Determining concentration and velocity profiles of non-Newtonian settling slurries using electrical resistance tomography

by A.P.N. Sutherland*, T.M. Long†, E.W. Randall‡, and A.J.Wilkinson†

Synopsis

Pipelines for transporting high concentration slurries are often designed using methods applicable to homogeneous non-Newtonian laminar flow. Many industrial slurries though comprise coarse particles in a non-Newtonian carrier fluid, and existing experimental evidence from tests with such slurries demonstrates that they are not homogeneous. Under laminar shear conditions the coarse particles settle, even in statically stable carrier fluids. To understand this better, the Institute of Materials Science and Technology (IMST) at Cape Peninsula University of Technology (CPUT) is conducting ongoing research into the flow of these high concentration non-Newtonian settling slurries. Since pressure gradient and flow rate measurements alone are insufficient to adequately model the flows, an electrical resistance tomography (ERT) instrument developed by the University of Cape Town (UCT) was incorporated into the pipe test loop to determine pipe cross-section concentration and velocity profiles. The software developed to do this is modular and allows different image reconstruction and cross-correlation algorithms to be implemented and tested without significantly changing the rest of the application. At a total frame capture rate of 566 frames/sec a 2.67 GHz Intel Celeron processor with 500 MB of RAM is fast enough to calculate and display velocity profiles in 'real-time' at an update rate of one profile every two seconds. Examples of experimental concentration and velocity profiles obtained using the system are shown and some proposed improvements/extensions to the system are listed.

Introduction

Environmental and economic considerations have required the mining industry to increasingly use high concentration co-disposal systems. Existing design methods, however, are not always appropriate or adequate and improved modelling of pipe flow of these materials is needed. The flow of these high concentration settling slurries is often laminar, and they appear to be and are taken as pseudo-homogeneous fluids. In reality they form stratified flows with the coarse solids transported as a sliding bed, even when the carrier fluid yield stress is high enough to support these coarse particles under static conditions (Pullum and Graham, 1999; Cooke, 2002). Although homogeneous and non-Newtonian stratified flows look similar in terms of pressure gradient versus flow velocity behaviour, flow regimes (patterns) inside the pipe are very different (Pullum and Graham, 1999), and homogeneous prediction techniques based on small pipe data can greatly underpredict pressure gradients on scale-up (Pullum and Graham, 2002). Flowrate and pressure gradient measurements alone provide insufficient data and insight to adequately model the flows. Under some flow conditions velocity gradients are known to exist through the sliding bed, so existing two- and three-layer prediction models are inadequate for these conditions. These stratified models for settling slurries need to incorporate the physics of the flow mechanisms correctly. Knowledge of parameters affecting the flow and a better understanding of the mechanisms governing the flow are needed (Pullum and Graham, 1999; Matousek, 2004). Determining concentration and velocity profiles across the pipe cross-section, under specific known conditions, will contribute significantly to this knowledge.

Electrical resistance tomography (ERT) is a non-intrusive technique that is able to provide information required for improving laminar flow modelling of settling slurries. Several researchers (for example Lucas *et al.*, 1999; Deng *et al.*, 2001; Henningsson *et al.*, 2006) have used ERT to investigate multiphase flows to determine not only concentration, but also velocity profiles using cross-correlation of signals from adjacent electrode rings. However, the range of flow velocities investigated by these groups was limited by the data capture rate of their tomography

* *Department of Civil Engineering, Cape Peninsula University of Technology*

† *Department of Electrical Engineering, University of Cape Town*

‡ *Department of Chemical Engineering, University of Cape Town*

instruments.

As part of on-going research into the laminar flow of non-Newtonian settling slurries, the UCT instrument (Wilkinson *et al.*, 2005; Randall *et al.*, 2008), which is capable of capturing data at up to 1 000 frames/second, is used in the Institute of Materials Science and Technology (IMST) pipe test loop. Due to its high speed it has good velocity discrimination for the flowrates of interest. Image reconstruction can be performed faster than the raw data are captured, enabling velocity profiling to be implemented to operate in 'real time'. The development of the system is described briefly and some typical concentration and velocity profile results for sliding bed flow in a kaolin/silica sand slurry are shown.

Equipment

The pipe test loop

Figure 1 shows the IMST pipe test loop, which comprises two class 12 uPVC test pipes on both the out and return legs (nominally Ø56 mm and Ø80 mm), a 1 500 litre mixing tank (18.5 kW) and a Warman 4/3 AH pump (55 kW) with a variable speed drive. Instrumentation comprises Fuji FCX-CII differential pressure transmitters, a Fuji MAG600 Ø80 mm electromagnetic flowmeter and a Ramsey γ-ray density gauge in the Ø80 mm vertical return pipe, thermocouples, paired conductivity probes to measure bed surface velocity via cross correlation and the ERT system. A glass viewing piece (Figure 2) is also included. The pressure, flow, density and temperature information is acquired via a 6014 National Instruments DAQ card.

ERT equipment

ERT electrode rings

Each pipe has a single ERT spool piece with three 16-electrode rings spaced at 50 mm then 100 mm, enabling dual plane data to be acquired at cross-correlation distances of 50, 100 or 150 mm. The vertical return pipe has a single ERT ring located above the density gauge, used to monitor the homogeneity of the mixture where the density is measured (Figure 11). RG174U coaxial cable is currently used for the ERT cables.

