Four-dimensional space-time correlations in wall turbulence

S. Giussani, M. Quadrio

Dipartimento di Ingeneria Aerospaziale, Politecnico di Milano

Abstract

Space-time correlations are often considered a perfect candidate for the meaningful description of wall turbulent flows, but their complete determination, either from laboratory experiments or Direct Numerical Simulations (DNS), has proven to date to be a challenging task. We have recently developed an innovative implementation for a low-cost, parallel DNS of turbulent wall flows, and we prove in this work that this approach can be extended to deal with the demanding tasks of computing the full correlation tensor, which implies a computational load significantly larger than the DNS itself.

To build the complete, spatio-temporal correlation tensor, a numerical simulation is purposely carried out, through which a suitably defined database of velocity fields is stored. Its optimized format makes possible an efficient access when, during the postprocessing stage, the correlation functions are computed. The database is stored in a distributed manner, so that a truly parallel I/O during the DNS is realized.

We present in this work a short review of typical quantities that can be extracted from such a database, with emphasis on the analysis of the correlations in the time domain. The main result is however the feasibility of such a demanding computation with commodity hardware. After some further optimizations, the present method will be able to compute the complete correlation tensor for higher-Reynolds-number turbulent flows.

1 Introduction

The description of turbulent flows over a wall is still a challenging problem. Theoretical difficulties and practical interest are such that during the last century a continuous effort has been devoted towards the objective of describing such flows in a way simple enough to be useful in applied problems, yet accurate and general enough to yield acceptable results.

Raw and accurate, multidimensional and time-dependent information on a turbulent flow field are now available, thanks to the significant advances recently made by experimental techniques (for example [5]) and numerical simulations, mainly the Direct Numerical Simulation of the incompressible Navier–Stokes equations [9]. To gather such a large amount of information and to process it in such a way that a compact and meaningful description of the flow is readily available represents a significant increase in the difficulty of the experiment or the numerical simulation.

Perhaps the first reliable description of wall turbulence that meets these often conflicting requirements has been the statistical description of the flow, based on the knowledge of the full correlation tensor for the velocity components. Correlations contain indeed a wealth of information: spectral energy density, spatial and integral scales where turbulent motions are significant are only a short list of quantities that can be extracted from the correlation functions. This explains the reason why experimental measurements of perhaps a single component of the tensor as a function of only *one* (spatial or temporal) coordinate can be traced back to the early '50s (see for example Favre, Gaviglio & Dumas, [3]). In more recent years, Kreplin and Eckelmann [8] in an oil channel explored the dependence of the correlation tensor on the various spatio-temporal separations, one at a time, and highlighted the convective nature of the flow in the viscous sublayer, educing space-time features of the supposedly dominant near-wall coherent structures, and opening the way to the full understanding of the concept of convection velocity in the near-wall region. Notwithstanding its fundamental interest, the complete measurement of the correlation tensor has been to date prohibitively difficult. According to the recent paper by Phillips [11], reliable estimates of the full tensor are still missing: even DNSs typically encounter significant problems when dealing with such quantities, since the acquisition of space-time correlations is time- and storage-intensive. Among the available data computed from a DNS, we mention those by Kim and Hussain [6], and the recent work by Quadrio and Luchini [13], where the wall-normal direction is not considered. Similarly, Jeon *et al.* in [4] focused on the spatio-temporal characteristics of friction and wall pressure, but disregarding the wall-normal direction.

The complete evaluation of the tensor requires a computational load roughly an order of magnitude higher than a standard DNS, and it is currently believed [11] that the latter can be afforded only when a supercomputer is available. Recently, Quadrio and Luchini [12], [15], [14] developed an ingenious method for the parallel DNS of turbulent wall flows, which is designed to run efficiently on low-cost commodity hardware. Its key property is the use of high-accuracy finite differences schemes for the discretization of the differential operators in the wall-normal direction: this reduces the communication among the computing machines to a minimum, so that the expensive, ultrafast networking hardware which connects the CPUs of a supercomputer is not needed, and the simulation can be run on simple commodity Personal Computers connected with standard Ethernet cards.

The aim of the present work is to investigate whether the same low-cost approach can be extended from the DNS itself to the complete calculation of the correlation tensor, which is an even more challenging problem from the computational point of view.

