1. INTRODUCTION

1.1 Motivation

Application of optical diagnostics to the field of fluid mechanics has proven to be invaluable in advancing the state of the art. Optical diagnostics are ideally suited to examination of many fluid flows, offering the advantages of noninvasive probing of flow fields, obtaining full-field in additional to point measurements, utilizing computational-optics techniques like optical Fourier transforms, and applying the wealth of readily available, well-developed image-processing techniques. From the early days of flow visualization in transparent fluids, optical diagnostics have become increasingly quantitative and now include such widely accepted techniques as laser Doppler velocimetry (LDV) and holographic interferometry. However, there are still needs for advanced diagnostics, especially for measurements in multiphase and turbulent flows. These are very important flow categories, as they encompass most environmental and industrial flows. In addition, multiphase and turbulent flows are difficult to simulate computationally, so high-quality measurements are in great need, both for fundamental understanding of the flow physics and for developing and validating advanced computational techniques. The noninvasive nature of optical techniques is particularly important in turbulent and multiphase flows, where probes may interact with the turbulence to change the flow properties or may interact differently with different components of a multiphase flow.

1.2 Program Objectives

The goal of this work was to develop optical diagnostic techniques to provide Sandia with advanced flow-measurement capabilities. The first technique developed and evaluated is Particle Image Velocimetry (PIV), in which a flow seeded with small particles is photographed repeatedly at short time intervals. Image-processing techniques are applied to find particle velocities from particle positional changes in successive frames, which permits determination of the instantaneous flow field and associated turbulence quantities. Standard PIV uses a laser light sheet to illuminate a plane of the flow for 2D measurements, while volume illumination and multiple cameras are used for 3D PIV. The second technique developed and evaluated uses nonlinear photorefractive optical crystals to measure turbulence quantities more directly. A laser beam passed through the flow acquires a small perturbation signal related to the turbulence, which is superimposed on a large steady informationless signal, rendering interpretation difficult. The laser beam is then passed through one or more crystals to remove most of the steady portion, thereby greatly improving the signal-to-noise ratio. Since turbulence models are at best of questionable validity and often difficult to extend to new flow situations, this capability will enable us to assess and potentially improve the accuracy of numerical turbulence models by directly measuring turbulence quantities to compare with model predictions.

The flows of interest for application of these new diagnostics are typically turbulent, multiphase, and/or nonuniformly heated, making traditional point measurements difficult and unreliable. These new diagnostics provide unique capabilities to make such measurements in a noninvasive fashion.

2. PARTICLE IMAGE VELOCIMETRY (PIV)

2.1 Introduction

Particle Image Velocimetry (PIV) is used to measure two-dimensional velocity fields by analyzing the motion of flow-tracing seed particles, recorded in photographic or video images (e.g., Adrian, 1991). This technique promises to become one of the most useful in experimental fluid mechanics. In a PIV experiment, a flow seeded with small particles is photographed repeatedly at short time intervals, typically using a laser light sheet to illuminate particles only in the plane of interest. Image-processing techniques are applied to find particle velocities from particle positional changes in successive frames, which permits determination of the instantaneous flow field and associated turbulence quantities. This near-instantaneous, full-field velocity mapping complements such point velocity measurement techniques as laser Doppler velocimetry and hot-wire anemometry. Where these techniques are capable of providing detailed velocity data at a point in the flow, they can only be used to map out the velocity field in steady-state flows. However, full-field velocity data are acquired very quickly using the PIV technique, allowing velocity field mapping in unsteady flows. Quantities such as vorticity and strain rates in fluids are very difficult to determine using point measurement techniques since they involve derivatives of the velocity field, which require measurements of very fine spatial resolution. PIV measurements allow determination of such quantities since they are of sufficiently fine spatial resolution.

PIV is often referred to as a "quantitative flow visualization" technique since the data record is typically a high-quality photograph or video of the flow. The human eye is very good at tracking particles between frames of a multiframe sequence, or even performing the autocorrelation to determine particle motion in multiple-exposure images. Velocity quantification in PIV is done by automating this process, using a computer to measure the displacement of individual particles in the flow field, over a known time period, using one of several different algorithms (see, e.g., Lourenco et al. 1989; Grant 1994). Two different PIV algorithms will be presented in this chapter, one based on frame-to-frame cross-correlation, the other based on frame-to-frame particle tracking.

Among its other advantages, PIV is especially useful in multiphase flows, at least for those cases where the separate phases can be tracked individually. For example, in a gas-liquid two-phase flow, small solid flow-tracing seed particles could be used to track the liquid velocity, while bubbles could be directly tracked to give gas velocity (Hassan and Blanchat, 1991a). In turbulent flows, the fast data acquisition of PIV holds promise for measuring rapid fluctuations throughout a flow field. However, turbulence statistics must be built up by working with a large number of frames.

At the time that this LDRD work was undertaken, there were no commercial sources for PIV systems. At present, there are at least three companies offering PIV hardware, software, and/or full systems. They include TSI Inc. (St. Paul, MN), Dantec Inc. (Mahwah, NJ), and Oxford Lasers (Acton, MA).

2.2 Two-Dimensional PIV

One of the first decisions to be made when setting up a PIV system is the data acquisition method, usually either a film camera or an electronic (usually CCD) camera. The primary advantage of a film camera system is better resolution. However, there is a trade-off between resolution and acquisition speed. High-speed motion picture cameras typically operate with smaller film sizes (e.g., 16 mm format), so resolution is lowered. For higher-resolution film cameras, the film advance times limit the data acquisition speed for sequential images. For this reason, multiple exposures on a single film frame are often used. An autocorrelation is then used to determine particle motion, but there is a directional ambiguity unless special image-shifting techniques are applied. The use of film also involves the time-consuming step of darkroom processing before image quality can be determined. PIV analysis of multiple-exposure film images is sometimes done using the Young's fringe analysis, where a probe laser beam traverses the image (usually the photographic negative), and the spacing and orientation of the diffraction fringe pattern is measured to determine local velocity (magnitude and direction).

Electronic image capture is usually limited to 30 frames per second for a conventional RS-170 video signal, at lower resolution than available with film. The main advantage of an electronic camera is immediate feedback, very useful for determining illumination levels, camera settings, seed density, etc. Other advantages are that the picture can be directly stored as a binary computer file available for image processing and measurement, and that there is no directional ambiguity when a sequence of frames is used. Both film and CCD cameras have been used in this work, although primarily the latter.

Light sheets for planar illumination can be formed either by expanding a laser beam with cylindrical optics, or by rapidly sweeping a focused laser beam through the flow field during exposure, forming an effective light sheet. Both planar illumination techniques have been applied for this work.

Two methods have been developed to obtain the displacement of multiple particles between flow images. The first is the cross-correlation tracking (CCT) algorithm, and the second is the multi-frame particle tracking (MFT) algorithm. These two methods are used both for 2D PIV as discussed herein and for the 3D PIV discussed in Chapter 3 (programs PST3D and TP3D do cross-correlation and particle tracking, respectively, for both 2D and 3D PIV). The programs are written in Fortran, and are based on the graduate work of two of the authors (TKB and JAH) at Texas A&M University (cf. Blanchat, 1991; Hassan and Blanchat, 1991ab)

2.2.1. Cross-Correlation Algorithm (CCT)

Cross-correlation tracking (CCT) is a dynamic particle tracking method that can be performed between two sequential digital images. The particle velocity is found by determining the correspondence between particles in the comparison areas of the two images. Every particle belongs to a characteristic group that has a specific local 2D distribution pattern in the first image and multiple

