Control of Intermittency in Near-wall Turbulent Flow

Peter N. Blossey and John L. Lumley

1 Introduction

The boundary layer exhibits intermittency at the interface between the turbulent fluid inside the boundary layer and the irrotational flow outside the boundary layer as well as internal intermittency. However, one finds in the boundary layer a third type of intermittency as well: the intermittent production of turbulence in the wall region of the turbulent boundary layer. The production of turbulence near the wall is dominated by a few, large-scale coherent structures which break down intermittently, resulting in peaks in turbulent production and the generation of small-scale turbulence — seen in experiments and simulations as bursts of Reynolds stress. Experimental work focusing on the breakdown of the coherent structures has led to the identification of a characteristic time scale, the mean inter-burst period. Lumley & Kubo (1985) collected the available data on the inter-burst period and found that the product of the bursting period with the turbulent skin friction at the wall was approximately constant. Thus, the shorter the bursting period becomes, the more active the coherent structures near the wall, and the higher the skin friction. If the bursting period is prolonged, the coherent structures are less active, and the skin friction will decrease. In this paper, we construct a control algorithm which identifies when the coherent structures are tending towards bursting and intervenes to suppress the burst, thereby increasing the inter-burst duration and decreasing the drag at the wall. Our control relies on the identification of the coherent structures in the boundary layer and on the prediction of their strength based on available information, namely the shear stress at the wall. However, some insight into the dynamics of the large scales of the turbulence in the wall layer — along with an estimate of their strength — is necessary to recognize the signs of bursting in the coherent structures.

Since the wall region is dominated energetically by just a few coherent structures, it is tempting to model the dynamics there including only these large-scale structures, while parameterizing the smaller scales of the turbulence. Many different techniques have been used to identify coherent structures in turbulent shear flows, from visualization and pattern matching to conditional sampling. Each of these techniques requires some degree of subjectivity in the identification of coherent structures, for example, in the choice of an initial pattern in pattern matching or in the selection of a threshold when using conditional sampling. However, Lumley (1967) proposed an objective technique for the identification of coherent structures in turbulent flows, the Proper Orthogonal Decomposition (POD). The POD, also referred to as the Karhunen-Loève decomposition, selects the most energetic feature from an ensemble of (random) turbulent velocity fields. In the wall region, the first and most energetic eigenfunction of the POD captures 60% of the turbulent kinetic energy there. The eigenfunctions of the POD are referred to as empirical eigenfunctions because they are derived from data (obtained either experimentally or through simulations) rather than from the governing equations. The empirical eigenfunctions form a basis for the turbulent velocity which possesses optimal convergence: when the velocity field is expanded in terms of the eigenfunctions, any truncation of the expansion will contain the most energy among all truncations of that order.

Using the empirical eigenfunctions as a basis, a model for the dynamics of the large-scales may be constructed from the Navier-Stokes equations by means of Galerkin projection. The projection yields a coupled system of nonlinear ordinary differential equations. When the equations are truncated — including only a few modes — the effect of the unresolved modes and the mean velocity profile must be accounted for. The unresolved modes exert a Reynolds stress on the large, resolved scales through the nonlinear terms. The Reynolds stress may be approximated using a gradient diffusion model with an eddy viscosity that may be varied parametrically. The mean velocity profile may be expressed as an integral of the Reynolds stress from the resolved modes. This introduces cubic terms into the model equations which globally stabilize the dynamical system which results from the low-dimensional model. These models may be studied through simulations and by use of the tools of dynamical systems theory.

The coherent structures observed in the boundary layer — the streaks and rolls — appear as fixed points in the models, and for certain values of the eddy viscosity parameter the models display intermittent dynamics with the trajectories in the phase space of the model following heteroclinic orbits connecting fixed points which correspond to the structures. The jumps in the phase space of the model correspond roughly to the bursts observed in simulations and experiments in the boundary layer. The updraft between the rolls strengthens into an ejection as energy is transferred from/to the smaller scales represented in the model. (The severe truncation in the models limits their ability to capture the energy transfer fully.) The energy transfer leads to the weakening and breakdown of the structures. Following their breakdown, the structures re-form, and the process repeats.