The UCT instrument

The ERT system developed at the University of Cape Town (Wilkinson *et al.*, 2005; Randall *et al.*, 2008) is used in the IMST pipe test loop. It was chosen because of its high speed, access to the developers, and cost. The instrument is able to capture data from up to eight 16-electrode rings at a maximum total rate of 1 000 frames/second for a minimal independent set of 104 measurements using the adjacent pair

(a)

(b)

Figure 1—IMST pipe test loop: (a) the test pipes; (b) flow meter, density gauge and single plane ERT ring in Ø80mm vertical return pipe

(a)

(b)

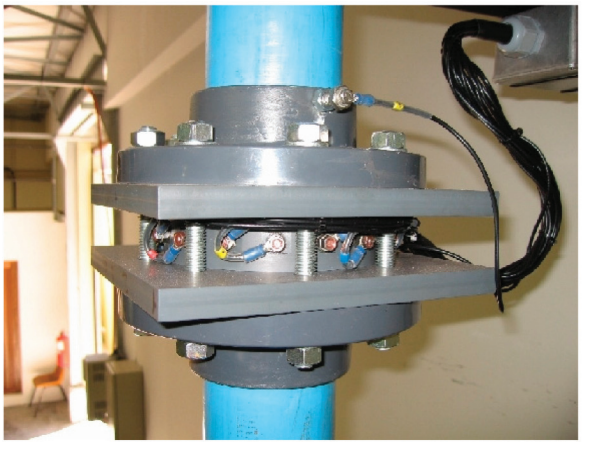

Figure 2—ERT electrode rings in IMST pipe loop: (a) Ø80 mm; (b) Ø80 mm vertical return pipe

'current injection' strategy (Barber *et al.*, 1983; Dyakowski *et al.*, 2000). In practice, modifications were made to the data capture sequence to improve the instrument's noise performance and the data set size was increased to 208 measurements, which reduced the data capture rate. For the results presented here the total capture rate was 566 frames/second, thus for two sensing rings as used for the velocity profiling application, data were recorded at 283 dual frames/second.

Concentration and velocity profile estimation

Reconstruction algorithms

2D algorithm

Several different ways have been proposed to perform tomography reconstructions—see for example Hua and Woo (1990) or Wang (2002). Multi-step algorithms produce accurate results, but are slow. The primary requirement of the UCT system was speed, so the 2D reconstruction method used in the UCT software is similar to the NOSER (Newton's One-step Error Reconstructor) algorithm (Cheney *et al.*, 1990), using only the first step of a Newton-Raphson type algorithm, expressed as a linear problem. The method is very fast because precomputation can be done. Reconstruction then requires only a simple matrix multiplication with the measured voltage set. Reconstruction time increases linearly with the number of elements in the reconstruction finite element model. As the emphasis in this application was fast generation of velocity profiles, qualitative conductivity results suffice since the cross-correlation algorithm effectively operates on relative changes in the signals. Speed is traded for accuracy, but the one-step algorithm is successful if difference imaging or a form of calibration is used. This system uses the calibration described by Wilkinson *et al.* (2005), whereby measured voltages are scaled by a set of factors that would force a measurement set from a homogeneous (constant conductivity) medium to equal that predicted by the forward model.

3D algorithm

The EIDORS (Adler and Lionheart, 2006) suite of Matlab functions implements different algorithms to enable comparison of new techniques with existing methods, and use was made of this package. A 3D version was implemented in the UCT software, based on the algorithm developed by Polydorides and Lionheart, (2002) (also a Newton-Raphson method). Calibration of the data for this algorithm was again done using difference imaging.

Meshes used in 3D algorithms appear coarse compared to those of the 2D algorithms, since the elements are spread over three dimensions and the number of elements in a cross-section is reduced. The velocity profiling algorithm only uses elements in the plane of the electrodes, so once the pre-computations are done, rows that calculate the resistivity of elements outside this plane can be dropped. This means that the 3D mesh can have the some cross-sectional resolution as the 2D mesh without a loss of speed, even though the forward model that generated it had many more elements. The matrix multiplication operation of the 3D one-step algorithm is almost identical to the 2D version, and reconstruction time varies linearly with the number of elements as in the 2D version (Long, 2006).

Comparison of 2D and 3D algorithms

Both algorithms have similar computational times, so they were compared on the basis of the reconstructed images. Although it is not ideal to compare images generated from different meshes, the 2D algorithm seems to discriminate better between the conducting and non-conducting regions. To illustrate this, 2D and 3D reconstructions from (the same) data captured from a tank containing a saline solution with a non-conducting rod placed off centre are shown in Figure 3, which also shows the effect of increasing the number of elements.

Velocity estimation via cross-correlation of reconstructed images

Using cross-correlation to estimate velocity assumes that a particular disturbance in the flow sensed by the first electrode ring (sensing plane) (x) produces a unique signature that can be identified by the second electrode ring (y) a time τ_d later at a distance L further down the pipe.

