upstream side of the water retaining outer wall of the cavity (figure 3). At this location, the surface variations due to the gravity waves are significantly larger than the measurement uncertainty of the pressure sensor.

Similar to section 4.1, a small region (30×30 mm) at the upstream wall of the cavity is reconstructed, located close to the pressure sensor inlet. To avoid the ambiguity between a change in water depth and inclination as discussed in section 3, a (straight) plane quantified by a single water depth is used to describe the selected region. Although this neglects a varying surface inclination, it is reasonable to assume that the local water depth in a small reconstruction area does not vary significantly and can be approximated with a single averaged value (a_0). This simplification of the optimisation problem also facilitates the comparison with the single-point pressure readings, as the algorithm directly computes the best average water depth in the reconstructed region close to the pressure sensor inlet.

The first dynamic test (test 4) employs images of camera 2, positioned at the same position as during the tests on still water (figure 2). Camera 2 is positioned such that 1 cm of the reconstructed surface area corresponds with approximately 30 pixels in the image plane. The camera position in the second dynamic test (test 5) is indicated as camera 3 in figure 2.

Table 4. Overview of the flow characteristics of the two dynamic tests.

	h_m (mm)	U_m (m s^{-1})	Re (—)	Fr (—)
Test 4	83.2	0.42	3.5 10^4	0.47
Test 5	97.1	0.40	3.9 10^4	0.41

Table 5. Summary of the theoretical maximum expected accuracy of the two tests on a dynamic water surface. Test 4 is done with images of camera 2, while test 5 employs images of camera 3.

	h_m (mm)	$\theta_{air}(°)$	Δh_{min} (mm)
Test 4	83.2	43	0.40
Test 5	97.1	38	0.01

Because camera 3 is positioned closer to the surface, camera 3 has a higher pixel cm^{-1} ratio of 53.

Table 5 presents a summary of the set-up characteristics that determine the theoretical magnitude of the image disparity shifts, in which the theoretical maximum accuracy (appendix) is indicated by Δh_{min}. Since $\Delta h_{min} \approx 0$ mm for test 5, the simplified theoretical approach does not longer seem valid due to the simplifications and/or assumptions made during the derivation. However, it indicates that camera 3 is able to register even very small water level differences.

In test 4, camera 2 is positioned so that the viewing direction approximately corresponds with the (transverse) *y*-direction in the chosen frame of reference. To account for a varying inclination of the surface in the (longitudinal) *x*-direction, the component n_x of the surface normal is therefore included in the optimisation. This results in a better match with the disparity shifts seen by the camera, which is perpendicular to the viewing rays' direction, and improves the reliability of the reconstruction result. The other component (n_y) of the surface normal is only computed once, using a single, averaged image of the entire image sequence. This component enables to account for a non-perfectly horizontal bottom plane relatively to the flat water surface, but is kept constant while optimising each individual image frame. This constant value ($n_y = a_2$) during reconstruction of the entire image sequence avoids ambiguity between a variation in average water depth or inclination in the *y*-direction of the reconstructed surface. The resulting parametrised surface model therefore becomes

$$\begin{aligned} z &= \eta(x, y) = a_0 \text{ and} \\ \overrightarrow{n} &= (a_1, a_2, 1) \end{aligned} \tag{10}$$

in which a_2 is a constant during the optimisation of the individual images; and a_0 and a_1 are the two parameters that need to be fitted.

