Tracking bed load particles in a steep flume: methods and results.

J. Heyman & C. Ancey
Laboratory of Environmental Hydraulics, School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne, Switzerland

ABSTRACT: In this paper, we present an experimental study about the motion of bed load particles in water. Recent developments in the stochastic theory of bed load transport call for precise experimental data to validate the models. We set up the experiment in a tilted narrow flume, where we could control water discharge. The erodible bed was made of natural rounded particles of mean diameter 8 mm. Two high speed cameras, disposed next two each other, recorded the transport process at 200 frames per seconds over an observation window of approximatively 1 meter and during periods of 150 seconds. Contrasting with previous experimental study that aimed to track moving particles, we imposed ourselves three constraints: (i) all moving particles have to be tracked, (ii) tracking should be fully automatic, (iii) bed load transport rates have to be measured independently. The former constraint is mandatory to correctly compare to theoretical predictions while the second is a condition to be able to get large samples of trajectories in relatively short time periods. The latter is obviously needed to validate the algorithm. After briefly introducing the numerical treatment that allow the particle tracking, we present some interesting results about the bed load particles dynamics.

1 INTRODUCTION

Recently, the stochastic approach to the problem of the transport of bed load particles has regained the interest of the scientific community (Parker et al. 2000, Ancey et al. 2008, Singh et al. 2009, Turowski 2010, Ancey 2010, Martin et al. 2012, Furbish et al. 2012, Roseberry et al. 2012, Heyman et al. 2013, Furbish and Schmeeckle 2013). This might be due to the apparent failure of the majority of deterministic based formulas in predicting bed load transport rates, both in the laboratory and in the field, within descent precision (Recking 2013). Indeed, the relative error between predicted and observed transport rates rarely drops beyond 5 in laboratories, and often reaches several orders of magnitudes in natural rivers. In the latter, this surprising unpredictability is generally attributed to the lack of precise knowledge of boundary and initial conditions (topography, grain size, shear stress) and to complex effects such as armoring (Chiari and Rickenmann 2010), segregation, dunes and bars migration (Dinehart 1992, Nelson et al. 2010). In the laboratory, that is in a precisely controlled environment, the failure of averaged bed load formulas is harder to justify. Meanwhile, time fluctuations in the transport rates are strongly interfering in the sampling of the average rates (Carey 1985, Bunte and Abt 2005, Ancey et al. 2006). The amplitude of those fluctuations is known to be large, often several time larger than the mean transport rates (Ancey et al. 2006).

The stochastic approach to bed load transport aims thus to take into account the intrinsic fluctuating nature of the transport of particles in order to derive statistically consistent averages as well as their expected fluctuations. In the recent contributions, (Furbish et al. 2012) comes up with a macroscopic definition of the solid flux based on statistical arguments. Contrarily to classical bed load flux definitions, a diffusive flux appear in the macroscopic equations, due to particle velocity fluctuations. They also presented few experiments that confirmed their findings. In the same time, (Ancey and Heyman 2014) generalized a previously published stochastic model (Ancey et al. 2008) to a spatio-temporal stochastic theory of bed load transport. In a similar manner, a diffusive flux emerges from particle velocity fluctuations. The authors showed that the particle activity, defined as the number of moving particles per unit bed area, is responsible for the large fluctuations of the bed load flux. From their stochastic model, (Heyman et al. 2013) showed that correlation in time and space could emerge from particle motions, and that the theoretical predictions compared well with previous experimental studies.

Meanwhile, the refinement of theory calls for new experimental techniques to measure the dynamics of individual particle. Indeed, the knowledge of instan-

taneous positions and velocities of moving particles is particularly relevant to describe completely the phase-space domain. In the past, computer imaging techniques have been rarely used to monitor bed load transport. Since recently, many authors have preferred to treat images manually. The obtained samples were thus extremely restricted in time and space. For instance, (Drake et al. 1988) analysed an analogue film to reconstruct particle motions. In total, only a few seconds of film could be treated. It spanned a small region of the bed, for about 200 particles trajectories. (Drake et al. 1988) pointed that "[...] these areas and times [were] not large in comparison with the spatial and temporal variation in bed surface material and bed load transport [...]". Other authors, such as (Roseberry et al. 2012) and (Lajeunesse et al. 2010) recently used manual tracking to reconstruct the particle trajectories over relatively short time and space scales. Some researchers (Martin et al. 2012) succeed in tracking automatically a sample of a few particles, called tracers. However, extrapolating results to the whole solid mass in motion is not obvious. Few years ago, (Böhm et al. 2004) used an extremely narrow flume, filled with glass beads of equal diameter and reconstruct trajectories automatically over 30 seconds for a 25 cm observation window. In total, a few thousand trajectories could be obtained via an automatic tracking algorithm. However, turbulence and grain dynamics were severely affected by the strong experimental constraints. Extrapolating their results to unconstrained bed load transport is thus not obvious.

