Experiments in Fluids 33 (2002) 743–751
DOI 10.1007/s00348-002-0498-4

Accuracy and time performance of different schemes of the local field correction PIV technique

A. Lecuona, J. Nogueira, P.A. Rodríguez, D. Santana

Abstract Local field correction particle image velocimetry (LFC-PIV) has become an established alternative among high-resolution PIV techniques. Previous works by the authors introduced its implementation by means of simple algorithms. In these works the initial limitation of the method, which was related to the mean distance between particles, was removed. Comparison with other contemporary high-resolution techniques indicates that it offers advantages in robustness and accuracy. The trade-off for this better performance is a heavier computing load. Until now, the computing time that this load requires has not been characterized in detail, since this computing time could be substantially reduced by accepting a reduction in accuracy. This paper focuses on the characterization of the trade-off between time and accuracy, thus offering a new perspective to PIV. In this field, LFC-PIV offers a wide range of possibilities that are described in the paper. Several alternative schemes for LFC-PIV are tested, together with an analysis of the influence of the number of iterations. Performance figures for both accuracy and expended time are given. Metrological evaluation is carried out over synthetic images. A test of coherence between these results and the performance on real images is also presented. The paper shows that even for a limited number of iterations this technique offers advantages.

1 Introduction

Local field correction particle image velocimetry (LFC-PIV), because of its unique approach, is the only correlation PIV technique able to resolve flow structures much smaller than the interrogation windows employed (Nogueira et al. 1999, 2001). This allows for an iterative high-resolution system using large window sizes.

There are other solutions to obtain high resolution below the size of an initial interrogation window size. Two examples are hybrid PIV systems and multigrid PIV systems. Hybrid PIV systems switch to particle-tracking velocimetry (PTV) once a first guess of the velocity is obtained by the usual correlation PIV method. To do this, each particle has to be identified and isolated from the surroundings, with the consequent advantages and disadvantages (Keane et al. 1995; Cowen and Monismith 1997, among others). In multigrid PIV systems a reduction of the interrogation window size allows high resolution. This reduction is possible due to an iterative calculation of the window shift (Soria 1996; Scarano and Riethmuller 1999; Lecordier et al. 1999; Hart 2000). Again, the advantages and disadvantages of the reduction of the interrogation windows have been already studied.

The LFC-PIV does not need to identify the particles nor is it impaired by the loss of robustness related to small interrogation windows. Using large windows makes the system very robust was well as very accurate (Nogueira et al. 1999, 2001). Nevertheless, the use of large windows results in longer computing time. This last issue can become relevant if the number of iterations is large. The iterations can be stopped at any time, yielding results that are monotonically better until the optimum. This offers an additional degree of freedom to conventional PIV. The user (expert) can decide the level of accuracy needed, even on-line, stopping at a certain point for fast results and continuing the processing at a later time to extract the ultimate possibilities of the images. This is especially important for industrial applications, where operating costs of the facilities are very high.

In addition, there are alternative schemes to implement the technique that can also affect the trade-off between processing time and accuracy. This work arises from the need to characterize this balance.

Nogueira et al. (2001) and Rodríguez et al. (2001) documented the basic procedure of LFC-PIV. Although this paper does not discuss the LFC-PIV technique itself, it offers some notes about the technique as an introduction and to establish a framework. LFC-PIV can be briefly summarized as an iterative method in which previous measurements are used to deform the original images (including translation) at a subpixel level. The deformation compensates for the particle pattern deformation

Accepted: 25 April 2002
Published online: 14 August 2002

A. Lecuona (✉), J. Nogueira, P.A. Rodríguez, D. Santana
Department of Mechanical Engineering,
Universidad Carlos III de Madrid,
Butarque 15, 28911-Leganés, Madrid, Spain
E-mail: lecuona@ing.uc3m.es

This work was partially funded by the Spanish Research Agency grant AMB1999–0211, DPI2000–1839-CE, Feder 1FD97–1256 in a collaboration with the Universidad de Las Palmas de Gran Canaria and under the EUROPIV 2 project. A Joint Program to Improve PIV Performance for Industry and Research is a collaboration between LML URA CNRS 1441, Dassault Aviation, DASA, ITAP, CIRA, DLR, ISL, NLR, ONERA, DNW and the universities of Delft, Madrid (Carlos III), Oldenburg, Rome, Rouen (CORIA URA CNRS 230), St Etienne (TSI URA CNRS 842) and Zaragoza. The project is managed by LML URA CNRS 1441 and is funded by the CEC under the IMT initiative, Contract no. GRD1–1999–10835.

caused by the continuous displacement field between both images, thus increasing the signal-to-noise ratio for further processing. This idea was proposed by Huang et al. (1993) and Jambunathan et al. (1995). Unfortunately, its direct use leads to two difficulties: instabilities along the iterations and an incorrect deformation for spatial wavelengths smaller than the interrogation window size. Both difficulties precluded high-resolution results. LFC-PIV is a unique method that overcomes these two strong limitations by the use of a proprietary weighting function (Nogueira et al. 1999, 2001).

A more detailed picture of the operations performed by LFC-PIV at each iteration is:

- Processing of the images: the usual cross-correlation PIV processing of the two particle images delivers vectors at the grid nodes. These images are the original images, a and b, in the first iteration and corrected ones, $a^\star$ and $b^\star$, in consecutive iterations. In this processing the mentioned weighting function υ is applied to the windows. The use of this weighting allows for compensation of the particle pattern deformation in successive steps towards the right direction, even with large windows. This gives the technique its advantages in terms of accuracy and robustness. Nevertheless, the application of a weighting function introduces a small error. Selecting the vectors to be corrected for each iteration partially avoids this error. Further correction of the error can be performed through a modification of the scheme of application of the function (Sect. 3.2).
- Compensation of the particle pattern deformation: with the displacement information obtained from each iteration, the deformation of the particle pattern between images a and b is compensated, obtaining images $a^\star$ and $b^\star$ through pixel interpolation. These images are processed in place of a and b, respectively, after the first iteration. The result allows for correction of the previously calculated displacement field. Further correction of the particle pattern deformation of images a and b gives new images $a^\star$ and $b^\star$, closing the iterative loop.