In test 5, the adopted camera (camera 3) is positioned differently, such that the viewing direction approximately corresponds with the (longitudinal) *x*-direction. Similarly to the first dynamic test (test 4), only the surface normal component perpendicular to the viewing rays is included in the optimisation algorithm. Therefore, the same surface model (equation (10)) is used, but for test 5 characterised by a constant a_1 and a variable a_0 and a_2 that are fitted to the refractive deformation.
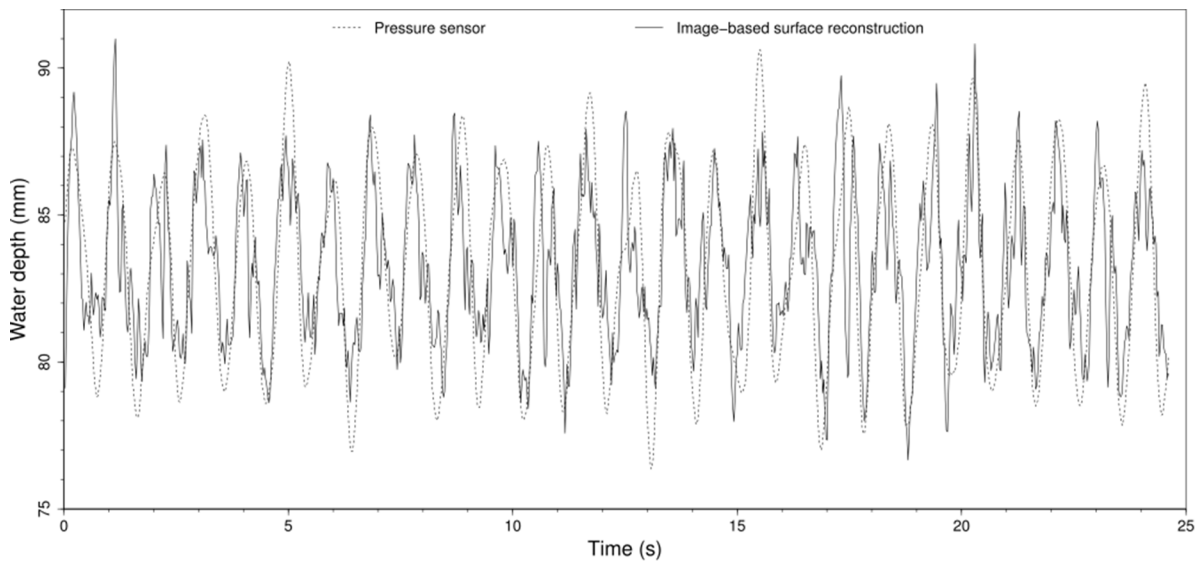
Figure 9. Dynamic test 4: Comparison between the measured water depth variation using the pressure sensor (dotted line) and the image-based surface reconstruction (solid line).
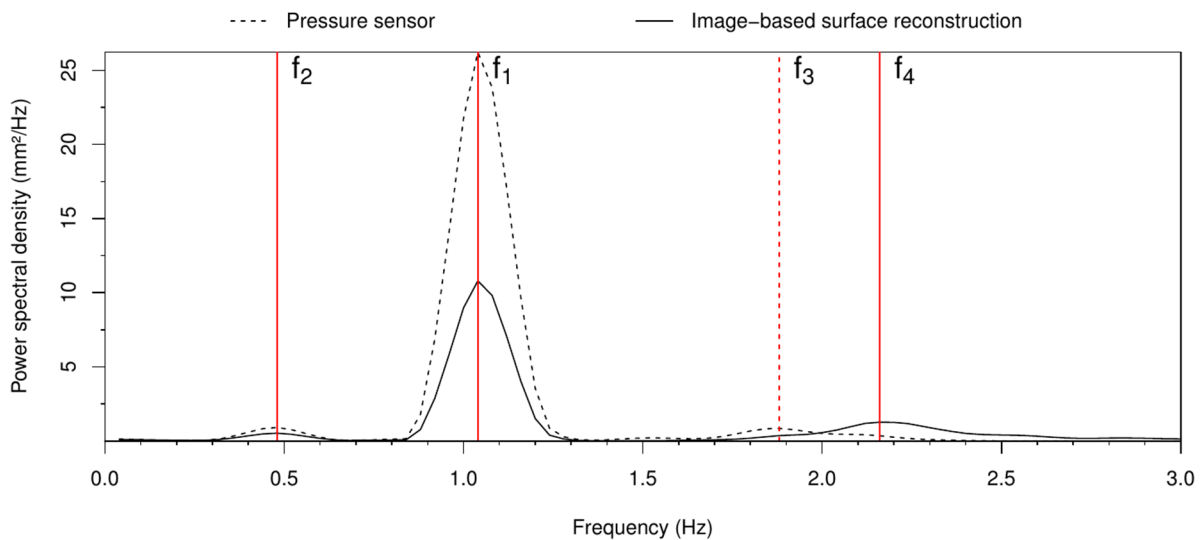

Figure 10. Dynamic test 4: Smoothed power spectral density of the water depth variation using the pressure sensor (dotted line) and the image-based surface reconstruction (solid line). Peaks in the spectral plots are indicated by a vertical line at frequency f_i.

4.2.2. Validation results. Figure 9 illustrates the variation of the local water depth for the first dynamic test (test 4), in which the dotted line corresponds with the pressure measurements and the solid line with the image-based results. This shows that in general, the data measured with the pressure sensor and the image-based time series (the variation of a_0) follow the same sinusoidal trend. The sub-millimetre accuracy that was achieved in section 4.1 is nonetheless not retained, and the image-based time series suggests a much more irregular variation of the water depth. Although both time series oscillate approximately in phase, the amplitude of the image-based surface oscillations sometimes differs from those measured by the pressure sensor. Plausible reasons for this discrepancy will be suggested in section 5.

Because the surface oscillates periodically in time, the spectra of the time series should reveal the most dominating frequencies. Figure 10 compares the (smoothed) power

spectral density of the image-based time sequence (solid line) and that of the pressure readings (dotted line), both computed with a fast Fourier transform of the original signal. Both power spectra show a clear maximum at a peak frequency f_1 of 1.04 Hz, indicating that the water oscillations are dominated by periodic surface gravity waves oscillating at that frequency. Other modes of oscillation are also visible, although the energy corresponding to these frequencies is much smaller. Some of these secondary peaks even only appear in one of both spectra, which will further be discussed in section 5.

Although also present in test 4, a large amount of small-amplitude surface ripples (see section 5) disturbed the water surface in test 5. This made it difficult to use the raw camera images, being heavily distorted by these high-frequency surface effects. As will be explained in section 5, averaging over multiple images enables to deal with these disturbances and obtain a reliable reconstruction result. The power spectra
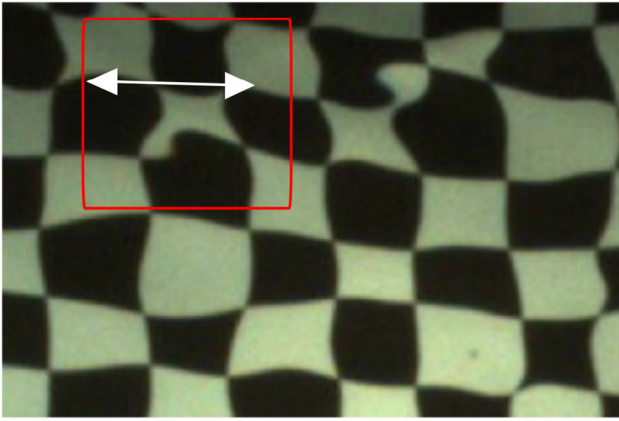
Figure 11. Local surface ripples cause a strong curvature and inclination of the surface shape and result in significant distortion of the adopted chequerboard pattern. The reconstructed surface area is indicated in red, while the approximate wave length of the disturbing surface ripples is indicated in white.

corresponding to test 5 are not presented in this paper, but also showed the same peak frequencies in the pressure and image-based time series.