To overcome this limitations, we present an automatic algorithm that tracks bed load particles through the frames of a film. The method borrows widely used concepts to image and signal processing science (Yilmaz et al. 2006). First, we present the experimental setup and the acquisition hardware. Second, we draw the main lines of the processing algorithm and show some typical output. Finally we present several statistical results about bed load particle dynamics under water.

2 EXPERIMENTS & METHODS

2.1 *Experimental setup*

We carried out experiments in a 2.5 m long, 3.5 cm wide tilted flume. We imposed a constant water discharge at the inlet with a pump. The erodible bed was made of natural gravels of narrow grain size distribution ($d_{50} = 8$ mm). Natural particle colors were sufficiently bright to obtain a good contrast in images. Constant sediment feeding at the flume inlet was also maintained with a conveyor belt. As our problematic concerns sediment transport in mountain rivers, the slope was adjusted between 2 and 4 degree to reproduce a supercritical flow. In this regime, the water depth is small and surface waves strongly perturbs the free surface. Filming particle motions from above is thus impractical. Indeed, the use of a glass lid (to avoid light refraction effects) is not possible without perturbing the flow. As a consequence, we took the image sequences by the side transparent wall. The channel width being relatively small ($\sim 4d_{50}$), moving particles were rarely occluding each other. Experimental conditions are resumed in table 2.1.

Table 1: Experimental parameters. θ, slope in degrees; Fr, Froude number; Re, Reynolds number; τ, Shields stress; $\bar{h}$, mean water depth (m); R_h, Hydraulic radius; $\bar{u}$, mean velocity (m/s); $\bar{q}_s$, mean bed load flux (particle/s).

θ	Fr	Re	τ	$\bar{h}$	R_h	$\bar{u}$	$\bar{q}_s$
1.99	1.52	40800	0.098	0.034	0.012	0.89	1.16

We used 2 cameras Basler© A504k of resolution 1280×256 with 8 bits gray scale pixels. The cameras were placed next to each other so that their fields of vision overlap by a few centimetres. We used 28 mm lenses, a fair compromise between image deformation and angle of vision. In total, a one-meter long observation zone of the bed was achieved (Fig. 1). The frame rate was fixed to 200 fps so that the quickest particles could not move further than one diameter away between two frames. This parameter was found to be crucial for the particle trajectories reconstruction. A set of 4 halogen lights and diffusers were placed above the flume so that the camera exposure time could be decreased down to 0.02 seconds for a 4 mm lens aperture. As a consequence, even the fastest particles were sharply delimited on frames (Fig. 2). 328 Kb are required to store each frame, so that, with 20 Gb available random access memory on our computer, 30000 frames (2.5 minutes) could be acquired in once. These video sequences were taken repetitively after saving the frames in the disk.

Simultaneously, we also monitored the bed load discharge at the outlet of the flume by an indirect acoustic technique described in (Heyman et al. 2013). This technique allows for a high temporal resolution (from 0.1 seconds to days) and was systematically controlled by weighting the cumulative sediment collected. This simultaneous measurements were necessary to infer the accuracy of the processing algorithm used on video frames as shown in §2.4 .

2.2 *Image processing*

The image processing consists in three main tasks: (i) Background subtraction (ii) Centroid detection and (iii) Trajectory reconstruction. Each of these tasks can be obtained by different techniques. Without pretending to be exhaustive, we will try to review the different options that offered to us.

2.2.1 *Background subtraction*

In bed load transport, particles in motion have generally the same appearance of those resting on the erodible bed. Indeed, no distinction between "moving" or

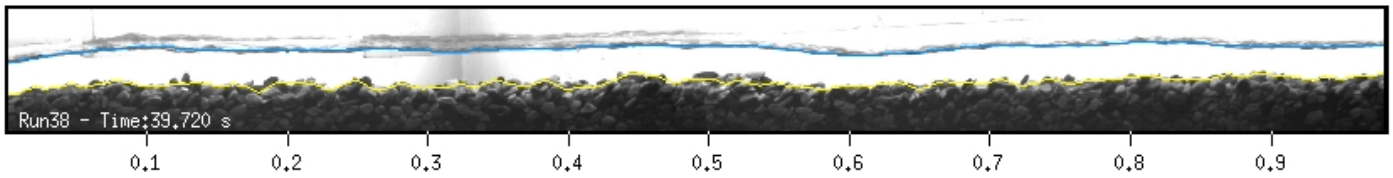

Figure 1: The whole observable window. Estimated water and bed elevation were superposed in blue and yellow.