The ERT hardware captures electrical measurement sets in each sensing plane at a rate of f frames/second. The reconstructed conductivity at each position in an image fluctuates as a function of time, with a bandwidth that depends on the flow speed and the electrode geometry. The frame rate of the system should be high enough to satisfy the Nyquist criterion over all flow speeds of interest.

To determine the speed of the fluid at a particular position (say element i) in the pipe, the relative delay must be estimated from the two sampled time waveforms $x[m]$ and $y[m]$ (where m is the sample index), extracted from the images at corresponding locations. To improve accuracy, the sequences are resampled to a finer sample spacing (typically 1/8th of the original sample spacing). The resampled data sets are then cross-correlated, an operation that involves sliding a windowed portion of one sequence over the other, to obtain the relative time delay τ_d at which the correlation measure is a maximum. A suitable measure is the normalized

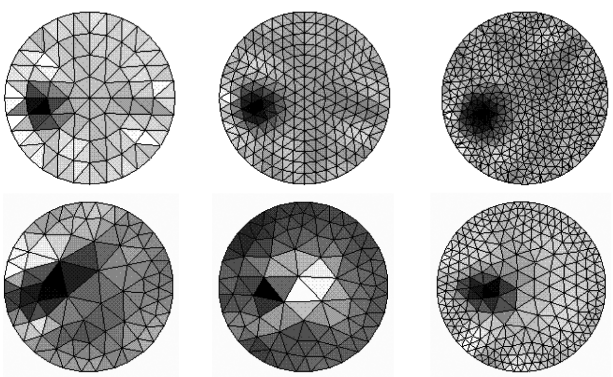

Figure 3—1st row: results of 2D image reconstruction with meshes of 128, 384 and 836 elements; 2nd row: results of 3D image reconstruction with meshes of 910 (126), 1265 (136) and 8749 (440) elements (numbers in brackets are the number of elements in the cross-section) (Long, 2006)

cros

$$cc_{xy}(p) = \frac{1}{N}\sum_{m=0}^{N-1}\left(\frac{x[m]-\bar{x}}{\sigma_x}\right)\left(\frac{y[m+p]-\bar{y}}{\sigma_y}\right) \quad [1]$$

$$p = p_{\min},\ldots,p_{\max}$$

where p is the offset (relative time shift), N is the number of samples in the correlation window, and the interval $[p_{\min}, p_{\max}]$ defines the range of the cross-correlation search; $\bar{x}$ and $\bar{y}$ are the mean values and σ_x and σ_y are the standard deviations, (ideally) calculated over the overlapping window of N samples. Subtracting the means and dividing by the standard deviations results in a correlation measure in the range $-1 \le cc_{xy} \le 1$, where a value of $cc_{xy} = 1$ corresponds to perfect correlation. The normalized correlation coefficient is a useful absolute measure for comparison purposes, as the normalization removes the scale dependence on slurry concentration and carrier fluid salinity, and allows a simple threshold to be used for accepting or rejecting correlation results.

In the practical implementation (Long, 2006), the means and standard deviations are precalculated and used to normalize the sequences to zero-mean, unit standard deviation sequences prior to execution of the correlation operation, i.e. $x[m]$ is replaced by $(x[m] - \bar{x})/\sigma_x$ and $y[m]$ by $(y[m] - \bar{y})/\sigma_y$. In the case of $y[m]$, $\bar{y}$ and σ_y are estimated from the batch of samples, and are not recalculated for each shift (which saves computational time). The correlation operation is then implemented as the standard signal processing cross-correlation operation

$$R_{xy}(p) = \frac{1}{N}\sum_{m=0}^{N-1} x[m]y[m+p] \quad [2]$$

$$p = p_{\min},\ldots,p_{\max}$$

where $R_{xy}(p)$ produces the normalized correlation coefficient $cc_{xy}(p)$ for each offset p. The result is numerically very close, but not identical to that of Equation [1], as the $\bar{y}$ and σ_y are not recalculated for each offset p. The exact calculation using Equation [1] can, however, be carried out once the correlation peak in $R_{xy}(p)$ has been found, with relatively small additional computational time.

The location of the peak $p = p^*$ is used to calculate the time delay $\tau_d = p^*/f$ seconds, and hence the fluid velocity

$$v = \frac{L}{\tau_d} = \frac{L}{(p^*/f)} \quad [3]$$

A velocity profile across the pipe cross-section can be generated by repeating the calculations of Equations [1] to [3] for each element of the reconstructed image (i.e. pixel by pixel).

Interpretation of the correlation coefficient

Ideally the cross-correlation calculation gives an easily identifiable peak so that the offset p^* is determined with confidence. In practice though, results are degraded by measurement noise, image resolution and non-axial movement of particles in the pipe. Results will also be poor in regions where the 'structure' is insufficient to generate a fluctuating conductivity signature, for example near the top of the pipe in this application of sliding bed flow (few coarse particles are present in supernatant, which is essentially homogeneous). Such factors can lead to poor correlation and/or errors in the location of the correlation peak. In this case the indicated maximum in the profile is likely due to a random fluctuation and not a particular disturbance in the flow, and will be significantly lower than a genuinely correlated peak. Thus some measure of the accuracy of the offset is needed.