2 Numerical method and computational parameters

We describe in this Section the numerical method employed for the numerical simulation of the flow, while the computational procedures related to the construction of the database and its postprocessing for the calculation of the correlation tensor will be discussed in §3.

We solve the Navier-Stokes equations for the flow in a plane channel with DNS. A pseudo-spectral solver for an incompressible fluid written in cartesian coordinates [12] is employed. At the wall, the usual no-slip and no-penetration conditions are applied. As demonstrated by [7], periodic boundary conditions can be adopted in the homogeneous directions, provided the computational domain is large enough. Fourier expansions are hence used for these directions, whereas fourth-order accurate, compact finite-differences schemes over a variable-spacing mesh discretize the differential operators for the wall-normal coordinate.

The calculations are performed at a value of the Reynolds number of $Re_\tau = 180$, being Re_τ based on u_τ, the friction velocity, and on h, half the distance between the channel walls.

The streamwise (x) length of the computational domain is discretized with $N_x = 321$ Fourier modes, and $N_z = 129$ of them are used for the spanwise direction; the number of collocation points in the wall-normal (y) direction is $N_y = 129$. The dimensions of the computational box are $L_x = 4\pi h$ and $L_z = 4/3\pi h$, so that the spatial resolution in the reference case is $\Delta x^+ = 12$, $\Delta z^+ = 6$ and $\Delta y^+ = 0.7 - 4.7$ (the $+$ superscript indicates quantities made dimensionless with inner variables, i.e. with the friction velocity u_τ of the reference case and the kinematic viscosity ν of the fluid). The dimensions of the periodic box and the spatial resolution are comparable to those currently used in the literature for similar values of the Reynolds number, see for example [7] and [10]. The time step size is $t^+ = 0.075$.

The incompressible Navier–Stokes equations are written in terms of two scalar independent equations for the wall-normal component of velocity and vorticity vectors, with pressure removed from the equations, as done for example by [7]. The non-linear terms are evaluated in a pseudo-spectral way, and the related aliasing error (in the homogeneous directions) is removed by expanding the number of Fourier modes by a factor of at least 3/2 before transferring from Fourier space into physical space. Time integration of the equations is performed with the classic partially implicit approach, using a third-order, low-storage Runge-Kutta method for the convective terms and a second-order implicit Crank-Nicolson scheme for the viscous terms.

The code is able to exploit the computing power of shared-memory SMP machines, as well as connect together multiple machines for distributed parallel computing. An innovative parallel strategy is designed so that the overall amount of required memory can be subdivided among the computing nodes and the communication is reduced to a minimum: during the computations a complete transpose of the data among the computing nodes is never required. Each machine computes a wall-parallel slice of the whole computational domain; since the computing machines are identical, the whole channel height is subdivided in slices of equal size.

The code is designed for the highest computational efficiency, both in terms of RAM and CPU requirements. When run with a time-integration scheme that requires one past time level (like the third-order Runge-Kutta used in the present work), the code requires a storage space of the order of $5N_xN_yN_z/2$, while similar codes described in the literature require at least $7N_xN_yN_z/2$. For a 128^3 case, it needs 95 MB of RAM (double precision) and, when executed on a single Pentium III 733 MHz CPU, it takes 44 seconds for computing a full time step with a 3-substeps Runge-Kutta scheme.

The code has been run on a dedicated computing system, designed and built on purpose, made by 8 commodity SMP Personal Computers, each one equipped with two Intel Pentium III CPU of 733 MHz, 512 MB RAM and two cheap Fast Ethernet cards. Most of the machines are equipped with 19GB hard disks, but two of them have 14GB hard disks. The machines are connected together in a dedicated, ring-like connection topology which avoids any hub or switch, and replicates the way data are stored in the machines during the computations. Notwithstanding the commodity networking hardware, the calculations are never communication-bound.

3 Calculation of the correlation tensor

The generic component i, j of the correlation tensor between the velocity components u_i and u_j is defined as

$$R_{ij}(t_1, t_2, x_1, x_2, y_1, y_2, z_1, z_2) = \frac{\langle u_i(t_1, x_1, z_1, y_1) u_j(t_2, x_2, z_2, y_2)\rangle}{u_i^{rms}(y_1) u_j^{rms}(y_2)} \qquad i, j = 1, 2, 3$$

and from this definition it is evident that the tensor R_{ij} is simmetric.