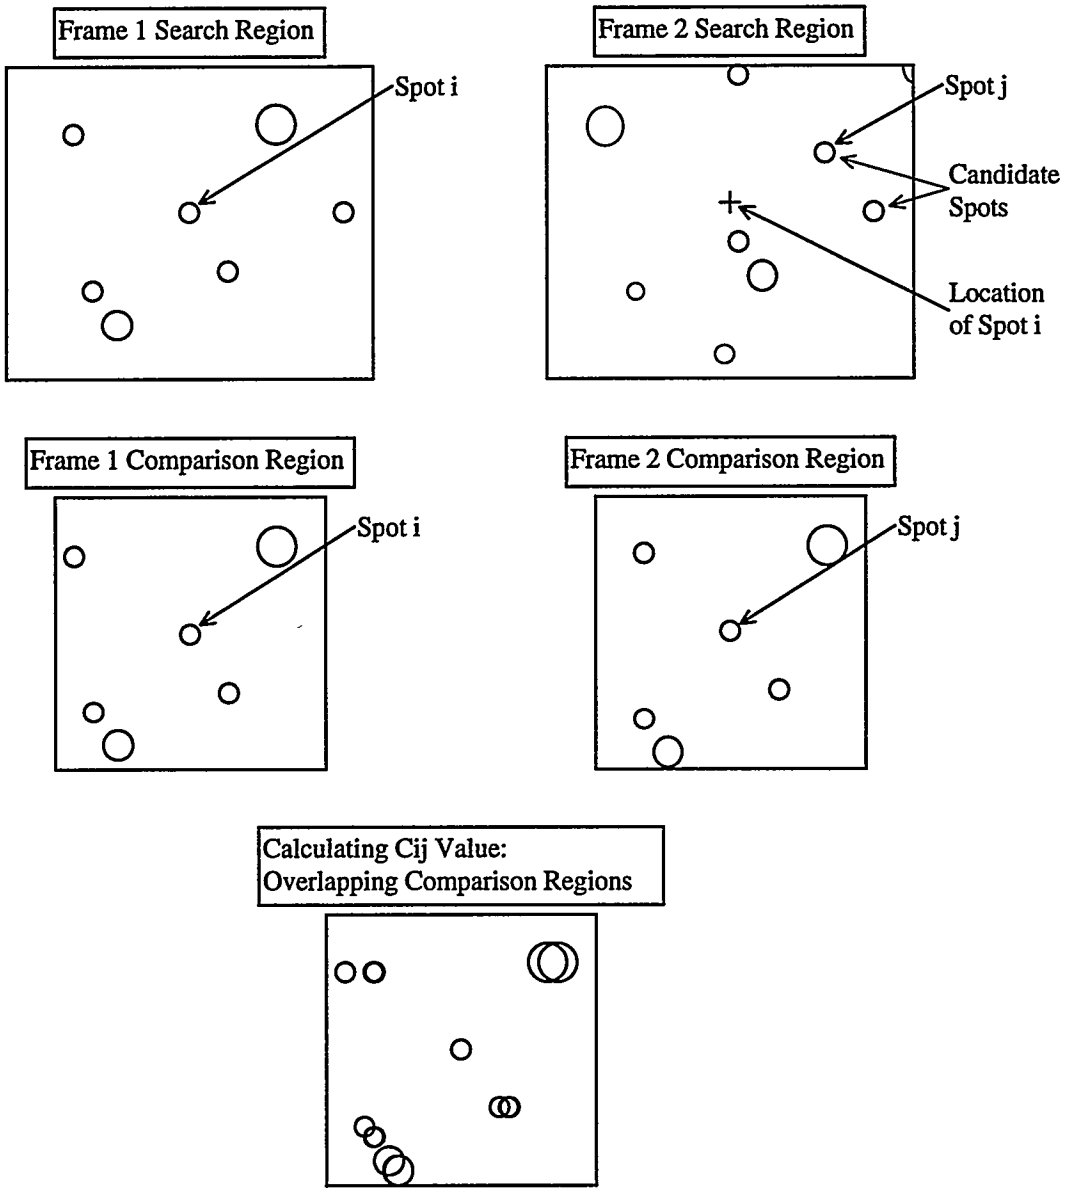

Figure 2.1. Calculating C_{ij} value for Cross Correlation Tracking (CCT) technique.

candidate patterns in the second image. Each second image candidate's comparison area is shifted over the first image comparison area such that the original particle's center is positioned over the candidate particle's center. The particle in the first image corresponds to the particle in the second image that maintains the most similar pattern within the comparison area, providing that the local pattern of the distributed particles changes little within the time that the original images were taken. This method is particularly useful when only two sequential images are available, and the multi-frame tracking method (see Section 2.2.2) cannot be performed.

The CCT algorithm correlates between two frames of data. A search area is formed in the second frame, centered on the location of the original particle in the first frame. The size of this area is determined by estimating the maximum possible movement for particles between frames. Each Frame 2 particle in this area is a candidate for being the original particle as moved by the flow. A comparison region is then created around each candidate, encompassing at least ten other particles in Frame 2. A comparison area of the same size is created in Frame 1, centered on the original Frame 1 particle. The two areas are then compared by shifting the center of the candidate particle over the center of the original particle. Figure 2.1 demonstrates this procedure. This comparison is characterized by the correlation coefficient, C_{ij}, which measures the overlap of the particles in the two areas. Figure 2.2 shows how the overlap of the diameters of a Frame 1 and Frame 2 spot is measured. A C_{ij} value of 1 indicates a perfect correspondence between two areas; the closer the value is to 1, the more likely that it is the correct particle match. A C_{ij} value of 1 is possible only if each spot is the same size in both comparison areas, and each spot is centered on the same point when shifted to overlap the two areas. Equation 1 is used to calculate the correlation coefficient between the comparison areas for frames 1 and 2:

$$C_{ij} = \frac{\sum_{i=1}^{P_1} \sum_{j=1}^{P_2} D_{i,1} \cap D_{j,2}}{\sqrt{\left(\sum_{i=1}^{P_1} D_{i,1}\right)\left(\sum_{j=1}^{P_2} D_{j,2}\right)}} \qquad (1)$$

where the numerator is the number of overlapping pixels between the two frames, and the denominator is the geometric mean number of total particle pixels in the two frames. The maximum value of C_{ij} is therefore equal to one. In other words, C_{ij} is simply the ratio of the overlapping "on" pixels (i.e., particle present) in the two frames to the total possible number (Blanchat, 1991). The particle pair with the highest C_{ij} value is then identified as the best particle pair match. Given a purely translational flow and a sufficiently small comparison region, C_{ij} values obtained for the correct particle match are much greater than those of incorrect matches. In areas of high vorticity, or if the comparison region is too large, the difference in C_{ij} values can be much smaller, and the identified matches are not always correct. Another means of checking is required to remove erroneous particle pairs. Therefore, two more checks are performed. The first check calculates the sum of the distances between all corresponding centroid pairs. For perfect overlap of all particles in the Frame 1 and Frame 2 search areas, this value will be zero.

The second check determines a reliability index for a possible pair based on the number of particles that overlap, N_{ij}, and the amount the diameters of those particles overlap, D_{ij}, which occur when checking the correspondence between spot i and j. This overlap was also used in calculating the C_{ij} value and is displayed in Figure 2.2. When all correspondences are completed, the pair reliability index R_{ij} is calculated as follows:

$$R_{ij} = N_{ij} \times D_{ij} \qquad (2)$$

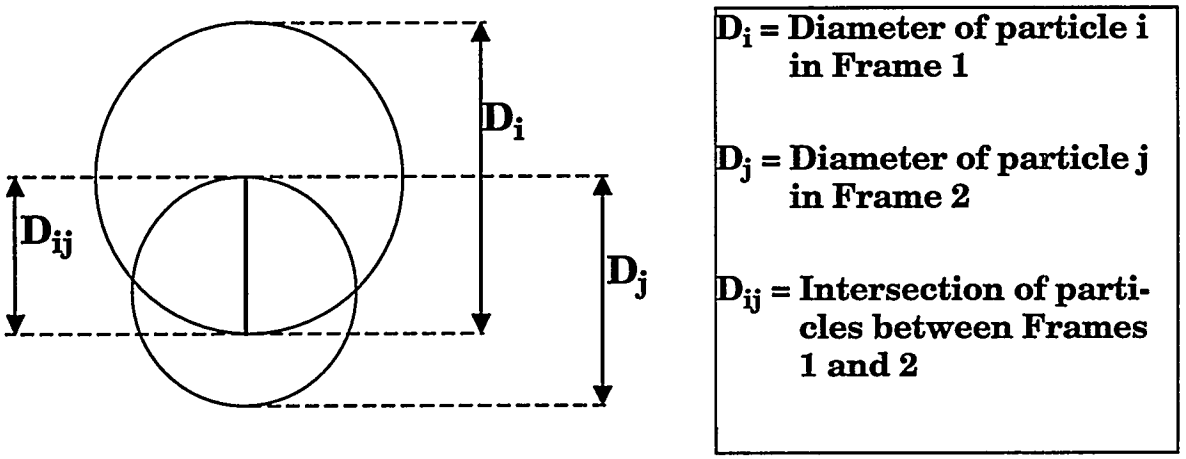

Figure 2.2. Intersection of particle images i and j between frames 1 and 2.

The possible pair with the largest C_{ij} value, the largest R_{ij} value, and the smallest sum of the distances between centroids is generally found to be the correct match. The velocity for a particle is then determined by dividing the particle movement by the known time between frames.

The Appendices provide step-by-step instructions (Appendix 1) and detailed descriptions of the programs (Appendix 2).

2.2.2. Multi-Frame Tracking Algorithm

The purpose of the Multi-Frame Tracking (MFT) algorithm is to track each particle through multiple time steps. This technique requires a minimum of four sequential images. A direct point-by-point matching of particles from one image to the next is performed. The MFT code outputs the best particle tracks through all the sequential space images.

The sequential images contain the information necessary for the two-dimensional tracking scheme. The maximum number of images that can be tracked through at one time is limited only by the memory of the computer on which the code is executed. The MFT routine tracks particles through four frames directly, and then it combines the four frame tracks to indirectly track through all the frames available. Several possible tracks exist for each particle in the first interrogation area. Possible tracks are evaluated based on the variance of length and angle from a path of constant curvature, referred to here as the sigma value of the possible track. The possible track with the lowest sigma value is assumed to be the best track for that particle. Changes in size and gray level are measured between frames and can also be used to evaluate possible tracks. The tracking is accomplished by prediction of the displacement and direction of the particle through the four consecutive time steps. Figure 2.3 shows the tracking procedure in two dimensions. The search area in the sec-