5. Discussion

The validation with still water proved sub-millimetre accuracy, but this level of accuracy could not be achieved in a validation with a dynamic surface. Figure 9 indicates that, although the image-based results show a good resemblance with the pressure readings, the amplitude of the image-based water depth oscillations sometimes differs from those measured by the pressure sensor. Moreover, the image-based water depth is characterised by a much more irregular and peaked variation in time. In the following, plausible reasons for this behaviour are explored and it is discussed how the results can be improved in further research.

A first reason is the simplified surface model adopted during the dynamic validation, in which the (average) water depth within the reconstructed zone is quantified by a single value a_0 and the corresponding surface normal can only be fitted by a component in the longitudinal direction of the cavity. Consequently, this surface model may not be appropriate to describe the more complex surface shape during the dynamic test with sub-millimetre accuracy.

Secondly, it was visually observed that the dynamic cases are affected by small and high-frequency surface perturbations. Generated at the downstream corner next to the main channel, these small surface waves are transported within the cavity due to the circulatory flow field. The surface ripples disturb the surface locally and cause a significant curvature of the surface shape, resulting in a heavily distorted view of the pattern. Figure 11 shows that the wave length of these small-amplitude waves, indicated in white, is of the same order of magnitude as the dimensions of the reconstructed area, indicated in red. Such distortion makes the optimisation problem less clearly defined and makes it more difficult to fit a straight surface to the complex local surface shape. The optimisation algorithm then sometimes stops its search close

to the starting position of the iterative optimisation, in this paper being an estimate based on the coefficients found in the previous time-step. Although including more parameters in the model theoretically enables to better describe the irregular surface, the simplified surface model given by equation (10) is considered as a suitable choice considering the objectives of this paper. The flat surface makes an easy and straightforward comparison with the pressure readings possible, and is also more generally applicable because it does not require any prior knowledge about the expected surface shape.

Compared to the measurements using the pressure sensor, the image-based methodology seems much more sensitive to these small disturbances that are superimposed on the large-scale surface oscillations. While having a small influence on the local water depth as measured with the pressure sensor, their strong local curvature causes a significant refraction which affects the image-based time sequence.

One possible solution to deal with these local surface ripples is to use the average intensity over multiple images as input for the reconstruction algorithm. This averaging filters out the high-frequency deformations, since it can be assumed that the average refractive distortion is dominated by the slower and large-scale gravity waves within the cavity. However, averaging over too much images is also not advised because it reduces the temporal resolution of the time-dependent reconstruction. Taking the median over four images gave for the test case presented in this paper a reliable result, in which the large-scale dynamics could still be distinguished. Figure 12 provides a comparison between the image-based time series, using the median over four subsequent images, and the median pressure reading within the time span of the four images. This decreases the influence of the surface ripples and moderates the peaked behaviour of the original image-based time series (figure 9). On the other hand, it also attenuates the large-scale dynamics of the surface oscillations. This attenuation is much more pronounced in the time series of the image-based results, of which the amplitude is smaller than that of the smoothed pressure readings. Depending on the time-scale of the studied phenomenon and the presence of undesired surface ripples, averaging the image sequence seems useful but should be done with care. A compromise might thus be necessary to smooth out noise induced by these disturbances while minimising the reduction in amplitude and temporal resolution of the reconstructed surface dynamics.

A second solution is to use a more extensive surface model to describe the surface shape in the area influenced by surface ripples. This is not attempted in this paper because the single-parameter model allows a straightforward comparison with the single-point measurements of the level gauge. Additionally, the simplified optimisation problem avoids ambiguity between a change in water depth and a variation in inclination as discussed previously. This ensures a reliable solution while using only a single view to reconstruct the surface.