Figure 4 shows the relationship between the correlation coefficient and the variation of the position of the peak for varying window sizes (number of samples used in the correlation). These results were generated by adding increasing levels of noise to two identical signals, cross-correlating them, and noting the maximum correlation coefficient and corresponding offset. For each noise level, this was repeated 100 times (using different random seeds to generate the noise). The standard deviation of the peak positions was then plotted against the mean of the correlation coefficient for each noise level. Based on peak offset, a higher correlation coefficient gives a better velocity estimation. Increasing the correlation window size reduces the variation of the peak position for a given correlation coefficient cc_{xy} of signals x and y, which suggests that if the magnitude of the peak is low, the velocity estimation can be improved by increasing the window size. The effect of measurement noise on the correlation of two signals is shown as follows (Long, 2006):

Taking $n_c(t)$ as the variation of the flow, two signals $x(t)$ and $y(t)$ to be correlated are represented as

$$x(t) = A + n_c(t) + n_x(t) \quad [4]$$

$$y(t) = A + n_c(t-\tau) + n_y(t) \quad [5]$$

where A is a DC component, $n_c(t-\tau)$ is a delayed version of $n_c(t)$ (both with zero mean) and $n_x(t)$, $n_y(t)$ are independent, zero mean random noise additions with standard deviation σ_n. If the standard deviations of $x(t)$ and $y(t)$ are σ_x and σ_y respectively, and that of the common component, $n_c(t)$, is σ_c, then since

$$\sigma_x^2 = \sigma_y^2 = \sigma_c^2 + \sigma_n^2 \quad [6]$$

the correlation cc_{xy} of these signals is

$$cc_{xy} = \sigma_c^2/(\sigma_c^2 + \sigma_n^2) \quad [7]$$

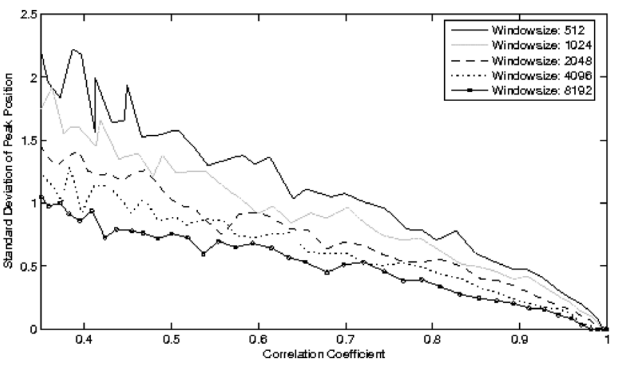

Figure 4—The relationship between the position of the peak and normalized correlation coefficient cc_{xy} (Long, 2006)

Equation [7] shows that if the standard deviation of the noise equals that of the changes due to the flow, then 0.5 is the maximum correlation coefficient attainable. The derivation is given in more detail in Long (2006).

Dealing with bad correlations

Simply finding the maximum in a set of very low correlation values is unlikely to produce the correct offset, so bad correlations need to be properly handled. One option is to discard pixels for which the peak correlation value is below a set threshold. Another is to infer the likely value from surrounding pixels and the velocity history of the pixel. This is valid if there is spatial and temporal coherence in the flow. If these conditions are met, the replacement value of a pixel can be found as the mean of the surrounding pixels, giving each surrounding pixel equal influence on the estimated value. Alternatively, the influence of neighbouring pixels can be weighted by their correlation coefficients, but this won't work if they too are low, in which case the pixel should be excluded from the velocity profile. The threshold for determining whether a pixel should be discarded depends on the flow and measurement conditions. Henningsson *et al.* (2006), for example, took correlation coefficients less than 0.45 as noise. The UCT software shows interpolated pixels in green and discarded pixels in red when visualizing the computed velocity profiles. See for example Figure 8.

Cross-correlation algorithm

Time domain

The most apparent way to implement Equation [1] is directly in the time domain, with an inner loop (m) from 0 to N-1 and an outer loop (p) from P_{min} to P_{max}. This is equivalent to 'sliding' a window of the most recent values from the second measuring plane pixel over the values from the first measuring plane pixel, and multiplying and summing the overlapping values at each step to produce the cross-correlation profile. Resolution can be improved by quadratic interpolation between the values surrounding the correlation peak (Yang and Beck, 1998). The range that the correlation window is 'slid' over can be different from the window size, so the technique of 'auto pre-delay' (Yang and Beck, 1998) can be used if the location of the correlation peak does not vary significantly. The cross-correlation range need only cover the expected variation of peak position, which can greatly improve calculation speed. The time domain algorithm also allows for incremental updates to the correlation profile. Wang *et al.* (2005) note that an extra data point added to the cross-correlation, without repeating the entire calculation, potentially reduces calculation time and improves velocity estimation over time. However, this means the cross-correlation window effectively becomes the entire data set, and variations in flow velocity will not be identified. This optimization should be used only if the velocity is expected to be constant.