Compensation of the particle pattern deformation of both correlating images in the right direction increases the correlation peak. Additionally, iterative compensation of the deformation eliminates the low-pass effect of the large windows, thus offering high-resolution capability. Under this circumstance, it is the grid distance but not the size of the window that limits the available resolution.

Overlapping is usually very high in order to obtain vectors at distances short enough to satisfy the Nyquist criterion on the small scales of the flow. It is important to consider that the advantages that high-resolution methods offer are effective only when the density of the seeding in the flow is large enough to give information at these small spatial wavelengths. Large windows avoid the lack of information related to the use of small windows at points where there is high noise, thus increasing robustness.

It is worth mentioning that vector validation and interpolation at intermediate iterations are useful to accelerate convergence. The iterative character of the method and its robustness make it capable of cancelling false measurements to a certain extent, at the expense of prolonging the iterations. The result is that vector validation is not critical. The progressive compensation of the particle pattern deformation increases the signal-to-noise ratio, thus vectors that have been interpolated in the initial steps can lead to valid measurements in consecutive ones.

If no error is present, the process would converge to a situation where images $a^\star$ and $b^\star$ are identical. The real situation is different; at a certain point, the iterations do not further increase the accuracy of the measurement. The above-mentioned authors' works introduced ways to estimate this state and hence stop the iterations. This point will define the last iteration in the present work, although results for intermediate points are given.

If not specified, the interrogation window size F will be set to 64 pixels. This minimizes the error introduced by the weighting function. Since the system can extract high spatial frequency features if the seeding density allows for it, the grid node spacing Δ will be set to 4 pixels. The above-mentioned works give details in all these aspects.

2 Synthetic images

Synthetic images have been used to evaluate the performance of LFC-PIV. The displacement field s measured in these images corresponds to one-dimensional single-frequency harmonics.

$$s = 2\sin(2\pi x/\lambda_x);\,(\text{pixels}). \tag{1}$$

Shifting the wavelength λ_x gives the response of the technique to different spatial wavelength features. This test is considered valid to give an idea of the performance, since the good performance of the technique does not come from the periodicity in the flow field (Nogueira et al. 2001). The results of the recent worldwide cooperation project for the assessment of PIV performance (www.pivchallenge.org) also sustain this.

The remaining parameters that define the images divide them into two categories, type 1 and type 2. Type 1 images are designed to exclusively test the performance of the technique in respect to the spatial wavelength. The images are generated as an example of good conditions for testing high resolution. They contain no noise except for the spatial discretisation of the CCD and the usual 8-bit grey level. The mean distance between particle images is high, δ=2 pixels, i.e. $4/(\pi\delta^2)\sim$0.3 ppp (particles per pixel area). The e^{-2} diameter of the Gaussian particle images is d=2 pixels. This monodispersivity arises because the diffraction limit of the lenses defines the optical diameter for small particles, instead of the geometrical diameter. Their shape is integrated with unity fill factor over each pixel surface (Westerweel 1998). Where particles overlap, the corresponding intensities are added. No out-of-plane displacement was considered. Figure 1 depicts a zoomed-in portion of these images. This type of image is comparable with those obtained by seeding oil droplets in air in large wind tunnel facilities.

It is also important to consider the effect of noise sources. To do this and to allow direct comparison with previous work (Nogueira et al. 1999, 2001), a different type

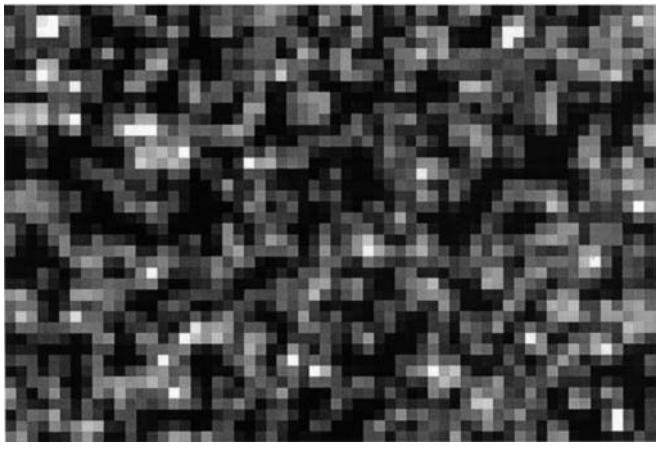

Fig. 1. Close-up zone of a synthetic image of type 1 (60 by 40 pixels)

of images has been devised. In these type 2 images the mean distance between particle images is larger, δ=4.5 pixels, i.e. $4/(\pi\delta^2)\sim$0.06 ppp (particles per pixel area). The diameter of particle images ranges from d=2 pixels to d=4 pixels. Their intensity ranges from 0 to 92% of the maximum grey level. Both intensity and diameter are allowed to randomly differ up to 10% within each particle pair. Their intensity distribution is Gaussian and is integrated with unity fill factor over each pixel surface. Where particles overlap, the corresponding intensities are added. Five percent of particles have no second image to correlate, simulating out-of-plane displacement. A continuous background grey level reaches up to 20% of the maximum range and is distributed over large zones ($\sim$60 pixels). Figure 2 depicts an example of a zoomed-in portion of these images. This type of image is similar to those obtained in small facilities where water is seeded with pollen, spores or plastic particles, leading to larger diameters.