Finally, the spatial resolution of the camera influences the magnitude of the image disparity shifts, i.e. the pattern deformation due to a changing surface shape. As explained in section 3, this affects the maximum accuracy that can be achieved. Given the characteristics of camera 2 (table 5), the theory of appendix predicts a minimal water depth change of 0.40 mm
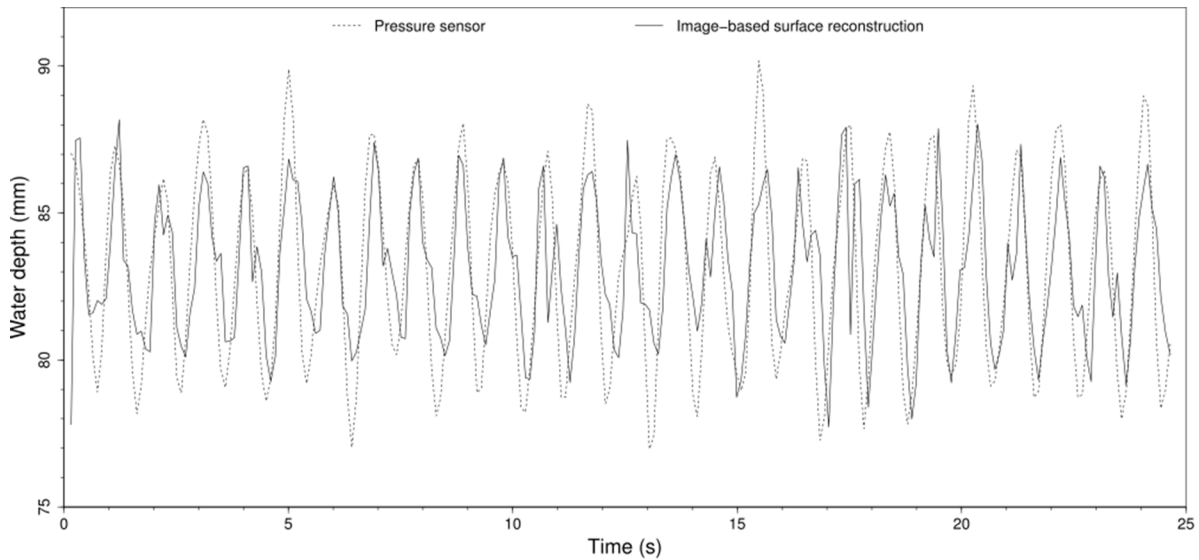
Figure 12. Dynamic test 4: Comparison between the measured water depth variation using the pressure sensor (dotted line) and the image-based surface reconstruction (solid line) using the median pressure reading over four samples or the median intensity over four images.
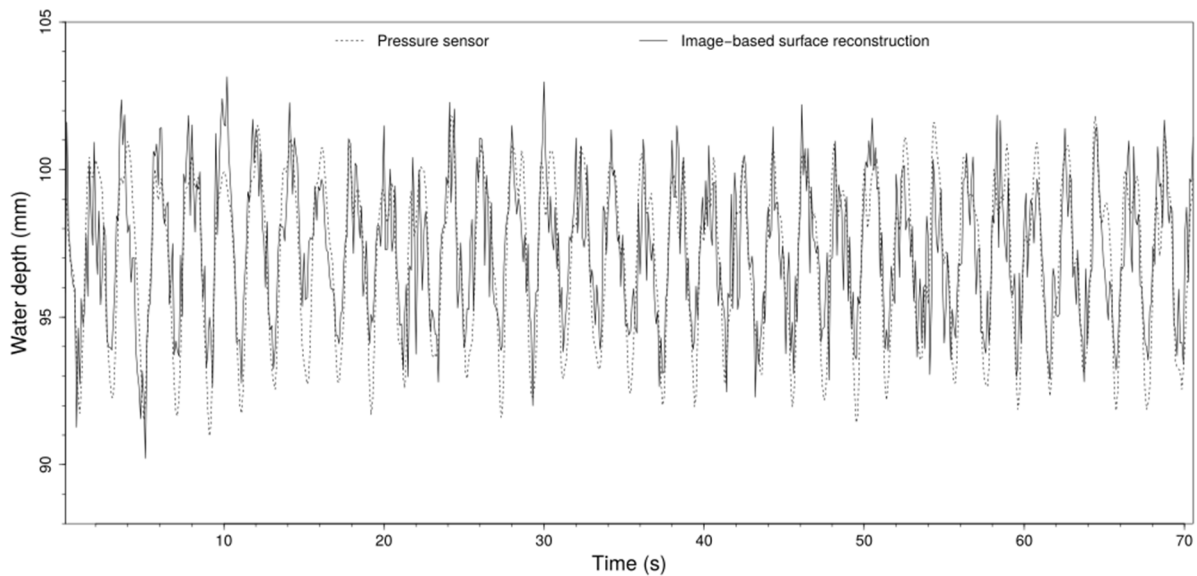
Figure 13. Dynamic test 5: Comparison between the measured water depth variation using the pressure sensor (dotted line) and the image-based surface reconstruction (solid line) using the median pressure reading over four samples or the median intensity over four images.

for an image disparity shift of minimal 1 pixel. Moreover, rapid pattern deformations complicate a pixel-accurate distinction between black and white squares. In practical applications, image disparity shifts of even more than 1 pixel are thus preferable to improve the reliability of the results. In case the adopted camera has a higher pixel cm^{-1} ratio, the same movement of surface points p correspond with more pixels in the image plane. Hence, increasing the pixel cm^{-1} ratio of the camera should improve the performance of the methodology.

To verify this statement, camera 3 is positioned closer to the surface (see figure 2 and table 2). This increases the pixel cm^{-1} ratio for camera 3 (test 5) to 53, almost twice as large as the pixel cm^{-1} ratio = 30 for camera 2 (test 4). Because test 5 was even more affected by small surface ripples, a similar procedure as for the first dynamic test is followed. The median intensity over four images is used as input for the

surface reconstruction, and is subsequently compared with the median over 4 pressure readings. Figure 13 depicts the (averaged) pressure time series and image-based results for test 5. The more peaked behaviour of these (averaged) image-based results indicates that a larger pixel cm^{-1} ratio increases the sensitivity of the methodology and results in less smoothing compared to figure 12.

