

# A parallel digital processor system for particle image velocimetry

C D Meinhart†, A K Prasad‡ and R J Adrian†

† Department of Theoretical and Applied Mechanics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

‡ Department of Mechanical Engineering, University of Delaware, Newark, DE 19716, USA

Received 11 September 1992, in final form 23 December 1992, accepted for publication 1 February 1993

**Abstract.** A high-speed, high-resolution, digital computer-based interrogation system for particle image velocimetry (PIV) has been developed. It utilizes advanced parallel array processing technology to achieve computational speeds comparable to those of current supercomputers, and high-resolution image acquisition equipment to achieve the data input rates necessary for high-speed analysis of PIV images. Advanced cross correlation techniques are incorporated into the algorithm, which take advantage of the high-speed capabilities and improve the quality of the measurements. The architecture, algorithms, and performance characteristics of the system are described.

## 1. Introduction

### 1.1. The requirement for speed

Interrogation systems for PIV are currently able to analyse image fields, either photographic or video-graphic, at rates that make analysis of high image density recordings using the Young fringe method or the auto-correlation method lie between 0.2 vectors per second and 5 vectors per second, depending upon the pixel array size of the interrogation spot and the speed of the computer. One vector per second is a fair number for reference. While these rates are satisfactory for many purposes, PIV will not realize its full potential for contribution to fluid mechanics until the interrogation rate is improved by several orders of magnitude. To justify this assertion, consider that interrogation of a  $100 \times 100$  field of vectors requires around 3 h at one vector per second.

There are three situations in which one wishes to analyse hundreds, or even thousands, of such fields: cinematic PIV, holographic PIV and accumulation of a data base of two-dimensional fields that is large enough to provide statistically stable ensemble averages. In cinematic PIV one can readily imagine acquiring of the order of 1000 frames of PIV images, each containing  $10^3$ - $10^4$  vectors. In holographic PIV, it appears feasible to acquire 100 planes of data, each plane containing  $10^4$  vectors. Lastly, for the purposes of measuring flow statistics in turbulence, the minimum number of two-dimensional fields needed is of the order of 100 to 1000. In this situation, and in the case of cinematic PIV and holographic PIV, the total number of vectors is of the order of  $10^6$ - $10^7$ , corresponding to between 10 days and 100 days at one vector per second.

The purpose of this paper is to describe a new computer system for PIV interrogation that implements parallel processing technology to achieve speeds of the order of 100 vectors per second using a robust interrogation algorithm. Here, the parallel processing strategy involves using a small number of relatively powerful computers, as opposed to a strategy in which a very large number of relatively simple processors are used. The system is based on eight processors whose central component is an Intel i860 vector processor operating at 80 Mflops. The total peak speed of this system is 640 Mflops, although it can, in principle be expanded to 32 processors for a peak speed of 2.56 Gflops.

To make full use of such high-speed processing power, it is necessary to design the total system so that the major elements are mutually compatible. The most important elements are the data acquisition device, the image array sizes, the data communications from the external device to the parallel processors and between parallel processors, the algorithm used for analysis, and the execution speed of the algorithm on each processor. We have used FFT routines which are highly vectorized and therefore efficiently executed on the i860 processors, and we have developed a pipelined process which parses image data to each processor in such a way as to keep each processor operating close to its peak efficiency. Specifically, image data are acquired in a  $1024 \times 1024$  video camera array which is divided into eight sub-arrays that are passed to each of the eight array processors. Communication requirements between array processors are minimized, and the use of a large camera array makes the system suitable for direct video capture of the images (for example 'digital' PIV). Further details are presented next.

1.2. Cross correlation versus Autocorrelation

Previous work by Keane and Adrian (1992) suggests that double-exposure single-frame *cross correlation* offers many advantages over conventional double-exposure single-frame *autocorrelation*. Following the notation and theory in Adrian (1988) and Keane and Adrian (1992), the cross correlation function  $C(s)$  of two particle image fields  $I_1(X)$  and  $I_2(X)$  is defined by

$$C(s) = \int I_1(X)I_2(X+s) dX \quad (1)$$

where  $X$  is the spatial variable on which the image intensities are defined and  $s$  is the correlation separation vector.  $C(s)$  is defined as a single-frame cross correlation function when the two particle image fields  $I_1(X)$  and  $I_2(X)$  are subsets of a larger single-image field. The autocorrelation function, denoted by  $R(s)$ , refers to the case in which the two particle image fields  $I_1(X)$  and  $I_2(X)$  are identical. Single-frame cross correlation is computationally more expensive than autocorrelation: it requires three FFTs versus two FFTs for autocorrelation.