In general, the correlations depend on 6 spatial coordinates and on two temporal coordinates. In the case of the turbulent flow in an indefinite plane channel, however, we can take advantage of the hypothesis of ergodicity and homogeneity in the streamwise and spanwise directions. Due to ergodicity, correlations do not depend on t_1 and t_2 separately, but on the temporal separation $\tau = t_1 - t_2$ only. Due to homogeneity in the spanwise direction, correlations do not depend on z_1 and z_2 separately, but on the spanwise separation $\zeta = z_2 - z_1$ only. Due to homogeneity in the streamwise direction, correlations do not depend on x_1 and x_2 separately, but on the streamwise separation $\xi = x_2 - x_1$ only. The wall-normal direction is inhomogeneous. In conclusion we have:

$$R_{ij} = R_{ij}(\tau, \xi, \zeta, y_2 - y_r; y_r) \qquad i, j = 1, 2, 3 \tag{1}$$

i.e. each component of the correlation tensor is a *4d function which depends on an additional parameter* (the reference distance from the wall y_r).

In principle, the calculation of the correlation tensor can be carried out following its definition (1), after a suitable database of velocity fields has been precomputed and stored during the DNS. In practice, however, such an approach can be very time-consuming, since this definition requires the evaluation of many convolution integrals, whose computational cost is very high. Moreover, since the size of the database is typically so large that the database cannot be kept in the main memory of the computer, the convolutions become even more costly, since they require repeated accesses to a database stored on disk, and the typical access times to disks are order of magnitudes higher than the access times to RAM. This is why efficiency plays a key role even in the postprocessing of the database to produce the correlation tensor.

3.1 Database design and layout

In this preliminary study, we use the somewhat old computing hardware described in the preceding Section, and the limited availability of disk space has been a primary design constraint. The database of velocity fields is distributed across the computing machines, each of whom stores as a function of time the values of the velocity components in its own wall-normal slice. In this way the limit for the maximum size of the database is around 100GB. The simulation is run for approximately $t^+ = 1000$, and whole velocity fields are stored on disk every $\Delta t^+ = 0.75$, i.e. every ten computational time steps. A total number of $N = 1280$ global flow fields are thus stored on disk for later postprocessing, in the form of Fourier coefficients of velocity components, as a function of time and of the distance from the wall.

For efficiency reasons R_{ij} cannot be computed after its definition, but it has to be evaluated by taking advantage of the Wiener-Kinchin theorem, which states that the correlation of a periodic function is the Fourier Transform of its spectrum. Since the spectrum in Fourier space can be easily calculated with a simple multiplication, the computational advantage of evaluating a convolution integral as the (fast) Fourier Transform of a product is readily evident. While in the database the velocity components are natively expressed as a function of streamwise α and spanwise β wavenumbers, and while the wall-normal direction is inhomogeneous and calls for the calculation of R_{ij} according to its definition, a significant saving can be attained if the time coordinate is treated according to the Wiener-Kinchin theorem. Since the first operation to be performed on the database will be a transform from the time domain into the frequency domain, we have optimized the database format so that it stores in contiguous positions values of velocity components of a given α, β, y as a function of time.

Of course, this is the most efficient format when reading the database, since the contiguous data are ordered with the most rapidly varying index in the innermost position. On the other hand, it represents the most unefficient format for writing, since the write operation happens at a given time for any α, β, y value. This latter disadvantage is however negligible, since ten DNS time steps have to be computed between two subsequent write operations, and this requires a computing time of the order of the order of a minute, a time interval during which the data can be safely trasnferred from the virtual memory of the computing units to their hard disks. In this way the I/O load is effectively parallelized, since all the disk controllers are used at the same time. Moreover, a buffering techinique has been employed, so that a few tenths of flow fields are written to disk at once after temporarily storing then in RAM in the meanwhile. We recall moreover than writing the database takes place only once, while the calculation of correlations can be repeated many times. Figure 1 is a sketch of our computing system (for simplicity only a limited number of computing machines is represented), where the ring-like connection topology among the nodes is evident. Each computing box (which may contain more than one CPU) own a wall-parallel slice of the channel, and stores to the local disk its part of the global database.