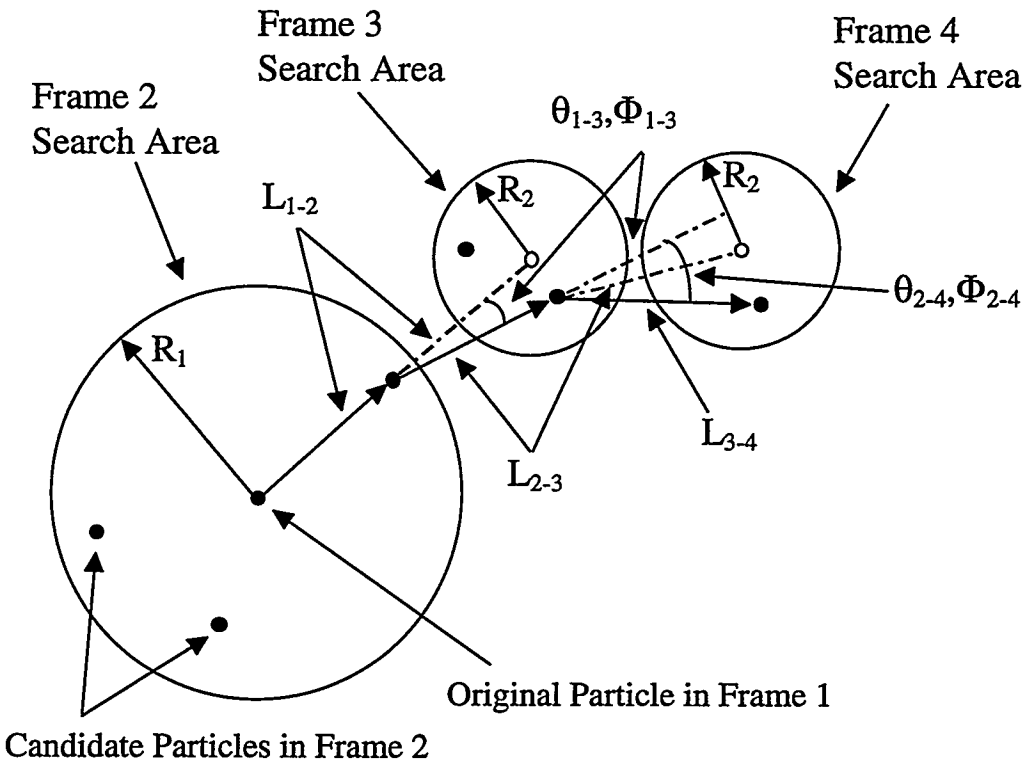

Figure 2.3. Determination of possible track for Multiframe Tracking (MFT) technique.

ond frame for a particle starting in the first frame is a circle of radius, R_1, determined by a rough estimate of the maximum velocity of the flow field. The search areas in the third and fourth frames have their radius, R_2, determined by a rough estimate of the maximum acceleration of the flow field. The center of each search area in Frame 3 is found by projecting out along the straight line, L_{1-2}, of a possible track for each particle found in the search area of Frame 2. For each particle found in the search area of Frame 3, the center of a search area in Frame 4 is determined by using the length of the track from Frame 2 to 3, L_{2-3}, and the angles created by the track from Frame 1 to 3, θ_{1-3} and Φ_{1-3}. This results in several different possible tracks for each particle in Frame 1, the best of which is used as the actual track for that particle, as determined by the lowest sigma value. If different tracks originating in the same frame share a particle, then the one with the highest sigma value is discarded as incorrect. The following calculations are performed to determine the sigma value for each possible track:

$$\sigma_l^2 = \frac{\left[L_{1-2} - \bar{L}\right]^2 + \left[L_{2-3} - \bar{L}\right]^2 + \left[L_{3-4} - \bar{L}\right]^2}{3} \qquad (3)$$

$$\sigma_\theta^2 = \frac{\left[\theta_{1-3} - \bar{\theta}\right]^2 + \left[\theta_{2-4} - \bar{\theta}\right]^2}{2} \qquad (4)$$

$$\sigma_{\Phi}^2 = \frac{\left[\Phi_{1-3} - \bar{\Phi}\right]^2 + \left[\Phi_{2-4} - \bar{\Phi}\right]^2}{2} \tag{5}$$

$$\sigma^2 = \sigma_l^2 + \sigma_\theta^2 \bar{l}^2 \cos^2\bar{\Phi} + \sigma_\Phi^2 \bar{l}^2 \sin^2\bar{\Phi} \tag{6}$$

Where:

σ_l	Standard deviation from the mean for the lengths of the vectors
σ_θ	Standard deviation from the mean for the angles in the x-y plane
σ_Φ	Standard deviation from the mean for the angles from the z axis (for 3D tracking, see Ch. 3)
σ	Total standard deviation
l_{1-2}	Length between particles in the 1st and 2nd frames
l_{2-3}	Length between particles in the 2nd and 3rd frames
l_{3-4}	Length between particles in the 3rd and 4th frames
θ_{1-3}	Angle in the x-y plane between l_{1-2} and l_{2-3}
θ_{2-4}	Angle in the x-y plane between l_{2-3} and l_{3-4}
Φ_{1-3}	Angle from the z axis between l_{1-2} and l_{2-3}
Φ_{2-4}	Angle from the z axis between l_{2-3} and l_{3-4} (for 3D tracking, see Ch. 3)
$\bar{l}$	Average displacement of particles between frames = $(l_{1-2} + l_{2-3} + l_{3-4})/3$
$\bar{\theta}$	Average θ of the track = $(\theta_{1-3} + \theta_{2-4})/2$
$\bar{\Phi}$	Average Φ of the track = $(\Phi_{1-3} + \Phi_{2-4})/2$

The total standard deviation can be made a dimensionless value as follows:

$$\sigma_{\text{total}}^2 = \frac{\sigma^2}{\bar{l}^2} = \frac{\sigma_l^2}{\bar{l}^2} + \sigma_\theta^2 \cos^2\bar{\Phi} + \sigma_\Phi^2 \sin^2\bar{\Phi} \tag{7}$$

Tracks can start in any frame of a multiframe sequence except the last three, so that each track can extend over four frames. After all four-frame tracks are found, any tracks that start in different frames but have the same particles in all their common frames are combined to form longer tracks, possibly spanning all the frames. For example, consider two four-frame tracks, the first through Frames 1-4 and the second through Frames 2-5. If the tracks share the same particles in frames 2, 3, and 4, then they are combined into a single five-frame track. On the other hand, if the two four-frame tracks shared the same particles in frame 2, but had different particles in the other frames, then the track with the lowest sigma value would be kept, and the other track would be discarded.

The MFT algorithm is generally preferable to the CCT algorithm for multiphase flow measurements, where the separate phases can be tracked individually (Hassan and Blanchat, 1991a).

The Appendices provide step-by-step instructions (Appendix A) and detailed descriptions of the programs (Appendix B).

2.3 2D PIV Validation Experiments

This section presents two validation experiments that were performed using 2D PIV. The first is the flow over a square cavity, where both PIV and LDV were applied to the experiment and a computational simulation was also performed. The second flow is a thermal convection experiment in which PIV measurements were compared to a computational simulation.

2.3.1. Cavity Flow Experiments

The flow in a cavity provides a well-characterized separated, recirculating flow in a closed domain. There have been a number of previous experimental and computational examinations of the shear-driven flow in a cavity (e.g., Koseff and Street, 1984abc; Iwatsu et al., 1989), and this geometry is often used for computational model validation. However, most previous investigations have been concerned with the lid-driven cavity, in which the flow is driven by a belt which imposes a constant uniform velocity at the cavity upper surface. The cavity flow in the present study is driven by the laminar channel flow passing over the top of the cavity (cf. Reiman and Sabersky, 1968). This configuration is challenging for the PIV technique since it extends the required velocity range: the recirculating flow in the cavity has a peak velocity several times lower than that of the driving channel. Previous studies of the lid-driven cavity have shown complex three-dimensional effects, although many of those studies were performed at Reynolds numbers (Re) in the range 3000-10000, much higher than those considered here (100<Re<2758). For Reynolds numbers greater than approximately 3000, the flow becomes unsteady (Koseff and Street, 1984abc). Thus, the flows examined in this study are all expected to be steady, except possibly for the highest case (Re = 2758). Laser Doppler velocimetry (LDV) was applied to measure point velocities throughout the cavity and channel, and PIV was then applied to measure full-field velocity maps in the same region.

2.3.1.1. Experimental Setup

There are a number of reasons that the cavity flow was chosen for this benchmark experiment. The extent of the recirculating region is independent of Reynolds number, thus easing the PIV imaging requirements. The flow is interesting because unusual three-dimensional recirculating regions are expected, although the present experiments are limited to examination of the two-dimensional centerline symmetry plane. Finally, there are a number of computational simulations of this flow for comparison, at least for the idealized case in which a fixed velocity condition is imposed at the cavity upper boundary.

The cavity flow geometry is set up in a small recirculating water tunnel (Fig. 2.4). The flow is pressure-driven, with the driving pressure supplied by a constant-head tank (62 by 62 cm cross-section, 114.5 cm deep). The water from the tank passes through a flow conditioning section consisting of a 62 cm long PVC pipe (9 cm I.D.) containing a packed bed of 3 mm glass beads. A 14 mesh stainless steel screen is placed at the downstream end of the flow conditioning tube. The flow passes through the tube into the 7.76 by 1.56 cm rectangular duct leading to the cavity. The cavity upstream corner lies approximately 193 cm (76") downstream of the head tank (74 hydraulic diame-