The autocorrelation function,  $R(s)$ , consists of five components,

$$R(s) = R_C(s) + R_F(s) + R_{D^+}(s) + R_{D^-}(s) + R_p(s) \quad (2)$$

Here,  $R_C(s)$  is the convolution of the mean image intensity,  $R_F(s)$  is fluctuating noise due to random correlation of non-paired particle images,  $R_{D^+}(s)$  and  $R_{D^-}(s)$  are the positive and negative displacement peaks due to correlation between image pairs, and  $R_p(s)$  is the self-correlation peak due to images correlating with themselves. Single-frame cross correlation significantly improves the signal-to-noise ratio over autocorrelation by (1) increasing the positive displacement peak without significantly changing the fluctuating noise, (2) reducing the self-correlation peak, and (3) reducing the negative displacement peak (Keane and Adrian 1992).

Figure 1(a) shows a digitized (128 × 96 pixel) double-pulsed particle image field of poor quality. The two-dimensional autocorrelation of this image field is shown in figure 1(b). The lack of particle image pairs in figure 1(a) leads to a weak correlation in which the displacement peak  $R_{D^+}(s)$  is barely distinguishable from the noise peaks. By using single-frame cross correlation instead of autocorrelation, the signal strength is significantly enhanced (Keane and Adrian 1992). Figure 2(a) shows the particle image field in figure 1(a), with the addition of a second correlation window offset 36 pixels to the right of the first correlation window. Figure 2(b) is the cross correlation of the image contained in window 1 with the image in window 2 of figure 2(a). The correlation function in figure 2(b) has a much stronger signal than that in figure 1(b). Clearly, interrogations of image fields with marginal particle pairing benefit from single-frame cross correlation.

Single-frame cross correlation is most successful when the window parameters are chosen such that (1) the relative offset between the first and second correlation window roughly equals the local mean particle displacement, (2) the first correlation window is large enough to enclose at least 10 first exposure particle images, and (3) the second window is large enough to capture all the second images of particles which have first particle images in the first window. This allows for efficient in-plane particle pairing, which increases the signal-to-noise ratio and eliminates the in-plane velocity gradient *bias* (Keane and Adrian 1992).

Asymmetric correlation windows are used to increase spatial resolution in one particular direction, while maintaining an adequate number of particles in the window. By choosing the first window to be asymmetric, say 128 × 64 pixels, twice the resolution is achieved in one direction than in the other. The size of the second window is not critical, assuming that it is large enough

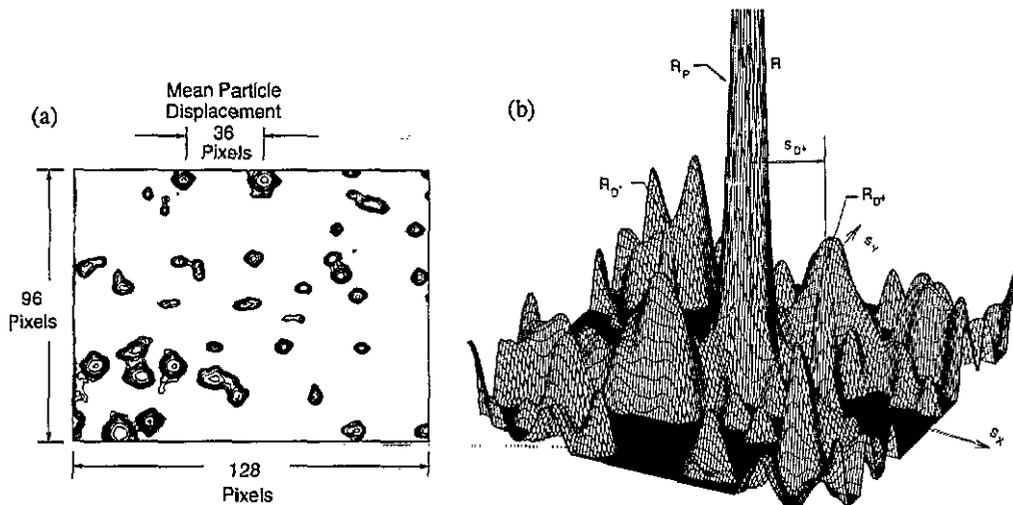
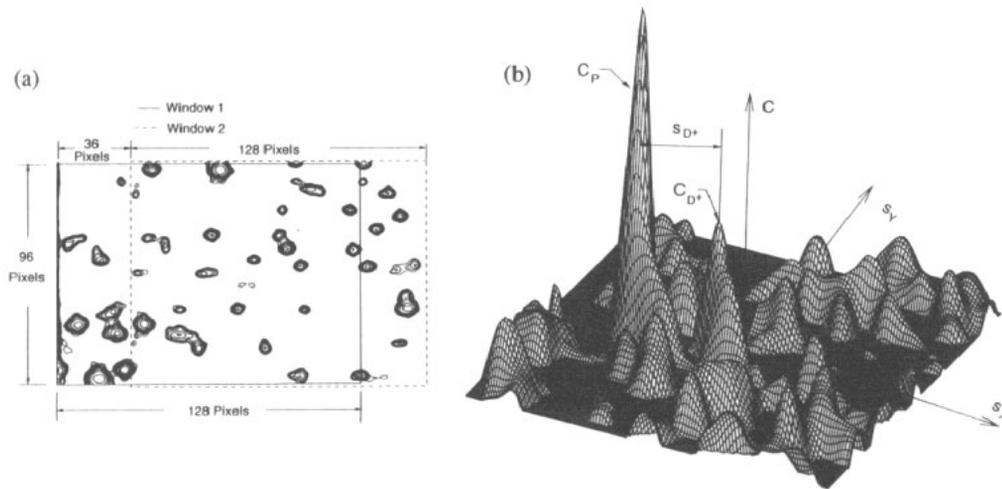


