

Low-dimensional dynamics of a turbulent wall flow

By JAVIER JIMÉNEZ^{1,2} AND MARK P. SIMENS¹

¹School of Aeronautics, U. Politécnica, 28040 Madrid, Spain

²Centre for Turbulence Research, Stanford U., Stanford, CA 94305, USA

(Received 1 December 2000 and in revised form 1 February 2001)

The low-dimensional dynamics of the structures in a turbulent wall flow are studied by means of numerical simulations. These are made both ‘minimal’, in the sense that they contain a single copy of each relevant structure, and ‘autonomous’ in the sense that there is no outer turbulent flow with which they can interact. The interaction is prevented by a numerical mask that damps the flow above a given wall distance, and the flow behaviour is studied as a function of the mask height. The simplest case found is a streamwise wave that propagates without change. It takes the form of a single wavy low-velocity streak flanked by two counter-rotating staggered quasi-streamwise vortices, and is found when the height of the numerical masking function is less than $\delta_1^+ \approx 50$. As the mask height is increased, this solution bifurcates into an almost-perfect limit cycle, a two-frequency torus, weak chaos, and full-fledged bursting turbulence. The transition is essentially complete when $\delta_1^+ \approx 70$, even if the wall-parallel dimensions of the computational box are small enough for bursting turbulence to be metastable, lasting only for a few bursting cycles. Similar low-dimensional dynamics are found in somewhat larger boxes, containing two copies of the basic structures, in which the bursting turbulence is self-sustaining.

1. Introduction

The physical mechanisms of wall-bounded turbulent shear flows are considerably less well understood than those of most other ‘classical’ turbulent flows. There are two competing conceptual models. In the first one, wall turbulence is just a modification of ordinary shear turbulence, occurring when the latter approaches a wall, and is therefore dependent on the prior existence of an outside turbulent flow. In the second, it is an independent phenomenon which coexists with the outer flow and merges into it when the distance from the wall is large enough.

In this paper we present some evidence supporting the latter view, examining the dynamics of the structures of the viscous and buffer layers in simplified situations in which their interactions with the outer flow are severely restricted. Even natural flows scale in this region approximately on wall units, defined in terms of the kinematic viscosity ν , and of the friction velocity $u_\tau = (\nu \partial_y U_w)^{1/2}$, where $\partial_y U_w$ is the derivative at the wall of the mean velocity profile. In that approximation, and if we admit that near-wall turbulence is not just a modification by the wall of the outside turbulent flow, only local quantities such as the dimensionless distance to the wall, $y^+ = u_\tau y / \nu$, should matter, while global parameters such as the Reynolds number of the outer flow should be irrelevant. In the reduced systems considered in this paper the outer

flow is effectively removed, and the scaling in wall variables should be strict. No bulk Reynolds number may be relevant because no bulk turbulent flow exists.

Since the local Reynolds numbers of the near-wall layer are low, this region corresponds to the Kolmogorov viscous range of isotropic turbulence, but it has been known for some time that, while that range is a sink for the energy cascading from the larger scales, the near-wall structures actually export energy to the rest of the flow (see for example the discussion by Jiménez 1999).

Because of this influence on the global energy balance, and consequently on the magnitude of the wall drag, the region below $y^+ \approx 100$ has been subjected to intensive study. Its dominant structures are streamwise velocity streaks and quasi-streamwise vortices. The former are an irregular array of long ($x^+ \approx 1000$) sinuous alternating streamwise jets superimposed on the mean shear, with an average spanwise separation of $z^+ \approx 100$ (Smith & Metzler 1983). The quasi-streamwise vortices are slightly tilted away from the wall and stay in the near-wall region only to $x^+ \approx 200$ (Jeong *et al.* 1997). Several vortices are associated with each streak, with a longitudinal spacing of the order of $x^+ \approx 400$ (Jiménez & Moin 1991). Most of them merge into disorganized vorticity after leaving the immediate wall neighbourhood (Robinson 1991).