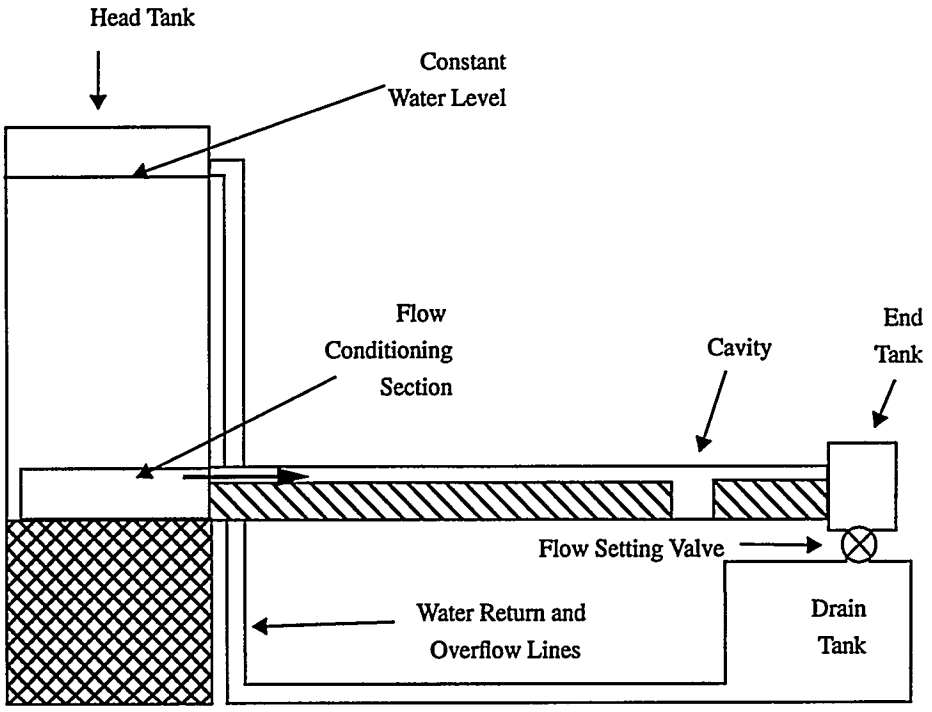

Figure 2.4. Schematic diagram of water channel/cavity flow.

ters). The downstream continuation of the channel, after the cavity, extends an additional 47.8 cm (18.8") before exhausting into an end tank. The water flows from the end tank to a drain tank through the flowrate-controlling valve. The water from the drain tank is then pumped back to the head tank, where a weir-type overflow connection back to the drain tank is used to maintain a constant head level. All walls are made of plexiglass to provide optical access.

Figure 2.5 is a schematic layout of the cavity flow geometry. A Reynolds number for this flow can be defined based on the mean incoming velocity $\bar{u}$ and the upstream hydraulic diameter, which for an infinitely wide channel is simply equal to twice the flow channel height h, yielding

$$Re = \bar{u}(2h)/\nu$$

where ν is the kinematic viscosity of the water. Although the channel here is not infinitely wide, this Reynolds number definition was chosen to allow comparison with 2D computations. The cavity depth and streamwise length are approximately 2.36 times the upstream channel height (3.36:1 expansion), so Reynolds numbers based on the cavity depth or length are approximately 18% higher than the Reynolds number defined above.

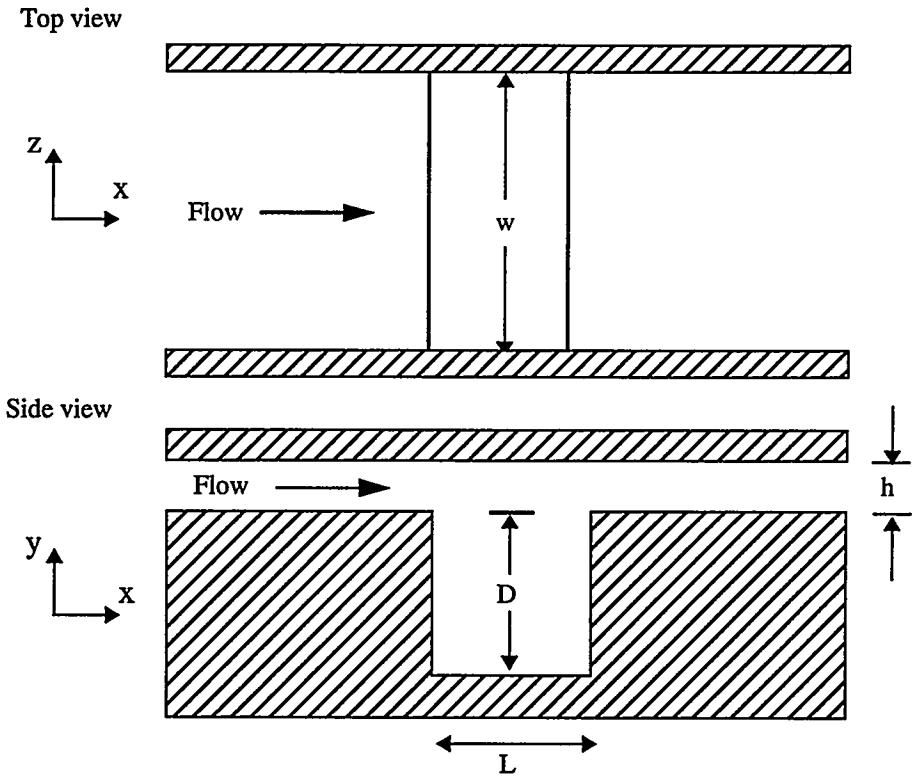

Figure 2.5. Schematic layout of cavity flow system. D = 3.68 cm, L = 3.70 cm, h = 1.56 cm, w = 7.76 cm.

2.3.1.2. Experimental Techniques

2.3.1.2.1. Laser Doppler Velocimetry (LDV)

The LDV system used in this study is a two-component TSI (St. Paul, MN) fiber-optic probe system using a 5 W argon ion laser (Spectra-Physics Model 2020). The green (514.5 nm) and blue (488 nm) lines of the Ar+ laser are used for measurement of the streamwise (x) and vertical cross-stream (y) velocity components, respectively. Frequency shifting is used on both channels. Data acquisition and analysis are performed using Model 1980B counter type processors, a Model MI990 data interface, and the TSI FIND software (Ver. 4.03) running on a 486-based computer.

The fiber-optic probe is positioned by a computer-controlled traverse. The light collection optics operate in the back-scatter mode. A 121.1 mm focal length probe lens (10.78° half angle) is used to focus the four laser beams at the measurement volume and to collect the Doppler-shifted light scattered by particles passing through the measurement volume.

Tap water, passed through a 10 μm filter, is used to fill the tank. Since the fiber-optic LDV probe operates in the backscatter mode, it is necessary to add seed particles to the flow. Titanium dioxide

(TiO_2) powder, ultrasonically dispersed in water, is added to provide seed particles. The TiO_2 mean particle diameter is approximately 2 μm, and its specific gravity is approximately 4. The TiO_2 suspension was added until the water was slightly cloudy and the LDV data rate was approximately 1000 Hz. Although TiO_2 particles are too heavy to follow the flow with perfect fidelity, their small size gives them a very low settling velocity and allows them to remain suspended and to follow all but high frequency flow fluctuations (Durst et al., 1981).

2.3.1.2.2. Particle Image Velocimetry (PIV)

A CCD-based, particle-tracking PIV system was developed for this study, based on previous work by one of the authors (Blanchat, 1991; Hassan and Blanchat, 1991ab). In this setup, particle centroids are mapped out in each frame of an image sequence, then the particle images are paired up through the correlation algorithm (CCT and MFT) to determine particle displacements and velocities. The PIV system uses a 5 W argon ion laser for light sheet generation. The laser beam is shaped into a narrow line (approximately 1 mm thick) using conventional spherical and cylindrical lenses. The light sheet is directed downward through the flow from the top wall to the floor and spans the entire cavity width in the streamwise direction. Because of the continuous nature of the illuminating laser light sheet, care was taken not to heat the cavity floor. If the light sheet was left on for more than about 5 s, localized floor heating created plainly visible regions of upwelling from the floor. For this reason, the light sheet is on for only one or two seconds during each PIV data acquisition period. Images are recorded with a CCD camera located perpendicular to the light sheet. The camera is a Pulnix Model TM745, with a 720 (x) by 480 (y) by 8 bit pixel format. Images can be recorded as quickly as every 1/30 s (33 ms apart). An image separation period of 1/30 s was used for all of the present experiments. Since the laser illumination for these experiments is continuous, the exposure period is determined by the camera shutter rate. The camera has an adjustable shutter, variable from 1/60 to 1/10000 s exposure times. Exposure times of 1/60 to 1/250 s are used for the present tests, which are sufficient because the flow velocities are low enough to prevent particle image blurring. The camera is connected to a frame grabber board (Epix 4MEG VIDEO Model 10 with 4 MByte memory), which is controlled by the Epix 4MIP frame grabber/image processing software running on a 486-based personal computer. This setup allows sequential acquisition of video images until the frame grabber memory buffer is filled. Eleven images are recorded in a typical flow sequence (approximately 3.8 MBytes).

Particle tracking from one image to the next has the advantage that there is no directional ambiguity, i.e., the direction as well as the magnitude of the velocity vector is directly measured. The video technology currently used limits data acquisition to approximately 30 frames/s. However, images of higher-speed flows could be made using two or more cameras or by double exposure of a single image with a pulsed light source.

The practical use of PIV requires the use of fast, reliable, computer-based methods for tracking the numerous particles suspended in a fluid flow. In some cases, image processing is performed on each frame prior to analysis by the PIV software. First, a background removal is performed by averaging together the eleven video frames, then subtracting the average frame from each of the other

frames. Then the histogram of the pixel values in the image is plotted to determine the need for contrast enhancement. If the histogram shows the majority of the pixel intensities are at one end of the range (0 to 255), then image contrast enhancement is needed. Typically, most of the pixel values fall into the 0 to 100 range, so the contrast is enhanced by scaling all pixel values with intensities in the 0 to 100 range to the expanded 0 to 255 range, while all pixel values of 100 to 255 in the original image are set to 255 (white). After the contrast enhancement, a Laplacian filter is applied to sharpen the edges in the image. Finally, to increase the visibility of the bright particles in the image, all pixels below a determined threshold, usually 100, are set to 0 (black). These steps convert a somewhat low-contrast original image to a sharp field containing well-defined particle images suitable for particle tracking with the PIV software.

