

Rachel Cassidy · Philip J. Morrow · John McCloskey

A machine vision system for quantifying velocity fields in complex rock models

Received: 10 May 2005 / Accepted: 12 September 2005 / Published online: 9 December 2005
© Springer-Verlag 2005

Abstract In this paper we describe a machine vision system capable of high-resolution measurement of fluid velocity fields in complex 2D models of rock, providing essential data for the validation of the numerical models which are widely applied in the oil and petroleum industries. Digital models, incorporating the properties of real rock, are first generated, then physically replicated as layers of resin or aluminium (200 mm × 200 mm) encapsulated between transparent plates as a flowcell. This configuration enables the geometry to be permeated with fluid and fluid motion visualised using particle image velocimetry. Fluid velocity fields are then computed using well-tested cross-correlation techniques.

Keywords Fluid velocity fields · PIV · Cross-correlation · Porous fractured rock · Model validation

1 Introduction

Imaging-based approaches for the unobtrusive quantification of fluid velocity have grown in popularity in recent decades. Particle image velocimetry (PIV) (reviewed in [1]) emerged in the early eighties and has since evolved and been deployed in a range of challenging fluid problems. The basis of the technique rests in seeding a fluid with reflective, neutrally buoyant particles so that fluid motion can be visualised. Using a camera and illumination system to record

the displacement of the moving particles in temporally separated images, the 2D velocity of the fluid is determined. Included among the many applications of this technique are the measurement of accelerations in water waves in a laboratory wave tank [2]; wind tunnel testing of racing car aerodynamics [3] and the measurement of velocity fields in microfluidics devices [4]. Common to all these applications are closely controlled laboratory conditions where the experimental apparatus is designed to maximise the visibility of the fluid and the camera and lighting systems are integrated into the system. Field-based applications are less common and the potential for geophysical fluid applications, where flow is located deep in the sub-surface and access is restricted to boreholes, is particularly and understandably limited. In this applied project, however, which focuses on the validation of numerical models of flow in rock, it has been possible to develop a laboratory-based system that investigates flow in materials designed to incorporate the properties of real rock in a way that has not previously been possible.

The motivation for this research rests in the use of numerical modelling to investigate different scenarios and test different hypotheses in the earth sciences and, more specifically, in the characterisation of inaccessible sub-surface aquifers and oil reservoirs. Numerical modelling is commonly applied to such tasks and suites of models have been developed dealing specifically with the movement of geofluids. Originally models tended to be based on a continuum approximation and reliant on defining some average properties of the material but more recently there has been a growth in the application of discrete modelling techniques and in particular the Lattice Gas [5] and Lattice Boltzmann [6] schemes, variations of which are now widely available in both 2D and 3D. These models are considered more appropriate for predicting flow in realistic rock geometries, as field and laboratory measurements indicate that many properties of the material are inherently complex and scale invariant (reviewed in [7]), and thus lack the characteristic length scale which is requisite to the continuum approximation.

R. Cassidy (✉)
Geophysics Research Group, School of Environmental Sciences,
University of Ulster, Coleraine BT52 1SA, N. Ireland
Tel: 0044 2870324041, Fax: 0044 2870324911
E-mail: ri.cassidy@ulster.ac.uk

P. J. Morrow
School of Computing and Information Engineering, University of
Ulster, Coleraine BT52 1SA, N. Ireland

J. McCloskey
Geophysics Research Group, School of Environmental Sciences,
University of Ulster, Coleraine BT52 1SA, N. Ireland

For these and other models the quality of the predicted velocity fields is wholly dependent on the realism of the numerical model, which must be established through thorough validation. At present however, due to the difficulty in obtaining reliable measurements of flow in well-defined rock geometries, modellers can only assume that, if validated for simple geometries, their models are equally accurate in more complex geometries, even though the boundary conditions required for these models are very much more complicated.

Motivated by the inadequacy of existing data for the validation of numerical models, this research uses a machine vision-based measurement system to fully determine 2D velocity fields in precisely defined complex, physical geometries. Following the definition of the numerical geometries as binary image maps and their translation to a physical format, the imaging system is used to analyse the velocity field in a fully automated process requiring a precision lighting and image capture system, an automated positioning system and image analysis software. A stationary flow field is imperative and is maintained by a fluid system that guarantees a constant, controlled flow through each model for the duration of each experimental run.

In this paper we will specifically concentrate on describing the machine vision aspects of our work; further details on the geophysical aspects of the research and more detailed discussions on the results can be found in [8]. We describe the system as follows. In Sect. 2, we first briefly outline the definition of numerical geometries and their translation into precise physical copies, encapsulated in a Hele–Shaw cell and able to be permeated with fluid. In Sect. 3, the experimental hardware is described, comprising a fluid system, an imaging system and a stepper-motor controlled drive system. Section 4 considers the software which automates the entire analysis process, controls the hardware and measures fluid velocity. In Sect. 5, we give a number of results and finally, in Sect. 6 some conclusions are presented.

2 Rock models

The first step towards generating a physical model for analysis is the definition of a 2D binary image map, a boolean image of solid and void space, incorporating the features of a rock geometry requisite to a particular investigation. These investigations may focus on any structure of geophysical interest; from the degree to which fluid flow is controlled by the roughness of fracture surfaces, the flow interaction between a tightly packed rock matrix and open fracture channels or the complexity of flow through a complex fracture network, where connectivity produces preferential pathways for flow.

A software package has been produced to control the construction of the various 2D numerical geometries, integrating the different numerical algorithms that are used to produce the features required for each investigation. Rock matrix, for example, is represented by distributions of discs, compressed using a Discrete Element Method [9]

into a set area, until a specified porosity is attained. Fracture networks are represented by distributions of line segments with fractal or lognormal distributions of length (e.g. [10]) and various statistical controls on connectivity and aperture. Single fractures can also be defined with a roughness controlled by self-affine fractal statistics, and corresponding with measurements of fracture roughness from field and laboratory studies (e.g. [11–13]). Additionally there is provision for the conversion of actual images of the pore space from rock thin sections to a binary image map which can then be translated to a physical model. The output image file (holding the binary map) serves both as the template for the production of the physical model using computerised machining and rapid prototyping techniques and as input for simulation using a numerical model. On inspection the binary image maps (two of which are shown in Fig. 1) may seem a crude approximation of actual rock. This is necessarily so, as they are constrained by what is resolvable in the production of the physical geometry (for the current dimensions of the model, 200 mm × 200 mm, we cannot exceed three orders of magnitude in detail), yet within these limits the algorithms used in their generation incorporate the complex, fractal properties measured in field and laboratory studies, resulting in complex flow patterns which are ideal to thoroughly assess the performance of numerical models.