Frequency domain

For large window sizes it is computationally efficient to calculate

$$R_{xy} = F^{-1}\{X^*Y\} \quad [7]$$

where F^{-1} is the inverse discrete Fourier transform, X^* is the conjugate of the fast Fourier transform (FFT) of the time series of pixel resistivities in the first measurement plane and Y is the FFT of the time series of resistivities of the corresponding pixel in the second plane. The speed of this method is significantly better than that of the direct time domain correlation method, and it was implemented using the FFTW (http://www.fftw.org) software libraries. The resolution of the correlation function is limited by the rate at which frames are captured, but zero padding the result of the multiplication in the frequency domain before calculating the inverse transform has the effect of oversampling in the time domain (a form of interpolation). This simulates a higher sample rate and is used to improve the estimate of the location of the correlation peak.

Algorithm speed comparison

Figure 5 compares the algorithm speeds for different window sizes when calculated using a 2.67 GHz Intel Celeron processor with 500 MB of RAM. If the expected variation of the peak position is low so the range of the time domain algorithm is small, the time domain algorithm is significantly faster, but for larger correlation ranges the fast correlation algorithm (FFT method) is faster.

Software implementation

Details of the software implementation can be found in Long (2006, 2007). Broadly though, the aim of this work was to develop a real-time application, so speed was important and a compiled language (C++) was chosen to implement the algorithms, using the wxWidgets library (http://www.wxWidgets.org). The software was designed to be modular, with the interfaces between modules explicitly defined. In this application there are two classes of operations—those done irregularly and relatively infrequently, e.g. user interface and display of the velocity profiles, and those which must be done continuously as fast as possible, such as processing the captured data, image reconstruction and cross-correlation. These two classes were implemented as two control loops or threads, resulting in an efficient arrangement in which the application can respond fast to user input without having to wait for a particular data processing step to complete.

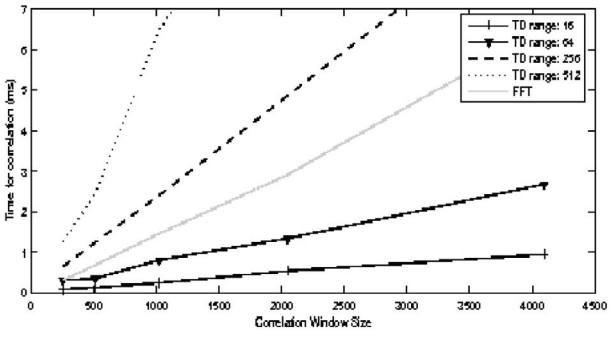

Figure 5—Time domain and fast correlation algorithm speed comparison (Long, 2006)

Results

Real-time operation

In this application 'real-time' means the system updates information at least as fast as it receives raw data—sufficiently fast to track fluctuations in the flow. So, for example, when using a correlation window size of 4096 and a dual frame capture rate of 283/s, data processing must take less than 4096/283 = 14.46 seconds for the velocity profiling system to be real-time and keep up with the data acquisition rate. The time taken to generate a complete velocity profile is the sum of the image reconstruction times, the cross-correlation time and any additional overhead times. Figure 6 shows overall computation time for increasing correlation window sizes for two mesh reconstruction sizes, as well as the real-time limit (time taken to capture the number of samples in the correlation window), using a 2.67 GHz Intel Celeron machine with 500 MB of RAM. Time to visualize the velocity profile must also be accounted for, but is not shown in Figure 6.

Velocity profiles

Measurements with a 5% kaolin carrier fluid with 10% coarse sand flowing in the Ø56 mm pipe were used to initially evaluate the ERT system. Figure 7 shows the results of calculations as described above (2D reconstruction) for a flow condition in which a fast sliding bed was present. Correlation at the bottom of the pipe where there is a high concentration of solid particles (better 'structure') is good, and velocity is correctly shown to increase from bottom to top through the sliding bed, supported by observation of the sliding bed. At the top of the pipe, where there was little or no sand, correlation is bad and the estimated velocities were discarded (elements shown as black). Figure 8 shows velocity profiles (presented as height maps) of the flow for the same mixture in the Ø56 mm pipe at low, medium and high superficial velocities. Velocity profiles derived from the 3D reconstruction were similar, but like the images (see Figure 3) were limited by the number of 3D elements used.

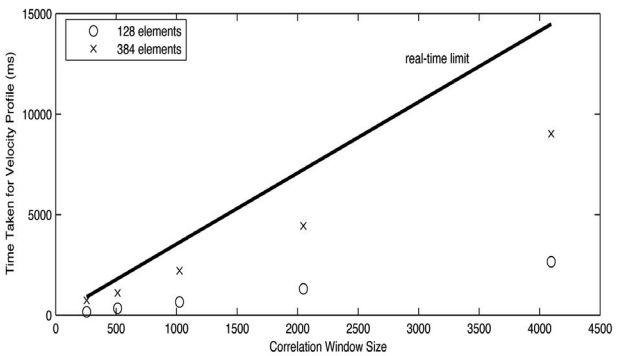

Figure 6—Times to generate velocity profiles using different window sizes for the cross-correlation algorithm (meshes 128 and 384 elements) (Long, 2006)