It was proposed by Kim, Kline & Reynolds (1971) that streaks and vortices are parts of a cycle in which the vortices are the result of an instability of the streaks, while the streaks are caused by the advection of the mean velocity gradient by the vortices (Swearingen & Blackwelder 1987). While there is still some discussion on the mechanism by which the vortices are generated, it was shown by Jiménez & Pinelli (1999 referred to as JP from now on) that disturbing the streaks inhibits their formation, resulting in the eventual decay of the turbulence. The manipulation is only effective below $y^+ \approx 60$, and fails if it is applied only below $y^+ \approx 10$. This suggests that it is predominantly between those levels that the streaks are involved in the generation process. That paper should be consulted for a fuller discussion of the different models that have been proposed for the near-wall region in natural turbulent flows, and for additional references.

A slightly different point of view is that the cycle is organized around a nonlinear travelling wave, a fixed point in phase space, which would represent a non-uniform streak. This is actually not too different from the previous model, which essentially assumes that the undisturbed streak is a fixed point in phase space, and that the cycle is an approximation to a homoclinic orbit running through it. Candidate nonlinear waves have been computed by Waleffe (1998) and others, and identified as part of a particular path to turbulent transition by Toh & Itano (1999). Reduced models based on this approach have been formulated by Sirovich & Zhou (1994) and Waleffe (1997).

A difficulty with most of these studies, with the possible exception of the one by Toh & Itano (1999), is that there is no clear connection between the objects being studied and full-scale turbulence. The models behave in a way that is reminiscent of turbulence, but are derived from the full equations by some drastic approximation, such as projection over a few selected modes, without intermediate steps. While the results are very suggestive, they are open to the same criticism as paleontological trees with ‘missing links’. The similarity in the final behaviour might be a stray coincidence between unrelated phenomena, rather than a proof of kinship. We intend in this paper to show that a turbulent flow can be continuously modified into one of those reduced objects.

A way of doing this was proposed by Jiménez & Moin (1991), who substituted a full turbulent channel by an array of identical computational boxes, periodic in

the streamwise and spanwise directions, x and z , while retaining its full wall-normal extent. The idea was to substitute the complexity of the mutually interacting turbulent units near the wall by a ‘crystal’ of identical structures executing synchronously the hypothetical turbulence regeneration cycle. The ‘unit cell’ of the crystal was adjusted to the smallest dimension that would maintain turbulence, which turned out to contain a single wavelength of a wavy low-velocity streak and a pair of quasi-streamwise vortices. These structures went through a complex evolution that was still difficult to analyse, while the statistics of the near-wall fluctuations were essentially identical to those of fully developed channels. This ‘minimal’ system has often been used as a surrogate for real wall turbulence, and most of the reduced models mentioned above refer to it.

That experiment suggests that the high dimensionality of natural wall turbulence does not derive from the interaction among neighbouring structures, all of which are in that case identical and cannot contribute additional degrees of freedom, but says nothing about the importance of its interaction with the core flow. A step in clarifying that question was taken by JP, who removed the vorticity fluctuations in the outer flow of a channel by using a numerical mask acting as a variable viscosity increasing away from the wall. They were able to show that the wall cycle was ‘autonomous’ in the sense that it could run independently of the outer flow as long as it was not disturbed below $y^+ \approx 60$. Again the statistics were similar, although no longer identical, to those of full turbulence.

In the present paper we describe the results of simulations which are ‘minimal’, in the sense that their wall-parallel size is small enough to contain a single copy of each structure, and ‘autonomous’ in the sense that they have no turbulent external flow.

The numerical experiments used to simplify and isolate the wall region are described first. The results are then presented and discussed, with emphasis on the low-dimensional behaviour of the structures in the simplified flows, and on how they evolve into a fully turbulent flow once the constraints are gradually removed.

2. The numerical experiments

The numerical scheme used for the simulations is similar to the one described by JP. The Navier–Stokes equations are integrated in the form of evolution equations for the wall-normal vorticity ω_y and for the Laplacian $\nabla^2 v$ of the wall-normal velocity. We use a pseudospectral code with dealiased Fourier expansions in the two wall-parallel directions, and Chebychev polynomials in y without dealiasing, as in Kim, Moin & Moser (1987). At each time step the right-hand sides of the two evolution equations are multiplied by a damping mask

$$F(y) = 1 \quad \text{if } y \leq \delta_1, \quad F(y) = F_0 < 1 \quad \text{if } y \geq \delta_2. \quad (2.1)$$

The two limits are connected smoothly by a cubic spline. In all our experiments $\delta_2 \approx 1.5 \delta_1$. Because the mask is applied at every time step, it is very effective in damping the vorticity fluctuations, and it was shown by JP that the complete suppression of the fluctuations only requires that $1 - F \gg \Delta t^+$, where Δt is the computational time step. Because the time step is limited by the velocity at the centreline, which is large when expressed in wall units, Δt^+ is small, and the damping is effective midway between δ_1 and δ_2 for the chosen value $F_0 = 0.95$. The evolution equations for the wall-parallel instantaneous mean values of the streamwise and spanwise velocities, which cannot be expressed in terms of ω_y and $\nabla^2 v$, are not modified by the mask.