The physical model is based on a Hele–Shaw cell, and comprises a pair of rigid, transparent plates (200 mm × 200 mm) separated by a layer of material identical to the binary image map, ensuring that fluid motion throughout the void space of the material can be quantified using imaging techniques. The solid layer, identical to the numerical model, is generated using either stereolithography (SL) or wire electric discharge machining (EDM), in resin and aluminium, respectively. The thickness of the layer was optimised at 2 mm following preliminary testing which showed this to be the minimum aperture which allowed saturation without trapping air pockets in the interstices.

Complex geometries consisting of many small elements, that could not be arranged manually between the plates, are produced using SL in which a laser is used to reproduce the binary image map of the geometry, layer by layer, in photo-sensitive monomer on a glass base plate according to a 3D representation of the object in a CAD system [14]. Such a model, a porous fractured rock with a fracture network, is shown in Fig. 1(a).

To generate the physical model, a glass base-plate (200 mm × 200 mm) is first placed on a platform and submerged in liquid resin to a depth of 0.1 mm (Fig. 2(a)). This provides a base on which the medium is built and subsequently forms one side of the flowcell. Following the binary image map of the medium a laser then moves across the plate, selectively curing sections corresponding to the solid areas of the model. Once a layer is complete, the model is again submerged in the resin to coat the previous layer and the process is repeated, continuing until a 2 mm layer has been generated. The model is then cleaned thoroughly

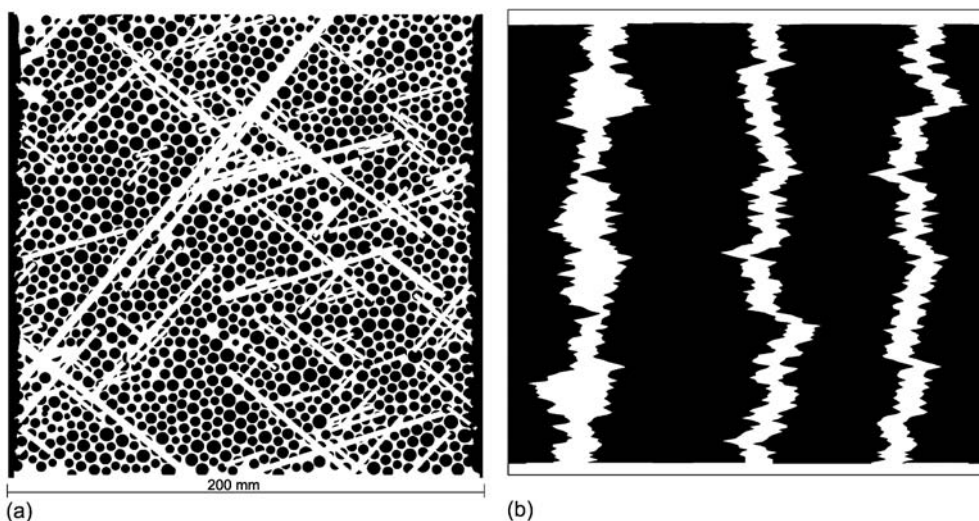


Fig. 1 2D digital models (200 mm × 200 mm) are generated for production of the physical model and as templates for simulation using the numerical models. Two examples are shown: **a** An idealised fracture set superposed on a rock matrix, comprising many small elements and suited to production using stereolithography and **b** A model of fracture roughness consisting of fewer elements and produced using wire EDM

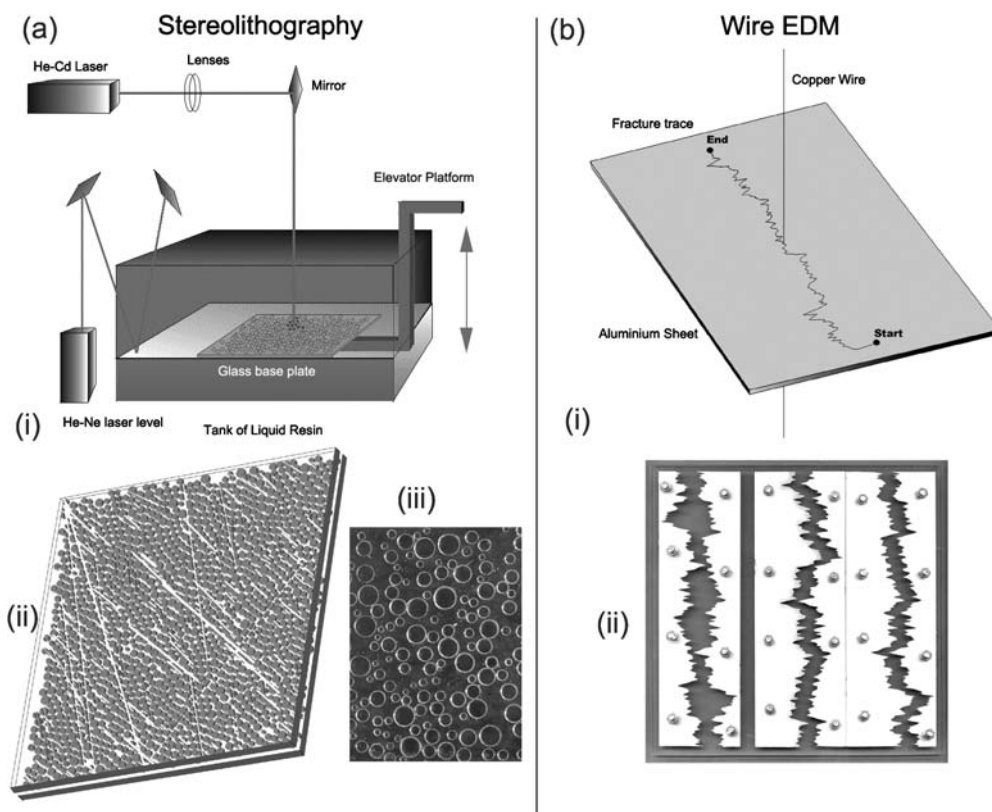


Fig. 2 Production of the physical models using stereolithography (SL) and Wire EDM. **a** Models composed of many small elements are produced using SL (i), where a laser selectively hardens photosensitive resin following the digital image of the model. (ii) 3D drawing of a completed physical model (iii) enlarged section (30 mm × 40 mm) showing the resin elements. **b** Models composed of fewer elements are produced in aluminium sheet by wire EDM (i). Photograph of a completed model (ii)