To quantitatively compare the two dynamic tests, the rms error is calculated for every time series (test 4: raw image sequence and averaged image sequence; test 5: averaged image sequence). For all three cases, the rms error is determined by comparison between the image-based time series and the corresponding simultaneous pressure time series. For test 4, this results in a rms error of 2.20 mm and 2.00 mm for the raw images and averaged image sequence, respectively. For test 5, the rms error equals 1.68 mm. This better

agreement between the (averaged) image-based and pressure time series of test 5 confirms that increasing the pixel cm^{-1} ratio of the camera improves the reconstruction accuracy. In practice, this can be achieved by using more advanced (high-resolution) optical equipment or positioning the camera closer to the imaged surface.

The power spectral densities that are presented in section 4.2 (figure 10) prove that the dominant frequency f_1 of 1.04 Hz is the same for both time series of test 4 (raw images). Also a second peak, indicated as f_2 at 0.48 Hz, is in both spectral plots clearly distinguished. However, the pressure data shows a third peak f_3 at 1.88 Hz ($\approx 4 \times 0.48$ Hz), while the image-based results are characterised by a peak f_4 at 2.16 Hz ($\approx 2 \times 1.04$ Hz). Formulas to predict the frequency of the (bi-directional) oscillations within a lateral cavity (see [29]) suggest that for test 4, the first and second harmonic of longitudinal standing gravity waves in the cavity occur at approximately f_1 and f_4. This indicates that the chosen surface model, adapted to include inclinations of the surface in the longitudinal direction, is able to discern these surface oscillations reliably. The transverse oscillations in the cavity are on the other hand expected at a frequency f_2 (or higher harmonics at a multiple of f_2). Transverse modes of a higher order (multiples of f_2) are nonetheless less energetic and influential on the the surface dynamics. Moreover, they are more difficult to register using the image-based methodology because of the chosen surface model (n_y constant) and viewing direction of the camera. Their smaller amplitude and the smaller refractive disparity shift caused by these standing waves might explain why the peak at f_3 is not present in the image-based spectrum.

6. Conclusion

An image-based and non-intrusive surface reconstruction approach is presented, compatible with (subsurface) image-based velocity measurement techniques. It only requires a pattern to be visible through the water column, and a single camera view is sufficient due to the parametric description of the surface shape. This single-view approach makes a more flexible experimental set-up possible, and improves the robustness and reliability because no feature tracking is needed in a twofold manner. Firstly, it makes the optimisation of a single image less sensitive to these disturbing surface effects. Secondly, it ensures that surface reconstruction for a single image is not affected by loss of feature points in previously processed image frames. Compared to the original method presented by [21], this ensures that reconstruction of the entire surface can be maintained in case of high-frequency surface perturbations. Moreover, the technique remains reliable in presence of seeding particles that (partly) obstruct a clear view of the pattern below.

A validation on still water, which is valid due to the image-based reconstruction methodology, shows that the presented measurement technique is able to reconstruct steady, shallow flows with sub-millimetre accuracy. Due to the integrated approach, in which the deformation of the entire pattern is used to fit the (limited set of) surface parameters, the experimentally achieved accuracy is significantly better than theoretically predicted. Additionally, the static validation confirms that a larger water depth improves the reconstruction accuracy, as was also expected based on the simplified theoretical derivation.

Furthermore, the methodology is validated by two tests on a dynamic surface phenomenon, more specifically to surface oscillations in an open-channel lateral cavity driven by the flow past the cavity. Although the accuracy of the static validation cannot be achieved, the reconstructed time series show close correspondence with simultaneous local pressure measurements. Additionally, a comparison of the power spectra of the first test indicates that the pressure and image-based time series are characterised by the same dominant peak frequency, corresponding to surface gravity waves in the longitudinal cavity direction.

Because the presented technique relies on the detection of a regular pattern, projected on the water surface, small surface perturbations complicate the image-based reconstruction. In this paper, these surface ripples are mitigated by averaging the image-intensity over subsequent images in order to filter out the high-frequency surface effects. This reduces the difference between the pressure and image-based time series, although also the large-scale dynamics of the water oscillations are (partly) attenuated. Finally, two tests on a dynamic surface are compared, in which the camera is positioned differently with respect to the surface. This indicates that increasing the pixel cm^{-1} ratio of the camera enhances the sensitivity to small changes in water depth and improves the reconstruction accuracy.

Although the presented validation is based on a comparison of only a small reconstructed surface area with alternative measurement techniques, good agreement for even non-smooth surface conditions is achieved. Moreover, the proposed averaging procedure allows to cope with high-frequency surface perturbations, which makes the image-based methodology useful for laboratory experiments of numerous hydraulic research fields. This methodology can easily be extended to a larger mesh composed of independently reconstructed areas. Further research will combine the presented technique with simultaneous (PTV) velocity measurements, enabling detailed understanding of the cavity flow field.

Acknowledgments

Appendix. Theoretical estimate of the image disparity shift

In the following, a theoretical approach is presented to estimate the image-disparity shifts of the depicted pattern due to a changing water surface. Herewith, some simplifying assumptions are made, although we believe that the results can still be used to gain more understanding in the reconstruction

conditions that affect the accuracy of the methodology. As such, we intend to provide a guideline to future users to obtain the best possible results.