Image data are analyzed in several steps. First, the particle centroid locations (gray level averaged), size, and average gray levels for each image recorded in a sequence are determined. Second, a dynamic, binary, spatial cross-correlation tracking method is used to analyze the PIV images. This method can be performed quickly if the particle tracer image information is first converted to binary data. That is, the 8-bit (0 to 255 gray level) pixel data defining a particle are converted to 2-bit (value 0 or 1) pixels through image processing techniques, e.g., thresholding and connectivity algorithms. The particle velocity is found by determining the correspondence between particles in two sequential video frames. Correspondence is obtained using the CCT technique described in Section 2.2.1.

The PIV analyses were run on an Hewlett-Packard Model 720 RISC Workstation. Typical data processing times were approximately 1-2 hours for 1000 particle tracks in 10 image pairs.

An auto-cleaning validation method is used to eliminate manual operator assistance in identifying and removing erroneous vectors. This method is an iterative process involving an interpolated field produced from the most reliable vectors, which allows fast analysis and presentation of sets of PIV image data. The interpolation is based on the Hardy multiquadratic equations, and produces a vector at any given point in the flow field from the sparse, scattered velocity data (Blanchat, 1991).

2.3.1.3. Experimental Approach

LDV and PIV data were recorded at Reynolds numbers of 100, 300, 600, 900, and 2758. The experimental sequence was to make LDV scans one cavity depth upstream of the cavity entrance to check the incoming velocity profiles (in y and z directions) and the flow steadiness. The velocity profiles were used to determine the mean velocity. When the desired Reynolds number was achieved, the flow was allowed to stabilize approximately one hour prior to data acquisition. At that time, the velocity profiles were checked, an LDV cavity scan was performed, followed by several sequences of PIV data acquisition, and finally a post-test LDV cavity scan and check of upstream velocity profiles to check flow steadiness. There were 458 LDV-velocity measurement locations within and above the cavity (see Figure 2.8). While the PIV images were captured in 1/60 to 1/250 s, the LDV cavity scan data are time-averaged in the sense that each scan took approxi-

mately two hours (1000 samples at each of 458 points). Both LDV and PIV data were recorded only on the flow centerplane.

2.3.1.4. Computational Approach

The computational fluid dynamics code FIDAP (Fluid Dynamics International, 1993) has been used to simulate the flow over and through the cavity-channel geometry. The flow is taken to be incompressible, isothermal, laminar, and steady, with constant material properties. The penalty method is used to enforce the incompressibility constraint, where a value of 10^{-5} is used for the penalty parameter. The experimental flow channel has a spanwise width of 7.763 cm, and a height of 1.560 cm, with a cavity in the lower side having a depth of 3.675 cm and a length of 3.700 cm (see Figs. 2.5 and 2.6). However, in these simulations, the following normalization is used: the den-

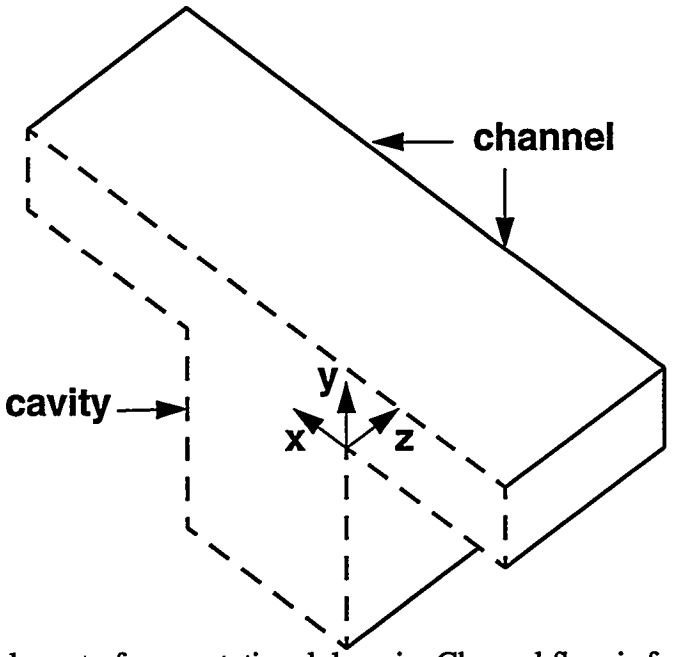

Figure 2.6. Schematic layout of computational domain. Channel flow is from lower right to upper left. Symmetry plane (dashed) is exposed to view.

sity ρ and the channel mass-flow velocity u_{av} are set equal to unity, and the channel height h is taken to be 0.5. With this choice, the channel has a width of 2.488, and the cavity has a depth of 1.178 and a length of 1.186. The Reynolds number, here defined as $Re = \rho(2h)u_{av}/\mu = 1/\mu$, is thus specified by selecting a value for the viscosity μ. A fully developed velocity profile with an average speed of unity is applied as the boundary condition at the inflow portion of the computational domain, natural boundary conditions (Fluid Dynamics International, 1993) are applied at the outflow boundary, symmetry conditions are applied on the channel midplane, and no-slip conditions are applied on all remaining surfaces (walls). The inflow and outflow boundaries in the computational domain are placed one cavity length upstream and downstream, respectively, from the cavity. Figure 2.6 shows a schematic diagram of the domain, and Figure 2.7 shows the two computational meshes computational meshes, composed of 27-node brick elements, on which solutions were obtained (Torczynski and O'Hern, 1994). Comparisons of results on these two meshes indicate that

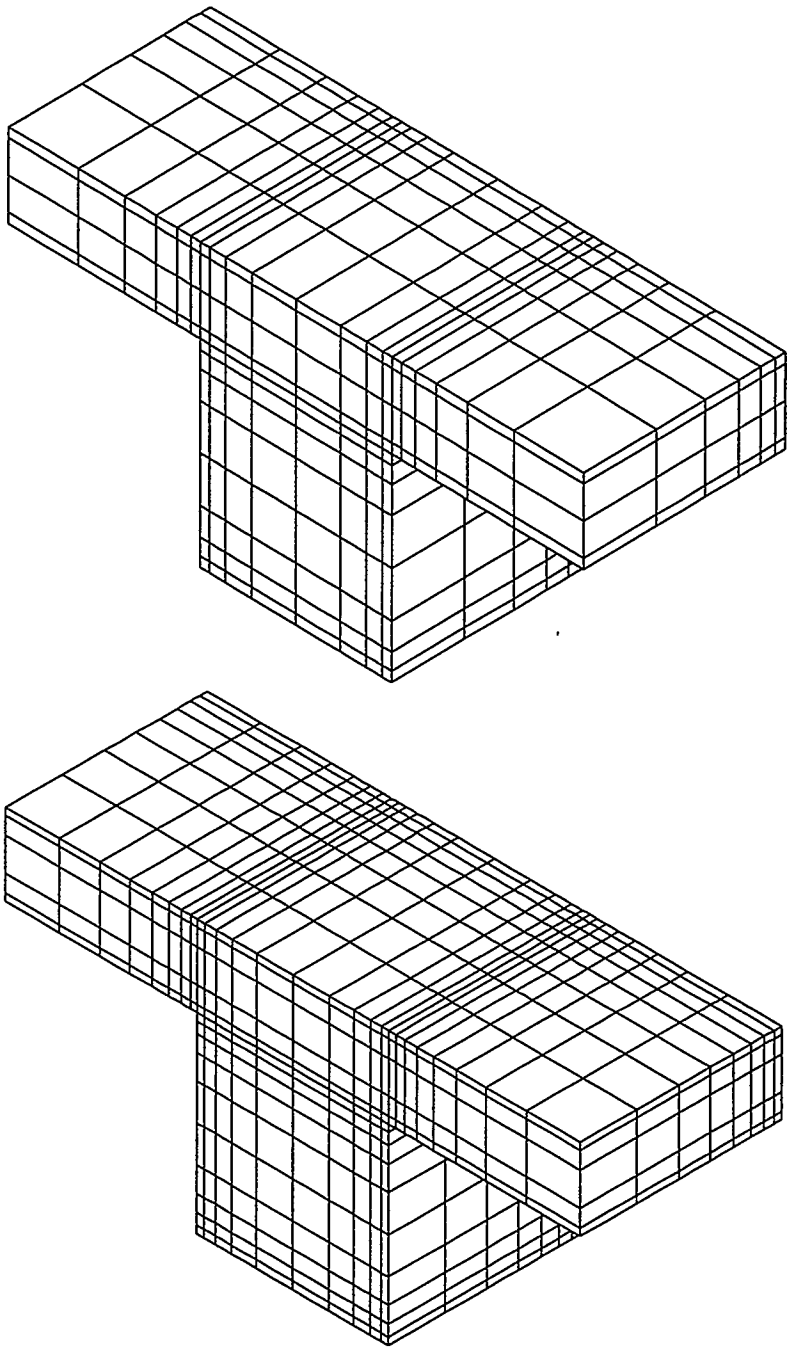

Figure 2.7. Computational meshes (27-node brick elements): top, coarse; bottom, refined.

the second mesh has adequate resolution, so results are reported only for the second mesh. All post-processing results are generated using the FIDAP post-processor FIPOST, which determines streamlines by integrating particle paths through the flow.

2.3.1.5. Results of Cavity Flow Study

2.3.1.5.1. Laser Doppler Velocimetry (LDV)

Figure 2.8 shows the LDV-measured planar centerline velocity fields recorded at Re = 100, 300, 600, 900, and 2758. A primary recirculating region essentially fills the cavity. A small secondary recirculation region in the bottom upstream corner was evident for all Reynolds numbers examined. However, the LDV system was not well configured to examine the corner flows in much detail. The fairly short probe lens focal length and the fact that measurements were performed on the flow centerline led to beam clipping in the corner regions.