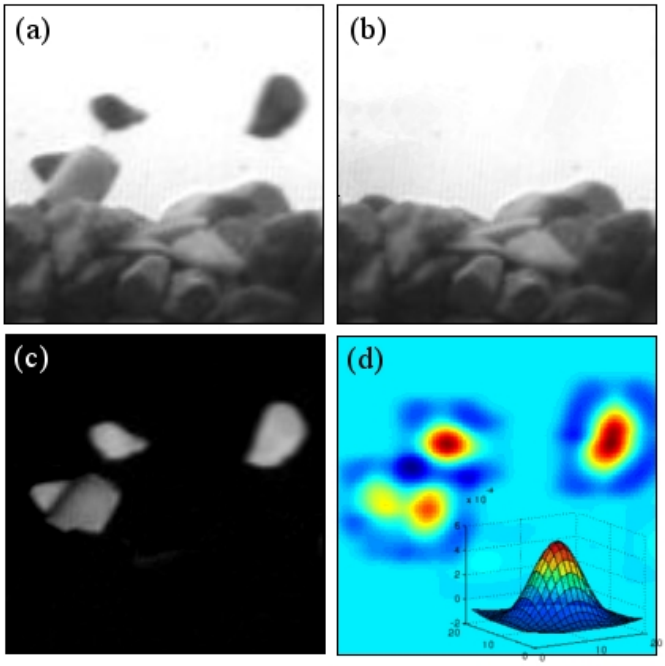

Figure 2: Example of the algorithm steps: (a) Original image (b) Background (c) Foreground (d) Convolution product.

"resting" particles can be a priori made when looking at an individual frame. Thus, we need to consider an algorithm to split the image into moving and still regions. This is what we call "Background subtraction" (or its complementary "Foreground selection"). There are (at least) 3 methods to achieve this.

The simplest way is to subtract pixel to pixel the current frame from the immediate previous one. The resulting image is thus formed of only the pixels values that change in time (particles that moved). Although being extremely fast, this basic method raises issues when particles travels less than their own diameter, resulting in "crescent moon" particle shapes. A better result is obtained by the technique of median filtering. A background image $\mathcal{B}$ is updated sequentially by comparing it to the current frame $\mathcal{I}$ according to the rules:

$$\begin{aligned} &\text{if} \quad \mathcal{B}_{ij} > \mathcal{I}_{ij} \text{ then } \mathcal{B}_{ij} \leftarrow \mathcal{B}_{ij} - 1, \\ &\text{if} \quad \mathcal{B}_{ij} < \mathcal{I}_{ij} \text{ then } \mathcal{B}_{ij} \leftarrow \mathcal{B}_{ij} + 1. \end{aligned} \tag{1}$$

These simple rules insure that the background image converge to the median image, that is an image where only the resting particles would appear. Furthermore, as the background of the scene evolves through time (because of erosion and deposition of particles), the median background automatically adapts itself to the changes. The foreground image is simply obtained by subtracting the median background to the current frame. The last possibility, more technical, is called mixture-Gaussian background subtraction. It consists in representing the value of each pixel as a superposition of n Gaussian distributions. At each frame, a pixel is set to foreground if the probability to observe its value fall beyond a certain threshold. The gain in performance of this method compared to median filtering is not obvious, so that we chose the latter for our purposes.

2.2.2 *Centroids detection*

Once the foreground image has been appropriately separated from the background image, positions of the particle centroids have to be determined. One expect the foreground picture to be noisy so that image filtering has often to be apply. Two options are described in the following.

The first one consists in thresholding the foreground image and successively apply morphological operations on pixel (erosion, dilatation, closing...) to remove isolated pixels. At the end, only groups of pixels delimiting moving particles should remain. By studying the properties of these regions, the center of mass, the area, the eccentricity and orientation of the particles can be obtained. This option assumes the choice of a threshold, that has to be retrospectively adjusted. The disadvantage of this method is that when two particles collide, the thresholding operation may result in a large unique region, precluding the detection of two particles. This drawback is partially solved by using the alternative method presented below.

The second option involves a convolution filter. Recall that the convolution product is a "comparator" of two functions such that:

$$(f \star g) = \int f(x) g(y - x) \mathrm{d}y. \tag{2}$$

In our case, f is the two-dimensional foreground image while g is a two-dimensional kernel that has approximately the shape of a particle. The Laplacian of Gaussian kernel is particularly suited for our purpose (Fig 2(d)). It is symmetric (rounded particles), it allows the definition of a mean (mean particle diameter) as well as a variance (variance in particle diameter). Particle centroids are obtained by finding local maxima in the convolution product. In other words, the maxima are located where the kernel fits particularly well the foreground picture, that is when a moving particle is present. The advantage of the Laplacian of Gaussian kernel (compared to a simple gaussian for

instance) lies in the presence of negative values on the edges of the filter, that result in a negative "halo" over the boundary of particles. Thus, when two particles collide, the halo prevents them to appear as a unique particle, so that two local maximums are still detected. Note that, unlike the thresholding method, none of the other properties (area, eccentricity...) can be obtain from the convolution product so that the best method would be a combination of both.