It is assumed that the water surface is a flat, horizontal plane, of which the water level changes from h_i to $h_{i+1} = h_i + \Delta h$. Given that the presented technique is used to capture smooth surface phenomena, in which the inclination of the surface remains relatively small compared to the viewing angle, this simplification does not influence the general applicability of the results presented hereafter. For example, the dynamic tests presented in this paper are dominated by large-scale oscillations of the surface (not considering the disturbing surface ripples), of which the maximum surface slope varies between approximately $+35\,\text{mm m}^{-1}$ and $-35\,\text{mm m}^{-1}$. This inclination is of the same order of magnitude as the variability of the inclination angle within the small reconstruction domain, and is therefore considered negligible for the geometry of the refractive problem.

Figure A1 shows that a single viewing ray intersects the initial water surface, characterised by its water depth h_i, with an angle $\theta_{\text{air},i}$ relative to the surface normal. The direction of the refracted viewing ray, defined by $\theta_{\text{water},i}$, is then calculated using Snell's law:

$$\theta_{\text{water},i} = \arcsin(\sin(\theta_{\text{air},i})/r_w) \qquad (A.1)$$

with r_w the refractive index of water, equal to 1.336 for a wavelength of 520 nm (section 2.4) and water at 20 °C (4). The horizontal distance x_i between the surface point $\boldsymbol{p}_i$ and corresponding point $\mathbf{f}$ on the bottom plane is then calculated as

$$x_i = \tan(\theta_{\text{water},i}) \cdot h_i. \qquad (A.2)$$

In case the surface level rises to h_{i+1}, the same point $\mathbf{f}$ on the bottom plane will be viewed at surface point $\boldsymbol{p}_{i+1}$ instead of $\boldsymbol{p}_i$. Figure A1 illustrates that a water level rise results in a horizontal shift from $\boldsymbol{p}_i$ to $\boldsymbol{p}_{i+1}$, directed towards the camera in the plane of incidence of the viewing ray. Additionally, the water level rise causes an increase of the angle of incidence with respect to the surface normal from $\theta_{\text{air},i}$ to $\theta_{\text{air},i+1}$. This change in viewing direction, denoted as $\Delta\theta$, depends both on the camera position with regards to the surface and the change in water level Δh. This makes it difficult to derive a general rule of thumb to quantify $\Delta\theta$.

Therefore, we hypothesise that $\Delta\theta$ will result in a movement of 1 pixel in the image plane of the camera. The cameras adopted in this paper have a horizontal angle of view of 30° and corresponding resolution of 1920 pixels, which is representative for a typical (low-cost) optical set-up. As a consequence, a movement of 1 pixel in the image plane corresponds to a change in viewing direction of $\Delta\theta = 30°/1920 = 0.0156°$. As such, the new angle of incidence of the viewing ray becomes $\theta_{\text{water},i+1} = \theta_{\text{water},i} + \Delta\theta$. The refracted viewing direction $\theta_{\text{water},i+1}$ is then found similarly as $\theta_{\text{water},i}$:

$$\theta_{\text{water},i+1} = \arcsin(\sin(\theta_{\text{air},i+1})/r_w). \qquad (A.3)$$

The corresponding horizontal distance x_{i+1} between $\boldsymbol{p}_{i+1}$ and $\mathbf{f}$ is measured in the same plane of incidence as x_i:

$$x_{i+1} = \tan(\theta_{\text{water},i+1}) \cdot h_{i+1}. \qquad (A.4)$$
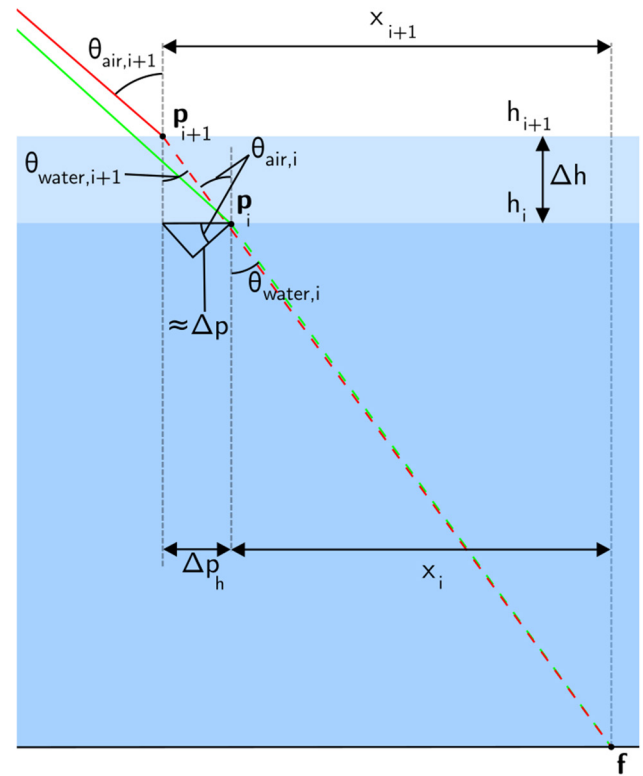
Figure A1. Schematic overview of the theoretical approach to estimate the image disparity shift for a given change in water depth Δh and initial viewing direction $\theta_{\text{air},i}$.