2 Motivation for Control

Our success in applying open-loop control to turbulent channel flow gave us confidence to take the same approach to closed-loop, feedback control with the aim of sustained drag reduction. The open-loop control (described in Lumley & Blossey (1998*a*)) interfered with the coherent structures near the wall, draining energy from the rolls and weakening momentum transport away from the wall. As a result, the drag was substantially reduced. Although the reduction in drag lasted much longer than the application of the control, the drag eventually rose again as the coherent structures strengthened and burst. The success of the open loop control led to two conclusions: that channel flow may be successfully controlled by focusing the control on only the large-scale structures, and that the effect of the control may last much longer than the duration of the control. This suggests an approach for applying control in a closed-loop, feedback control framework. The control may be applied as relatively short pulses whose purpose is to weaken the cross-stream coherent structures, or rolls. These pulses may be repeated when the coherent structures strengthen to prevent the bursting of these coherent structures. To be a realistic approach, the control must rely only on information available in a practical implementation. This information includes pressure and shear stress measurements at the wall, but not any information about the velocity field above the wall. Thus, our heuristic procedure of choosing the position of the forcing in the open-loop control — from observations of the high-speed streak — must be replaced by an objective method which relies, most likely, on the shear stress at the wall.

This control strategy is descended from work originating from our extended group at Cornell. A low-dimensional model was constructed for the wall region of the turbulent boundary layer by Aubry et al. (1988). Their model successfully mimicked the bursting process observed in experiments and simulations of near-wall turbulence, and was found to possess heteroclinic connections between fixed points in the phase space of the model. The fixed points — which were saddle points for realistic parameter values — represented rolls in the fluids systems. The jumps along the heteroclinic connections resembled the bursting process with a strengthening of the updraft between the rolls and a transfer of energy to the higher-wavenumber modes. If we accept the dynamical system as a reasonable model for the bursting process, the problem of control may be approached from this point of view. The first method to explore is stabilizing the fixed point: setting up linear feedback control with the aim of turning the saddle point into a stable fixed point. The stabilization of these saddle points would require sustained and substantial control input. This is not realistic in a practical setting. The difficulty of estimating the state of the coherent structures based on partial information available at the wall and the susceptibility of such a strategy to the dynamics of unmodeled modes would further complicate such an approach. In designing a control strategy for near-wall turbulence, our hope is that limited but intelligent application of our control effort would lead to efficient techniques for the reduction of drag.

Bloch & Marsden (1989) suggested another strategy for controlling the heteroclinic orbits that was more in line with the thinking outlined above. Their strategy relies on the identification of a controllable region around the stable manifold of a given saddle point. Once a trajectory enters the controllable region, the control is applied to prevent the escape of the trajectory along the unstable direction. In the absence of noise, the trajectory can be directed right to the fixed point and will sit in the neighborhood of the fixed point for all time. For a more realistic setting which includes noise, the trajectory will fall in towards the fixed point along the stable direction — with the control keeping the trajectory close to the stable manifold — until the controllable region becomes small enough that the noise may bump the trajectory out of the controllable region, leading to the escape of the trajectory along the unstable manifold. The controllable region usually takes the form of a cone which encloses the stable manifold and has its vertex at the fixed point. Therefore, the control can be expected to be most susceptible to noise near the fixed point. However, the goal of a practical control strategy is not to delay the heteroclinic jumps or bursts indefinitely, but only to increase the period between bursting events (the inter-burst time, T_B). By increasing the time between jumps/bursts, fewer turbulence-generating events will occur, less turbulence will be generated and the momentum transport away from the wall by the turbulence (skin friction drag) will be weakened.

Our strategy for control in the minimal flow unit, which contains a single set of coherent structures, is to apply the control directly to the coherent structures. Earlier control work (Coller et al. 1994*a*, Coller et al. 1994*b*) relied on an adjacent pair of vortices as an actuator for control, but the minimal flow unit is not large enough to permit the introduction of another set of structures without directly interfering with the naturally occurring coherent structures. The form of the control in the minimal flow unit will be similar to that used in our experiments with open-loop control and will be applied with a body force whose form is a Gaussian in the wall-normal direction which is largest near the wall (at $y^+ = 10$). The body force is applied to the first two spanwise Fourier modes, and its position is determined by our estimation technique which predicts the location of the rolls. The strength of the control is determined by our control algorithm which is outlined in section 4. The control is chosen to be invariant in the streamwise direction. Although the structures themselves are not uniform in the streamwise direction, streamwise invariance is a reasonable assumption in a small box like the minimal flow unit except when the structures break down. The control focuses on the prevention of the instability and bursting of the near-wall coherent structures. To this end, we will attempt to suppress the rolls when they grow strong. Weakening the rolls reduces the Reynolds stress generated by the near-wall structures and prevents the ejections which lead to instabilities and bursts. However, before we can describe the control strategy in detail, the problem of estimating the state of the flow based on measurements at the wall must be addressed.