2.2.3 *Trajectories reconstruction*

Once a list of particles centroids through space and time is obtained, the last task is to link them between two frames to construct a unique trajectory. The issues are numerous:
- Several particles have to be tracked at the same time.
- The number of particles to track is varying in time, due to erosion, deposition and migration of particles in and out the image.
- Two particles can collide or/and overlap for a few images, leading to only one centroid detection. Although this phenomena is mitigated by using the Laplacian of Gaussian convolution kernel, it is still likely to happen. This is even more probable when the images are taken from the side wall, as in our application.
- The observation of the true position of centroids is not guaranteed over all frame, as some measurement noise can interfere in the process.
- The motion of particles is chaotic due to repeated impacts and turbulent forces.

Several tools have to be used conjointly to reduce at a minimum the error in trajectory reconstruction. We present some of them in the following. The simplest way to reconstruct a trajectory is to link a centroid to its nearest neighbour in the next frame. However, when multiple trajectories are simultaneously tracked, this does not guaranty the best combination. One would prefer an optimisation algorithm (also called *Hungarian algorithm*) that minimize a global cost function, defined as the sum of the distance between present particle positions and their potential candidates in the next frame. This insure that the constructed trajectories are the best in average.

To tackle the issue of errors or failures in centroid detections, the use of a Kalman filter is opportune. The Kalman filter consists in two successive step: a prediction step where the new positions of particles are predicted according to some kinematic model, and a correction step where these predictions are compared to the observations. The kinematic model can be made as complex as wanted, but in the absence of precise knowledge about the forces involved, the simplest model is to assume a constant acceleration (constant forces acting on particles). That is, given the position x_i^t of the i-th particle at frame t, the predicted position at frame $t+1$ is given by:

$$\tilde{x}_i^{t+1} = x_i^t + v_i^t \Delta t + a_i^t \Delta t^2, \tag{3}$$

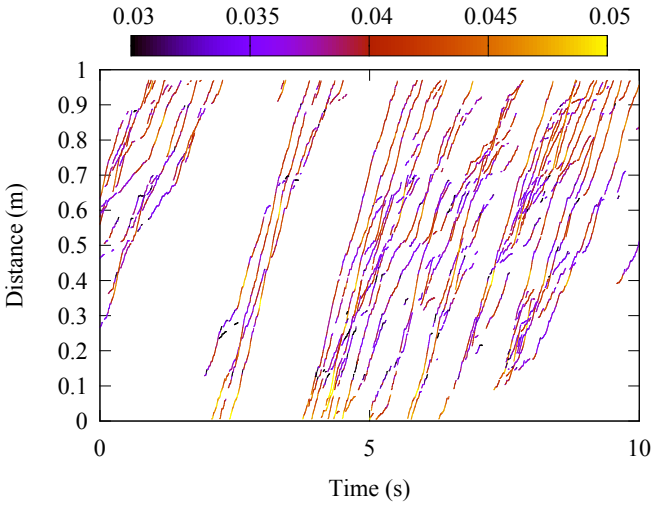

Figure 3: Particle trajectories in the $x-t$ plane. The color encodes the vertical component (in meter).

where v_i^t and a_i^t are particle velocity and acceleration respectively. We then compare the centroids detected in frame $t+1$ with the previous estimations, for instance by mean of the *Hungarian algorithm* mentioned above. Finally the corrected position x_i^{t+1} of particle i is obtained by a weighted combination of the observed and the predicted position. The relative weight is given as a function of the noise expected in measurements and the expected departure from the kinematic model. This prediction-correction algorithm allows the continuous tracking of particles. When a particle centroid is missing for a few frames, the kinematic model continues to estimate its position, and eventually reconnects it with a future observation.

More general methods use Bayesian statistics to predict and update the particle positions (Schikora et al. 2011). The kinematic model is now defined as a *prior* probability function. The ultimate position is then obtained by maximizing the *posterior* function (maximum likelihood).

Fig. 3 shows a sample of particles trajectories obtained by the algorithm. Each trajectory consists in a continuous line. Some trajectories appear to be broken in several pieces, when the Kalman filter couldn't predict accurately enough their positions. Note that, at this point, still no manual treatment has been involved.