3
Implemented LFC-PIV schemes

As commented Sect. 1, several alternatives in the LFC-PIV technique modify the trade-off between accuracy and computing time. The following subsections define the more relevant options. To introduce them, this section delivers some performance examples. The measurement error is evaluated by its root-mean-square value rms(e) normalized by the root-mean-square value of the signal rms(s). This can be easily translated into absolute values, observing that rms(s)=1.41 pixels. The average of the bias error over the whole image is negligible, even in the presence of peak locking, thus justifying the use of the proposed single metrological parameter.

In this section we decide which schemes should be evaluated later in detail. Section 5 discusses results from the schemes selected and gives error measurements and time performance on synthetic images for the most balanced options. Coherence with the behaviour in real images is studied in Sect. 6.

3.1
Number of iterations

As happens in many iterative techniques, the criterion to decide the number of iterations in LFC-PIV is an open subject. In previous works by the authors, the criterion was based on an estimation of the maximum accuracy obtainable by the technique. In this work, the importance of the computing time is also considered.

Nogueira et al. (1999) showed that the rms error in the measurement rapidly diminishes monotonically in the initial iterations, but decreases more slowly as the number of iterations grows. Consequently, good measurements are possible far before the best one is reached, resulting in a reduced computing time (Fig. 3). Figure 3 depicts the behaviour for several wavelengths of the system described in Nogueira et al. (2001) on type 1 images. The curves stop when the algorithm finishes as indicated above.

Before translating the number of iterations into computing time, some considerations have to be raised. They concern the type of iteration from the computational point of view. The initial iterations must take into account the possibility of large displacements. For this reason the whole correlation plane is calculated by means of a fast Fourier transform (FFT). After a few iterations, only corrections to the displacement field in the range of

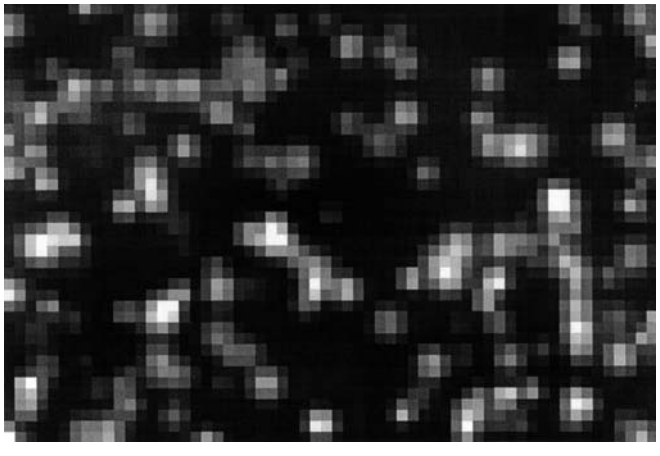

Fig. 2. Close-up zone of a synthetic image of type 2(60 by 40 pixels)

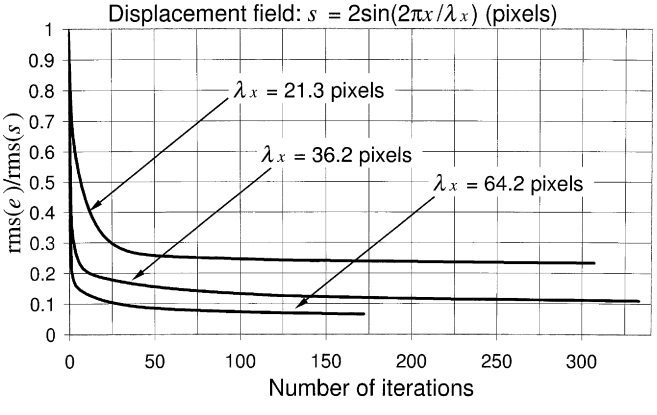

Fig. 3. Behaviour of basic LFC-PIV scheme based on F=64 pixels for several wavelengths

–0.5 to 0.5 pixels have to be measured. This can be performed by directly calculating only the five central values in the correlation plane, which is faster than the FFT method. In a generic LFC-PIV, the jump from one type of correlation to the other is decided depending on the measurement from the previous iteration. This can happen after a single iteration for easy cases, and after up to five iterations for difficult ones. Here it is locked to three iterations. This restriction allows for a uniform performance comparison of different schemes working over different test images.

Nogueira et al. (2001) also established two other types of iterations: either all the grid nodes are allowed to evolve or only some of them. This brings a further discrepancy between number of iterations and expended time, which hardly ever amounts to more than 5% difference in computing time for different runs. Where this happens the maximum time obtained is given.

As Fig. 3 depicts, for large wavelengths a good accuracy is obtained after a few iterations. However, the smaller the wavelength is, the larger becomes the number of iterations needed to achieve good accuracy. There are other details to consider before the trade-off between accuracy and computing time is characterised in Sect. 5.

3.2
Weighting function implementation

LFC-PIV generally uses large windows and small grid distances. In this scenario, a weighting function is needed to avoid instabilities in the iterative processes of compensation of the particle pattern (Nogueira et al. 2001). This means calculating the correlation coefficients C_{lm}

$$C_{lm}=\frac{\sum_{\xi,\eta=-F/2}^{F/2}\upsilon(\xi,\eta)f(\xi,\eta)\cdot\upsilon(\xi,\eta)g(\xi+l,\eta+m)}{\sqrt{\sum_{\xi,\eta=-F/2}^{F/2}\upsilon^2(\xi,\eta)f^2(\xi,\eta)\sum_{\xi,\eta=-F/2}^{F/2}\upsilon^2(\xi,\eta)g^2(\xi+l,\eta+m)}}. \quad (2)$$

Here f and g are the grey values of the two correlated images, a and b, respectively; l and m are the displacements associated with each calculation of the correlation coefficient; ξ and η are coordinates with origin at the center of the interrogation window and F is the length of its sides. The recommended weighting function υ is

$$\upsilon^2(\xi,\eta)=9\left(4\left|\frac{\xi}{F}\right|^2-4\left|\frac{\xi}{F}\right|+1\right)\left(4\left|\frac{\eta}{F}\right|^2-4\left|\frac{\eta}{F}\right|+1\right). \quad (3)$$