FIGURE 1. Mean profiles for three autonomous computational boxes: —, permanent wave, $L_x^+ = 145$, $L_z^+ = 180$, $\delta_1^+ = 49$, $\delta_2^+ = 71$; - - - -, torus, $L_x^+ = 150$, $L_z^+ = 185$, $\delta_1^+ = 55$, $\delta_2^+ = 86$; — · —, chaos, $L_x^+ = 155$, $L_z^+ = 190$, $\delta_1^+ = 63$, $\delta_2^+ = 94$. The line with symbols is the full channel from Kim *et al.* (1987), included for comparison. The solid triangle on each profile marks the lower limit δ_1 of the mask. (a) Mean velocity. (b) Fluctuation intensity of the spanwise vorticity.

The flow is integrated in a channel which is periodic in the two wall-parallel directions, with wavelengths L_x and L_z . No-slip impermeable boundary conditions are imposed at $y = 0$ and $y = 2$, and the constant volumetric flux is adjusted so that the Reynolds number of an undamped turbulent channel would have been $Re_\tau \approx 200$. Because there are no Reynolds stresses in the region where the mask is active, the mean velocity profile is parabolic in that region (see figure 1a), sequestering much of the total mass flux, and the final Reynolds number is closer to $Re_\tau \approx 120$. As previously mentioned, however, such Reynolds numbers are irrelevant because no turbulent flow extends to the far wall. The relevant Reynolds number is the height of the damping function $\delta_1^+ = u_\tau \delta_1 / \nu$, which determines the size of the largest turbulent scales.

It was stated by JP that the damping function was approximately equivalent to a viscous term proportional to the velocity, instead of to its Laplacian. This is only partly correct. While the effect of (2.1) is equivalent to a zeroth-order viscosity for the two evolution variables, the velocities are obtained from them using the continuity equation and the definition of the vorticity. Potential velocity fluctuations penetrate unimpeded into the damped region. Even vorticity fluctuations are not completely suppressed, since only ω_y is directly damped while the other two vorticity components have to satisfy the solenoidality condition (i.e. the vortex lines have to close).

The flow below $y = \delta_1$ satisfies the unmodified Navier–Stokes equations. JP showed that the profiles of the mean velocity and of the fluctuations were very similar in that layer to those in regular channels. This was true even when the mask was set to an approximately minimum value of $\delta_1^+ \approx 50$. Much lower masks led to laminarization. Those results were obtained in a computational box whose length and width were much larger than the mask height, and which was essentially turbulent in the sense of displaying strong chaos both spatially and temporally. We will present here results for much smaller boxes, some of which contain ordered flows in which there are no chaotic fluctuations to smear the effect of the individual structures, and whose mean profiles therefore differ substantially from those of natural turbulent channels. Some examples are given in figure 1.

The three boxes used for that figure contain respectively a fixed point in phase space,

corresponding to a permanent traveling wave, a two-frequency torus, and a mildly chaotic flow. All of them are too small to sustain for long times flows undergoing the strong chaotic bursts which characterize wall turbulence and, although they are wider than the mean distance between streaks in natural turbulence, each of them is shown by direct visualization to contain a single low-velocity streak (see figure 5 below).

The mean velocity profiles have plateaux which can be interpreted as abortive logarithmic layers, in which the velocity decreases slightly before joining the parabolic profile of the irrotational region. The local maxima of the vorticity profiles below the mask height are real, and correspond to individual features which, in these relatively ordered flows, stay at a constant distance from the wall instead of wandering as in natural turbulence, but the strong peaks located just above δ_1 are artifacts. They mark the reconnection of the vortex lines after they are truncated by the masking.