Figure 1. Autocorrelation of double-pulsed particle images, with marginal particle pairing: (a) particle images, (b) autocorrelation function,  $R(s)$ .



**Figure 2.** Cross correlation of double-pulsed particle images, with marginal particle pairing: (a) particle images, (b) cross correlation function,  $C(\mathbf{s})$ .

to capture all the second particles which have pairs in the first window.

### 1.3. Pixel resolution

*Bias* errors occur when the particle images are not adequately resolved by the CCD camera. Prasad *et al* (1992) state that in the case of photographic PIV, the bias error is negligible when the ratio between the particle image diameter,  $d\tau$ , and the CCD camera's pixel centre-to-centre distance (when projected back onto the photograph),  $d_{\text{pix}}$ , is  $d\tau/d_{\text{pix}} \approx 4$ . Increasing the pixel resolution will greatly increase computational time, without significantly reducing the *bias* error.

Typically, the particle image diameter (on photographs) is  $30 \mu\text{m}$ , which suggests an optimal value of  $d_{\text{pix}} \approx 8 \mu\text{m}$ . A typical correlation window of  $1 \text{ mm} \times 1 \text{ mm}$  ( $d_1 \approx 1 \text{ mm}$ ), implies an interrogation resolution of  $128 \times 128$  pixels ( $d_1/d_{\text{pix}} \approx 2^7 = 128$ ). Therefore, the present system has been designed to correlate image fields with resolutions between  $64 \times 64$  and  $128 \times 128$  pixels.

## 2. System hardware

The particle image velocimeter's acquisition and processing hardware consists of five primary components: (1) Videk Megaplus CCD camera, (2) Imaging Technology Model VSI-150 frame grabber, (3) two Mercury Computer Systems Model MC860VS array processing boards, (4) SUN SPARCstation Model 370 host computer and (5) Unidex X-Y-Z translation stage. The system's hardware components are schematically displayed in figure 3.

### 2.1. The CCD camera

A Videk Megaplus CCD camera provides  $1320 \times 1035$  pixels of resolution. The pixels are square

( $6.8 \mu\text{m} \times 6.8 \mu\text{m}$ ) with a fill ratio of unity. The camera provides either an analogue or an eight-bit digital signal to the frame grabber. The camera acquires up to 6.9 frames per second.

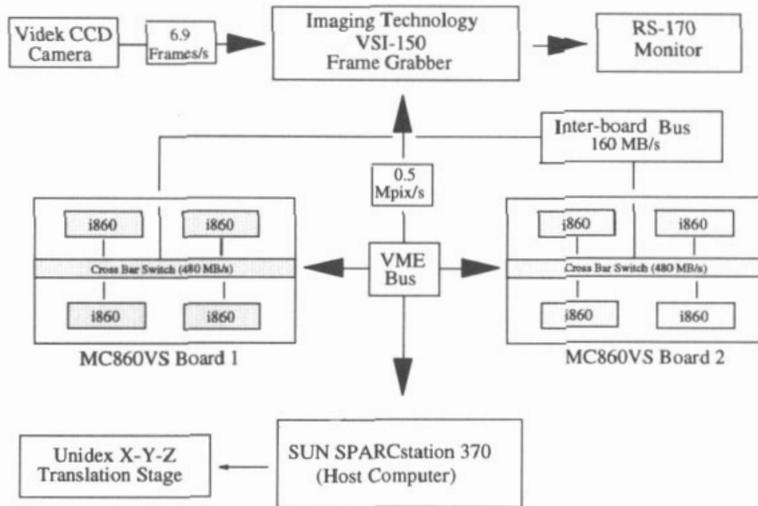
### 2.2. Frame grabber

An Imaging Technology Model VSI-150 frame grabber board digitizes the analogue signal sent from the Megaplus camera at a frequency of 20 MHz. The VSI-150 is equipped with double input buffers ( $1024 \times 1024 \times 16$ -bit) and includes an RS-170 display controller. (Only  $1024 \times 1024$  pixels out of the  $1320 \times 1035$  pixel Videk output are used.)