Fig. 3 The experimental rig. The numbers refer to the components listed as follows: (1) flowcell unit, (2) horizontal and vertical trolleys, (3) camera, (4) lighting (hidden), (5) fluid lines, (6) HPLC (high-pressure-liquid-chromatography) pump, (7) reservoir and (8) pressurised eluent flask (water and flow marker particles). All electronic components are connected to a computer located to the right of the rig

and a top plate is bonded in place by applying epoxy resin to the surface of the resin layer using a silk-screening procedure. The model is then complete and is left to cure for 48 h at room temperature ($\sim 16^{\circ}\text{C}$) before use.

Simpler geometries with fewer elements are produced by wire EDM, which uses spark erosion to cut defined traces in aluminium sheet. These are then assembled by hand between Perspex sheets to match the digital model. A model produced using this technique is shown in Fig. 2(b), consisting of three fracture apertures with self-affine roughness. Each trace is produced as a self-affine trace [15], normalised to the scale of interest and a 300 mm section extracted and output as a script file to AutoCAD where the model is constructed as a 2D image and prepared for cutting in 2 mm aluminium using wire EDM. This technique uses a copper wire electrode to generate an electric discharge and erode a trace (0.1 mm wide) through 2 mm aluminium sheet. The resolution of the resulting model is on the order of 0.1 mm ($\sim 2\%$ of the minimum amplitude). When complete, the aluminium sheet is removed from the machine, trimmed to size and anodised to prevent the aluminium oxidising when saturated with fluid. The fracture sections are then bolted in place between transparent plates to match the binary image map of the model.

3 Experimental hardware

The experimental hardware (shown in Fig. 3) comprises three key components:

1. The fluid system
2. Imaging and illumination
3. Automated positioning of the flowcell

The components of the system are housed on a steel frame designed to minimise vibrations. The flowcell is attached to a pair of moveable trolleys on the central support of the rig that allow the position of the flowcell to be adjusted by controlled stepper motors. The trolleys move the flowcell relative to the camera and lighting which are mounted on stands to the front and rear of the rig and ensure the systematic measurement of fluid velocity in small sub-areas to produce a full, high-resolution map of velocity throughout each model. With the exception of the fluid system all hardware is interfaced with the measurement software as a complete application that manages the measurement of the velocity field in each model.

3.1 The fluid system

Fluid is circulated through the flowcell at a constant volumetric flow rate for the duration of each analysis. Foam-lined clamps, fitted to a 10 mm Perspex plate, compress against the edges of the flowcell and seal it. The clamps at the top and base of the unit are fitted with connectors to the inlet and outlet pipes and all contacts are sealed with a silicone-based sealant. The clamped unit bolts vertically to trolleys at the front of the rig. Fluid is pumped from the base




Fig. 4 Side-view of the imaging system showing the cross-focussed illumination and the camera co-focussed on an 8 mm sub-area of the flowcell

of the system through the flowcell, allowing any gas bubbles to rise with the fluid and leave the system. Maintaining a stationary flow field throughout the flowcell for the duration of each analysis is imperative. This is achieved with a pumped fluid system comprising a high pressure liquid chromatography (HPLC) pump, a pressurised storage container, fluid lines and a reservoir. The HPLC pump ensures a fully adjustable flow rate with zero pulsation which can, if required, be software controlled. Fluid is extracted from the reservoir and passed through the pump to a pressurised eluent bottle filled with a suspension of micro-particles in distilled water. These hollow borosilicate glass spheres (TSI model 10089), with a nominal mean diameter of 8–12 μm , are almost neutrally dense in water ($1.05\text{--}1.15\text{ g cc}^{-1}$) and contrast strongly when backlit in water. The particles are initially added to 100 ml water and ~ 1 ml of dilute detergent solution (to break surface tension) and stirred vigorously to produce a well-mixed suspension which is then gradually added to another 1.5 l of water in the pressurised eluent flask. The seeded fluid is driven from the pressurised bottle through an input line to a manifold at the base of the flowcell unit. The same arrangement at the top of the unit terminates at an outlet pipe to the reservoir where filters ensure the flow seed cannot be re-circulated through the pump. Particle density is monitored as the concentration is increased to ensure it is optimal for analysis ($\sim 30\text{ particles mm}^{-2}$).

3.2 Imaging and illumination

The experimental rig ensures the movement of the fluid through the flowcell is visible and quantifiable using a particle image velocimetry (PIV) technique. (An optical flow approach was considered to be too computationally expensive

for this particular application). The precise specifications of such PIV systems must be customised to the application, but in general, one of two standard approaches is adopted for image acquisition; either multiple exposures of an area within a single frame using a pulsed light source or a single exposure per image frame using a continuous light source and a precisely triggered camera. In this application, fluid velocities are low and laminar (the range in a standard medium is generally $0\text{--}15\text{ mm s}^{-1}$) making it feasible to employ a single frame capture approach with a camera acquiring up to 30 images of an $8\text{ mm} \times 8\text{ mm}$ sub-area of the flowcell per second. This approach obviates the directional ambiguity associated with a single frame approach as the direction is easily determined by registering the particle images according to the order in which they are captured.

Following a lengthy test period during which many different configurations were evaluated, the optimal set-up for measuring fluid velocities involves indirect cross-focussed backlighting of an 8 mm sub-area of the flowcell (see Fig. 4). This is imaged by a high-resolution camera focused on the plane central to the flowcell and the positions of the flow marker particles at intervals are recorded. The region of interest for measurements is a narrow ($\sim 50\text{ }\mu\text{m}$) plane at the midpoint of the 2 mm space between the plates of the flowcell. Under laminar flow conditions this midpoint represents the peak of the parabolic velocity profile which develops between the plates. While most PIV applications illuminate a two-dimensional plane with a light sheet, the configuration of the flowcell in this application precludes this and consequently, in common with [4], the entire volume is illuminated and the quasi-two-dimensional measurement region is determined by focusing the camera to minimise the depth of field. Particles outside this narrow region are out of focus,

appear larger and darker and can, therefore, be excluded from analysis using image processing techniques.