3.2 Database postprocessing

After the calculation of the full database, in the format described in the preceding Section, one has to compute from it the correlation tensor, which is a function of 4 spatio-temporal separations once the reference position y_r has been chosen. Since the statistical sample is somewhat limited, in the procedure of time averaging we have decided to double the available statistical sample by considering both walls: a further reference position $y_r^{(s)} = h - y_r$ is considered, and the correlations over the two channel halves are averaged together.

The minimum temporal separation τ_{min} considered in the present work cohincides with the time interval at which the velocity database is constructed. The maximum temporal separation τ_{max} has been selected after examining the longest significant temporal scales in the flow (see for example [13], and in the present work is $\tau^+_{max} = 112.5$. The averaging period (total length of the simulation, i.e. the temporal extent of the database) is approximately 10 times τ_{max}.

The first step to compute the correlation tensor is to read from the database the quantities $\hat{u}_i(\alpha, \beta, y)$ as a function of time, where $\hat{u}_i$ denotes the Fourier coefficient of the i-th component of the velocity vector. Then we evaluate $\hat{u}_i(\alpha, \beta, y, \omega)$ with a Fast Fourier Transform algorithm. The FFT of nonperiodical functions in the time domain is carried out with the usual technique of zero-padding [1]. Then the spectral energy density is computed, as

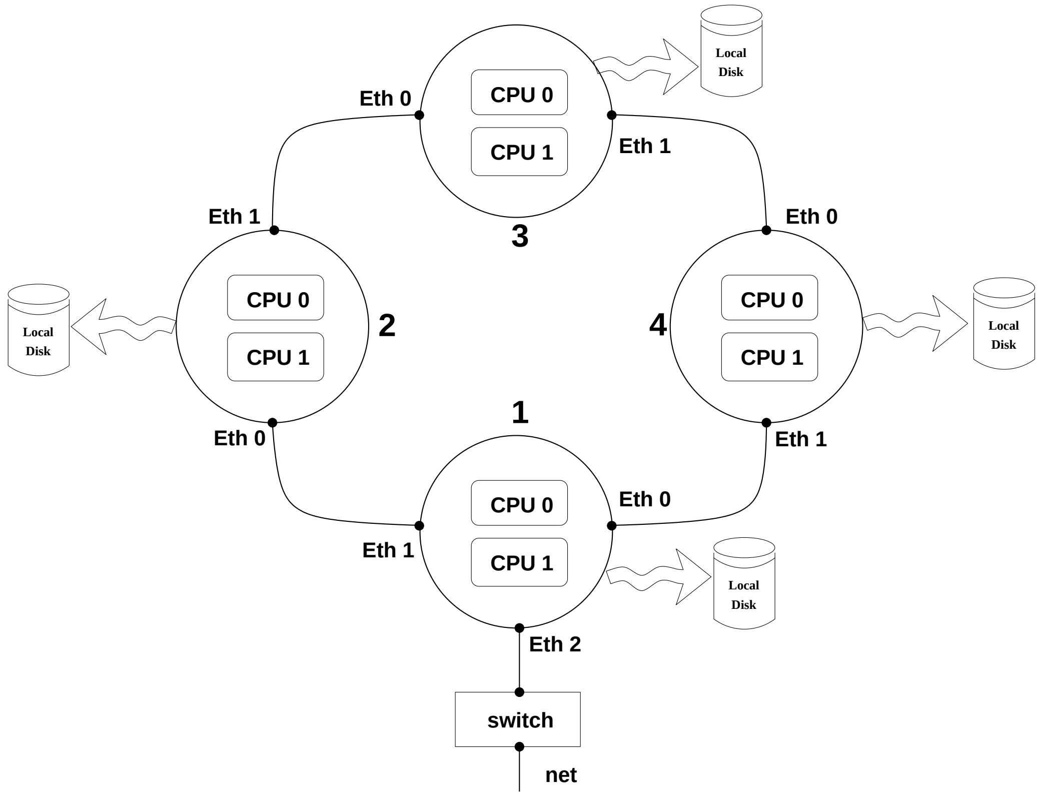

Figure 1: Sketch of the computing system (only a limited number of computing nodes are shown). It is evident the ring topology of the internode connections, and the use of local disks for the distributed storage the database.

$$S_{ij}(\alpha, \beta, \omega, y_r, y_2) = \hat{u}_i(\alpha, \beta, \omega, y_r)|\hat{u}_j(\alpha, \beta, \omega, y_2)$$

and averaged with the simmetrical position (indicated with the superscript s):

$$S_{ij}^{(s)}(\alpha, \beta, \omega, y_r^{(s)}, y_2^{(s)}) = \hat{u}_i(\alpha, \beta, \omega, y_r)|\hat{u}_j(\alpha, \beta, \omega, y_2)$$

The calculation of the tensor R_{ij} is carried out on a single machine, where the disk space does not allow the simultaneous presence of the whole database and the correlation tensor. The necessary parts of the database are thus dynamically transferred to the computing machine, and erased after use, while the database itself resides permanently in the distributed disk space across the computing system.