Two types of numerical boxes are studied, both of which have widths $L_z^+ \approx 180$ when the mask height is set to its minimal value of $\delta_1^+ \approx 50$. The longer ones, which contain two streamwise copies of the basic structure, have $L_x^+ \approx 300$, and were run using $42 \times 97 \times 31$ spectral modes in the streamwise, wall-normal and spanwise directions, after dealiasing. The shorter boxes, which are only half as long in the streamwise direction, were run using 20 streamwise modes. The resolution, $\Delta x^+ \approx 7$, $\Delta z^+ \approx 6$, with the first collocation point at $y^+ \approx 0.07$, is therefore excellent.

As the mask height is raised, the effective Re_τ increases, and the grid becomes coarser by about 50% in wall units. The resolution of those simulations, which are essentially full-height minimal channels, is therefore correspondingly worse, but still adequate.

3. Results

In a first series of numerical experiments the vertical mask was set to $\delta_1^+ \approx 50$, as suggested by the results of JP, and the wall-parallel box size was changed. The simulations were then run until they achieved statistically steady state. This usually required discarding an initial transient of length $t^+ \approx 5000$, and collecting statistics during times of the order of twice as long. The fluctuation profiles collected during the first and second halves of this latter period were visually indistinguishable. The boxes discussed here are those small enough to contain recognizable simple structures and a single low-velocity streak. Each simulation was initialized from a statistically converged instantaneous field from another box of roughly similar dimensions. The transition from long to short boxes was made by zeroing the odd streamwise Fourier coefficients of one of the long boxes, essentially averaging the two wavelengths contained in it. Because the basin of attraction of some of the solutions is fairly small, this initialization procedure was essential for their survival.

For two selected box lengths the height of the mask was varied systematically. As the mask height increases, especially in the cases in which it is comparable to the channel half-height, the extent of the parabolic part of the profile decreases, and Re_τ increases. Although some efforts were made to compensate for this effect by changing the length and width of the computational boxes, there is some systematic variation of both quantities with the mask height, when expressed in wall units. The lowest possible mask height was not explored in detail, but masks lower the $\delta_1^+ \approx 35$ laminarize the flow.

The time evolution of the flow was characterized in two ways. The first one is the evolution of the instantaneous skin friction, defined as the gradient at the wall of the plane-averaged streamwise velocity. Some representative time histories are shown in




FIGURE 2. Time evolution of the instantaneous friction coefficient: (a) —, $L_x^+ = 150$, $\delta_1^+ = 52$; ----, $L_x^+ = 152$, $\delta_1^+ = 55$; (b) —, $L_x^+ = 154$, $\delta_1^+ = 63$; ----, $L_x^+ = 161$, $\delta_1^+ = 78$. Note the different scales of the two figures.

figure 2. They are normalized with the skin friction of a parabolic laminar profile with the same mass flux, to separate them vertically for visual clarity but, because of the ambiguities mentioned above in the use of bulk variables, the only relevant normalization of the skin friction history is by its own mean value,

$$c_f^+ = \frac{\partial_y \bar{u}}{\partial_y U}. \quad (3.1)$$

The overline stands for instantaneous averaging over wall-parallel planes, and U is the long-time average of \bar{u} .

All the simulations in figure 2 are ‘short’ boxes containing a single wavelength of the low-velocity streak. For $35 < \delta_1^+ < 50$ the flow in these boxes is a permanent wave moving with a celerity $U_c^+ \approx 15$. One such case is the solid profile in figure 1(a), and the celerity is very close to the plateau in the mean velocity. Increasing slightly the mask height results in the two traces in figure 2(a), both of which can be shown to be two-frequency tori by spectral analysis of their evolution over a long time. Both have a short period $T^+ \approx 60$ and a long one $T^+ \approx 400$. The relative importance of both periods switches with the mask height. While for $\delta_1^+ = 52$ the long period is only a weak modulation of the short one, invisible at the scale of figure 2, the opposite is true when $\delta_1^+ = 55$.