The use of a weighting function distorts the grey levels of the original images, introducing error. It is small for the large recommended F values (64 or 32 pixels). The LFC-PIV method works accurately even if it has this error built in. The error can be reduced by "double weighting", which consists of averaging the result of Eq. (2) with that calculated by its symmetrical counterpart, indicated in Eq. (4)

$$C_{lm}=\frac{\sum_{\xi,\eta=-F/2}^{F/2}\upsilon(\xi,\eta)f(\xi-l,\eta-m)\cdot\upsilon(\xi,\eta)g(\xi,\eta)}{\sqrt{\sum_{\xi,\eta=-F/2}^{F/2}\upsilon^2(\xi,\eta)f^2(\xi-l,\eta-m)\sum_{\xi,\eta=-F/2}^{F/2}\upsilon^2(\xi,\eta)g^2(\xi,\eta)}}. \quad (4)$$

The averaging can be performed either in the correlation plane or by just averaging both displacements. This last option is computationally more efficient, because part of the denominators in both expressions can be ignored.

With this option, the accuracy is improved at the expense of almost doubling the computing time. Thus, if the criterion were to obtain the most accurate result, this would be the option to choose. On the other hand, if the computing time has to be taken into account, the small increase in accuracy obtained has to be considered. Figure 4 compares the error in both cases. This figure shows an expanded scale in order to appreciate the changes in the measurement error. The maximum reduction is ~0.02 pixels for the same number of iterations. On the other hand, if the comparison is based on the same computing time, this figure is in favour of the double-weighting scheme only after ~80 single weighting iterations, and in the best cases the gain is ~0.01 pixels. Consequently, this option is only of value when there is plenty of computing time.

3.3
Smoothing of initial errors

Another relevant issue is the spatial smoothing of the displacement field calculated at initial iterations (Rodríguez et al. 2001). Obviously, this reduces the response of the system to small features in the flow for these initial iterations. However, Fig. 5 shows that the final error is clearly reduced. The question of how this procedure benefits a high-resolution system arises. The rationale is that the particle pattern deformation is highest when correlating the initial pair of images. This implies that, at this stage, the possibility of erroneous measurements is also highest. These errors can introduce small-wavelength false features that deviate the compensation in the wrong direction, thus reducing the signal-to-noise

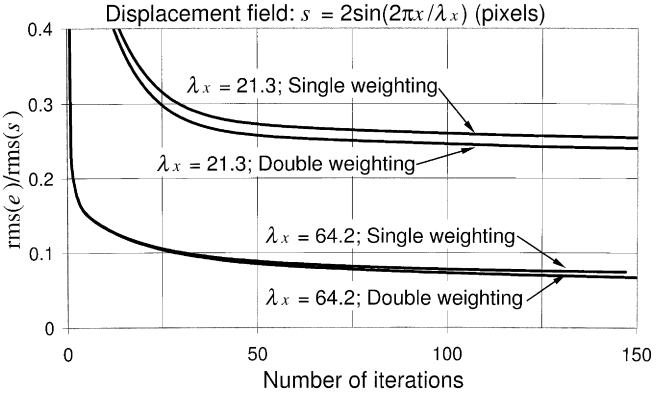

Fig. 4. Behaviour of single- and double-weighting LFC-PIV schemes on type 1 images for several wavelengths

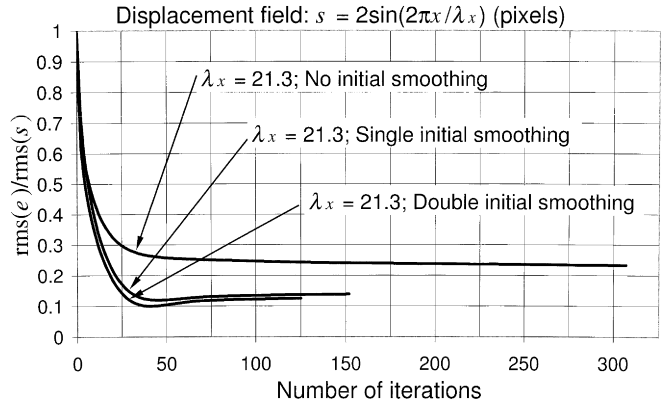

Fig. 5. Behaviour of unsmoothed, single-smoothed and double-smoothed LFC-PIV schemes for initial window with F=64 pixels

ratio. The response of the large windows to small wavelengths is small. Consequently, the system needs a large number of iterations to remove those errors.

A way to avoid this is to apply a moving average to smooth the displacement field obtained in the initial iterations. This smoothed displacement field still contains enough information to allow for reduction of the particle pattern deformation. Thus, the signal-to-noise ratio is increased in the following iterations and the possibility of introducing small-wavelength errors is reduced. Although the initial iterations have a reduced small-wavelength response in this method, the following iterations operate in a better signal-to-noise environment and cleanly extract the information in the small features of the flow after the smoothing.

Two alternative schemes based on smoothing initial iterations have been tested. One consists of smoothing the measured displacement in the first iteration with a moving average of 3 by 3 grid nodes (single initial smoothing). The other scheme additionally performs a second interrogation pass with smaller windows (F=32 pixels), which is also smoothed the same way (double initial smoothing). Figure 5 shows that this second scheme performs slightly better for the same number of iterations. Furthermore, the computing time is reduced as this second pass computes the correlation in a smaller interrogation window.

These results show that there is no drawback for the double-smoothing scheme in comparison with the others. Therefore, it would be the one to use in Sect. 5 for further analysis. Eventually the PIV images may be so noisy as to make the 32-pixel window yield a false result. Then the validation algorithm would refuse this measurement and the scheme would perform as the single-smoothed one.