Each component of the correlation tensor has a size of approximately 6GB, for a total occupancy of 35 GB for the complete correlation tensor. The computing time is approximately 10 hours when an AMD AThlon CPU is used on a Personal Computer equipped with a standard EIDE disk.

4 Results

In this Section we report some preliminary results. All of them have been obtained by setting $y_r = 0$.

Figure 2 reports a planar view of the correlation function at the wall, where the longitudinal and spanwise components of friction and pressure are defi ned. It is readily evident that the longitudinal friction possesses a very long spatial scale. There is also (not shown) a small region of negative correlation for pressure. The time evolution of these quantities will be shown at the Conference.

The time correlation of the longitudinal friction (at zero spatial separations) is shown in fi gure 3. The interesting feature of this plot is the secondary minimum evident at a temporal separation of $\tau^+ \approx 180$. As recently discussed in [13], this is an effect of the periodic boundary conditions applied to a computational box of fi nite length, and raises the question whether a streamwise length of $L_x = 4\pi h$, which is a commonly accepted value, is really enough to obtain turbulence statistics which are independent on the box size. It is interesting to note that this question has been raised essentially at the same time by [2] and [13], where the former authors looked at velocity spectra while the latter examined correlations functions. It is well known that spectra and correlations are a Fourier transform pair, but each of these quantities presents its own practical advantages and disadvantages.

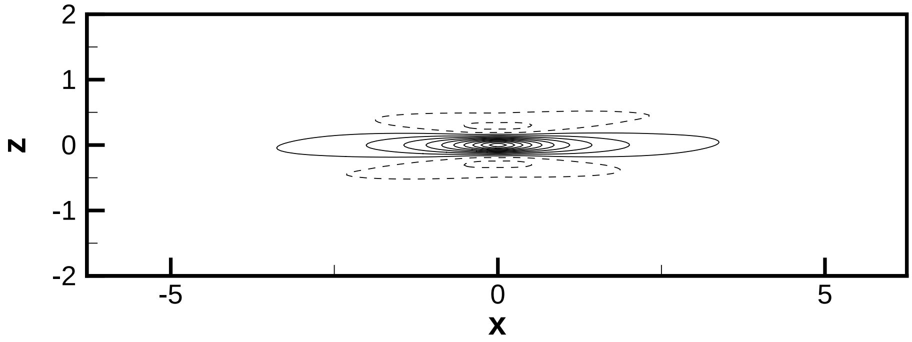

Figure 2: Correlation function for the longitudinal component of wall friction as a function of the longitudinal and spanwise separations, at $\tau = y_2 = 0$.