To maximise the contrast between the micro-particles and the fluid within each sub-area, the field of view is backlit obliquely by cross-focusing two uniform, cold light sources. Back illumination diffracts the incident light, and consequently each particle appears larger, occupying an area of 4–5 pixels, which is readily identifiable when imaged and, importantly, allows the use of smaller and neutrally dense marker particles. The lighting unit consists of a fan-cooled illumination source containing a 150 W 15v Tungsten Halogen lamp mounted in a dichroic mirror reflector, providing a near-uniform light source. A fibre-optic cable with a single input transmits light from the source unit and splits it into two intense outputs. This provides intense, cold illumination and the branched cable is easily positioned and fitted into the condenser system which is fixed at the rear of the system and centred on the optical axis of the camera. The lighting is focused to intersect and coincide with the field of view of the camera, illuminating the marker particles moving through that area.

The camera, an Adimec MX12P, is positioned at the front of the flowcell unit to focus on an 8 mm × 8 mm field of view along the central plane of the flowcell. A Computar M-55 telecentric lens with a 2X-extender lens sets the minimum field of view to 8 mm × 8 mm and ensures a constant, distortion-free magnification across the field of view. The camera charged-coupled device (CCD) array comprises 1024 × 1024 pixels, is 2/3 in. in diameter with a 1:1 aspect ratio and delivers up to 30 frames per second through a composite video output signal to the frame grabber memory. The minimum field of view of the camera determines the measurement resolution of the system, as the flow marker particles must be clearly visible within the 1024 × 1024 pixel image frame (a pixel resolution of 7.8125 μm).

Image acquisition is precisely controlled by an externally generated trigger input signal supplied from the controlling software which triggers the camera to acquire images at any required separation time up to a maximum of 30 frames per second. This ensures velocity measurements, made on the basis of comparing particle positions in temporally separate images, relate precisely to a velocity in mm s⁻¹ and ensures consistency among all image pairs. The triggering is synchronised with the high-speed variable-scan frame grabber (DT3162 Data Translation Inc), which connects to the PCI bus to allow rapid transfer of each monochrome image to CPU memory. The use of cross-correlation in the present system requires images to be acquired at precise separation times and written into labelled memory buffers in the order in which they are captured.

3.3 Automated positioning of the flowcell

Measurement of the flowcell is effected by subdivision into 8 mm sub-areas and moving the flowcell relative to the camera on a fixed reference system so that each sub-area lies within the field of view of the camera. The sub-areas are defined,

prior to analysis, as a grid on the binary image map, with an overlap of 1 mm. This overlap ensures the edges of each sub-area are adequately sampled when further sub-divided into 1 mm interrogation areas (IA) at 0.5 mm intervals. The IA are then compared with the solid/void map of the sub-area and those with >60% solid are excluded from the analysis, as IA with solid fractions above this cannot be measured accurately using cross-correlation. The co-ordinates of all sub-areas and the IA within these sub-areas are then stored in a look-up file and read into the control software at the start of each analysis along with a defined reference location.

The positioning system controls the analysis of the flowcell by systematically placing each sub-area within the field of view of the camera and engaging position correction software to ensure the field of view corresponds exactly to the required section of the flowcell. The flowcell unit is attached to a pair of trolleys resting on linear guideways and is driven horizontally and vertically by geared stepper motors. A digital I/O and counter timer controls positioning by generating pulsed commands to the stepper motors from within the software, each output pulse producing a shift of 9.26 μm in the position of either trolley.

To initiate the analysis, an 8 mm reference area is located in the binary image map of the medium and then co-located on the flowcell. The position of the flowcell is then adjusted to bring this area within the field of view of the camera. The co-ordinates of this area are then set as a reference and all movements to other sub-areas throughout the analysis are relative to this.

On each move to a new sub-area, the accuracy of the position relative to the binary image is determined. Each movement incurs errors due to hysteresis on the order of 0.05%, small enough to cause no significant loss in accuracy in a move to a single sub-area but becoming unacceptable over several movements. To address this, the software incorporates a routine to assess the goodness of fit between an image of the field of view after each movement of the flowcell and the corresponding section of the binary image.

An 8 mm × 8 mm section of the binary image, and an additional edge of 1 mm, showing the exact position of the sub-area, are extracted. An image of the field of view is taken, thresholded to produce a binary image and converted to the same resolution where an 8 mm section is represented by 82 × 82 pixels. This coarsened image (R^I) is then moved systematically across the section of the binary image map (R^M) and the sum (σ_n) of the product of the image and the section of the binary image on which it is overlain calculated for each shift n , as

$$\sigma_n = \sum_{i,j} R_{ij}^I R_{ij}^M$$

The position corresponding to the best fit maximises the sum (σ_n) of the product. The offset is then calculated and the flowcell repositioned accordingly. The position is then rechecked and the process continues until the error in position is less than 50 μm.




Fig. 5 Outline of the application structure

4 The measurement system—software

The preceding sections describe the hardware components of this machine vision-based measurement application, developed to optimise the acquisition of the images that are the basis of quantifying the velocity field in the physical models. Controlling the hardware components and the process of analysing the images to determine velocity is a single software application, the structure of which is shown in Fig. 5. The complete application consists of seven core modules. The main module is written in FORTRAN and controls the overall process from image capture to final analysis. The frame grabber is accessed through an Active X control and both the camera and Amplicon card (controlling the stepper motors) are accessed via Dynamic Linked Library modules (DLLs). The three bottom layer modules were supplied with the associated hardware. In order to link these modules to the main application separate ‘interface’ modules were written in C, since there was no simple mechanism to call each of the modules directly from FORTRAN. These higher level modules also provide a simplified and consistent interface to the libraries controlling the hardware. Note that the frame grabber module also makes calls to the camera module. This is necessary in order to allow the camera integration and separation times to be set and to ensure that the frame grabber and camera are synchronized during image capture (via triggering signals). There is also provision for controlling the fluid system via the HPLC pump but this has not been necessary in experiments to date.

Pseudo-code describing the structure of the complete application and the processes involved in measuring a velocity field is given in Fig. 6. We follow this with a description of the analysis of a single sub-area, concluding with the treatment of the output velocity data to remove erroneous vectors and a description of the measures we have taken to ensure the quality of the data.

4.1 Initialising the measurement run

Following initialisation and registration of a reference location in the flowcell, the stepper motor controls move the flowcell to the first sub-area for measurement. The accuracy of the position is checked against the binary image map of

```

// Structure of the measurement algorithm
Initialise flowcell position to a reference point
Initialise camera, frame grabber, automated positioning system
Read binary image map
// Analyse a 8mm sub-area
Do for each sub-area
  Move flowcell to position
  Check accuracy of position (adjust if necessary)
  While sub-area is incomplete
    Acquire an image pair with a set separation time
    Subtract images and threshold to isolate flow seeds
    // Analyse an 'interrogation area' (IA) within the sub-area (128x128 pixels)
    Do for each IA within the sub-area
      If the IA velocity measurement is non-convergent then
        Extract IA from each image
        Compute frequency domain cross-correlation
        Stack correlation fields with previous to improve signal-to-noise ratio
        If sufficient image pairs stacked then
          Find peak in correlation field
          Interpolate to sub-pixel scale
          Store location of peak as direction vector
        EndIf
      EndIf
    EndDo
  // Assess vector quality
  Do for each IA within the sub-area
    If vector significantly different from neighbours
      Tag as poor quality
      Mark IA as non-convergent
    Else If vector too small
      Tag as slow area and mark IA as non-convergent
    Else
      Store vector and mark IA as convergent
    EndIf
  EndDo
  // Check quality for sub-area
  If more poor quality IAs than slow IAs then
    Sub-area incomplete – continue sampling (stacking up to 10 image pairs)
  Else If more slow IAs than poor quality IAs then
    Store poor quality IA tagged accordingly
    Scale up image separation time
    Sub-area incomplete – continue sampling
  Else
    Mark sub-area as complete
  End If
End While
End Do
Finalise system
  
```

Fig. 6 Pseudo-code describing the structure of the application

the area and following any re-positioning the analysis of the area commences.

4.2 Image acquisition

Cross-correlation operates on pairs of images separated by an interval adjusted according to the velocity of the fluid. A particle population captured in the first of a pair of images (captured at time t) and also present in the second image (captured at time $t + \Delta t$), when compared, allows the offset corresponding to particle displacement during the interval to be quantified.

The definition of the size of the interrogation areas and the separation time are crucial to the success of this approach and must be optimised to maximise the number of particles common to both images while ensuring a significant particle displacement between images. Large interrogation areas allow higher velocities to be measured by ensuring a greater percentage of the particles are common to both images but

at considerable cost in terms of resolution. In terms of the separation time used, we undertook tests on both real and synthetic data, which established that a separation time in which the population has moved in excess of 40% of the interrogation area increases measurement error significantly, as an increased portion of the particle population captured in the first field will have left the area, reducing the signal-to-noise ratio.

On balance, therefore, it was decided to restrict the interrogation areas (IA) to 1 mm^2 and to account for velocity variations by applying an adaptive separation time optimised to the velocities in each sub-area. The complexity of the models, where a high velocity fracture can be juxtaposed with a near-stationary pore throat, means that over the space of a few millimetres there can be a significant variation in velocity. This requires each IA to be dealt with separately in terms of separation times to ensure the displacement between images for any IA is within a range of 10–40 pixels in each direction.

This is implemented by first setting the HPLC pumping rate after saturation to ensure the highest velocity areas of the model can be measured at the maximum rate of 30 frames s^{-1} , which for 1 mm interrogation areas, is an upper limit on measurement of $\sim 15 \text{ mm s}^{-1}$. All interrogation areas within each sub-area are then initially imaged at this rate. The velocity vectors in each IA are then assessed to locate those with minimal displacement and these are then resampled at separation times scaled-up to ensure a significant displacement. In sub-areas exhibiting a broad range of velocities, up to six different separation times may be necessary to analyse the velocity field accurately. As there is no physical limit on the length of separation time that can be used, the system is able to measure very low flow rates. A separation time of 60 s will, for example, accurately measure velocities in the range $0.001\text{--}0.005 \text{ mm s}^{-1}$ although for most of the models the measured velocity extends over three orders of magnitude as it is necessary to balance the duration of the experimental run with the potential for deterioration in model quality from prolonged saturation.

4.3 Image processing

Subtraction and thresholding are first applied to each image pair to remove noise and isolate the flow marker particles. The pixels in each 1024^2 image array are allocated a greyscale value from 0 (black) to 255 (white). Under illumination, the flow marker particles appear white (high-greyscale values) while the fluid appears grey and the solid areas either white or black (dependent on whether a resin or aluminium medium is used).

Direct subtraction of an image pair isolates all altered pixels between the first and second images and eliminates stationary values. The moving marker particles are responsible for the main disparity between images, creating a strong signal by their presence at a location in the first image and their absence in the next. Those particles which are in focus (within the depth of field) generate the strongest signal.

Slight changes due to shifting ambient light and noise in the camera CCD also arise, but have generally a smaller amplitude, which enables their exclusion by careful optimisation of a threshold. The threshold, which is set in an empirical procedure at the beginning of each measurement run, assigns a unit value to all subtracted pixel values above a limit and a zero to those below. This is performed for both the positive and negative returns from subtraction, the former corresponding to the first image in the pair and the latter to the second image. Careful optimisation ensures that only particles within the depth of field, which have the highest greyscale levels, are included in the analysis and output as a pair of binary images containing the isolated flow marker particles corresponding to the original images.

4.4 Cross-correlation

Following subtraction and thresholding, individual IAs corresponding to a section of 128×128 pixels (Fig. 7) in each of the output binary image pair are analysed using a frequency domain cross-correlation algorithm which computes the correlation between images by comparing them both directly superposed, and with one offset relative to the other.

The correlation (described in detail in [16]) between two (1D) continuous functions $g(t)$ and $h(t)$ is defined by

$$\text{Corr}(g, h) \equiv \int_{-\infty}^{\infty} g(\tau + t)h(\tau) d\tau$$

where the correlation is a function of the lag, t . The correlation will be large at some value of t if the first function (g) is




Fig. 7 An 8 mm sub-area showing an interrogation area (128^2 pixels) in red. Resin elements and marker particles appear white and the fluid-filled areas dark grey

a close copy of the second (h) but lags it in time by t , i.e. if the first function is shifted to the right of the second to match the shift in particle positions between images. Similarly, the correlation will be large for a negative value of t if the first function is displaced to the left of the second.

It therefore lies in the time domain, and is one member of the transform pair:

$$\text{Corr}(g, h) \Leftrightarrow G(f) H^*(f)$$

where $G(f)$ and $H^*(f)$ are the Fourier transform and the complex conjugate of the Fourier transform, respectively, where g and h are real functions. Multiplying the Fourier transform of one function by the complex conjugate of the Fourier transform of the other gives the Fourier transform of their correlation. The inverse of this product (rk) is a complex vector of length N . As the original datasets were real, however, the imaginary parts are zero. The components of rk are the values of the correlation at different lags, with positive and negative lags stored in wrap-around order. The maximum value in the resultant 2D correlation function, represents the physical displacement between image sections producing the best fit. This displacement is directly related to fluid velocity through the separation time used. A single analysis is not, in general, enough to procure an accurate determination of velocity for all interrogation areas in any sub-area. The density of marker particles in an interrogation area may be too high or low, or light diffracted from the solid areas of the model may interfere with image quality resulting in a spurious peak. To counteract this, multiple measurements are taken by acquiring and processing multiple image pairs and the return from each cross-correlation stacked by repeated addition of the 2D correlation functions. After the correlation fields for five image pairs are stacked the vector field for the sub-area is assessed to determine:

- Areas requiring analysis at longer separation times to account for lower velocity regions and
- Areas requiring further analysis to improve the signal peak due, in general, to contamination of the field of view.

For the majority of areas five iterations are enough to produce a strong signal and these values are stored. For areas with a low signal to noise ratio (due, in general to some contamination of the physical model) and in areas with a broad variation in velocity, further analysis is required. While a displacement in excess of 10% of the IA is considered acceptable in terms of measurement resolution, areas with a displacement less than this are marked for analysis at extended separation times ensuring a significant shift in particle position between images. Additionally areas deviating significantly from adjacent values are tagged as potential outliers and further images are analysed and stacked to allow the signal peak to stabilise by dampening extraneous noise. If additional stacking fails to improve the measurement, the value is stored but tagged to denote the irregularity.

The accumulated peak for each interrogation area is interpolated to a sub-pixel scale to provide a higher resolution

estimate of the location of the cross-correlation maximum. Several peak fitting algorithms were assessed and for accuracy and computational efficiency a procedure which fits a series of 1D polynomials was finally adopted.

The computed x and y displacements, in pixels, d_p , are converted to velocity, v , (in mm s^{-1}) by the relation $v = \frac{1000}{\Delta t} p d_p$ where Δt is the separation time and p is the size of a pixel in millimetres (corresponding to 0.0078125 mm for an 8 mm field of view captured on a 1024² CCD). These are written to output files, each referenced to the centre of the interrogation area in real co-ordinates. The positioning system then re-engages to move to the next sub-area and repeats this process, building up a completed velocity field for the physical model.

Figure 8 shows the measured velocity field in an area of 14 mm \times 10 mm for flow converging at the top of a disc in a straight channel, showing the individual vectors plotted at the centre of each IA.

Each physical model takes ~ 48 h to complete and results in a data set which typically contains 80,000 measured velocities all referenced to precise locations and with attached meta-data describing the quality of each measurement. The need for computational efficiency in the analysis means that detailed examination of potentially irregular measurements during the analysis is not possible. Instead they are tagged and dealt with during post-measurement filtering, when potentially erroneous vectors are excluded from the data set.

4.5 Measurement quality and accuracy

Various controlling steps have been implemented to ensure the quality of the data sets produced using this technique. These include the development of a checking procedure which assesses the clarity of the sub-areas in each model prior to processing and extensive sensitivity testing of the measurement algorithms during development.

4.5.1 Image quality

Measurement accuracy is reliant on image quality and all models are affected to some degree by areas that cannot be measured due to poor contrast between the marker particles and the fluid. The factors which contribute to this include the following:

- *Impurities on the surface of the transparent plates.* The clarity of the plates, whether made of glass or Perspex, is crucial. Back illumination highlights any scratches, dust or residual sealant and the resultant glare obscures the flow marker particles when imaged.
- *Residual resin.* The variable quality of the models produced by the stereolithographers has been a problem. A recent change of resin and technique has resulted in hardened deposits of resin in the interstices between model sections which must be scraped or, in extreme cases, cut out.




Fig. 8 Measured sub-area of a model, approximately 14 mm \times 10 mm. Each vector corresponds to a single measurement in an interrogation area. Areas without vectors are where clarity affected velocity measurement resulting in tagged values which were excluded following analysis

- *Illumination and backscatter.* Diffraction around the edges of the resin particles can obscure areas of the field of view. Even where the plates and the interstices between resin sections are completely clean, the configuration of certain sections diffract the incident light to obscure the flow marker particles.

To ensure measurement quality all interrogation areas within each model are examined and assigned a degraded IA removal threshold (DIRT) value according to the fraction of each image that is contaminated. The basis of determining the DIRT is that opaque areas reflect the incident light in the same manner as the flow seed and thus appear white in the image (a greyscale maximum of 255). By assessing the fraction of the void space in each IA with a high greyscale value out of a worst-case where the entire area is obscured, problem-areas are excluded from analysis.

A 1024 \times 1024 pixel image of the entire 8 mm sub-area is first captured. The equivalent section of the binary image map is then extracted and re-scaled to allow direct comparison with the image. This section acts as a mask array, where the solid area has a null value and the void space a unit value, and is used to eliminate the solid part of the image from calculations by obtaining the product of the image and the mask. This sets all pixels in the solid fraction of the image to zero. For each defined interrogation area within this sub-area the clarity of the image is now assessed.

The total void area, the number of pixels in the void space (n_v) and the total light in this void area, the sum of all pixels ($\sum P$), are then computed. Where the entire area is obscured, and appears white, the light level would be $255.n_v$.

The opaque fraction, or DIRT, for the area is obtained by

$$\frac{\sum P}{255.n_v}$$

and is stored with each velocity measurement. During measurement all IA with $>75\%$ DIRT are ignored and, when evaluating the data set following analysis, the DIRT value is used as a guide to eliminating erroneous measurements.

4.5.2 Sensitivity of the measurement algorithms

In addition to determining image quality for the measurement of each model, we also examined the accuracy of the cross-correlation technique in determining the averaged velocity within a defined area. The algorithm was extensively tested using synthetic data during development to ensure for a given displacement of particles and different levels of noise that the returned vector was identical to that input, and careful measurements of noise tolerance were made and these used to limit the measurement process.