The horizontal shift from $\boldsymbol{p}_i$ towards $\boldsymbol{p}_{i+1}$ is then calculated as

$$\Delta p_h = x_{i+1} - x_i. \qquad (A.5)$$

Finally, the distance between $\boldsymbol{p}_i$ and $\boldsymbol{p}_{i+1}$, measured perpendicular to the viewing rays' direction, can be approximated by

$$\Delta p \approx \Delta p_h \cdot \cos(\theta_{\text{air},i}). \qquad (A.6)$$

In case the distance Δp corresponds to more than 1 pixel in the image plane of the camera, the original assumption about $\Delta\theta$ is correct. As an example, figure A2 shows Δp in function of the initial viewing direction $\theta_{\text{air},i}$, for a varying initial water depth h_i and a given change in water depth $\Delta h = 0.5$ mm. The red, green and blue lines denote the (average) viewing angle and mm pixel^{-1} ratio of camera 1, camera 2 and camera 3, respectively.

Table A1 gives an overview of the set-up characteristics that determine the image disparity shifts of the five validation tests presented in this paper. For each test, the direction of the viewing angle in the middle of the reconstructed domain is calculated, based on the camera position and average water depth h_m during the test. In table A1, $\Delta p_{\min}$ indicates the required movement of a surface point for a 1 pixel-image shift, based on the pixel cm^{-1} ratio of the adopted camera. The minimum water level change $\Delta h_{\min}$ that can theoretically be reconstructed is then found as the water level change for which Δp (equation (A.6)) equals $\Delta p_{\min}$. Because camera 3 is positioned closer to the reconstructed surface, 1 cm of the reconstructed surface corresponds to 53 pixels compared to 30 pixels for camera 1 and 2. This results in a smaller required
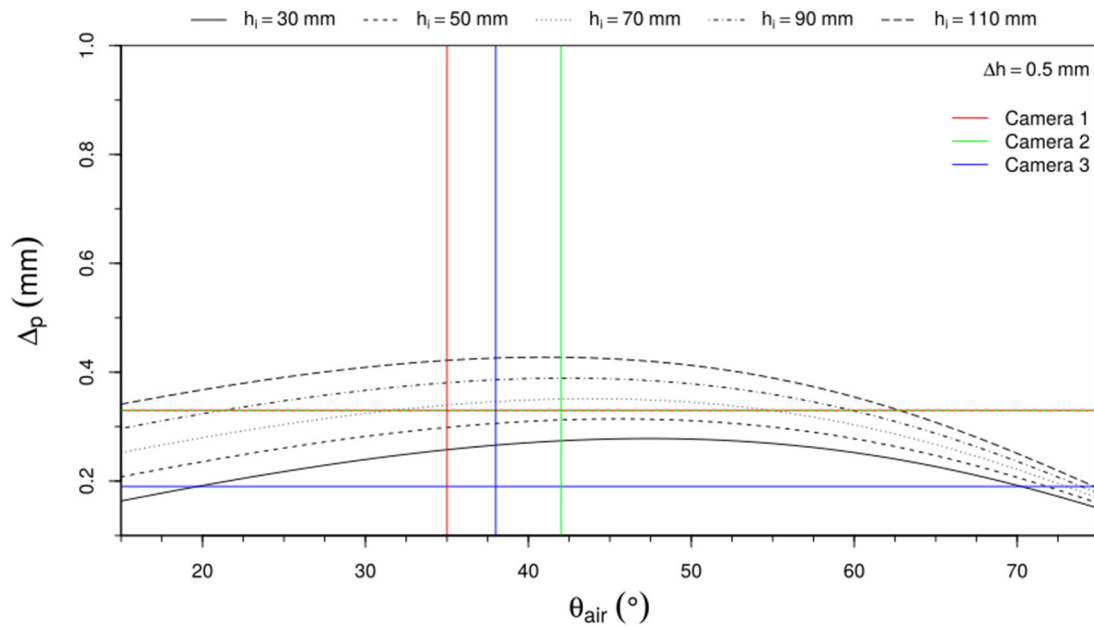
Figure A2. Δp in function of the initial viewing angle with respect to the surface normal for a varying initial water depth h_i and a given change in water level $\Delta h = 0.5$ mm. The viewing angle and required movement of the surface point **p** for a 1 pixel-image shift (Δp_{min} in table A1) are indicated in red, green and blue for camera 1, camera 2 and camera 3, respectively.

Table A1. Overview of optical and/or geometric characteristics of the tests on still water and dynamic water. The theoretical estimates of the required shift of surface point **p** (Δp_{min}) and the required minimal surface change (Δh_{min}) are given for each test and camera position.

		Camera 1			Camera 2		
	h_m (mm)	$\theta_{air,m}(°)$	Δp_{min} (mm)	Δh_{min} (mm)	$\theta_{air,m}(°)$	Δp_{min} (mm)	Δh_{min} (mm)
Still water test 1	38.0	35	0.33	0.64	42	0.33	0.59
Still water test 2	65.0	35	0.33	0.50	42	0.33	0.48
Still water test 3	102.4	36	0.33	0.31	43	0.33	0.31
	h_m (mm)	$\theta_{air,m}(°)$		Δp_{min} (mm)			Δh_{min} (mm)
Dynamic test 4 (camera 2)	83.2	43		0.33			0.40
Dynamic test 5 (camera 3)	97.1	38		0.19			0.01

Δp_{min} for a 1 pixel-image disparity shift. Based on the simplified approach, a value of $\Delta h_{min} \approx 0$ mm ($O(10^{-2}$ mm)) is found for test 5, which is considered to be the consequence of the simplifications made throughout the theoretical derivation. However, it indicates that even small surface variations can still be recorded by camera 3.