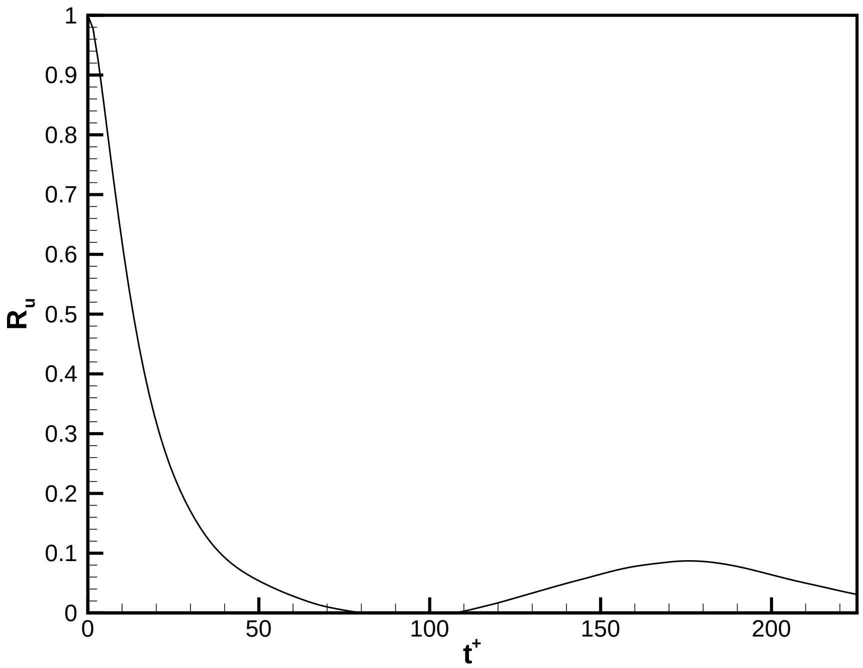

Figure 3: Time correlation for the longitudinal component of the wall friction, at $\xi = \zeta = y_2 = 0$. Note the secondary peak due to the artificial periodic boundary conditions.

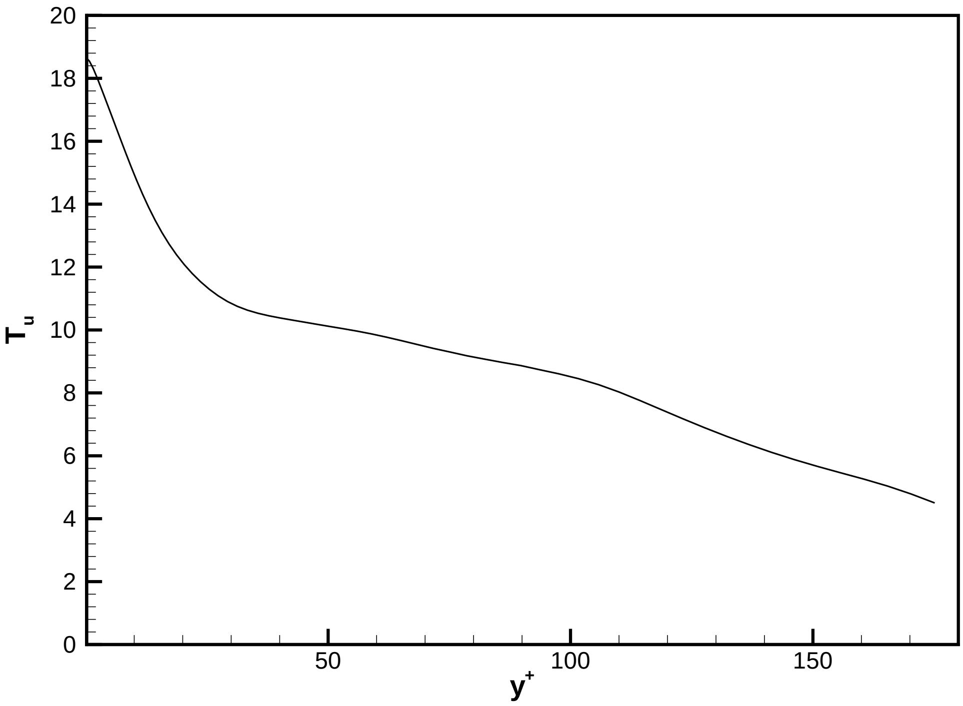

Figure 4: Variation with the distance from the wall of the integral longitudinal time scale T_u^+.

Correlation functions are often examined in terms of their integral scale. The integral scale is a dimensional quantity which expresses a typical (temporal or spatial) separation at which the signal decorrelates. For example the temporal integral scale for the correlation at a given reference distance from the wall y_r is a function of y_2, and it is measured as:

$$T(y_2) = \int_0^\infty R\,(\tau, 0, y_2 - y_r; y_r)\, d\tau$$

This quantity is shown in figure 4 for the longitudinal velocity component. The wall value appears to be in good agreement with the value of the integral time scale of the wall friction first computed with a DNS in [13]. The integral time scale decreases with the distance from the wall, but the decrease rate slows down at $y_2^+ > 30$ and at the channel centerline the value of T_u^+ is still a significant percentage of the value at the wall.

Figure 5 shows the spatial behaviour of the correlation for the longitudinal velocity component, at zero time and spanwise separation, as a function of the wall-normal and longitudinal separation. A slope of the contours towards the center of the channel can be appreciated, and the correlation is observed to remain significantly different from zero even in the central part of the channel. The time evolution of this quantity will be presented at the Conference.