Once the system was fully operational a further test was undertaken to examine the robustness of the algorithm in the working experimental system. This involved comparing the automated measurements returned by the cross-correlation algorithm for different locations and velocities in the flow-cell with measurements taken by hand from a live video of flow in the same areas. For the manual measurements, a live video image of a field of view was activated. The flow rate at the HPLC pump was set and allowed to equilibrate and the movement of the marker particles observed in real time.

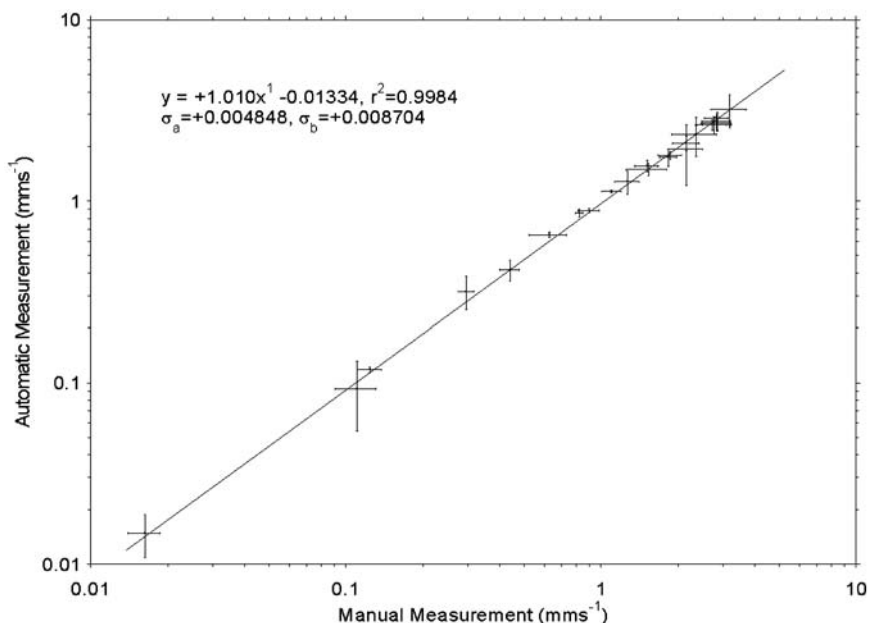


Fig. 9 Comparison of the automated measurements obtained from cross-correlation with those digitised manually in the image pairs for different velocities. Data are plotted on a log–log graph to show the lower velocities more clearly. Error bars are the standard deviation of the measurement in linear space

Lines were marked on the screen at 1 cm intervals normal to the direction of flow. A program was then activated which timed precisely the interval between successive key inputs. Marker particles were tracked and recorded as they crossed each line, storing the time taken for a particle to move the set distance between lines. Successive measurements, tracking multiple particles in an area ~ 1 mm wide, were taken and from the slope of a graph of time against distance, the velocity was determined. Simultaneously automated measurements were acquired using cross-correlation. The average and standard deviation of both sets of measurements were then calculated. This analysis was completed over a wide range of flow rates and the automated and measured velocities compared (see Fig. 9) Error bars on the log-log plot of the measured velocities are the standard deviations of the measurements in linear space. This exercise was the real test of the quality of the experimental system and demonstrates excellent agreement over almost three orders of magnitude, confirming the robustness and accuracy of the technique.

4.6 Filtering the output

The completed analysis produces an output data file containing on the order of 80,000 velocity measurements. Each measurement is referenced to the centre of the IA and has attributed to it information on the clarity of the IA and whether the vector was tagged as potentially erroneous. A fraction of incorrect measurements is inevitable in all data sets due to the automation of the measurement process and the impossibility of ensuring all IA are in optimum condition for measurement. The fraction is dependent on the quality and the clarity of the model.

While the estimation of a DIRT value for each IA prior to analysis is effective in excluding contaminated areas from analysis, it is not infallible and as such it is necessary at the conclusion of the measurement process to evaluate the output velocity field and filter erroneous vectors to maximise the quality of the output data sets.

Filtering initially involves identifying all tagged values in the data set and evaluating them through comparison with the computed mean and standard deviation of the x - and y -components of velocity for all neighbouring measurements. If the measurement exceeds one standard deviation of the mean for either component then it is excluded from the data set. This has been tested and is effective in eliminating the majority of outliers but fails where there is a cluster of incorrect measurements, as would occur when a large area of the field of view is obscured and in narrow interstices where there are few or no immediately adjacent measurements. After consideration of various methods and assessment of each, the most effective method for eliminating these values has proven to be by visual inspection. Remaining outliers are clearly identifiable in the vector displays by both colour and orientation and a program was produced to plot the vector field for a small area and to allow, using cursor control, erroneous values to be selected and removed by hand. This is not time consuming as the preceding assessment removed the majority and outliers are easily identified.

5 Experimental results

To date we have measured the velocity fields in a range of geometries. Each velocity field is produced by plotting the vectors, centred on each interrogation area and scaled




Fig. 10 **a** Map of velocity for a dual porosity resin model, shown in Fig. 1(a) and **b** an enlarged sub-area (red box on *left*) showing the vectors, with magnitude according to the colour scale

according to colour and length, on the digital image of the medium. There is no interpolation of the data, as this can imply measurements were taken in areas that could not be analysed. The white areas in each velocity field, therefore, correspond to areas for which there are no associated measurements. The absence of a measurement is indicative either of an area where a high DIRT fraction precludes measurement, where a measurement was taken and discarded during filtering or because the velocity in the area was too low for the range of separation times used, returning a value of zero.

Here, to maintain consistency and as it demonstrates the nature of the measured data as well as some limitations of the system, the velocity field for one of the models (Fig. 1(a)), produced using stereolithography is described briefly.

The completed velocity field for the model is shown in Fig. 10(a) and an enlarged section of the model showing the detail in a sub-area of the vector field is shown in Fig. 10(b). The measurement resolution with this technique allows the interaction between matrix and fracture flow to be clearly observed and recorded and such complex flow patterns provide a difficult test for numerical modelling schemes, testing the extent to which they are able to reproduce complex boundary conditions. The channelisation of the flow pattern is apparent as well as the intricacies of fluid interaction at pore junctions.