3.4 Window size

As commented in Sect. 1, the work presented here relies mostly on interrogation windows with F=64 pixels. The use of smaller windows reduces the computing time and gives a higher response to small wavelengths. This helps cut the number of iterations, but as the LFC-PIV resolution is not directly related to the size of the interrogation window, it does not increase the accuracy. Actually, accuracy is reduced due to a larger error introduced by the weighting function (Nogueira et al. 2001). These reasons define the selection of the interrogation window size as an issue, in order to balance the computing time and accuracy.

However, a particularity makes this consideration different from the ones offered in Sects. 3.1–3.3. Apart from the effect on accuracy and computing time, the robustness of the system is also reduced if only steps with F=32 pixels are implemented. To give an idea of the order of magnitude of this reduction, the authors' previous work allows to state that the robustness of a weighted interrogation window with F=64 is similar to that of an unweighted one with F between 32 and 64 pixels. Similarly, the robustness of a weighted window with F=32 can be identified with that of an unweighted one with F between 16 and 32. In summary, there is not a fixed rule to recommend the use of F=64 or F=32 pixels interrogation windows. In Sect. 5 time and accuracy are characterized for both options in order to help the designer.

Figure 6 plots the performance of both options over each type of synthetic image. For images of type 1 (those with almost no noise) it can be clearly observed that the number of iterations is reduced, and the accuracy is similar for F=32. However, the case of images of type 2 (with noise) shows some differences. The error from the weighting function accumulates faster, manifesting a positive slope in the curve. In addition, the algorithm that decides the moment to stop iterating (Nogueira et al. 2001) is much less effective than with the larger window, as it stops much later than the minimum.

For this reason, the ideal application of the smaller window size seems to be in the beginning of the processing, with a limited number of iterations in the range 5–10. After that the scheme should switch to the larger size. Nonetheless, a wider analysis seems necessary. Nogueira et al. (2001) and Lecuona et al. (2001) studied systems that change the sizes of the weighted windows when appropriate, named multigrid LFC-PIV.

The behaviour in respect to double weighting or initial smoothing for F=32 is similar to the cases with F=64 pixels. The only difference is that double smoothing is not considered here, as it is not advisable to use weighted windows with F=16 pixels.

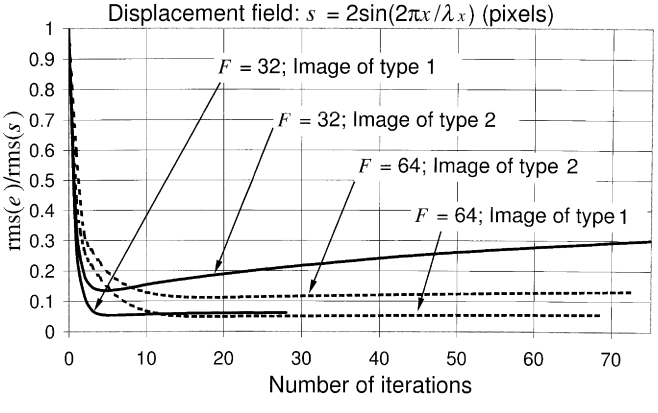

Fig. 6. Behaviour of LFC-PIV schemes with F=64 and F=32 over images of types 1 and 2; λ_x=36.2 pixels in all cases

Summarising, the study presented in Sect. 3 allows us to focus on only two schemes in Sect. 5 since they are good representatives of contemporary LFC-PIV methods: the single-weighted, double-smoothed scheme with F=64 pixels and the single-weighted, single-smoothed scheme with F=32 pixels.

4
Units to measure time performance

An appropriate time unit has to be defined to evaluate the computing time performance of the method. This unit is necessarily arbitrary due to the broad diversity of computing hardware and software. Here these time units correspond to milliseconds of computing time per vector in a 350-MHz Pentium II(r) computer. The authors consider that the time could be reduced to about one-quarter of this value by running on a 1.5-GHz Pentium IV computer. A further detail about time performance is that in the 350-MHz Pentium II(r) system, the program needs 14 ms for a 64 by 64 FFT (6 ms are needed to implement Eq. (2), 4 ms direct and 2 ms inverse). These data can be used to extrapolate results to other types of software and hardware just by evaluating the time spent for a 64 by 64 FFT.

A further breakdown of the computational cost is illustrative. The breakdown is different for each type of iteration detailed in Sect. 3.1. Figures are indicated for F=64 and Δ=4 pixels; F=32 is indicated within parentheses.

- FFT calculated vectors: ~96% (86%) of the computing time for FFT transforms, ~4% (14%) of the computing time for other computations (mainly control routine and image deformation). If only half of the vectors are allowed to evolve, the last part reaches 8% (25%).
- Direct correlation calculated vectors: ~ 35% (12%) of the computing time for direct correlation, ~ 65% (88%) of the computing time for other computations (mainly control routine and image deformation). If only half of the vectors are allowed to evolve, the last part reaches 77% (94%).

5
Balance between accuracy and time performance

Following the conclusions of Sect. 3, two schemes are tested here as relevant for an analysis of the balance between time and accuracy. The accuracy is evaluated by the error as a function of spatial wavelength and the number of iterations. Tables supply the correspondence between number of iterations and time units.

As an additional reference on accuracy, all plots show the best performance line for a conventional multigrid method using compensation of the particle pattern deformation and reduction of window size down to 16 pixels. The hatched areas indicate where the results of this method would fall. In applying this technique, care has been taken not to bring up the unstable frequencies. Results are consistent with those from other authors, as indicated in Nogueira et al. (2001) and Lecuona et al. (2001). Following these works, a certain frequency could be optimised for accuracy, showing lower errors than indicated by the hatched area. The enhancement at this frequency is at the cost of a strong deterioration at other frequencies, yielding a system unsuitable for measurement.