2.3 *Velocity threshold*

Once particle trajectories are obtained, the instantaneous particle velocity can be computed by finite difference:

$$u_k(t_i) = \frac{x_i - x_{i-1}}{\Delta t}, \quad v_k(t_i) = \frac{y_i - y_{i-1}}{\Delta t}, \tag{4}$$

where Δt is the inverse of the frame rate of the camera.

In order to discriminate between moving and resting particles, a velocity threshold needs to be defined.

This threshold is *arbitrary* in the sense that a particle is never totally resting on the bed but its position slightly oscillates due to the turbulent flow drag. It is worth noting that the choice of the threshold is likely to influence the results. Indeed, by including the slowest particles into the group of moving particles, more importance to transport by rolling or sliding will be given. On the contrary, by ignoring them, saltation will be the dominant mode of transport. In the following, we consider that a particle is moving if its instantaneous velocity magnitude is larger than 0.3 cm/s, which is approximately 1% of the bulk average particle velocity.

2.4 *Validation*

A detailed validation of the detection algorithm is difficult. A visual check of each single trajectories is the only way to insure the accuracy of our algorithm. Unfortunately, it is infeasible when a large number of trajectories is involved. However, visual check on a few video samples can help to get an idea of the obtained accuracy. To that end, graphic interfaces like the MtrackJ plug-in for the open-source imageJ software are convenient.

As we could measure the bed load flux at the outlet of the flume at the same time as acquiring images, we can also compare the resulting bed load flux obtained with both techniques. The instantaneous bed load flux defined as a volume average reads:

$$q_s(t) = \frac{1}{\Delta x}\sum_{i=1}^{N} v_i \text{ [Particles/s]} \tag{5}$$

where Δx is the length of the observation window (Ancey and Heyman 2014). The cumulative transported load between two times is thus:

$$Q = \int_{t_1}^{t_2} q_s(t)\mathrm{d}t, \tag{6}$$

and can be directly compared to the cumulative flux obtained at the outlet of the flume (Fig. 4). Note that we do not expect a perfect agreement between the curves because the observation window was not directly touching the outlet of the flume, but was ending approximately 10 cm before. Thus, the 10% discrepancy observed in Fig. 4 between the two cumulative fluxes at the end of a video sequence is likely to be caused by spatial variability in the bed load transport fluxes rather than an under estimation from the tracking algorithm. Despite this, the instantaneous fluxes of both methods compare well on Fig 4.

3 RESULTS

3.1 *Flux*

The Exner equation, or mass conservation equation, involves the spatial derivative of the solid flux. Thus,

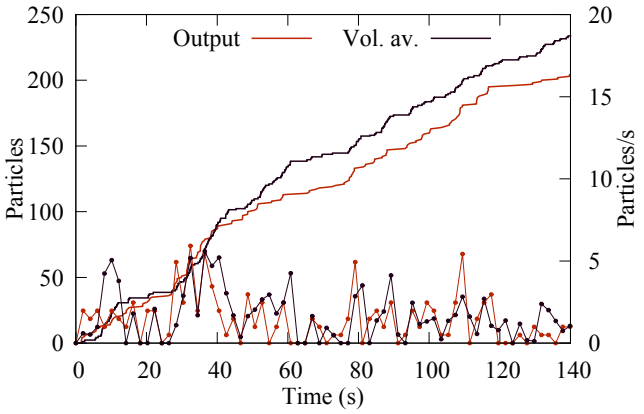

Figure 4: Comparison between tracking and acoustic techniques to measure bed load transport rates: cumulative flux (line) and instantaneous flux (line with points).

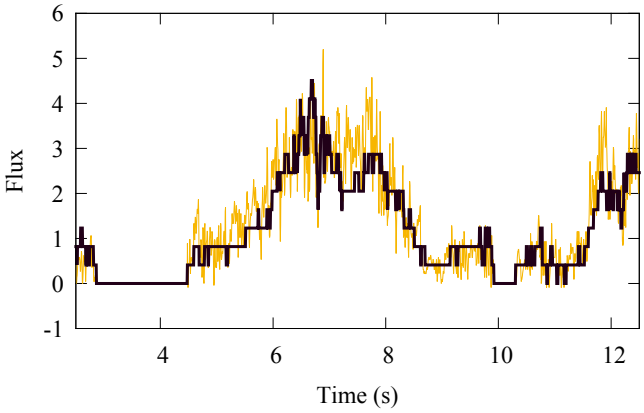

Figure 5: Solid flux in [particles/s]. Thin yellow line: q_s. Thick black line: $n\bar{v}/\Delta x$

its knowledge is of great importance for any morphodynamic model. The volume average instantaneous solid flux in an observation window of length Δx was defined in equation (5). It is tempting to decompose the bed load flux as:

$$q_s = \frac{1}{\Delta x} nv \tag{7}$$

where n and v are two random variables for the number of moving particles in the windows and their velocities. In a similar manner as a Reynolds decomposition in turbulence, we define $n = \bar{n} + n'$ and $v = \bar{v} + v'$ so that

$$q_s = \bar{n}\bar{v} + n'\bar{v} + v'\bar{n} + v'n'. \tag{8}$$

Fig. 5 shows both the variation of q_s and of n during a short time period. We can see that most of the fluctuations of q_s originate from the fluctuations of n. Indeed, in practice, $v' \ll (n', n)$, so that to first order, $q_s \sim \bar{v}\,(\bar{n} + n') + \mathcal{O}(v')$.

3.2 *Velocity distribution*

The velocity distribution of particles is shown in Fig. 6. It is fairly well described by a Gaussian probability law, of mean $\bar{v} = 0.34$ m/s and standard deviation $\sigma_v = 0.26$ m/s. This is in agreement with the

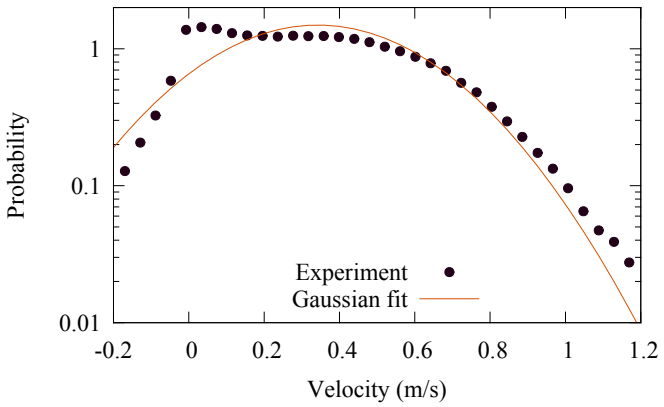

Figure 6: Probability density function of the streamwise velocity component of moving particles.

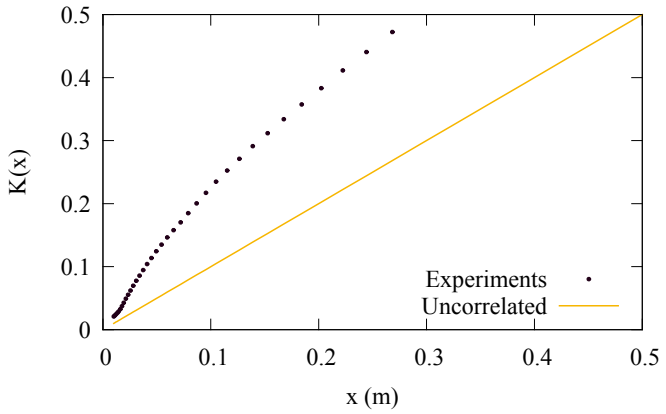

Figure 8: K-function for the experimental particle positions. The uncorrelated reference is shown.

model proposed in (Ancey and Heyman 2014). Other models suggest an exponential distribution for the particle velocities (Roseberry et al. 2012, Furbish and Schmeeckle 2013). Apart from the negative velocities and the high positive velocities, the experimental probability density function does not show an evident exponential shape (a line in semi-log plot).

3.3 *Diffusivity*

Unlike the movement of Brownian particles that undergo uncorrelated motion, the velocities of particles shows some non-vanishing correlation over time. Moreover, as particles start from rest and eventually return to rest after some time, the velocity of a single particle exhibits periodicity. According to (Furbish et al. 2012), the effective diffusivity can be obtained by calculating the variance of the particle velocity as well as the integral of its auto-correlation. However, because of the periodicity, the integral does not grow monotonically towards a constant value over long time scales.

An other method to determine the diffusivity makes use of the fact that the mean squared displacement should grow linearly with time $\langle X^2 \rangle \propto 2Dt$ (Taylor 1922). Finding D is thus equivalent to fitting the particle's mean squared displacement through time with a linear regression curve. For short time scales, the mean squared displacement shows a strong t^2 dependence, confirming the super-diffusive behavior due to particle velocity correlations. For scales greater than 0.5s, the mean squared displacement depends linearly on time, such that the diffusivity coefficient can be computed.

3.4 *Clustering*

The last interesting feature that can be highlight by our algorithm concerns spatial patterns in particles positions. It has been shown by several authors that, close to the threshold of motion, the transport of particle occurs intermittently. Indeed, (Drake et al. 1988) observed that "bursts" in particle activity were happening less than 9% of the time but were concentrating not less than 90% of the total load. He pointed that moving clusters of particles were responsible of those bursts in solid discharge.