Some typical pipe results

These developments with the UCT ERT instrument and data processing software were aimed at applying the instrument in the IMST settling slurry flow research work, to fulfil the needs as outlined in the introduction. Progress has been made in this regard, and much data were acquired during an extensive experimental program. Some typical outputs from the system are shown in Figures 9 to 11, obtained using an 836 element reconstruction mesh. Figure 9 shows a typical sliding bed and the corresponding vertical centre line (relative) concentration. Figure 10 shows similar information for different flow rates, where the concentration of coarse material varies from a settled sliding bed to an almost uniform distribution across the pipe cross section. Figure 11 shows the use of the ERT ring in the vertical return pipe to monitor homogeneity of the mixture where the density is measured. The vertical centre line conductivity distribution is overlaid on the image. As can be seen, the mixture in the

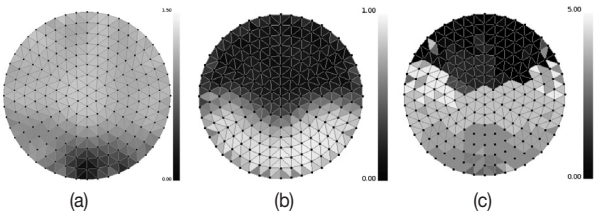

Figure 7—(a) Conductivity profile (darker regions are lower conductivity/higher sand concentration); (b) Corresponding correlation coefficient at each pixel (lighter regions indicate areas of better correlation); (c) Corresponding velocity profile (Long, 2006)

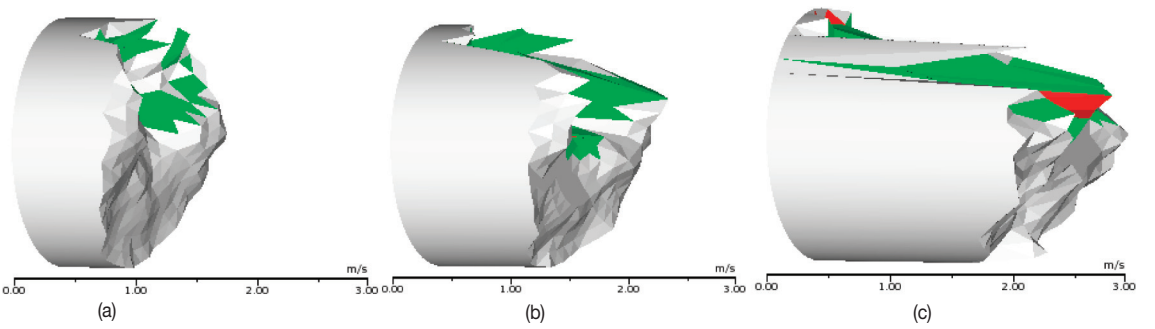

Figure 8—Height map representation of the flow velocity profiles with approximate scale, (a) low speed; (b) medium speed; (c) high speed (Long, 2006)

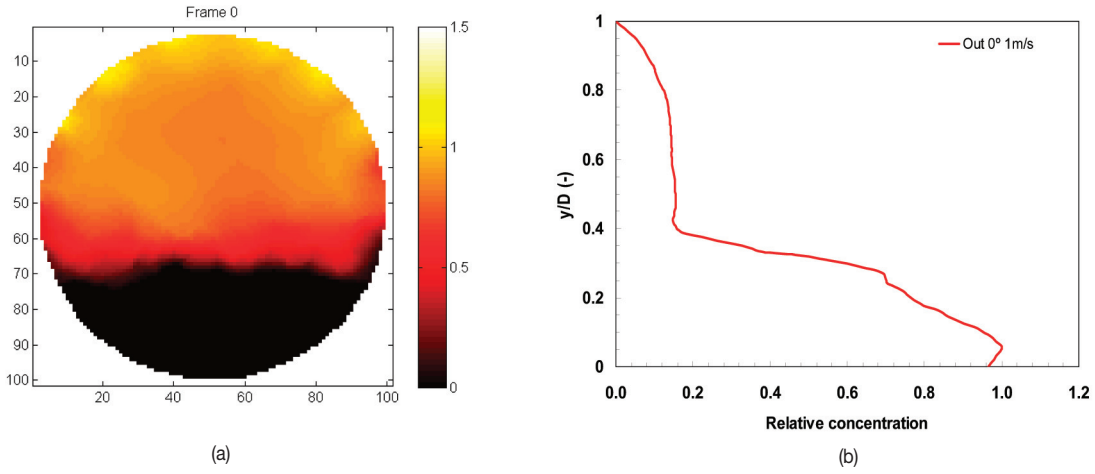

Figure 9—(a) Typical real-time display, settled sliding bed; (b) corresponding vertical centre line concentration distribution (kaolin carrier fluid with silica sand coarse particles)