3 Estimation

In our attempt to move from open-loop control to some form of feedback control, one key input to the control — an estimate of the state of the flow — is necessary. Our control strategy must incorporate information from the flow so that the control may adapt. The control will be applied to the flow selectively, i.e. only when the structures are tending towards bursting, in an attempt to maximize the effectiveness of the control. To this end, our estimate of the state of the flow will focus on the strength of the coherent structures. Our strategy for estimation draws from the work of Podvin & Lumley (1998), who applied the Proper Orthogonal, or Karhunen-Loève, Decomposition (POD) to velocity fields from the minimal flow unit and used the eigenfunctions of the POD to construct a low-dimensional model for the large-scale flow structures in the minimal flow unit. (For background on low-dimensional models and the POD, see Holmes et al. (1996) or Lumley & Blossey (1998*b*).)

The POD provides a complete expansion of the velocity field in terms of orthogonal eigenfunctions. The convergence of this expansion is optimal in the sense that any truncation contains, on average, the most kinetic energy of any truncation of that order. (The POD yields Fourier modes when applied to homogeneous directions — in this case, the streamwise and spanwise directions — so that the velocity field is expanded in Fourier modes in the homogeneous directions and eigenfunctions in the inhomogeneous directions.)

$$u_i(x,y,z,t) = \sum_{n=1}^{\infty} a^{(n)}_{k_1 k_3}(t) e^{2\pi i\left(\frac{k_1 x}{L_x}+\frac{k_3 z}{L_z}\right)}\ \phi^{(n)}_{i_{k_1 k_3}}(y) \tag{3.1}$$

The wall shear stress can be expanded in terms of the wall-normal derivatives of the eigenfunctions.

$$\frac{\partial u_i}{\partial y}(x, y=0, z, t) = \sum_{n=1}^{\infty} a^{(n)}_{k_1 k_3}(t) e^{2\pi i\left(\frac{k_1 x}{L_x}+\frac{k_3 z}{L_z}\right)}\ \frac{\partial \phi^{(n)}_{i_{k_1 k_3}}}{\partial y}(y=0) \tag{3.2}$$

We are interested in tracking the strength of the first two spanwise Fourier modes of the most energetic eigenfunction. These two modes combine to form the streaks and rolls that we have been talking about in our picture of turbulence generation in near-wall turbulent flow. The coefficients of these two modes, denoted by $a^{(1)}_{01}(t)$ and $a^{(n)}_{02}(t)$, can be estimated from the wall shear stress very simply if we assume that the wall shear stress in those two Fourier modes comes only from those two modes. Remember that the first eigenfunction carries the most energy among all eigenfunctions. Furthermore, the first eigenfunction in the wall region of the turbulent boundary layer carries more than half of the kinetic energy there. We truncate our eigenfunction expansion, including only the first eigenfunction, to determine the coefficients of the eigenfunctions from the shear stress. This truncation is a good approximation when the flow is dominated by the coherent structures. When the structures break down, we can expect higher eigenfunctions to play a larger role and perhaps degrade the accuracy of our estimate.

$$\frac{\partial \hat{u}_{i_{k_1k_3}}}{\partial y}(y=0,t) \approx a^{(1)}_{k_1k_3}(t)\frac{d\phi^{(1)}_{i_{k_1k_3}}}{dy}(y=0) \tag{3.3}$$

$$a^{meas}_{k_1k_3}(t) = \beta\left(\frac{D\hat{u}_{1_{k_1k_3}}}{D\phi^{(1)}_{1_{k_1k_3}}}\right)_{wall} + (1-\beta)\left(\frac{D\hat{u}_{3_{k_1k_3}}}{D\phi^{(1)}_{3_{k_1k_3}}}\right)_{wall} \tag{3.4}$$