As a last result, we report in figure 6 a different quantity, from the point of view of spectra energy density functions, computed from the same database. We have computed the premultiplied one-dimensional spectra of the longitudinal velocity component, for different distances from the wall. The two-dimensional spectral density function (for example for the longitudinal velocity component) is:

$$S_{uu}(\alpha, \beta; y) =< \alpha\beta\hat{u}(\alpha, \beta; y)|\hat{u}(\alpha, \beta; y) >$$

where α and β are the streamwise and spanwise wavenumbers, and $|$ indicates the multiplication for the complex conjugate. The one-dimensional spectrum is then obtained by summing over one direction.

A similar quantity has been recently reported by [2], who discussed the presence of turbulent motions of very large spatial scale. In particular they found that, when large spanwise scales are considered, their streamwise extension is very large and tends to exceed typical box sizes. We confirm this important result in figure 6. The plot represents the one-dimensional premultiplied spectrum of the u component, as a function of the distance from the wall, where only spanwise scales with spanwise wavelength $\lambda_z = 2\pi/\beta_0 > 0.7h$. It is seen that the longitudinal dimensions of the computational domain are not enough to accomodate the largest scales present in the turbulent flow.

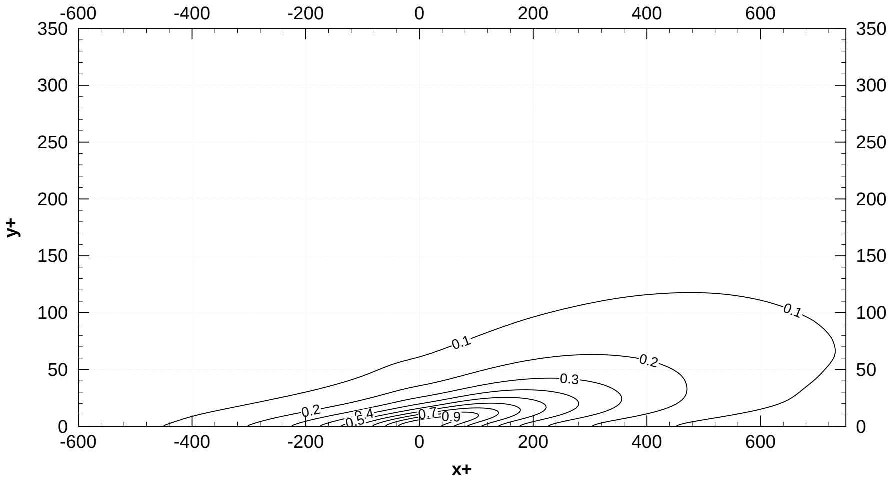

Figure 5: Contours of the correlation function, for the u velocity component, as a function of the streamwise and wall-normal separations, at $\zeta = 0$ and $\tau = 0$

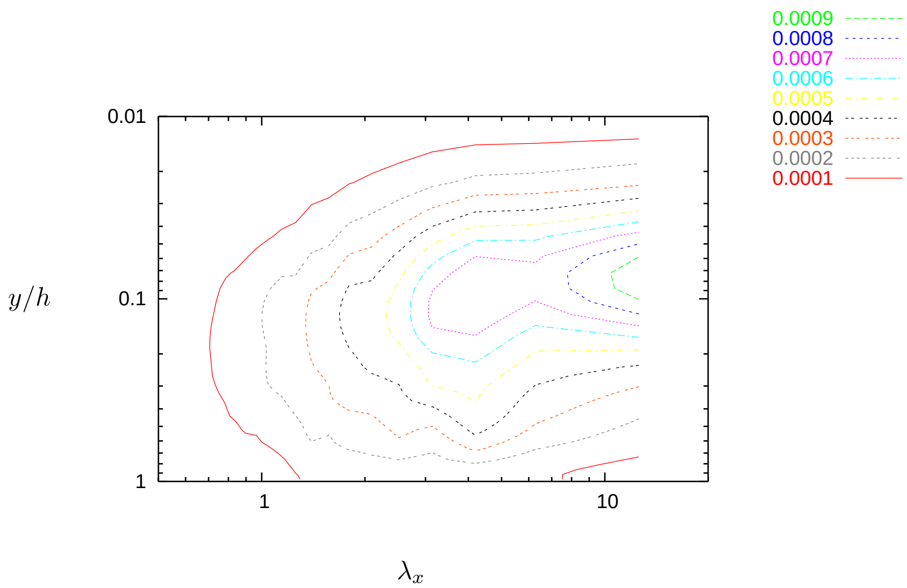

Figure 6: Premultiplied one-dimensional streamwise spectra for the longitudinal velocity component, as a function of the distance from the wall. Top: only spanwise modes with $\lambda_z > 0.7h$. Bottom: only spanwise modes with $\lambda_z < 0.7h$.