Figure 7 shows the performance of LFC-PIV with F=64 applied to images of type 1. As commented before, this scheme includes a step with F=32, but it was already observed that in Fig. 5 accuracy is very similar to the single-smoothing scheme without this step. It can be remarked that only five iterations give better results than the multigrid system ending with F=16. This highlights that although the low response of large windows does not match that of small windows, it is effective when appropriate weighting is applied, as in Sect. 3.2.

Table 1 gives the correspondence between computing time and the number of iterations. In these same figures, a discontinuous line shows the final error if the algorithm described in Nogueira et al. (2001) decides the number of iterations for stopping.

Figure 8 depicts the performance of the same scheme over images of type 2. For the required computing time,

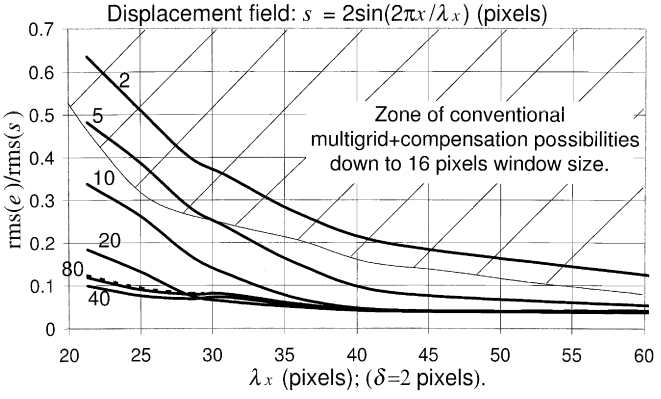

Fig. 7. Behaviour of LFC-PIV scheme with F=64 as a function of wavelength and for several numbers of iterations applied to type 1 images. Table 1 gives time performance

Table 1. Correspondence of time to number of iterations for Figs. 7 and 8

Number of iterations	2	5	10	20	40	80
Time units/vector	105	211	261	362	557	896

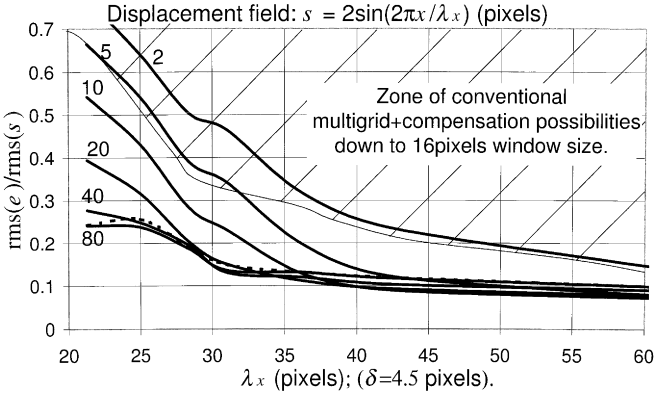

Fig. 8. Behaviour of LFC-PIV scheme with F=64 as a function of wavelength and for several numbers of iterations applied to type 2 images. Table 1 gives time performance

Table 1 is also valid. It shows the effect of higher noise on error. The low sampling of the displacement field adds further errors for small wavelengths because of the mean distance between particles.

In Sect. 3 we proposed studying the behaviour of a system based on F=32 pixels. Figure 9 shows the performance obtained for this LFC-PIV scheme, applied to images of type 1. Table 2 gives the correspondence between number of iterations and computing time. The application of the algorithm that stops the iterations, described in Nogueira et al. (2001), seems to work with these clean images, although it was not designed for this size of window.

Figure 10 shows the performance of the LFC-PIV scheme based on F=32, applied to images of type 2. Table 2 is also appropriate for this case. A remarkable fact is that error does not increase in respect to that of Fig. 8, although the time is noticeably reduced. The number of iterations could be clearly limited to something between 5–10.The algorithm devised for stopping the iteration around the optimum accuracy was designed for F=64 and does not seem to be valid here. Figure 10 shows results of its application as a discontinuous line, indicating larger errors.

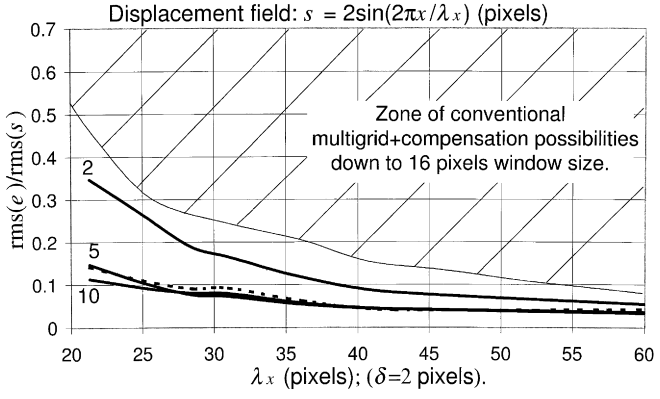

Fig. 9. Behaviour of LFC-PIV scheme with F=32 as a function of wavelength and for several numbers of iterations applied to type1 images. Table 2 gives time performance

Table 2. Correspondence of time to number of iterations for Figs. 9 and 10

Iteration number	2	5	10	20	40	80
Time units/vector	42	80	122	195	335	–

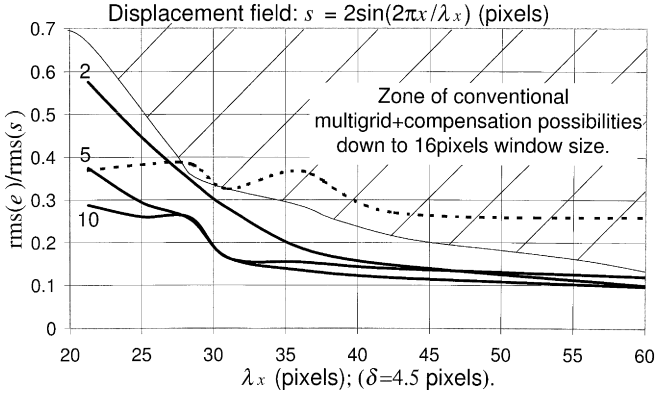

Fig. 10. Behaviour of basic LFC-PIV scheme with F=32 as a function of wavelength and for several numbers of iterations applied to type 2 images

6 Coherence between results from synthetic and real images

The exact displacement field is unknown in real images. This precludes direct evaluation of the errors. However, these images allow us to study the coherence between the results from synthetic images and those from real images.