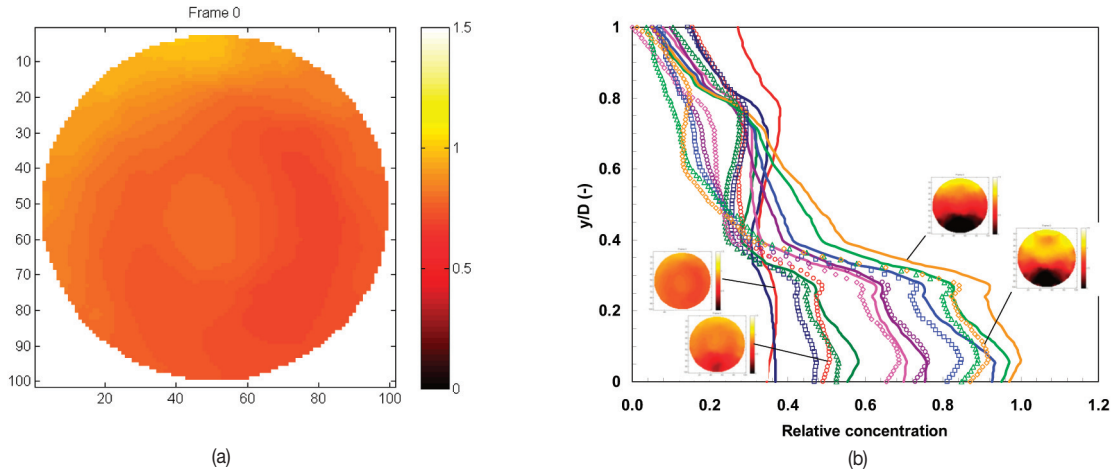

Figure 10—(a) Typical real-time display, particles largely suspended; (b) vertical centreline concentration distributions (kaolin carrier fluid with silica sand coarse particles) at different flow rates

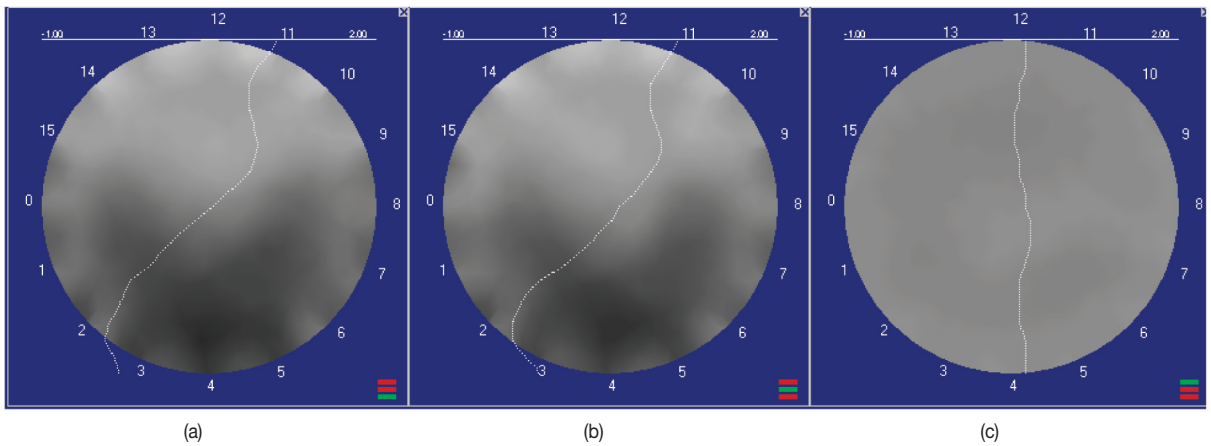

Figure 11—Typical real-time display – (a) out pipe, (b) return pipe, (c) vertical return pipe: Line plot of relative conductivity across diameter from 12 to 4 (rotated 90°) overlaid on image

vertical pipe is homogeneous, even though a settled bed is evident in the horizontal test pipes (Figure 11 (a) and (b)). The density gauge reading (single orientation) can therefore be used with confidence. As alluded to in the introduction the ultimate aim is to determine accurate concentration and velocity profiles over a sufficiently wide range of settling slurries and flow conditions, along with the flow parameters of interest, to enable correct modelling of these slurries for pipeline design.

Conclusions and future developments

A tool to visualize the flow of settling slurries using the UCT ERT instrument has been developed. The system is fast enough to allow real-time visualisation of velocity profiles for the flow rates investigated. The software design is highly modular and allows for different cross-correlation and reconstruction algorithms to be 'swapped' in and out of the program, using EIDORS open-source code or custom routines. Concentration and velocity profiles obtained so far (with kaolin carrier fluid and silica sand coarse particles), derived from data captured by electrode planes 100 mm apart, reflect the nature of the flow for superficial pipe velocities up to 4 m/s, although the accuracy of the velocity profiles has not yet been quantitatively verified.

During this experimental exercise the very real practical difficulty of electrical noise was found to be a major problem. The main sources of this noise are the variable speed drives of the pump and mixer. Efforts to fix this problem (within time/budget constraints) helped considerably, but did not entirely resolve it. This noise needs to be minimized before proceeding with further detailed testing and evaluation of the ERT instrument. In spite of this, significant progress has been made in incorporating the ERT instrument into the slurry pipe tests to acquire the additional information needed for improved settling slurry flow modelling.

Several ideas under consideration for further development of the system are:

- improvements in reconstruction routines and implementation of other methods such as the SCG (Wang, 2002)
- verification of ERT results using vertical and horizontal traversing γ-ray density gauges
- batch and automated post-processing of data
- using the 'best-correlated pixels' method (Mosorov *et al.*, 2002)
- the development of a low cost instrument (two 16-electrode rings) specifically for 'remote' pipe flow measurements and processing.