Here, $D \equiv \partial/\partial y$ and $a^{meas}_{k_1k_3}$ refers to our instantaneous estimate ("measurement") of $a^{(1)}_{k_1k_3}$. We drop the superscript in $a^{meas}_{k_1k_3}$ since we only estimate the strength of a single eigenfunction for a given Fourier mode (k_1, k_3). The scalar β determines the relative weight of the spanwise and streamwise shear stress in determining the coefficients. The two modes that we want to estimate have wavenumbers $k_1 = 0$ and $k_3 = 1, 2$. The streamwise and spanwise shear stress for modes of this form (with no streamwise variation) come from the streaks and rolls, respectively. The rolls are disturbances of the form $\mathbf{u}(x,y,z) = (0, u_2(y,z), u_3(y,z))$ which do not directly affect the streamwise shear stress, except indirectly through their nonlinear interaction with the mean velocity which supplies energy to the streaks. Similarly, the streaks do not affect the spanwise shear stress. For the purposes of control, we are more interested in the rolls and choose a value of β close to zero. In fact, Podvin & Lumley (1998) found that their estimation was more successful for such values of β. The importance of spanwise shear stress in the prediction of ejections was also highlighted by Lee et al. (1997), who employed neural networks to optimize the prediction of strong ejections and sweeps in the wall region based on wall shear stress measurements.

The procedure outlined above will give us instantaneous "measurements" of the coefficients, which will likely be aliased and noisy. Our estimate comes from the application of a simple filter to the time series of instantaneous measurements. The filter takes the form of a dynamic equation:

$$\frac{d}{dt}a^{est}_{k_1k_3} = \frac{a^{meas}_{k_1k_3} - a^{est}_{k_1k_3}}{T} \tag{3.5}$$

Here, T is a timescale. We find that values of T on the order of 40 wall time units are effective. The choice of T involves a balance between the smoothing of non-physical oscillations in measurements from aliased higher frequencies and the time lag inherent in choosing a large T. The performance of the estimation using an array of six rows of eight sensors each in the minimal

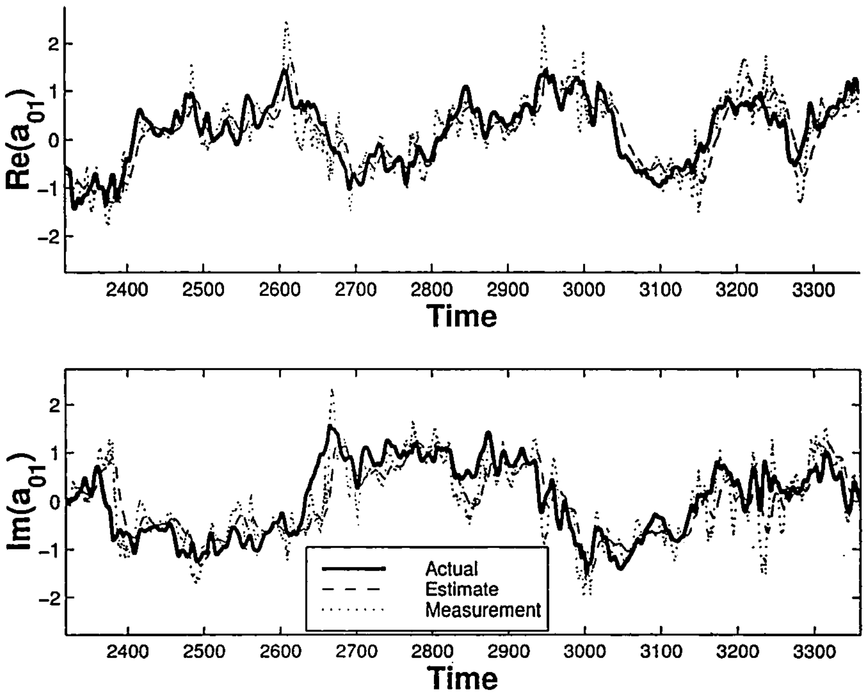

Figure 1: A comparison of our estimate of $a_{01}^{(1)}(t)$ with its actual value, computed by taking an inner product between the velocity field and the eigenfunction $\phi_{i01}^{(1)}$.

flow unit is displayed in Figure 1. The spacing of the sensors is approximately 22 wall units in the spanwise direction and 60 wall units in the streamwise direction. This array of sensors covers the entire wall in our minimal flow unit ($L_z = 184, L_z = 368$).