As the mask is raised further, the flow becomes chaotic, with the longest period of the torus acting now as a dominant fast frequency. Finally, for $\delta_1^+ > 70$, intermittent bursts appear which are typical of fully developed turbulent flows, and whose amplitude is substantially higher than in any of the other three cases (figure 2b). The shorter bursting events are roughly of the same length as the long period of the torus, while the longest ones are two or three times longer. These times are comparable to those identified for individual bursts in minimal channels by Jiménez & Moin (1991). The present boxes are too short to maintain bursting turbulence indefinitely. A lot of energy is dissipated in the bursts, and the velocity fluctuations become very weak in the dips of the skin friction history. One of the bursts eventually fails to regenerate and the flow laminarizes, as seen at the end of figure 2(b). A similar phenomenon was described in marginal minimal channels by Jiménez & Moin (1991).

Figure 3 shows the same data in a representation introduced by Toh & Itano (1999). Orbits are drawn in terms of the two-dimensional and three-dimensional energy of

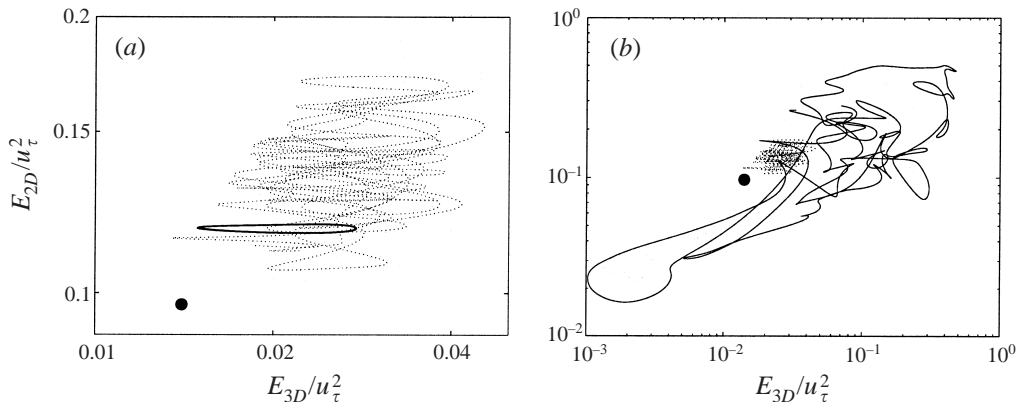


FIGURE 3. Orbits of the fluctuating energies. (a) \bullet , $L_x^+ = 145$, $\delta_1^+ = 49$; —, $L_x^+ = 150$, $\delta_1^+ = 52$; \cdots , $L_x^+ = 154$, $\delta_1^+ = 63$. (b) Symbols are as in (a) except for —, $L_x^+ = 161$, $\delta_1^+ = 78$. Note that the scales are logarithmic and different for the two figures.

the wall-normal velocity fluctuations. The two-dimensional energy is

$$E_{2D} = \frac{1}{H} \int_0^H \langle v_{2D}^2 \rangle dy, \quad (3.2)$$

where $v_{2D}(y, z, t)$ is the streamwise average of the fluctuation of v , and measures the intensity of objects which are long in the streamwise direction, essentially the wall-normal updraught responsible for the formation of the streaks. The three-dimensional energy is similarly defined for the remaining v fluctuations, $v_{3D} = v - v_{2D}$, and measures deviations from streamwise uniformity. We arbitrarily use $H = 0.56 \approx 70^+$, with the result that there is some inconsistency between energies from different simulations, as the top of the mask moves in and out of the integral in (3.2). This is however only an issue for fairly high masks, and does not affect any of the low-dimensional cases.

In this representation the permanent wave appears as a point, while the weaker torus of figure 2(a) appears as a slightly blurred limit cycle, wide in the sense of the three-dimensional energy and narrow in the two-dimensional one. A comparison of the behaviour of the energies during this cycle with the skin friction history in figure 2 shows that the skin friction changes in response to the variation of E_{2D} , while the faster oscillation of the three-dimensional energy does not appreciably affect it. The same is true for other cases.

Figure 3(b) underscores the amplitude difference between the simple flows which are obtained for low mask heights, and the bursts of the deeper channels. It also suggests that the simplifying effect of the mask is due to the damping of the large bursting excursions, while the low-activity periods of the deeper flows share with the constrained systems the same region of the energy space.