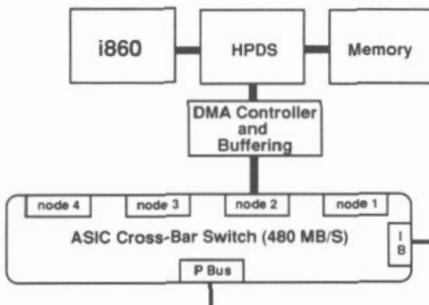
The frame grabber's architecture allows only one pixel in video memory to be read and transferred over the VME bus at any one time. Furthermore, only one array processor can map into its memory at a time. Thus, only 0.5 Mpixels can be transferred to an array processor per second. This low rate may or may not be a bottleneck, depending upon the computation time of the algorithm used on the i860s. A faster frame grabber is desirable for general purpose use.

### 2.3. Array processors

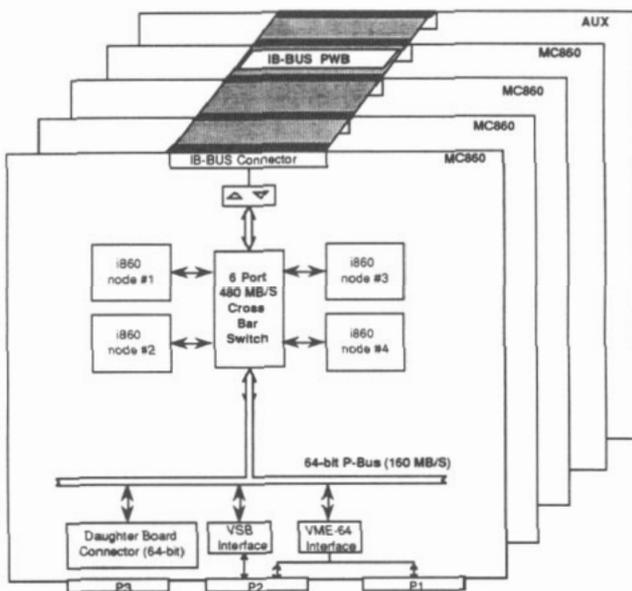
**2.3.1. MC860VS boards** The velocimeter utilizes two MC860VS boards for its high-speed array processing. Each board occupies a single 9U VME slot and mounts a cross bar switch and four i860 nodes. Each i860 node contains an Intel i860 microprocessor, a DMA controller, and 2 MB of DRAM memory (upgradeable to 16 MB), see figure 4. All four i860 nodes on a single MC860VS board communicate with each other at  $480 \text{ MB s}^{-1}$  via the cross bar switch. The cross bar switch is a six-port communication device which supports up to three channels of communication. Figure 5 shows four MC860VS boards being directly linked via the Inter-board bus



**Figure 3.** Schematic layout of interrogation hardware.



**Figure 4.** A single node on an MC860VS board (courtesy Mercury Computer Systems, Inc).



**Figure 5.** Parallel configuration of four MC860VS boards, each containing four i860 processors (courtesy Mercury Computer Systems, Inc).

(IBbus)<sup>†</sup> with a data transfer rate of  $160 \text{ MB s}^{-1}$ . Communication outside the boards is facilitated by the VME bus (Mercury Computer Systems 1991).

**2.3.2. i860 micro processor** The i860 is the processing element of each i860 node. It runs at a 25 ns (40 MHz) clock cycle and has nine functional units: (1) core execution unit, (2) floating-point control unit, (3) floating-point adder unit, (4) floating-point multiplier unit, (5) graphics control unit, (6) paging control unit, (7) instruction cache, (8) data cache and (9) bus and cache control unit (see figure 6). The i860 is capable of performing high-speed vector operations up to 80 Mflops per processor.

The core execution unit controls the overall operation of the i860 under an RISC core environment. The parallel ‘fetch/execute’ architecture enables instructional <sup>†</sup> Inter-board bus was under development at the time of publication.

Instruction Cache	Memory Management Unit	Data Cache
Bus Control Unit	RISC Core	Floating Point Control Unit
Integer Unit	FP Adder Unit	FP Multiplier Unit

- Single Package
- 40 MIPS
- 80 SP MFLOPS
- 60 DP MFLOPS
- 640 MB/S Internal
- 160 MB/S External
- Separate address and math units
- MMU
- 4 Kbyte I Cache
- 8 Kbyte D Cache
- 64-bit 3D Graphics Unit
- 128-bit data bus
- 64-bit instruction bus
- IEEE 754-1985 floating point

**Figure 6.** Specifications of the Intel i860 processor (courtesy Mercury Computer Systems, Inc).

execution every clock cycle, thus providing efficient use of the i860.

The floating-point control unit executes all floating-point instructions received from the core execution unit and maintains control of all the floating-point units (floating-point adder, multiplier and graphics units). The control unit also transfers data to the floating-point registers and subsequently from the registers to the adder, multiplier and graphics units. The instruction bus from the instruction cache allows both the floating-point control unit and the core execution unit to operate in parallel. In addition, the floating-point control unit has dual instruction capabilities for commanding both the adder and multiplier to run in parallel, generating two results every clock cycle.

The floating-point adder and multiplier units, using single (32-bit) precision, operate on a three-stage pipeline. One single-precision addition or multiplication is processed in three clock cycles. When a large series of floating-point operations are performed, the three-stage pipeline can perform three operations simultaneously. Thus, after the first two clock cycles of a series of operations are completed, a single result is produced every clock cycle. When the adder and multiplier operate in parallel, two results can be produced every clock cycle (Mercury Computer Systems 1991).

When all eight (four i860s on two MC860VS boards) i860 microprocessors operate in parallel, they have a combined rating of 640 Mflops (single precision). Each i860 microprocessor performs a  $128 \times 128$  complex two-dimensional FFT in 17.3 ms (real-to-complex FFTs in 8.65 ms). A  $128 \times 128$  single-frame cross correlation is calculated in 50 ms on one i860 microprocessor.

#### 2.4. Host computer

A SUN SPARCstation Model 370 computer, equipped with a 688 MB hard disk and 150 MB cartridge tape drive, serves as the host computer for the PIV system. The host computer houses the VSI-150 frame grabber board and the two MC860VS processor boards.

The host computer coordinates the camera, the frame grabber, the array processor boards, and the  $X-Y-Z$

translation stage. The stage traverses the PIV photograph with respect to the camera or (for the three-dimensional case) the CCD camera with respect to the PIV hologram.

### 3. System software

#### 3.1. Interrogation algorithm

During interrogation, the velocimeter digitizes a large ( $1024 \times 1024$  pixel) particle image field. This image, referred to as a 'camera frame', is divided into eight smaller sub-images called 'processor images'. The processor images are approximately  $512 \times 256$  pixels and usually overlap each other. A single processor image contains the pixels which are to be transferred to a particular array processor. Once a processor image is transferred to its array processor, it is further partitioned into several 'correlation blocks'. Each correlation block contains the necessary pixels required for a single cross correlation (that is, it encloses both cross correlation windows 1 and 2). Figure 7 shows the complete hierarchy of particle image data sets used in the algorithm.

The digital interrogation of a two-dimensional PIV photograph requires photographed particle images to be projected directly onto the CCD camera array by a uniform intensity incoherent white light source. The incoherent light source provides high contrasting images without the interference patterns caused by coherent light sources. The analogue image signal is sent from the CCD camera to the frame grabber, where it is digitized and stored as a  $1024 \times 1024$  pixel image in video memory.

A signal is then sent through the SUN's RS-232 serial port to move the  $X-Y-Z$  translation stage, so that another section of the photograph can be digitized. Simultaneously, the digitized image in the frame grabber's video memory is partitioned into the eight processor images; each processor image is sent via the VME bus to one of the eight i860 microprocessors for interrogation. During this transfer, only one i860 at a time can map into the frame grabber's memory and transfer its appropriate processor image. This often causes other

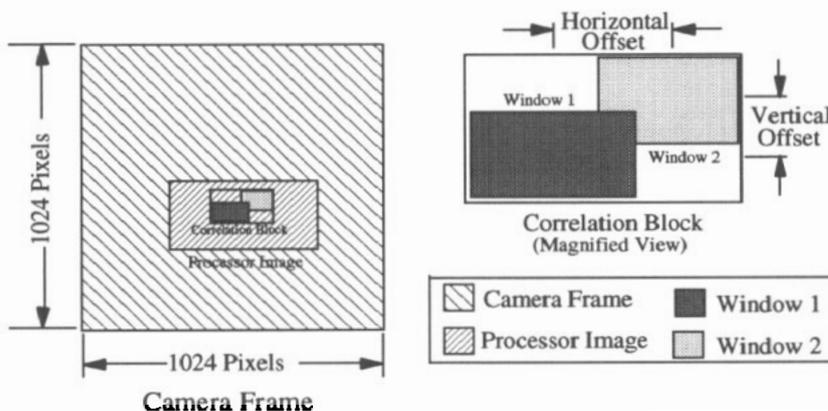


Figure 7. Hierarchy of particle image data sets.

i860s to remain idle until they transfer data and begin processing, which, in turn, causes system inefficiency.

Once each i860 receives its processor image, it further divides its image into several interrogation spots called correlation blocks. Typically there are between 8 and 100 correlation blocks in one processor image, depending upon the cross correlation windows' sizes and relative offsets. Neighbouring correlation blocks typically contain common pixels. The second window of one correlation block often overlaps with the first window of the adjacent correlation block.

During the interrogation, each i860 analyses individually every correlation block in its processor image. It transfers the particle images contained in windows 1 and 2 (of the correlation block in question) to two  $128 \times 128$  correlation buffers. When a window is smaller than  $128 \times 128$ , zero padding is used to completely fill the correlation buffer. The single-frame cross correlation function,  $C(s)$ , is then calculated by correlating the two buffers together.

The cross correlation function's positive displacement peak  $C_D+(s)$ , is located by (1) zeroing a predefined number of pixels around the self-correlation peak  $C_P(s)$ , (2) searching the correlation plane, on an *a priori* determined side of the self-correlation peak  $C_P(s)$  for the remaining three highest peaks, and (3) curve-fitting a parabola to the peak pixel value and its adjacent pixel values to obtain the peak locations with sub-pixel accuracy. This process is repeated until all the correlation blocks in the processor's image have been interrogated.

After the first i860 has interrogated all of its correlation blocks, the results are sent to the SUN computer, reassembled, and written to a binary file. The SUN computer then estimates the median particle displacement of all the valid measurements in the processor's image and sets the new relative window offset approximately equal to that displacement. (This allows the velocimeter to adapt continually to the changing flow field for better quality measurements.) By this time, the X-Y-Z translation stage has moved to a new position, and a new image has been digitized. Then, the first i860 receives its new processor image and begins a new interrogation. All of the remaining i860s (one at a time) repeat this process and begin their new interrogations. The SUN then sends another signal to move the X-Y-Z translation stage to the next region of the photograph. This interrogation cycle is repeated until the photograph is finished.

### 3.2. Adaptive windowing techniques

The velocimeter uses adaptive windowing techniques to optimize its performance. More specifically, it takes advantage of the benefits gained from using single-frame cross correlation instead of conventional autocorrelation. The velocimeter can be 'tuned' for a desired performance by selecting the proper shape, size, and relative offset of the correlation windows. For example, the velocimeter can be tuned for high spatial resolution

by decreasing window size or for high signal-to-noise ratio by increasing window size.

Many flow fields, such as near-wall shear flows, which have two or more velocity scales, require the velocimeter to have a large dynamic velocity range. When using autocorrelation, particle displacement should not be more than 30% of the interrogation window size ( $d_1$ ). However, by continually varying the relative window offset during the interrogation, this velocimeter can achieve dynamic ranges several times larger than the window size ( $d_1$ ). Similarly, if image shifting is used to eliminate velocity ambiguity, the effect of the image shift on the dynamic range can be completely eliminated by offsetting the second window by an amount equal to the image shift.

## 4. System performance

### 4.1. Processing speed

The velocimeter is efficient when all eight i860 array processors spend large percentages of time interrogating, instead of transferring images from the frame grabber or being idle. Since the frame grabber-to-i860 transfer time is fixed, efficiency is increased solely by increasing the number of correlation blocks (number of vectors) per processor (namely decreasing the window sizes and increasing their overlap). Figures 8(a) and (b) show the

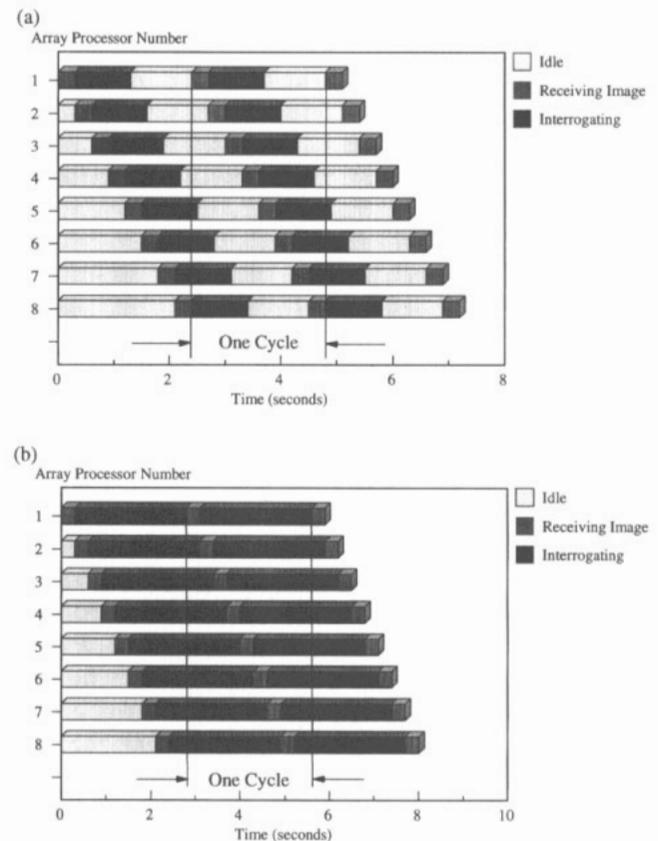


Figure 8. Array processor time-line: (a) 20 vectors per processor image, (b) 50 vectors per processor image.

activities of the eight array processors during a typical interrogation cycle as functions of time, when the number of vectors per processor image is 20 and 50, respectively. When there are only 20 vectors per processor image (see figure 8(a), each array processor spends one-half of its time idle, and the velocimeter only calculates 50 vectors per second. The efficiency is greatly increased when there are 50 vectors per processor image, see figure 8(b). Clearly, the i860s spend a larger percentage of time interrogating, and the velocimeter calculates 100 vectors per second.

If system overhead could be completely removed, the velocimeter would operate at a theoretical peak performance of 160 vectors per second using cross correlation. Because autocorrelation only requires two FFTs, the theoretical peak performance for autocorrelation is 200 vectors per second.

Arnold and Hinsch (1988) developed a holographic technique for optical parallel processing of PIV images. By using a holographically generated two-dimensional array of interrogation spots, multiple Young fringe patterns can be produced in parallel. Theoretically, this technique could be incorporated into our system by projecting multiple Young fringe patterns onto the CCD camera. The Young fringe patterns could then be digitally analysed in parallel by the i860 microprocessors. Because only one FFT would be required per correlation, the theoretical peak performance would be increased to 266 vectors per second.

In reality, the increases in theoretical peak performance gained by using autocorrelation and the Young fringe method instead of cross correlation are offset by system overhead caused by the frame grabber-to-i860 data transfer rate. Thus, the actual performance of the velocimeter would not be substantially increased by using autocorrelation or the Young fringe method instead of cross correlation. However, the benefits of increased flexibility and quality of measurement gained by using digital cross correlation make it the superior method of interrogation for this velocimeter.

#### 4.2. Application to a pipe flow experiment

The velocimeter's overall performance was tested by interrogating high-image-density double-pulsed photographs† taken in pipe flow with a Reynolds number  $Re_D = 50\,000$ . A light sheet spanning the pipe's radial streamwise plane illuminated  $1\ \mu\text{m}$  olive oil particles which were imaged onto  $4 \times 5$  inch photographic film with a magnification factor of six. The correlation windows for this flow field were chosen to be asymmetric  $128 \times 96$  pixels ( $1.6\ \text{mm} \times 1.2\ \text{mm}$  on the photograph), to provide higher spatial resolution in the wall-normal direction ( $y$ ) than in the streamwise direction ( $x$ ). After interrogation, the vector field was analysed to replace spurious vectors with alternatively measured vectors (the second or third correlation peaks, typically 5%), linearly interpolated vectors (typically 5%), or left blank (typi-

cally 5%). Figure 9(a) shows the instantaneous in-plane velocity vectors, with  $x$  and  $y$  being the streamwise and wall-normal spatial coordinates, respectively. Aberrations very close to the wall,  $y < 50\ \mu\text{m}$ , caused the particle images to blur, which made velocity measurements impossible. With the correlation windows overlapped 25% (which prevented aliasing of the velocity measurements), the wall-normal and streamwise spatial resolution was  $\Delta y = 160\ \mu\text{m}$  and  $\Delta x = 214\ \mu\text{m}$  respectively. By subtracting the line-averaged mean velocity from the total velocity, and increasing the velocity vector scale, the turbulence structure is revealed, see figure 9(b). The high spatial resolution ( $\Delta y = 160\ \mu\text{m}$ ) ensures that even the smallest scales of motion are resolved.

Significant out-of-plane loss of particle pairs, caused by out-of-plane motion in the pipe flow experiment, allowed only 85% of the velocity vectors to be correctly measured, during an interrogation using  $1.6\ \text{mm} \times 1.2\ \text{mm}$  (on the photograph) cross correlation windows. By keeping the interrogation spot size constant but using autocorrelation instead of cross correlation, the percentage of valid measurements decreased to 60%. When the interrogation spot size was increased to  $2.0\ \text{mm} \times 1.5\ \text{mm}$ , 85% of the vectors were measured correctly by autocorrelation.