Acknowledgements

The authors would like to thank the National Research Foundation and Cape Peninsula University of Technology for their ongoing support of this work.

References

Adler, A. and Lionheart, W.R.B. Uses and abuses of EIDORS: an extensible software base for EIT. *Physiological Measurement*, vol. 27 no. 5, 2006. pp. S25–S42.

Barber, D.C., Brown, B.H. and Freston, I.L. Imaging spatial distributions of resistivity using applied potential tomography. *Electronics Letters*, vol. 19, 1983. pp. 993–995.

Cheney, M., Isaacson, D., Newall, J.C., Simske, S. and Goble, J. NOSER: An algorithm for solving the inverse conductivity problem. *International Journal of Imaging Systems and Technology*, vol. 2 no. 2, 1990. pp. 66–75.

Cooke, R. Laminar flow settling: the potential for unexpected problems. *Proc: 15th Int. Conf. on Hydrotransport*, Banff, Canada, 2002. pp. 121–133.

Deng, X., Dong, L., Xu, L.J., Liu, X.P. and Xu, L.A. The design of a dual-plane ERT system for cross correlation measurement of bubbly gas/liquid pipe flow. *Measurement Science and Technology*, vol. 12 no. 8, 2001. pp. 1024–1031.

Dyakowski, T., Jeanmeure, L.F.C. and Jaworski, A.J. Applications of electrical tomography for gas-solids and liquid-solids flows – a review. *Powder Technology*, vol. 112, 2000. pp. 174–192.

FFTW (The Fastest Fourier Transform in the West), http://www.fftw.org.

Henningsson, M., Ostergren, K. and Dejmek, P. Plug flow of yoghurt in piping as determined by cross-correlated dual-plane electrical resistance tomography. *Journal of Food Engineering*, vol. 76 no. 2, 2006. pp. 163–168.

Hua, P. and Woo, E. Electrical Impedance Tomography. The Adam Hilger Series on Biomedical Engineering. Bristol: IOP Publishing Ltd. 1990. pp. 97–136.

Long, T.M. The development of an online velocity flow profiling system using electrical resistance tomography. Master's thesis, University of Cape Town. 2006.

Long, T.M., Wilkinson, A.J., Randall, E.W. and Sutherland, A.P.N. An Online Real-time Velocity Profiling System using Electrical Resistance Tomography. *Proceedings of 5th World Congress on Industrial Process Tomography*, Bergen. 2007.

Lucas, G.P., Cory, J., Waterfall, R.C., Loh, W.W. and Dickin, F.J. Measurement of the solids volume fraction and velocity distributions in solids-liquid flows using dual-plane electrical resistance tomography. *Flow Measurement and Instrumentation*, vol. 10 no. 41, 999. pp. 249–258.

Matousek, V. Research Developments in Pipeline Transport of Settling Slurries. *Proc: 12th Int. Conf. on Transport and Sedimentation of Solid Particles*, Prague (Czech Republic), 2004. pp. 19–34.

Mosorov, V., Sankowski, D., Mazurkiewicz, L. and Dyakowski, T. The 'best-correlated pixels' method for solid mass flow measurements using electrical capacitance tomography. *Measurement Science and Technology*, vol. 13 no. 12, 2002. pp. 1810–1814.

Polydorides, N. and Lionheart, W.R.B. A MATLAB toolkit for three-dimensional electrical impedance tomography: a contribution to the electrical impedance and diffuse optical reconstruction software project. *Measurement Science and Technology*, vol. 13 no. 12, 2002. pp. 1871–1883.

Pullum, L. and Graham, L.J.W. A new high-concentration pipeline test loop facility. *Proc: 14th Int. Conf. on Slurry Handling and Pipeline Transport*, Maastricht, The Netherlands, 1999. pp. 504–514.

Pullum, L. and Graham, L.J.W. Predicting fine particle suspension performance – the case for pipe tests. *Proc: 15th Int. Conf. on Hydrotransport*, Banff, Canada, 2002. pp. 109–120.

Randall, E.W., Wilkinson, A.J., Long, T.M. and Sutherland, A.P.N. The UCT Electrical Resistance Tomography Instrument and Its Applications. *The Southern African Institute of Mining and Metallurgy, Tomographic Symposium*, 25 July 2008.

Wang, M. Inverse solutions for electrical impedance tomography based on conjugate gradients methods. *Measurement Science and Technology*, vol. 13 no. 1, 2002. pp. 101–117.

Wang, M., Ma, Y., Holliday, N., Dai, Y., Williams, R. and Lucas, G. A high performance EIT system. *IEEE Sensors Journal*, vol. 5 no. 2, 2005. pp. 289–299.

Wilkinson, A.J., Randall. E.W., Cilliers, J.J., Durrett, D., Naidoo, T. and Long. T.M. A 1000-measurement frames/second ERT data capture system with real-time visualization. *IEEE Sensors Journal*, vol. 5 no. 2, 2005. pp. 300–307.

wxWidgets, http://www.wxWidgets.org.

Yang, W.Q. and Beck, M.S. An intelligent cross correlator for pipeline flow velocity measurement. *Flow Measurement and Instrumentation*, vol. 8 no. 2, 1998. pp. 77–84. ◆