Typical LDV velocity measurements of laminar flows are known to have precision of about ±1% or better (Durst et al., 1981). Measurement of the upstream peak velocity over 5 to 10 minute periods showed fluctuations on the order of 1% at Re = 100, and as discussed below, the upstream centerline velocity profiles (across both y and z directions) were typically constant to within the measurement accuracy over the several hour experiment run times.

Upstream velocity profiles were recorded at the start and end of each run at each of the Reynolds numbers of interest. Agreement was usually better than 1%, with a maximum variation of 1.7% for the center velocity at Re = 600. .

2.3.1.5.2. Particle Image Velocimetry (PIV)

Figure 2.9 shows the PIV-measured planar centerline velocity fields recorded at Re = 100, 300, 600, 900, and 2758. A primary recirculating region essentially fills the cavity. Again, the primary recirculation region essentially filled the cavity. At Re = 600, secondary recirculation zones became visible in the two bottom corners of the cavity.

The raw PIV data consisted of velocity maps with a random distribution of vectors, depending on the location of trackable particles during data acquisition. The Hardy multiquadratic interpolation was used to calculate velocity vectors on a regular grid corresponding to the LDV measurement locations.

As expected, the cavity/channel flow combination was a demanding one in terms of the PIV velocity measurement range. The operator has to input the size of correlation window to search for particle pairs from one frame to the next, which can usually be determined by visual observation of the video images. This is typically done by visually locating the maximum displacement of particle images between frames (maximum velocity) and choosing a correlation window 1.5 to 2 times this length. However, this parameter varied significantly from the cavity region to the channel region,

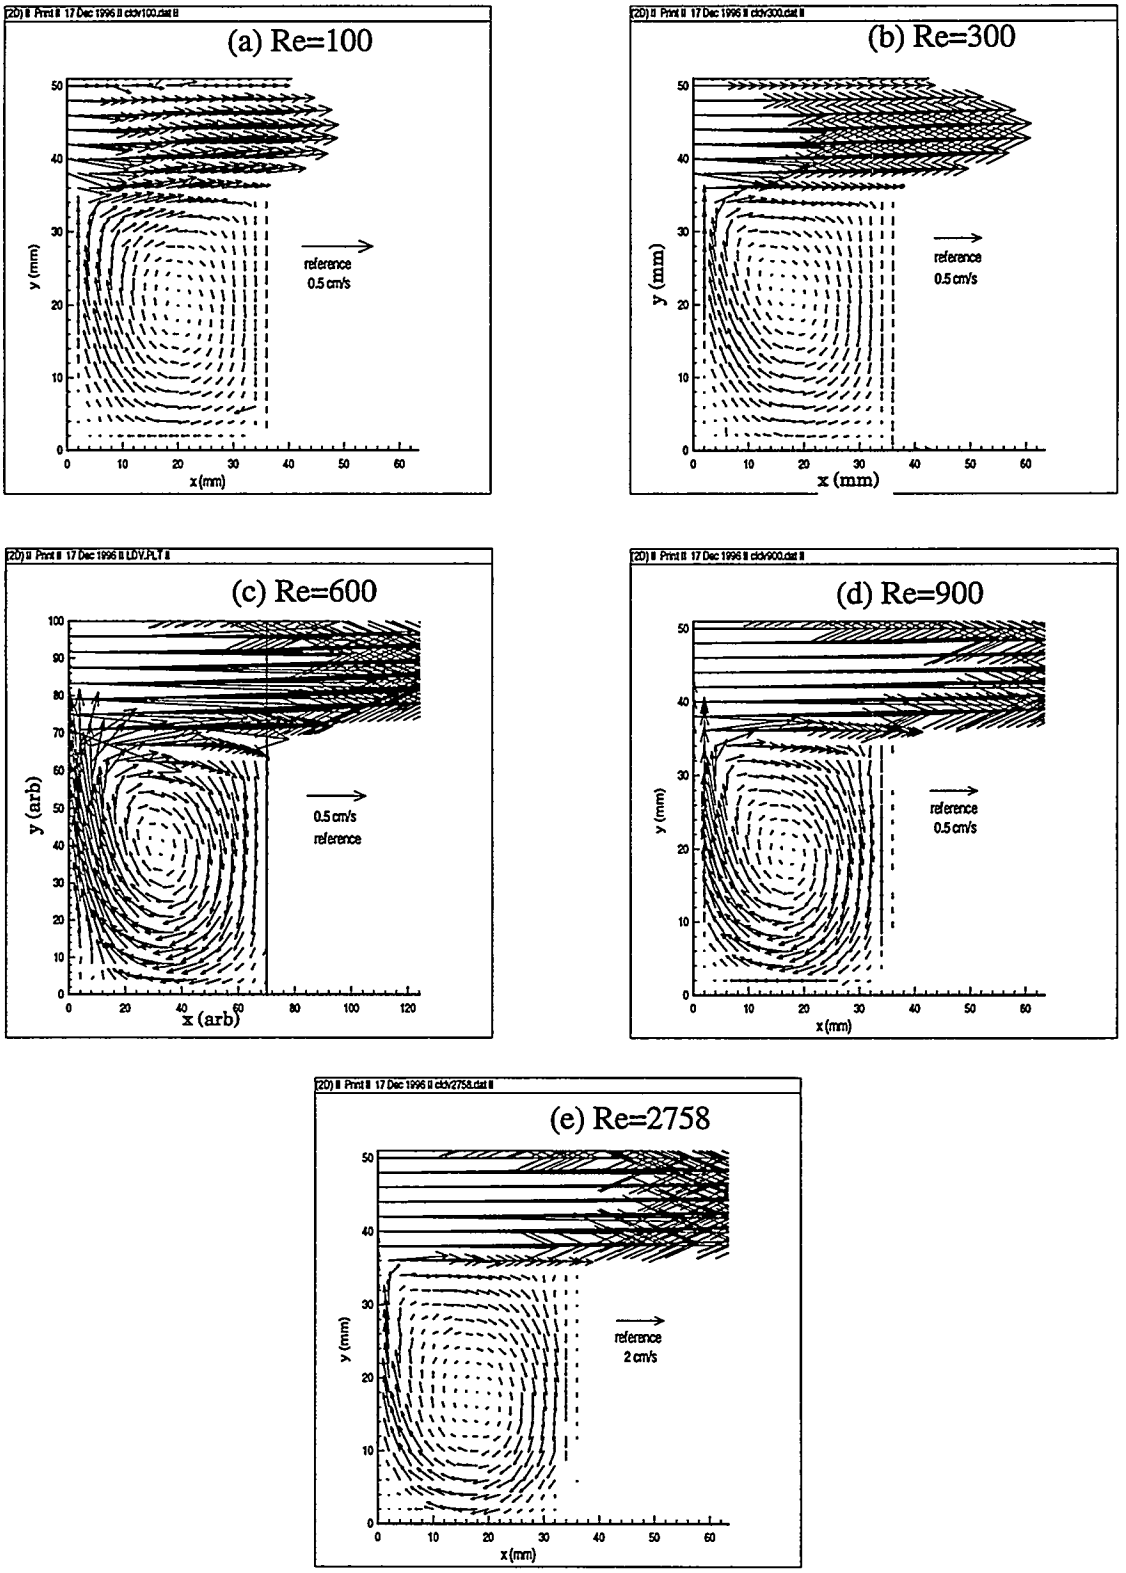

Figure 2.8. LDV results for cavity flow. Time-averaged centerplane velocity fields.