The real PIV images chosen correspond to the disperse phase motion in a two-dimensional fluidised bed under atmospheric operation in the bubbling regime. These images are very similar to the synthetic ones of type 2. The particular case depicted in Fig. 11 corresponds to the explosion of an air bubble over the surface of the bed, shown in black, projecting disperse phase to the freeboard, shown in grey. Diffuse back illumination has been used in this experiment.

In the image, a black square limits a measurement zone of 144 by 144 pixels. The magnification equals 16 pixels/mm, and the time between images is 4 ms. The disperse phase is formed by 0.2-mm natural silica particles. They do not act as air tracers; on the contrary, in this experiment the disperse phase velocity is being tracked. The fluidising air velocity is 10 cm/s.

As there is a strong effect from gravity, the mean velocity has been subtracted in order to more clearly depict the velocity gradients. Figure 12 offers a plot of the flow field measured on the squared zone of Fig. 11. Three-dimensional motion is evident, thus indicating a loss of particles.

Design of the coherence test is as follows. According to the results of Sect. 3, the most accurate result corresponds to processing with double initial smoothing and double weighting based on F=64 pixels. At about 40 iterations, the measurement would be good enough to take it as a reference. The deviation between this vector field and those calculated with other setups serves as an approximation of the error. It is marked with an asterisk rms(e)*. In this case the reference displacement field shows rms(s)=1.8 pixels, only slightly different from the amplitude chosen for the synthetic cases. Figures 13 and 14 depict the results of this evaluation.

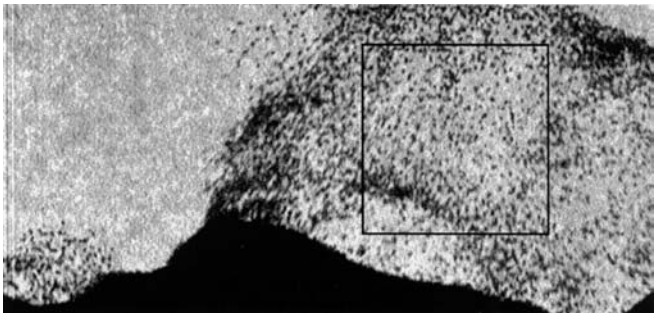

Fig. 11. PIV image of the disperse phase of an exploding bubble at the surface of a two-dimensional fluidised bed

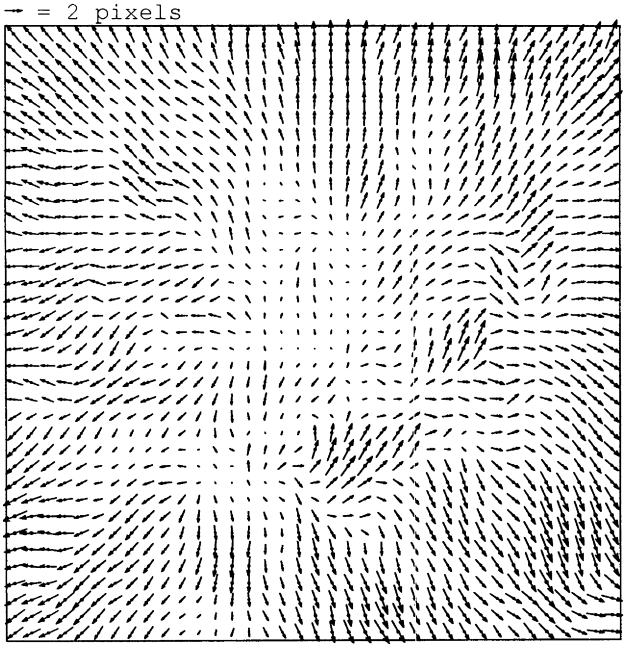

Fig. 12. LFC-PIV measurement corresponding to the squared zone in Fig. 11 (Δ=4 pixels). Mean velocity has been subtracted

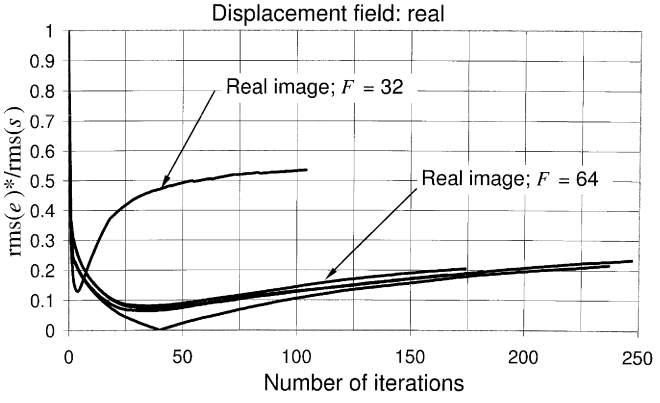

Fig. 13. Behaviour of different LFC-PIV schemes on the real image of Fig. 11

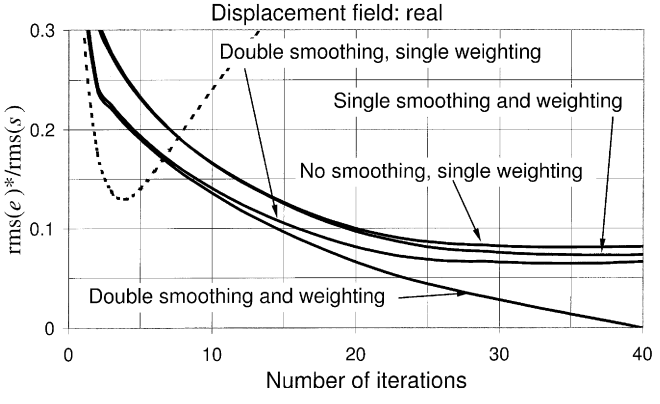

Fig. 14. Close-up zone of the graph in Fig. 13. The case with *F*=32 has been changed to a *dashed line*