Mathematically, this translates into correlations in particle positions. When moving particles are randomly distributed in the space, no correlation should exist between them. However, when particle clusters are present, this indicates a non-zero correlation between particles. The correlation originates from different processes, such that turbulent coherent structures or collective entrainment (Heyman et al. 2013).

A simple way to quantify this apparent correlation is the $K-$function (Ripley 1976), where $K(x)$ represents the expected number of moving particles found in a ball of radius x centred at a particle centroid divided by the mean process rate. If the process is uncorrelated, we expect $K(x)$ to grow linearly with x. If the process is positively correlated, $K(x)$ grows quicker than linearly (Fig. 8).

As expected, the experimental K-function shows a positive correlation for spatial scales smaller than 30 cm. Its behavior seems to become linear at longer scales, proving that long range correlations are not involved (e.g. the integral of the spatial correlation function exists).

4 CONCLUSIONS

In this paper, we presented an automatic algorithm to track the motion of bed load particles in water. Motion tracking is a vast research area and various methods have been proposed. Their performances depend strongly on the application, such that methods performing well for a soccer ball wont automatically apply to bed load transport. The latter presents several inherent difficulties (see 2.2), such as variations in the number of particles, particle occlusions... Some of them can be avoided by filming the transport process from above, if the free surface has no influence on the flow dynamics. For extremely shallow water flow, such as steep slope flows, filming from the side wall is unavoidable.

We showed that Kalman filters, and more generally all Bayesian filters, are or great help to deal with im-

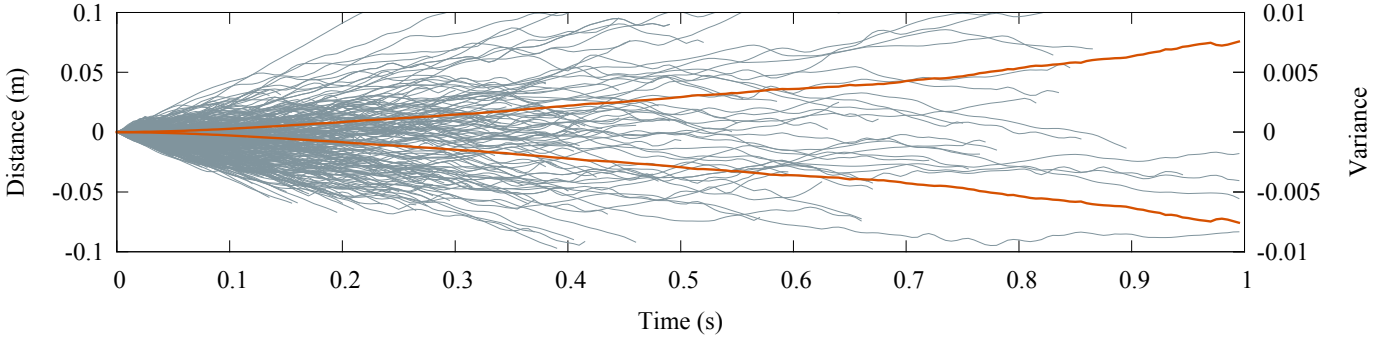

Figure 7: Dispersion of particles. Deviation from the mean particle displacement for a sample of trajectories (grey). Variance of the mean particle displacement (red).

age noise and measurements error. An *a priori* knowledge of some fundamental kinematics of the particle motion, like the constant acceleration model mentioned, allows a better reconstruction of particle trajectories. However, it is worst warning that those filters can also modify the apparent nature of particle motion (for instance, by smoothing strong variations) and thus, their impact on the statistical results have to be carefully assessed.

As an example, we presented an experiment carried out in a steep narrow flume. The whole observation window covered about 1 m. To the authors knowledge, this is the largest motion-picture experiment existing. Particle flux calculated from particle trajectories compared well with the flux measured at the outlet of the flume. This validates partially the tracking algorithm. The major limitation arise from the fact that particle trajectories are often broken in several pieces, so that estimation of the rate of erosion and deposition is difficult. Still, a lot of characteristics of the motion of particles can be extracted by our method such as particle flux, particle diffusivity, particle clustering...

Improving the algorithm is certainly possible, and even desirable. The use of advanced methods, such as optical flow or more complex Bayesian filters (Schikora et al. 2011) have to be considered in the future to overcome the present limitations.

REFERENCES

Ancey, C. (2010, April). Stochastic modeling in sediment dynamics: Exner equation for planar bed incipient bed load transport conditions. *J. Geophys. Res. 115*, F00A11.

Ancey, C., T. Böhm, M. Jodeau, & P. Frey (2006, Jul.). Statistical description of sediment transport experiments. *Phys. Rev. E 74*(1), 011302.

Ancey, C., A. C. Davison, T. Böhm, M. Jodeau, & P. Frey (2008). Entrainment and motion of coarse particles in a shallow water stream down a steep slope. *J. Fluid Mech. 595*, 83–114.

Ancey, C. & J. Heyman (2014, 4). A microstructural approach to bed load transport: mean behaviour and fluctuations of particle transport rates. *J. Fluid Mech. 744*, 129–168.

Böhm, T., C. Ancey, P. Frey, J. Reboud, & C. Ducottet (2004, Jun.). Fluctuations of the solid discharge of gravity-driven particle flows in a turbulent stream. *Phys. Rev. E 69*(6), 061307.

Bunte, K. & S. R. Abt (2005). Effect of sampling time on measured gravel bed load transport rates in a coarse-bedded stream. *Water Resour. Res. 41*(11), W11405.

Carey, W. P. (1985). Variability in measured bedload-transport rates. *J. Am. Water. Resour. As. 21*(1), 39–48.

Chiari, M. & D. Rickenmann (2010). Back-calculation of bedload transport in steep channels with a numerical model. *Earth Surf. Processes*, n/a–n/a.

Dinehart, R. L. (1992). Evolution of coarse gravel bed forms: Field measurements at flood stage. *Water Resour. Res. 28*(10), 2667–2689.

Drake, T. G., R. L. Shreve, W. E. Dietrich, P. J. Whiting, & L. B. Leopold (1988). Bedload transport of fine gravel observed by motion-picture photography. *J. Fluid Mech. 192*, 193–217.

Furbish, D., A. Ball, & M. Schmeeckle (2012). A probabilistic description of the bed load sediment flux: 4. fickian diffusion at low transport rates. *J. Geophys. Res.-earth 117*(F3), F03034.

Furbish, D., P. Haff, J. Roseberry, & M. Schmeeckle (2012, September). A probabilistic description of the bed load sediment flux: 1. theory. *J. Geophys. Res.-earth 117*(F3), F03031.

Furbish, D. & M. Schmeeckle (2013). A probabilistic derivation of the exponential-like distribution of bed load particle velocities. *Water Resour. Res. 49*(3), 1537–1551.

Heyman, J., H. Ma, F. Mettra, & C. Ancey (2013). A microstructural approach to bedload transport: spatio-temporal fluctuations of the particle activity. *J. Geophys. Res. submitted*.

Heyman, J., F. Mettra, H. B. Ma, & C. Ancey (2013). Statistics of bedload transport over steep slopes: Separation of time scales and collective motion. *Geophys. Res. Lett. 40(1)*, 128133.

Lajeunesse, E., L. Malverti, & F. Charru (2010). Bed load transport in turbulent flow at the grain scale: Experiments and modeling. *J. Geophys. Res. 115*(F4), F04001.

Martin, R., D. J. Jerolmack, & R. Schumer (2012). The physical basis for anomalous diffusion in bed load transport. *J. Geophys. Res.-earth 117*(F1), F01018.

Nelson, P. A., W. E. Dietrich, & J. G. Venditti (2010). Bed topography and the development of forced bed surface patches. *J. Geophys. Res.-earth 115*(F4), n/a–n/a.

Parker, G., C. Paola, & S. Leclair (2000). Probabilistic exner sediment continuity equation for mixtures with no active layer. *J. Hydraul. Eng. 126*(11), 818–826.

Recking, A. (2013). An analysis of nonlinearity effects on bed load transport prediction. *J. Geophys. Res.-earth 118*(3), 1264–1281.

Ripley, B. (1976). The second order analysis of stationary point processes. *J. Appl. Probab. 13*, 255–266.

Roseberry, J., M. Schmeeckle, & D. Furbish (2012). A probabilistic description of the bed load sediment flux: 2. particle activity and motions. *J. Geophys. Res.-earth 117*(F3), F03032.

Schikora, M., W. Koch, & D. Cremers (2011). Multi-object tracking via high accuracy optical flowand finite set statistics. In *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, pp. 1409–1412.

Singh, A., K. Fienberg, D. J. Jerolmack, J. Marr, & E. Foufoula-Georgiou (2009). Experimental evidence for statistical scaling and intermittency in sediment transport rates. *J. Geophys. Res. 114*(F1), F01025.

Taylor, G. I. (1922). Diffusion by continuous movements. *Proc. London Math. Soc. 20*, 196212.

Turowski, J. M. (2010, August). Probability distributions of bed load transport rates: A new derivation and comparison with field data. *Water Resour. Res. 46*(8), W08501.

Yilmaz, A., O. Javed, & M. Shah (2006, December). Object tracking: A survey. *ACM Comput. Surv. 38*(4).