Similar behaviours are found for the longer boxes, which generally contain two wavelengths of the wavy streak. A summary of cases is found in figure 4(a), which classifies the different simulations in terms of their temporal behaviours. Note the narrow range of mask heights for which the time evolution is not fully turbulent, but note also that the bursting appears as a continuous modification of the weakly chaotic flows. In the longer boxes the bursting turbulence is self-sustaining and does not decay spontaneously.

The root-mean-squared amplitudes of the skin friction fluctuations with respect to their long-term averages, normalized as in (3.1), are collected in figure 4(b). The




FIGURE 4. (a) Classification of the evolution of the different computational experiments, as a function of the box length L_x^+ , and of the mask height δ_1^+ ; (b) root-mean-squared temporal variation of the plane-averaged wall friction. \circ , Permanent wave; \bullet , cycle or weak torus; \triangle , simple torus; \blacktriangle , complex torus or chaos; \square , bursting turbulence. The cases for which $L_x^+ < 200$ have been overlaid with a cross in (b).

bifurcation parameter that best collapses the different experiments is δ_1^+ , and the fluctuations increase quickly above a minimum threshold $\delta_1^+ \approx 50$. There are two branches in this figure, representing the long and the short boxes. The steeper branch consists of short boxes, and the lower fluctuation amplitudes of the long ones reflects the fact that not all the structures burst simultaneously, and that the amplitudes measured in the long boxes represent averages of bursting and quiescent structures.

A three-dimensional representation of the flow field in one of the longer boxes is shown in figure 5(a). It shows a low-velocity streak with two waves, associated with each of which there is a pair of quasi-streamwise vortices. They undergo a numerically exact limit cycle, in which the amplitude of the lateral displacement of the streak grows and wanes, and so do the vortices, while their heights oscillate slightly between $y^+ \approx 40$ and 50. The streak is never uniformly straight. These are the fast oscillations of the three-dimensional energy E_{3D} , that were seen above not to influence the skin friction appreciably. The fluctuations of E_{2D} , which have a strong influence on the friction, correspond to a slower modulation of the intensity both of the vortices and of the streak. Careful tracing of the vortex lines, although visually complicated, suggests that the streak and the vortices are parts of the same vorticity structure, and are not independent of one another. The permanent wave mentioned above when discussing the time histories is shown in figure 5(b), and is essentially equal to one of the two waves in 5(a). It should be emphasized that this wave is both permanent, moving at a constant celerity without change of shape, and autonomous, in the sense that there are no vorticity fluctuations in the computational box except those seen in the figure. It has been obtained from a continuous deformation of a full turbulent flow, and is therefore one of the best representations that we have of the elemental ‘engine’ that powers near-wall turbulence.

Figure 5(a,b) is qualitatively very similar to the travelling waves obtained in Waleffe (1998), and to the structures deduced from larger channels by Stretch (1990), Jiménez & Moin (1991), Webber, Handler & Sirovich (1997) and JP, as well as to the saddle point identified by Toh & Itano (1999). They all share the structure of a low-velocity streak flanked by vortices at roughly the same vertical location as the present ones.

Careful observation of the differences between computer animations of these




FIGURE 5. Three-dimensional representations of two of the flows discussed in the text. The grey shaded surface is $u^+ - \bar{u}^+ = -3.5$. The coloured objects are $\omega_x^+ = \pm 0.18$. The cyan quasi-streamwise vortices are negative, so that the net effect of the vortex system is to induce a positive updraught along the low-velocity streak. The flow is from top-right to bottom-left, and only the region below δ_1 is represented. (a) A regular cycle with $L_x^+ = 520$, $\delta_1^+ = 51$. (b) The steady wave discussed in the text, $L_x^+ = 145$, $\delta_1^+ = 49$.

shallow boxes and of deeper ones suggests the root of the increased complexity of the latter. As the streak goes through the cycle, it ejects some vorticity into the outer flow, essentially an asymmetric vortex hairpin. In full-depth flows this vorticity evolves, becomes disorganized, and eventually modifies the next cycle of the streak. In autonomous flows the vorticity is damped by the mask as soon as it is ejected, and this randomizing mechanism is not present. The threshold mask height at which bursting first appears is that at which the vortices are just allowed to move far enough from the wall to cross over the top of the streak and interact with one another.

4. Conclusions

We have shown that simple low-dimensional systems can be obtained from fully turbulent channels by a continuous set of transformations, the most important of which is the distance from the wall above which vorticity fluctuations are damped.