For the purposes of our control, we are interested in the strength of the rolls in particular. However, in the interest of a robust estimation scheme, it is useful to include both streamwise and spanwise shear stress in the estimation routine. In fact, a value of $\beta = 0.1$ was most successful in providing an effective control for drag reduction. The inclusion of the streamwise shear stress can be seen as "contaminating" our measurement of the strength of the rolls, but the streaks (through the streamwise shear stress) do provide an indication of the past strength of the rolls, since the rolls give rise to the streaks through interaction with the mean flow. Nevertheless, the quality of our estimate suffers when the control is turned on, as can be seen in Figure 2.

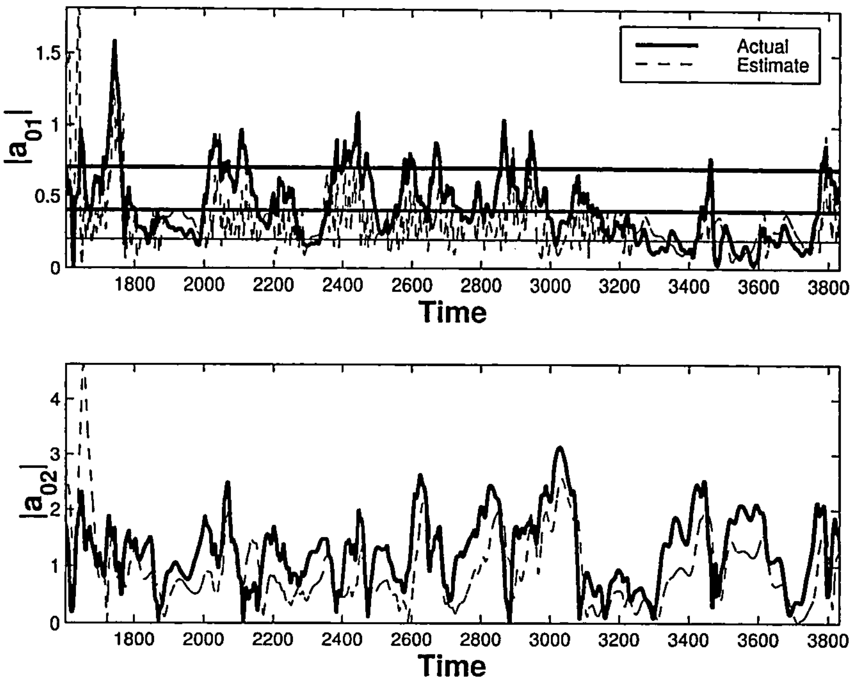

Figure 2: The estimation technique is sensitive to the control and is not as successful in predicting the strength of the first mode when the control is switched on. The horizontal lines correspond to the levels at which the control is switched on and off.

The estimate is much more sensitive to the control than the actual value of the coefficient: the control has the effect of driving the value of the estimate towards zero. After the control is switched off, the estimate increases and approaches the actual value of the coefficient until the control switches on again and drives the estimate back towards zero. The effect of the control on the estimate — driving it towards zero — apparently results from a decoupling of the shear stress at the wall, which provides the input for our estimate, from the coherent structures when the control is switched on. The causes of this will be explored in section 6, when we take a look at the effect of the control on the eigenfunctions of the POD. While the estimate is increasing and decreasing with the pulses of control in Figure 2, the shear stress and the actual value of the coefficient rise slowly but steadily. This seems to call into question the effectiveness of our control. However, if the control were switched off, the increases in shear stress and in the strength of the rolls would be much larger. In fact, turning off our control leads to a quick return of the shear stress to the levels of the uncontrolled flow. The motivation for our control strategy, which is based on a discrete switching algorithm, is described in the next section.