#### 5. Conclusions

Recent developments in array processing technology along with adaptive cross correlation interrogation tech-

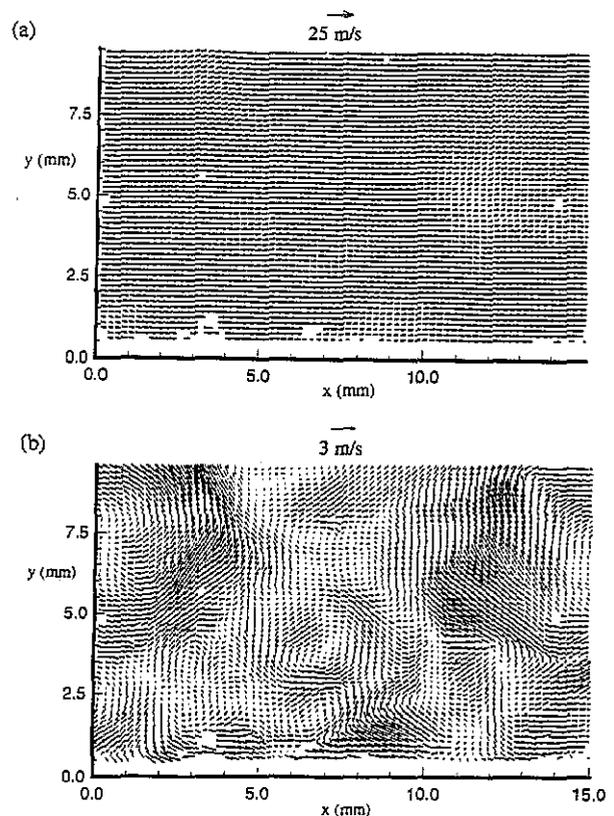


Figure 9. Instantaneous velocity vector field of pipe flow with  $Re_D = 50\,000$ : (a) total velocity, (b) fluctuating velocity.

† Photograph courtesy of T Urushihara, Nissan Motor Co.

niques have paved the way for a new generation of particle image velocimeters. The increase in speed and accuracy of this velocimeter makes it possible to process very large numbers of vectors.

### 5.1. Adaptive windowing techniques

Interrogation by single-frame cross correlation was chosen because previous work by Keane and Adrian (1992) shows that when analysing double-pulsed PIV image fields, single-frame cross correlation is superior to autocorrelation. It (1) increases the signal-to-noise ratio, (2) eliminates in-plane velocity gradient *bias*, and (3) increases the dynamic range of the velocity measurements. By selecting the window offset and size adaptively, it is possible to optimize the measurements. For example, high spatial resolution can be obtained in a preferred direction, while maintaining a necessary number of particle image pairs for strong signal strength by using asymmetric cross correlation windows. Flow fields, such as near-wall shear flows, which have large velocity gradients are ideal candidates for asymmetric windows because they require more spatial resolution in one direction than in the other. Using asymmetric ( $128 \times 96$  pixels) correlation windows, velocity vectors in pipe flow were resolved at  $160 \mu\text{m}$  increments in the wall-normal direction.

Conventionally, autocorrelation limited the dynamic range in which PIV could be used to measure velocities. With autocorrelation, the recommended maximum mean particle displacement was one-third of the interrogation window. Alternatively, single-frame cross correlation provides a larger dynamic range. By continuously varying the relative window offset to approximate the current median particle displacement, the velocimeter has a dynamic range of several windows. The velocimeter's digital windowing allows for quick and effective optimization of the interrogation parameters. For example, spatial resolution can be adjusted by merely inputting a different size correlation window into the host computer, thus avoiding the need to readjust or refocus optics.

### 5.2. System performance

The PIV interrogation system developed here, under normal operating conditions, can calculate over 100 vectors (correlations) per second. Thus, a 10000 vector double-pulsed two-dimensional photograph can be interrogated in less than 2 min, making it feasible to analyse large ensembles of photographs for the purpose of improving the statistical stability of averages. A  $35 \times 25$  vector field from a frame of 35 mm film can be analysed in less than 8 s, making cinematic PIV practical.

### Acknowledgment

This research was supported by the Office of Naval Research N00014-90-J-1415, and TSI Inc. The parallel computer system was acquired with the support of Department of Energy URI Equipment Grant DEFG05-89ER 75508. The authors would also like to thank Mr R D Keane for his helpful insight into the theory of single-frame cross correlation.

### References

- Adrian R J 1988 Statistical properties of particle image velocimetry measurements in turbulent flow *Laser Anemometry in Fluid Mechanics 3* (Lisbon: Instituto Superior Técnico) pp 115-29
- Arnold W and Hinsch K 1988 Purely optical parallel processing in particle image velocimetry and the study of flow structures *Proc. Conf. on Optical Methods in Flow and Particle Diagnostics, ICALEO '88* (Orlando, FL: Laser Institute of America) pp 1-9
- Keane R D and Adrian R J 1992 Theory of cross-correlation analysis of PIV images *J. Appl. Sci. Res.* **49** 191-215
- Mercury Computer Systems Inc 1991 *The MC860 Series - PRELIMINARY System Reference: Hardware, Configuration, & Instruction Set* Lowell, MA: Mercury Computer Systems, Inc.
- Prasad A K, Adrian R J, Landreth C C and Offutt P W 1992 Effect of resolution on the speed and accuracy of particle image velocimetry interrogation *Exp. Fluids* **13** 105-16