Direct comparison of these data with the predicted velocities at the same points using a numerical model allow a comprehensive and thorough investigation of the numerical scheme that has not to date been possible. The results of comparison between measurement and a modified lattice Boltzmann scheme can be found in [8].

6 Conclusions

Improvements in digital imaging technology, combined with falling camera and processor costs, have seen particle image velocimetry applied to problems outside of the engineering applications for which it was originally developed. This paper describes a geophysical application of this measurement technique to the investigation of fluid movement in geological materials, where a machine vision system has been developed to quantify fully velocity fields in 2D complex models of rock. The system, which integrates machine controls with measurement algorithms in a single software application to fully analyse a 2D rock model in a complete and automated process, provides, for the first time, a complete picture of flow in a complex geometry that can be directly applied to test and improve numerical models. The technique has been extensively tested and validated and we are confident in the accuracy of the output velocity fields.

The system was developed in response to the lack of detailed empirical data for fluid flow in realistically complex geometries, which limited the level at which model validation could be undertaken. Thorough assessment of existing numerical models is of fundamental importance if numerical modelling is to contribute to exploring possibilities for upscaling, particularly as a realistic representation of rock structure requires the incorporation of complexity. These data sets consist of up to 80,000 points across the void space of each medium, each with an associated velocity in mm s^{-1} . For any specified geometry fluid velocity can be simulated at an identical point in the numerical geometry and directly compared with measurement.

The system has much potential for further work in observing and measuring fluid flow. It is planned to combine, in future work, the particle image velocimetry approach with tracer experiments to explore the relationship between the material geometry, the velocity field and the movement of tracers in the sub-surface. The system is being retained as a research facility for interdisciplinary and inter-institutional use.

Acknowledgements This research was supported under the Micro-to-Macro Thematic Programme of the Natural Environment Research Council (NERC). We thank Harry Smyth and Terry Griffin of the University of Ulster Mechanical and Electronic Workshops for assistance in developing the experimental equipment. We would also like to thank the reviewers of this paper for their helpful comments.

References

- Grant, I.: Particle imaging velocimetry: A review. In: Proceedings of the Institution of Mechanical Engineers, vol. 211, no. C, pp. 55–76 (1997)
- Jensen, A., Svein, J.K., Grue, J., Richon, J.-B., Gray, C.: Accelerations in water waves by extended particle image velocimetry. *Exp. Fluids* **30**, 500–510 (2001)
- Eisele, K., Zhang, Z., Wildi, J., Müller, K.: The application of a particle tracking velocimetry system with a high speed video camera on racing cars. In: Proceedings of the 7th International Conference on Laser Anemometry, pp. 755–760 (1996)
- Santiago, J.G., Wereley, S.T., Meinhart, C.D., Beebe, D.J., Adrian, R.J.: A particle image velocimetry system for microfluidics. *Exp. Fluids* **25**, 316–319 (1998)
- Frisch, U., Hasslacher, B., Pomeau, Y.: Lattice-gas automata for the Navier-Stokes equation. *Phys. Rev. Lett.* **56**(14), 1505–1508 (1986)
- Benzi, R., Succi, S., Vergassola, M.: The Lattice Boltzmann Equation: Theory and Applications. *Phys. Rep. (Review section of Physics letters)* **222**(3), 145–197 (1992)
- Bonnet, E., Bour, O., Odling, N.E., Davy, P., Main, I., Cowie, P., Berkowitz, B.: Scaling of fracture systems in geological media. *Rev. Geophys.* **39**(3), 347–383 (2001)
- Cassidy, R., McCloskey, J., Morrow, P.J.: Fluid velocity fields in 2D heterogeneous porous media: Empirical measurement and validation of numerical prediction. In: Shaw, R.P. (ed.) *Understanding the Micro to Macro Behaviour of Rock-Fluid Systems*, vol. 249, pp. 115–130. Geological Society London Special Publications (2005)
- Cundall, P.A., Strack, O.D.L.: A discrete numerical model for granular assemblies. *Geotechnique* **29**, 47–65 (1979)
- Davy, P.: On the frequency-length distribution of the San Andreas fault system. *J. Geophys. Res.* **98**, 12141–12151 (1993)
- Brown, S.R.: Fluid flow through rock joints: The effect of surface roughness. *J. Geophys. Res.* **92**, 1337–1347 (1987)
- Power, W.L., Tullis, T.E., Brown, S.R., Boitnott, G.N., Scholz, C.H.: Roughness of natural fault surfaces. *Geophys. Res. Lett.* **14**(1), 29–32 (1987)
- Power, W.L., Durham, W.B.: Topography of natural and artificial fractures in granitic rocks: Implications for studies of rock friction and fluid migration. *Int. J. Rock Mech. Min. Sci.* **34**(6), 979–989 (1997)
- Bernhard, P., Hofmann, M., Schulthess, A., Steinmann, B.: Taking lithography to the third dimension. *Chimia* **48**, 427–430 (1994)
- Turcotte, D.L.: *Fractals and Chaos in Geology and Geophysics*, Cambridge University Press, Cambridge, pp. 73–94. (1993)
- Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T.: *Numerical Recipes in FORTRAN 77: The Art of Scientific Computing*, Cambridge University Press, Cambridge, pp. 992. (1992)



Rachel Cassidy is a Research Associate in Geophysics at the University of Ulster. Dr. Cassidy's research interests include percolation theory and its application to fluid flow in fractured rock, the fractal and multifractal properties of natural phenomena and the development of experimental techniques for investigating fluid flow in porous fractured media with realistic structure and exhibiting scale invariance. She is currently involved in the development of molecular tracer techniques for characterising reservoir heterogeneity.



Philip Morrow is currently a Senior Lecturer in the School of Computing and Information Engineering at the University of Ulster. Dr. Morrow has a BSc in Applied Mathematics and Computer Science, an MSc in Electronics and a PhD in Parallel Image Processing, all from the Queen's University of Belfast. His main research interests lie in image processing, computer vision and parallel/distributed computing. He has published over 65 research papers in these areas.



John McCloskey is Professor of Geophysics and Head of the School of Environmental Sciences at the University of Ulster. Prof. McCloskey's research interests are in the application of ideas of chaos and complexity to a variety of geophysical problems including earthquake dynamics and fluid flow in fractured porous rock. He has published over 100 articles and is a regular contributor to international press on matters connected with earth science.