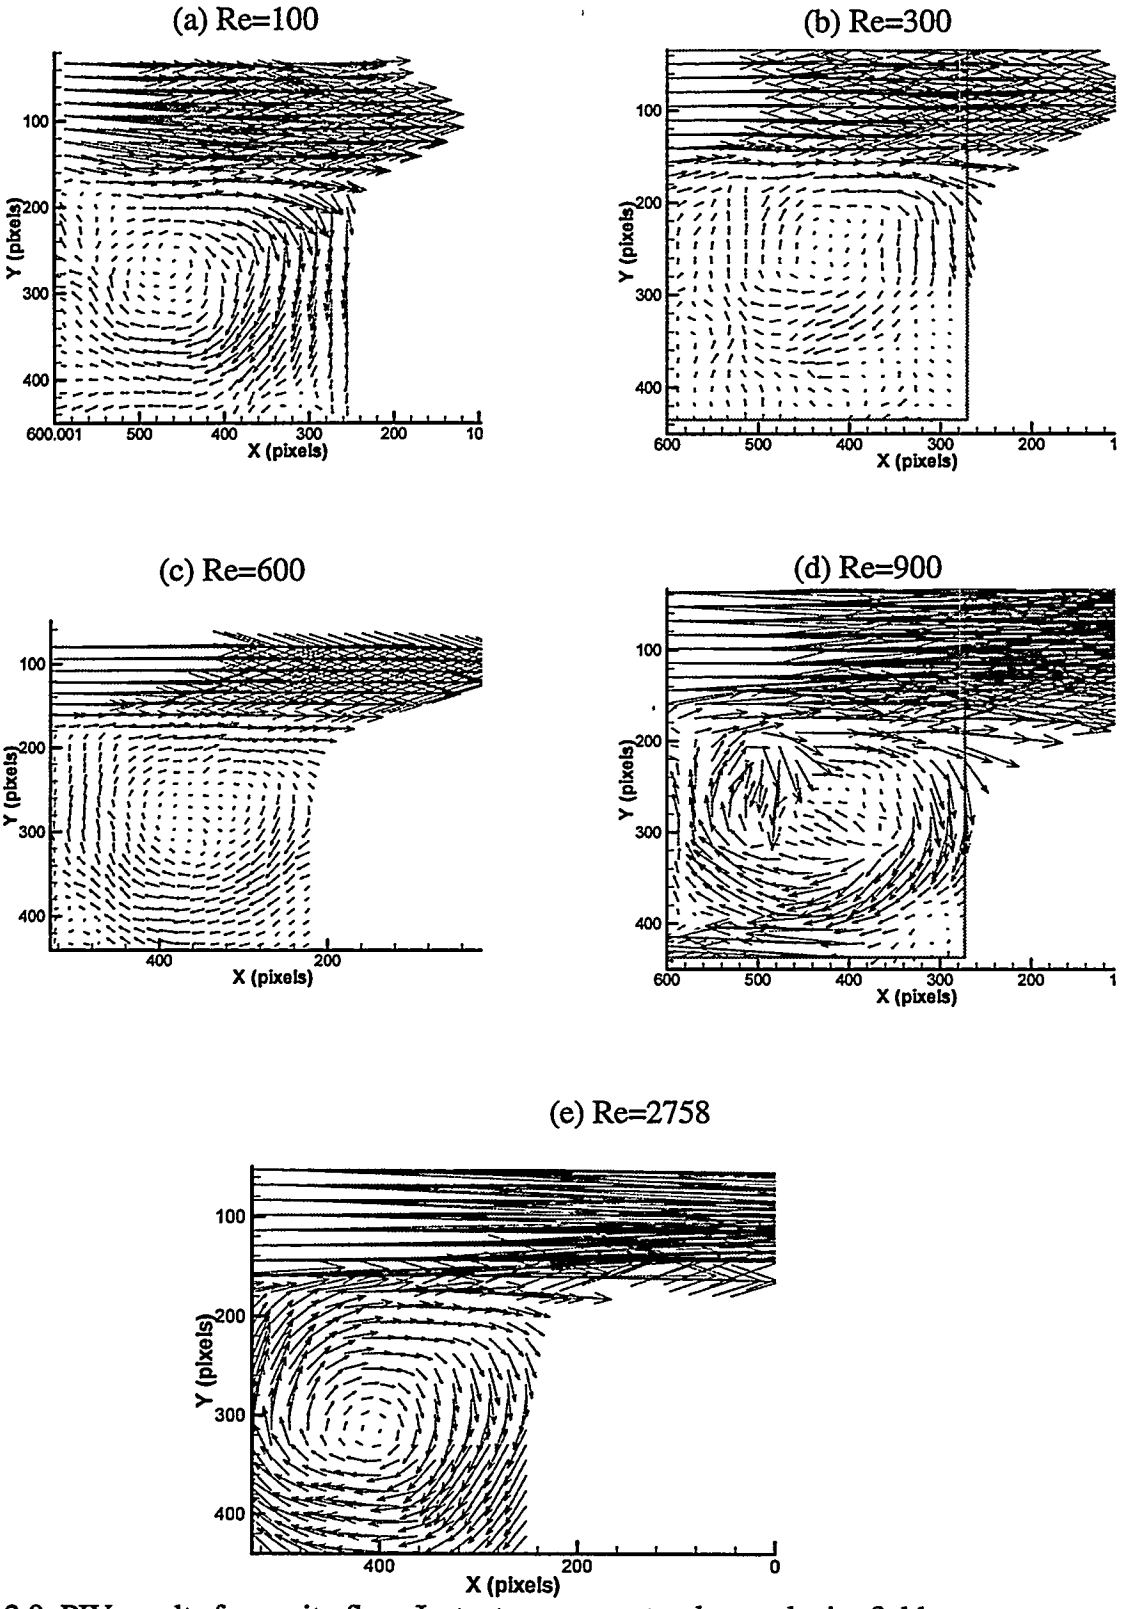

Figure 2.9. PIV results for cavity flow. Instantaneous centerplane velocity fields.

so two values were used. This led to some problems at the cavity/channel interface, where it was not always clear where to stop using the larger window needed for the channel flow and start using the smaller window appropriate for the cavity flow. In addition, the video timing was optimized for the low-speed cavity flow rather than the higher-speed channel flow above the cavity. In general, the number of velocity vectors found by PIV in the channel was much lower than the number found in the cavity.

Sources of experimental error during PIV data recording include the timing between the video capture periods and the measurement of the cavity dimensions (for converting pixel locations to actual displacement). Both of these are thought to be better than 1%. The main source of error in this technique is in the analysis, in particular in determination of which particles to match up for frame-to-frame tracking. This can be complicated by out-of-plane motions (not expected on the present experiments performed on the centerline symmetry plane) and nonuniformities in the light sheet.

2.3.1.5.3. Computational Simulations (FIDAP)

Flow simulations are performed for Reynolds numbers of 1, 3, 6, 10, 30, 60, 100, 300, 600, and 900. Figure 2.10 shows the centerplane velocity map, cf. Figs. 2.8 and 2.9 for LDV and PIV, respectively. Figure 2.11 shows how the location of the center of the main eddy varies with Reynolds number. Not evident in Figure 2.11 is the fact that the flow changes topology around $Re = 27$. In the lower Reynolds number cases, some streamlines from the upstream section of the channel are subducted into the cavity, swirl about a streamline emanating from an inward spiral point while being displaced laterally away from the symmetry plane and toward the side walls, and ultimately exit from the cavity into the downstream section of the channel. In the higher Reynolds number cases, the subducted streamlines swirl about a closed isolated volume of fluid, rather than a single streamline. The fluid within this volume is kinematically similar to two opposed vortex rings being rotated about their axes by the shearing channel flow passing above the cavity: each half possesses a core streamline analogous to a vortex core through which the fluid flows toward the symmetry plane and outside of which the fluid flows away from the symmetry plane. Thus, the spiral point is outward here. Figure 2.12 shows streamlines for Re = 3 and Re = 30 illustrating these differences, and Figure 2.13 shows the topologies that correspond to these cases. See Torczynski and O'Hern (1994) for a more detailed discussion of these results.

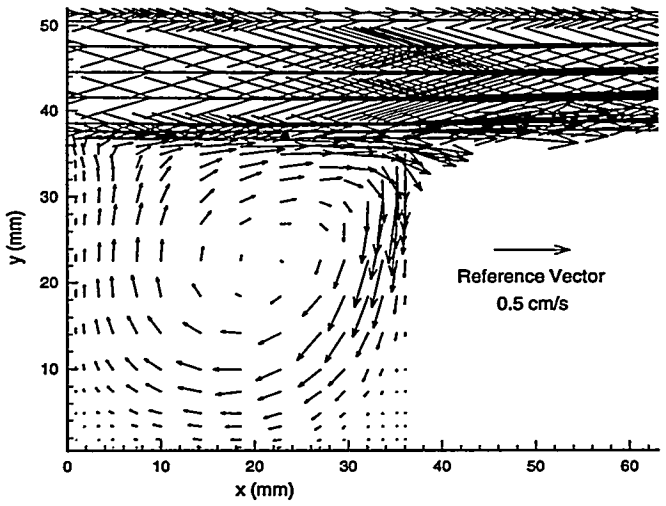

Figure 2.10. CFD simulations of velocity vectors in the cavity flow at Re = 600. Flow above the cavity is from left to right, with a peak velocity of about 2.8 cm/s.

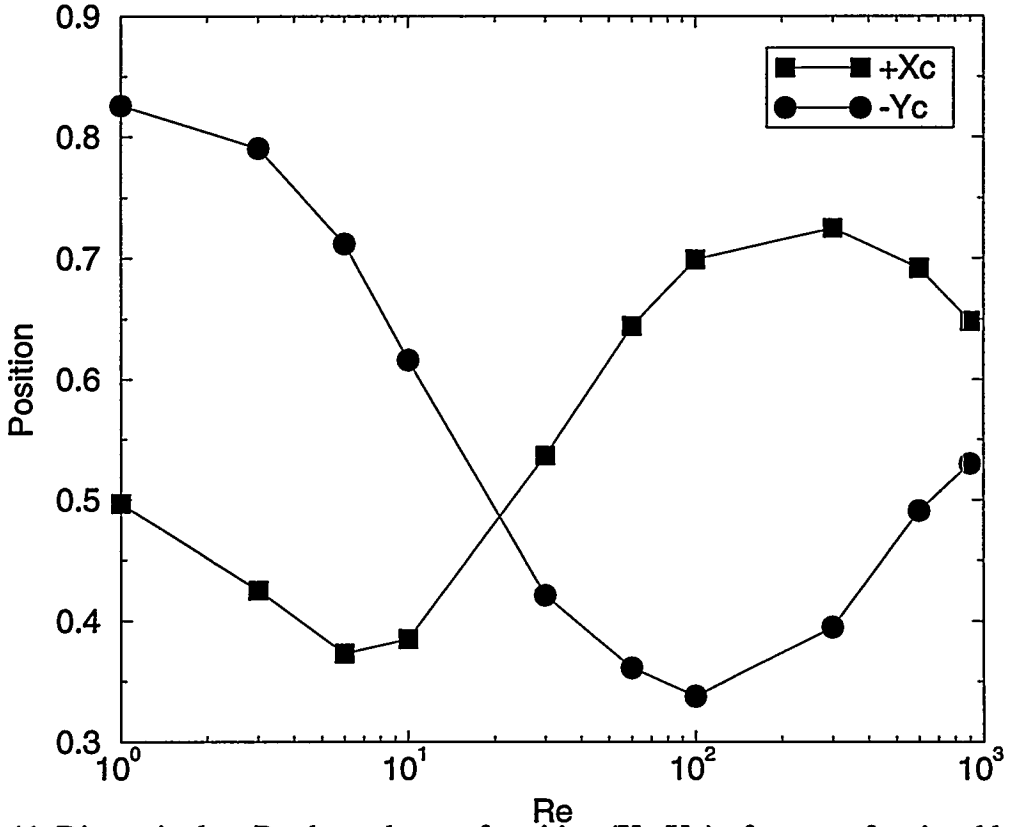

Figure 2.11. Dimensionless Re-dependence of position (X_C,Y_C) of center of main eddy in the symmetry plane z = 0.