The simplest solution identified in this way is a permanent travelling wave, similar to those found by other investigators in related flows. As the vertical constraints on the system are relaxed, this wave bifurcates into either an approximate limit cycle or a two-frequency torus, depending on the parameters. The longest of the two periods of the torus evolves into temporal chaos, and eventually into the bursts of full-scale turbulence, when the vertical restriction is further relaxed.

The critical bifurcation mask height is $\delta_1^+ \approx 50$, and the range of heights in which subsequent bifurcations occur is very narrow, with the flow being essentially turbulent when $\delta_1^+ \approx 70$. The bifurcation mechanism was identified visually as the interaction, across the top of the low-velocity streak, of the two counter-rotating vortices.

Even in the very short boxes used in these simulations the only streamwise-coherent structure is the low-velocity streak. The quasi-streamwise vortices, their associated wall-normal velocities, and the resulting high-shear regions near the wall, are shorter and do not span the box length.

It should be emphasized that the fact that the near-wall region can be described by autonomous dynamics does not mean that there are no inner–outer interactions in real turbulent flows. Such effects, in the form of Reynolds number dependence of near-wall quantities, have been demonstrated often (Gad-el-Hak & Bandyopadhyay 1994). The models derived from descriptions such as the one in this paper can at most be lowest-order approximations, useful in the same sense as the Kolmogorov theory is a useful first approximation to the turbulent energy cascade, even if Reynolds number dependent intermittency corrections are known to exist.

M. P. S. was supported in part by grant CT98-0175 of the TMR program.

REFERENCES

- GAD-EL-HAK, M. & BANDYOPADHYAY, P. R. 1994 Reynolds number effects in wall-bounded turbulent flows. *Appl. Mech. Rev.* **47**, 307–365.
- JEONG, J., HUSSAIN, F., SCHOPPA, W. & KIM, J. 1997 Coherent structures near the wall in a turbulent channel flow. *J. Fluid Mech.* **332**, 185–214.
- JIMÉNEZ, J. 1999 The physics of wall turbulence. *Physica A* **263**, 252–262.
- JIMÉNEZ, J. & MOIN, P. 1991 The minimal flow unit in near wall turbulence. *J. Fluid Mech.* **225**, 221–240.
- JIMÉNEZ, J. & PINELLI A. 1999 The autonomous cycle of near wall turbulence. *J. Fluid Mech.* **389**, 335–359 (referred to herein as JP).
- KIM, H. T., KLINE, S. J. & REYNOLDS, W. C. 1971 The production of turbulence near a smooth wall in a turbulent boundary layer. *J. Fluid Mech.* **50**, 133–160.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* **177**, 133–166.
- ROBINSON, S. K. 1991 Coherent motions in the turbulent boundary layer. *Ann. Rev. Fluid Mech.* **23**, 601–639.
- SIROVICH, L. & ZHOU, X. 1994 Dynamical model of wall-bounded turbulence. *Phys. Rev. Lett.* **72**, 340–343.
- SMITH, C. R. & METZLER, S. P. 1983 The characteristics of low speed streaks in the near wall region of a turbulent boundary layer. *J. Fluid Mech.* **129**, 27–54.
- STRETCH, D. D. 1990 Automated pattern eduction from turbulent flow diagnostics. *CTR Ann. Res. Briefs, Stanford University*, pp. 145–157.

- SWEARINGEN, J. D. & BLACKWELDER, R. F. 1987 The growth and breakdown of streamwise vortices in the presence of a wall. *J. Fluid Mech.* **182**, 255–290.
- TOH, S. & ITANO, T. 1999 Low-dimensional dynamics embedded in a plane Poiseuille flow turbulence. Traveling-wave solution is a saddle point? *Proc. IUTAM Symp. on Geometry and Statistics of Turbulence* (ed. T. Kambe). Kluwer (in press).
- WALEFFE, F. 1997 On a self-sustaining process in shear flows. *Phys. Fluids* **9**, 883–900.
- WALEFFE, F. 1998 Three-dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.
- WEBBER, G. A., HANDLER, R. A. & SIROVICH, L. 1997 The Karhunen-Loève decomposition of minimal channel flow. *Phys. Fluids* **9**, 1054–1066.