5 Conclusions

A computational procedure has been illustrated for the calculation of the full velocity space-time correlation tensor in a plane turbulent channel flow. The main result is that the full tensor can be evaluated with a limited computational effort and by using commodity hardware, if a database of velocity field is purposely built with suitable optimizations.

Correlation functions have been examined both in the space and time domain, illustrating how they can shed light on the spatial and temporal scales which characterize the turbulent motion. In particular the behavior of the integral timescale as a function of the distance from the wall has been examined. The same dataset has been moreover examined in terms of pre-multiplied one-dimensional streamwise spectra, to evidence how large-scale motions can be described in the turbulent plane channel flow.

After adding further optimizations to our computational procedure, we aim at the study of the correlation tensor (and its equivalent spectra density function tensor) at even higher values of the Reynolds number. In the present years an interesting debate is taking place in the scientific community, regarding whether inner, outer or both scaling can be applied to spectral quantities.

References

[1] J. S. Bendat and A. G. Piersol. *Engineering applications of correlation and spectral analysis*. John Wiley & Sons, 1980.

[2] J.C. del Álamo and J. Jiménez. Spectra of the very large anisotropic scales in turbulent channels. *Phys. Fluids*, 15(6):L41–L44, 2003.

[3] A. J. Favre, J.J. Gaviglio, and R.J. Dumas. Space-time double correlations and spectra in a turbulent boundary layer. *J. Fluid Mech.*, 2:313, 1957.

[4] S. Jeon, H. Choi, J. Y. Yoo, and P. Moin. Space-time characteristics of wall shear-stress fuctuations in a low-Reynolds-number channel flow. *Phys. Fluids*, 11(10):3084–3094, 1999.

[5] B.C. Khoo, Y.T. Chew, and C.J. Teo. Near wall hot-wire measurements. Part II: turbulence time-scale, convective velocity and spectra in the viscous sublayer. *Exp. Fluids*, 31(5):494–505, 2001.

[6] J. Kim and F. Hussain. Propagation velocity of perturbations in turbulent channel flow. *Phys. Fluids A*, 5(3):695–706, 1993.

[7] J. Kim, P. Moin, and R. Moser. Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.*, 177:133–166, 1987.

[8] H.-P. Kreplin and E. Eckelmann. Propagation of perturbations in the viscous sublayer and adjacent wall region. *J. Fluid Mech.*, 95(2):305–322, 1979.

[9] P. Moin and K. Mahesh. Direct numerical simulation: A tool in turbulence research. *Ann. Rev. Fluid Mech.*, 30:539–578, 1998.

[10] R. Moser, J. Kim, and N.N. Mansour. Direct numerical simulation of turbulent channel flow up to $Re_\theta = 590$. *Phys. Fluids*, 11(4):943–945, 1999.

[11] W. R. C. Phillips. Eulerian space-time correlations in turbulent shear flows. *Phys. Fluids*, 12(8):2056–2064, 2000.

[12] M. Quadrio and P. Luchini. A 4th order accurate, parallel numerical method for the direct simulation of turbulence in cartesian and cylindrical geometries. In *Proc. of the XV AIMETA Conf. on Theoretical and Applied Mechanics*, 2001.

[13] M. Quadrio and P. Luchini. Integral time-space scales in turbulent wall flows. *Phys. Fluids*, 15(8):2219–2227, 2003.

[14] M. Quadrio and P. Luchini. A low-cost parallel implementation of direct numerical simulation of wall turbulence. *In preparation.*, 2003.

[15] M. Quadrio, P. Luchini, and J.M. Floryan. A Parallel Algorithm for the Direct Numerical Simulation of Turbulent Channel Flow. In *Proc. of the XI Conf. of the CFD Society of Canada*. Vancouver (CAN), May 28-30 2003, 2003.