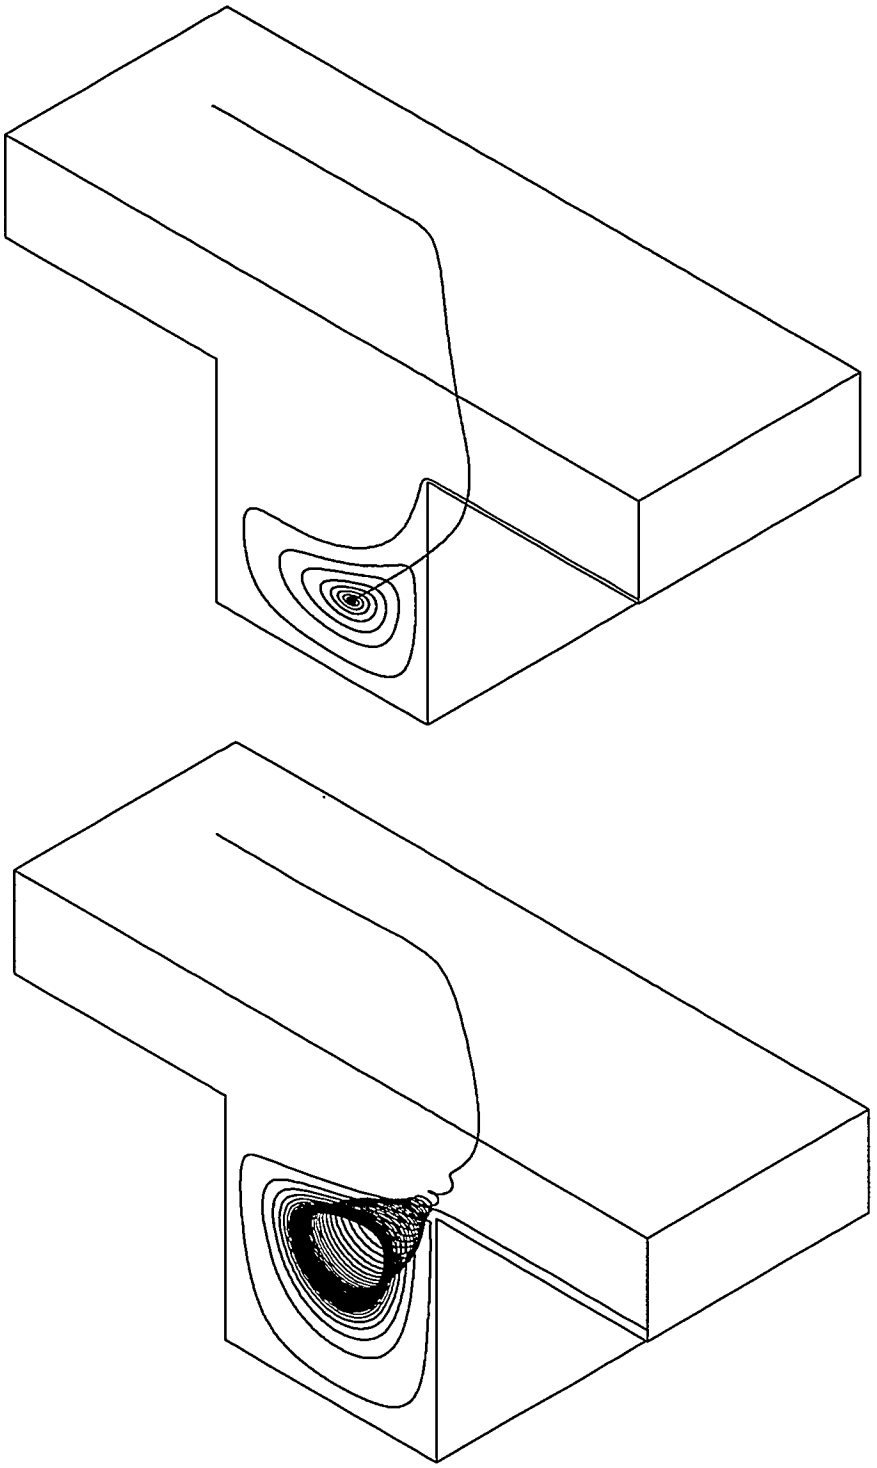

Figure 2.12. Streamlines: above, Re = 3 (no isolated region); below, Re = 30 (isolated region).

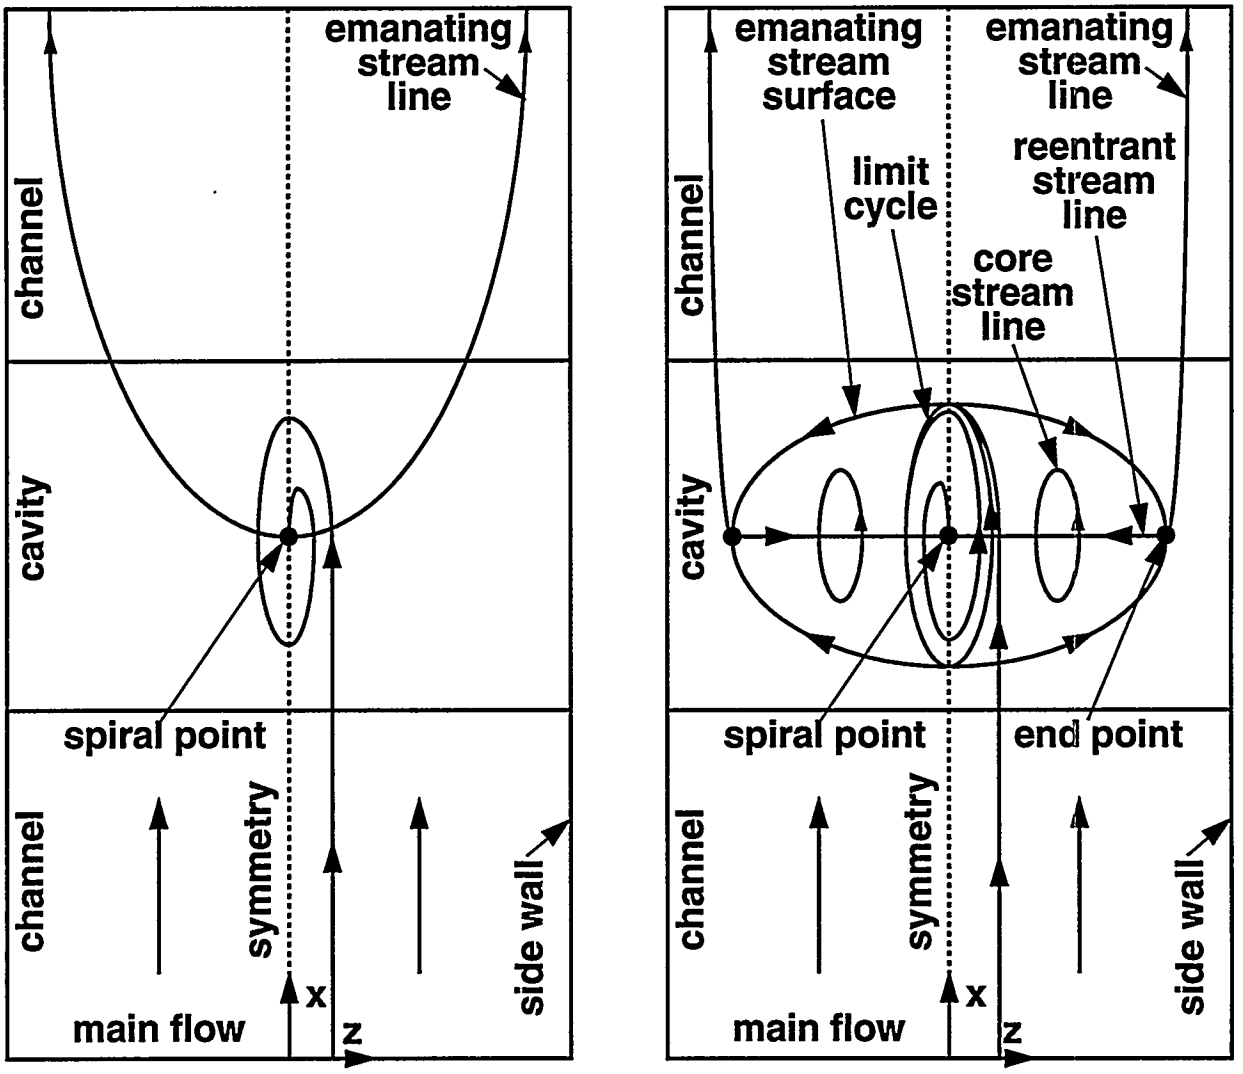

Figure 2.13. Topologies: left, low Re; right, high Re. Transition occurs at Re = 27.