ORCID iDs

L Engelen ⬤ https://orcid.org/0000-0003-1013-2181

References

[1] Akutina Y 2015 Experimental investigation of flow structures in a shallow embayment using 3D-PTV *PhD Thesis* McGill University

[2] ALGLIB 2016 Project: ALGLIB: a cross-platform open source numerical analysis and data processing library www.alglib.net/

[3] Aureli F, Dazzi S, Maranzoni A and Mignosa P 2014 A combined colour-infrared imaging technique for measuring water surface over non-horizontal bottom *Exp. Fluids* **55** 1–14

[4] Bashkatov A N and Genina E 2003 Water refractive index in dependence on temperature, wavelength: a simple approximation *Proc. SPIE* **5068** 393–5

[5] Bechle A J and Wu C H 2011 Virtual wave gauges based upon stereo imaging for measuring surface wave characteristics *Coast. Eng.* **58** 305–16

[6] Bradski G 2017 The OpenCV Library http://opencv.org/

[7] Brown D C 1971 Close-range camera calibration *Photogramm. Eng.* **37** 855–66

[8] Burkardt J 2010 MINPACK least squares minimization http://people.sc.fsu.edu/~jburkardt/cpp_src/minpack/minpack.html

[9] Chatellier L, Jarny S, Gibouin F and David L 2013 A parametric piv/dic method for the measurement of free surface flows *Exp. Fluids* **54** 1–15

[10] Cobelli P J, Maurel A, Pagneux V and Petitjeans P 2009 Global measurement of water waves by Fourier transform profilometry *Exp. Fluids* **46** 1037

[11] Dabiri D 2003 On the interaction of a vertical shear layer with a free surface *J. Fluid Mech.* **480** 217–32

[12] Evers F M and Hager W H 2015 Videometric water surface tracking: towards investigating spatial impulse waves *Proc. 36th IAHR World Congress (The Hague, The Netherlands)* (https://doi.org/10.3929/ethz-a-010630706)

[13] Fouras A, Hourigan K, Kawahashi M and Hirahara H 2006 An improved, free surface, topographic technique *J. Vis.* **9** 49–56

[14] Gomit G, Chatellier L, Calluaud D and David L 2013 Free surface measurement by stereo-refraction *Exp.Fluids* **54** 1–11

[15] Ihrke I, Kutulakos K N, Lensch H, Magnor M and Heidrich W 2010 Transparent and specular object reconstruction *Comput. Graph. Forum* **29** 2400–26

[16] Jähne B, Schmidt M and Rocholz R 2005 Combined optical slope/height measurements of short wind waves: principles and calibration *Meas. Sci. Technol.* **16** 1937–44

[17] Kutulakos K N and Steger E 2008 A theory of refractive and specular 3d shape by light-path triangulation *Int. J. Comput. Vis.* **76** 13–29

[18] Mignot E, Pozet M, Riviere N and Chesne S 2015 Bidirectional seiching in a rectangular, open channel and lateral cavity *Proc. 22ème Congrès Français Mécanique (Lyon, France)*

[19] Moisy F, Rabaud M and Salsac K 2009 A synthetic schlieren method for the measurement of the topography of a liquid interface *Exp. Fluids* **46** 1021–36

[20] Morris N J and Kutulakos K N 2011 Dynamic refraction stereo *IEEE Trans. Pattern Anal. Mach. Intell.* **33** 1518–31

[21] Morris N J W 2004 Image-based water surface reconstruction with refractive stereo *PhD Thesis* University of Toronto

[22] Ng I, Kumar V, Sheard G J, Hourigan K and Fouras A 2011 Experimental study of simultaneous measurement of velocity and surface topography: in the wake of a circular cylinder at low reynolds number *Exp. Fluids* **50** 587–95

[23] Rupnik E, Jansa J and Pfeifer N 2015 Sinusoidal wave estimation using photogrammetry and short video sequences *Sensors* **15** 30784–809

[24] Savarese S, Chen M and Perona P 2005 Local shape from mirror reflections *Int. J. Comput. Vis.* **64** 31–67

[25] Tsubaki R and Fujita I 2005 Stereoscopic measurement of a fluctuating free surface with discontinuities *Meas. Sci. Technol.* **16** 1894–902

[26] Turney D E, Anderer A and Banerjee S 2009 A method for three-dimensional interfacial particle image velocimetry (3d-ipiv) of an air–water interface *Meas. Sci. Technol.* **20**

[27] Uijttewaal W S 2014 Hydrodynamics of shallow flows: application to rivers *J. Hydraul. Res.* **52** 157–72

[28] Wolfinger M, Ozen C and Rockwell D 2012 Shallow flow past a cavity: coupling with a standing gravity wave *Phys. Fluids* **24** 104103

[29] Wüest A and Farmer D 2003 Seiches *Encyclopedia of Science and Technology* (London: McGraw-Hill)

[30] Zappa C J, Banner M L, Schultz H, Corrada-Emmanuel A, Wolff L B and Yalcin J 2008 Retrieval of short ocean wave slope using polarimetric imaging *Meas. Sci. Technol.* **19** 055503

[31] Zhang X and Cox C S 1994 Measuring the two-dimensional structure of a wavy water surface optically: a surface gradient detector *Exp. Fluids* **17** 225–37