4 Control Algorithm

Our approach to control is simply to limit the strength of the rolls in the minimal flow unit. The rolls supply energy to the streaks through nonlinear interaction with the mean velocity profile. In this way, the rolls cause an increase in momentum transport away from the wall and drag at the wall. In addition, by strengthening the streaks, the rolls promote instability and the generation of smaller-scale turbulence which gives rise to a further increase in momentum transport and shear stress at the wall. By acting to control the rolls, we attempt to intervene and prevent the production of turbulence and an increase in shear stress at the wall. This strategy is an attempt to mimic the success of a passive technique for turbulence control, polymer drag reduction. The polymers act only at the smallest scales of the turbulence by expanding in the fluctuating strain rate field just outside of the viscous sublayer and damping out the smallest eddies there. However, their action at the smallest scales has a significant effect on the large scales: the rolls are significantly weakened in polymer drag reduced flow, while the streaks grow stronger. We take this configuration of the coherent structures as a cue and focus on using the control to weaken the rolls when they grow strong while simply ignoring the strength of the streaks. Although the streaks are likely to burst sometimes while we focus on the rolls, our control effort will be most efficiently applied to the rolls which are less energetic than the streaks. Both in the polymer drag reduced case and with our control, the flow will sometimes burst — this is nearly unavoidable — but our goal is not the complete suppression of the turbulence. The rolls and streaks are part of the turbulence, and completely suppressing them would be nearly impossible in a practical situation. Rather, our control attempts to suppress the instability of the rolls and streaks, so that this instability is triggered less often, less small-scale turbulence is generated, and momentum transport away from the wall is weakened.

The full fluid system does not exactly possess the heteroclinic jumps of the low-dimensional models. The heteroclinic connections in the phase space of the model are partly a result of the combination of the roll and streak modes into a single mode in the dynamical system of Aubry et al. (1988). The behavior of the full fluid system does not seem to possess the simple geometry of a saddle point as the rolls and streaks give rise to a burst. It appears that the streaks do not change in strength very substantially during a burst. The rolls do strengthen in the time leading up to a burst, but the substantial signal of a burst lies in the streamwise-varying ($k_1 \neq 0$) modes which grow significantly in strength as the burst develops. (Hamilton et al. (1995) provide a nice description of the breakdown and regeneration of these structures in turbulent Couette flow.) However, for the purposes of our control, these streamwise-varying modes would be very difficult to track and become strong only when the structures have started to break down. Our control aims to prevent the breakdown of the structures and, as a result, focuses on the rolls. We attempt to limit the strength of the rolls when they appear to be growing and tending towards a burst. This control strategy draws from the work of Guckenheimer (1995) who constructed a class of hybrid (discrete switching plus linear feedback) controllers for problems with two unstable eigendirections. The linear feedback controllers have a limited domain of attraction and may be susceptible to disturbances and unmodeled dynamics. Guckenheimer proposed that the domain of attraction of these linear controllers might be increased substantially by applying discrete switching control to trajectories which escaped from the domain of attraction of the linear controllers. The amplitude and direction of this control would be piecewise constant over stretches of phase space, presumably increasing in amplitude as the trajectories moved farther away from the fixed point. The discrete control does not attempt to direct the trajectory gracefully and efficiently back to the fixed point, rather the control switches on and drives the system directly towards the linear controller's domain of attraction. When the domain of attraction is reached, the discrete control switches off and the linear feedback controller takes over. As a demonstration of his concept, Guckenheimer (1995) constructs a hybrid control strategy for the inverted double pendulum whose domain of attraction is considerably larger than that of a linear feedback controller alone.

For our fluid system, we wish to borrow from the hybrid control strategy outlined above. However, we do not have a clear-cut, low-dimensional dynamical system underlying our fluid simulations. The rolls have a stable fixed point at zero strength when the flow is laminar (since our simulations are subcritical), but the domain of attraction is apparently small and would be difficult to attain. We will not attempt to apply linear feedback control when the strength of the rolls is small.[1] However, we will apply discrete switching control to the rolls when they grow strong. Our "hybrid" strategy relies solely on the discrete switching control to limit the strength of the rolls. One point must be addressed here before describing the control explicitly: the control reacts in response to our estimate of the strength of the rolls, not to the rolls themselves. The estimate is susceptible to aliasing, a finite response time, and other difficulties, and should not be accepted without skepticism. If the estimate is varying wildly, the estimate should probably not be used to guide the control. Either the estimate is responding to higher modes which have been aliased down to the wavenumber of the coherent structures — or the coherent structures are drifting quickly and will not respond well to the ap-

[1]Some groups apply linear feedback control to channel flow, with feedback laws based on suboptimal control theory (Lee et al. 1998), neural networks (Lee et al. 1997), or arguments about the suppression of vorticity flux (Koumoutsakos 1997).