Figure 13 shows that the LFC-PIV schemes based on *F*=64 pixels quickly approach the reference measurement during the initial iterations, separating slowly for the following ones. This highlights the reliability of the method. It also shows that the case of *F*=32 should be limited to just a few iterations. This coincides with the observations previously mentioned.

Figure 14 shows that the difference between the reference case and the remaining ones with *F*=64 pixels increases when single weighting is performed and when the initial smoothing iterations are removed. Both figures show results that are consistent with those obtained with synthetic images.

7 Conclusions

LFC-PIV successfully handles the dynamic range and spatial resolution limitations of conventional correlation PIV because of the built-in iterative correction of the average displacement of the interrogation window and the compensation of the distortion caused by the local displacements within the window. The subpixel correction further contributes to higher accuracy of LFC-PIV than conventional PIV, also reducing the impact of peak locking.

The use of large windows in conjunction with the compensation for particle pattern deformation and the application of a proprietary weighting window allows for a high robustness. The size and separation of the particle images can be very small, thus offering the potential for high spatial resolution. This allows including more information on the images, which is of paramount importance for the limited resolution of CCD cameras.

This advanced algorithm increases the computing time to reach its full potential; thus a study of the loss when computing time is reduced has been undertaken. The results obtained in this study offer the possibility of prematurely stopping the iterations with the confidence that intermediate results do not contain any serious limitation caused by the appearance of instabilities or lack of spatial resolution.

The balance between time and accuracy performance for LFC-PIV schemes has been analysed. For this purpose the response to sinusoidal displacement fields has been used as a test, considering both clean images and images with realistic noise content. Although this methodology does not include the full richness of real flows, it has been established as a valid tool in PIV development.

Tests included several alternatives in the design of the method. The outcome can be summarised as follows:

- The difference in accuracy between 64- and 32-pixel windows is not crucial for clean images, but the former performs more robustly with noisy images.
- LFC-PIV with a small number of iterations delivers results that are comparable in accuracy with those obtained with multigrid PIV.
- A large number of iterations improves accuracy beyond contemporary multigrid PIV, especially for small wavelengths.
- Windows of 64 pixels per side offer the best compromise when an initial double smoothing is implemented, which includes the use of a 32-pixel window.

- Double weighting offers a marginal increase in accuracy at the expense of doubling the computing time.

The time performance assessment of LFC-PIV method was completed with an experimental contrast using images from a two-phase flow. The conclusions from this test corroborate the general conclusions obtained from the synthetic images.

The window weighting effect does not imply a severe reduction of the effective window size, as the robustness of the method demonstrates (more details can be found in Nogueira et al. 2001). The use of large windows offers a distinctive robustness, in contrast to the much smaller ones used in the final steps of multigrid methods. The impression is that such large windows effectively guide the processing towards increasing the signal-to-noise ratio; however, the trade-off is a longer computing time. Besides this, the error introduced by the weighting function limits its application to large windows. The alternatives presented offer the possibility of stopping the iterations without losing confidence in the results, making LFC-PIV a very flexible tool.

References

Cowen EA, Monismith SG (1997) A hybrid digital particle tracking velocimetry technique. Exp Fluids 22:199–211

Hart DP (2000) Super-resolution PIV by recursive local-correlation. J Visualiz 3:187–194

Huang HT, Fiedler HE, Wang JJ (1993) Limitation and improvement of PIV. 2. Particle image distortion, a novel technique. Exp Fluids 15:263–273

Jambunathan K, Ju XY, Dobbins BN, Ashforth-Frost S (1995) An improved cross-correlation technique for particle image velocimetry. Meas Sci Technol 6:507–514

Keane RD, Adrian RJ, Zhang Y (1995) Super-resolution particle-imaging velocimetry. Meas Sci Technol 6:754–768

Lecordier B, Lecordier JC, Trinité M (1999) Iterative sub-pixel algorithm for the cross-correlation PIV measurements. In: Proc 3rd Int Workshop on PIV'99, Santa Barbara, USA

Lecuona A, Nogueira J, Rodríguez PA (2001) Proposals on the design of 2D multigrid PIV systems: application of dedicated weighting functions and symmetric direct correlation. In: Proc 4th Int Workshop on PIV'01, Göttingen, Germany

Nogueira J, Lecuona A, Rodríguez PA (1999) Local field correction PIV: on the increase of accuracy of digital PIV systems. Exp Fluids 27:107–116

Nogueira J, Lecuona A, Rodríguez P A (2001) Local field correction PIV, implemented by means of simple algorithms and multigrid versions. Meas Sci Technol 12:1911–1921

Rodríguez PA, Lecuona A, Nogueira J (2001) Modification of the local field correction PIV technique to allow its implementation by means of simple algorithms. In: Proc 4th Int Workshop on PIV'01, Göttingen, Germany

Scarano F, Riethmuller M (1999) Iterative multigrid approach in PIV image processing with discrete window off-set. Exp Fluids 26:513–523

Soria J (1996) An investigation of the near wake of a circular cylinder using a video-based digital cross-correlation particle image velocimetry technique. Exp Thermal Fluid Sci 12:221–233

Westerweel J (1998) Effect of sensor geometry on the performance of PIV. In: Proc 9th Int Symp on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal