Channel 2000: Computing Turbulence at Experimental Reynolds Numbers

Sergio Hoyas[1] & Javier Jiménez[1,2]

1. School of Aeronautics, Madrid.
2. Centre for Turbulence Research, Stanford.

Funding:
DEISA, BSC,CICYT,PIC

Why 6e6 CPU-Hours?

Why we needed 6e6 CPU-Hours and 25TB of storage to run a simulation?

Energy Cascade

Kolmogorov (1941)

Big whorls have little whorls
that feed on their velocity,
and little whorls have lesser whorls
and so on to viscosity.

Richardson

Energy flux in homogeneus flows

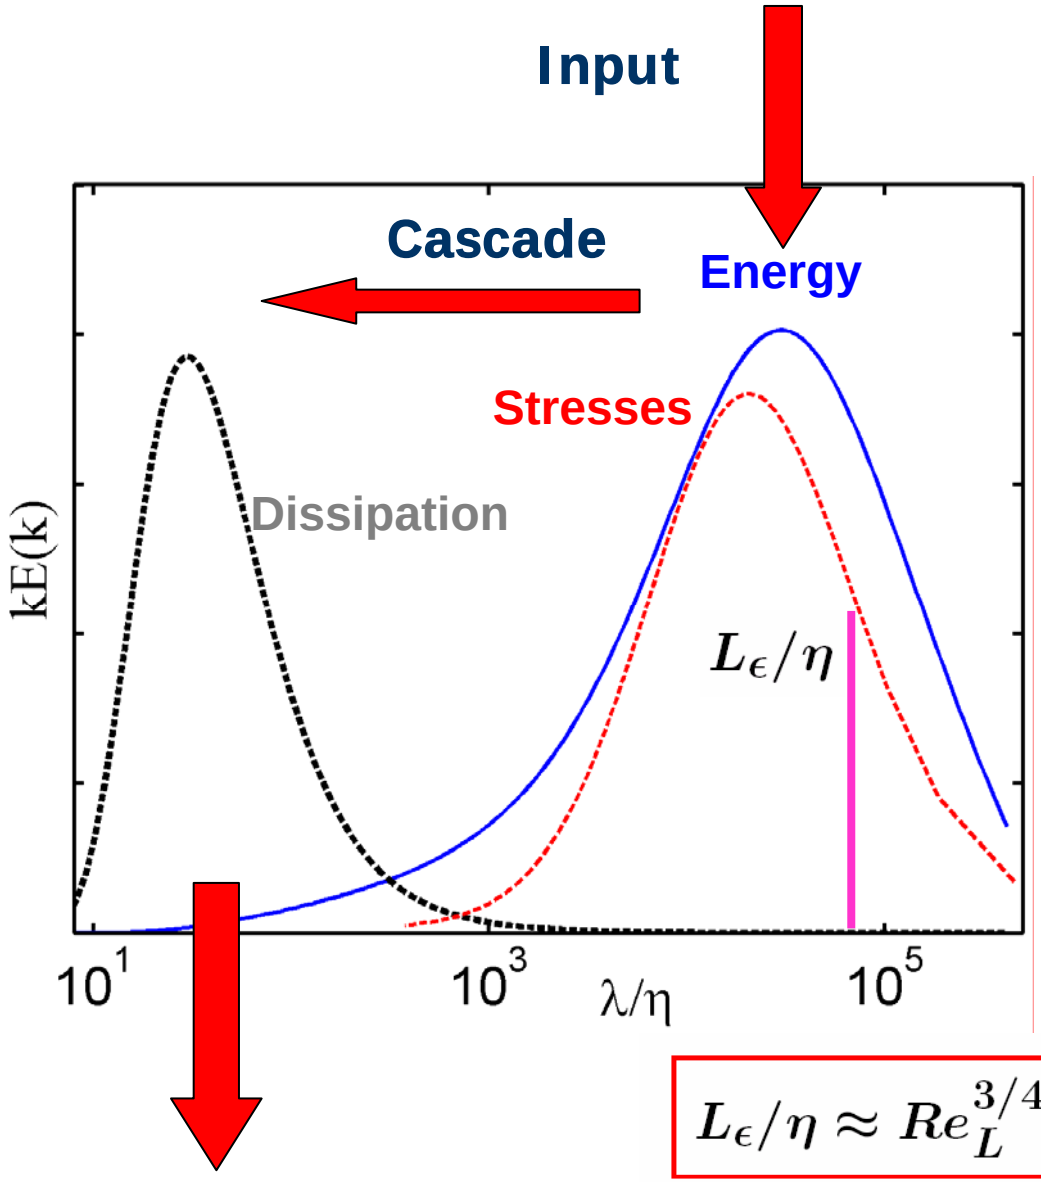

$$L_\epsilon/\eta \approx Re_L^{3/4}$$

$$\text{Dissipation: } 5 - 100\eta$$
$$\text{Energy: } 0.01 - 50L_\epsilon$$
$$\text{Stress: } 0.1 - 20L_\epsilon$$

Typical industrial flows:

$$Re_L = 6000 - 1000000$$

Wake of a person walking

$$Re_L = 16000$$

Boundary layer of a commercial aircraft

$$Re_L = 6000000$$

What is a turbulent channel?

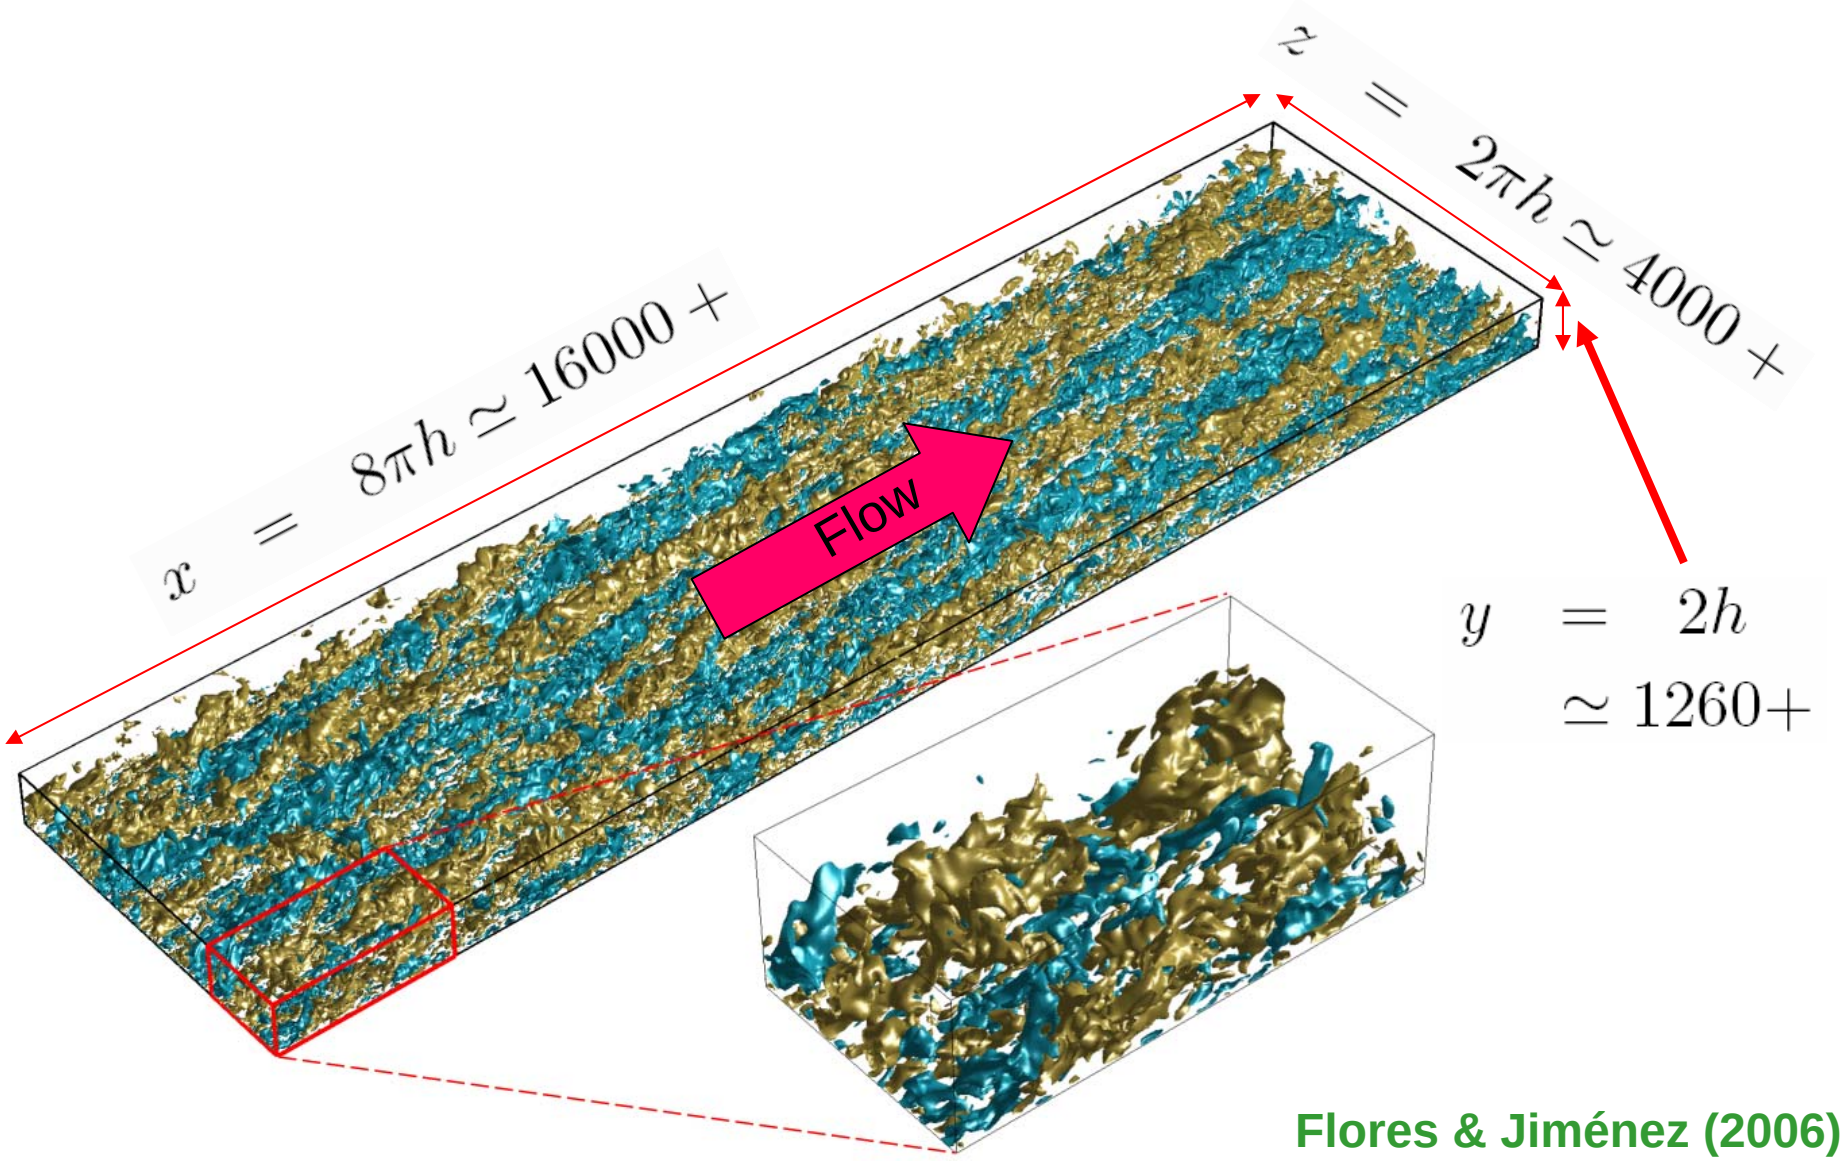

Flores & Jiménez (2006)

Cascades in wall turbulence

$Re_\tau = u_\tau h/\nu = 2000$

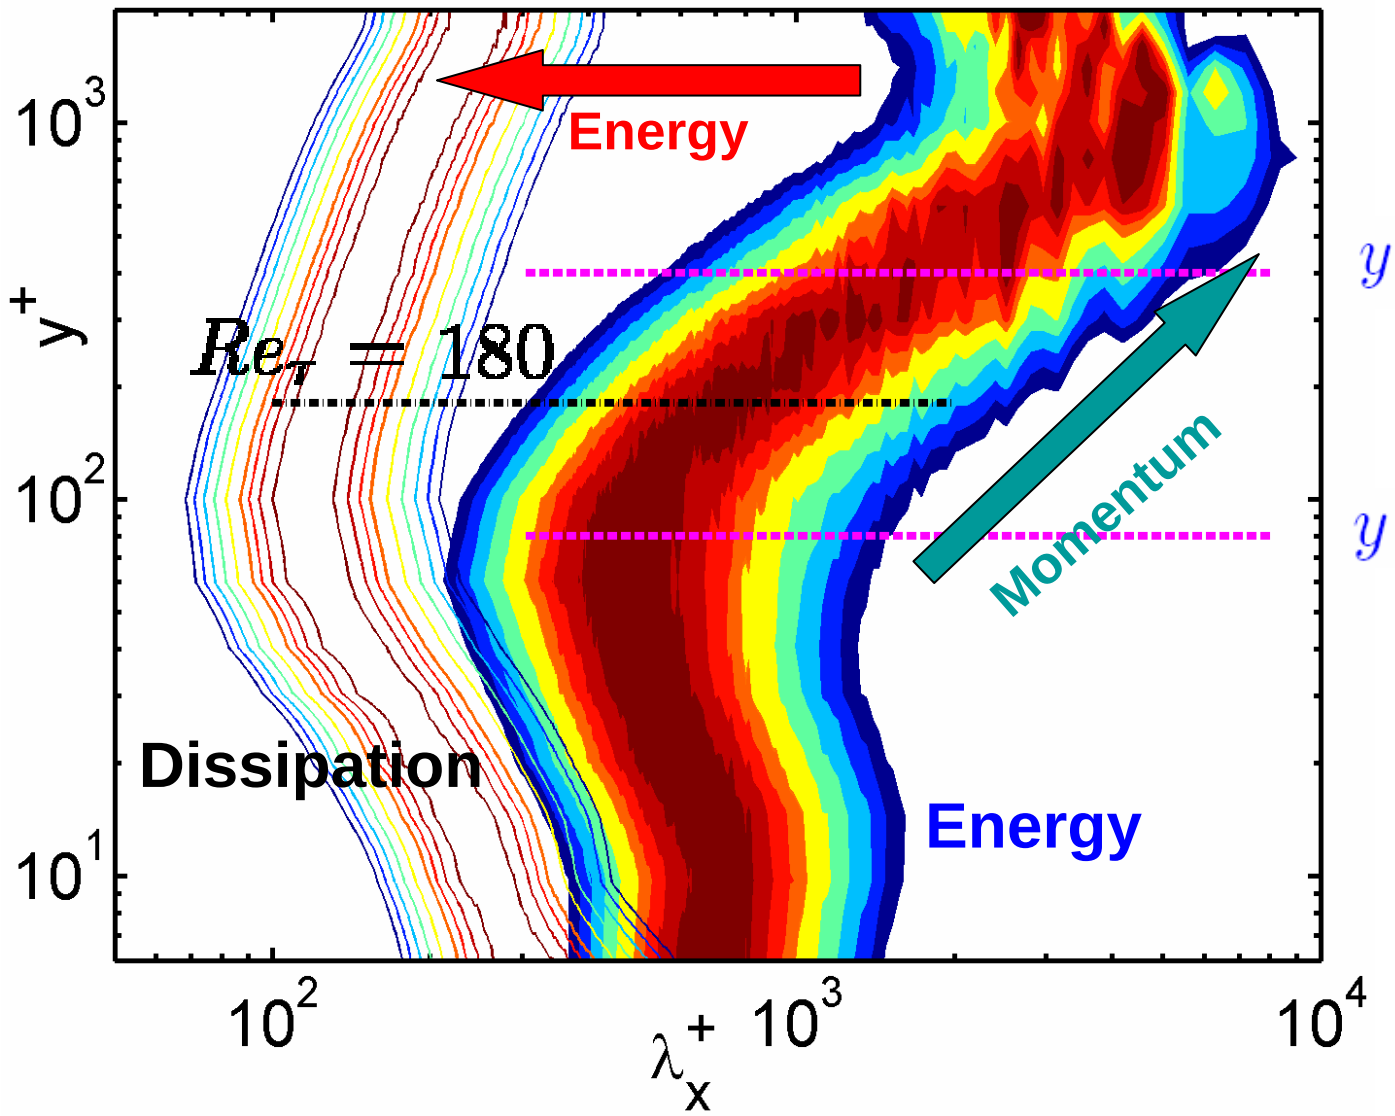

Outer

$y < 0.2h = 400^+$

Log

$y = 80^+$

Viscous and Buffer

Domain and grid size

Grid size

$$N_x = 6144$$
$$N_y = 633$$
$$N_z = 4608$$

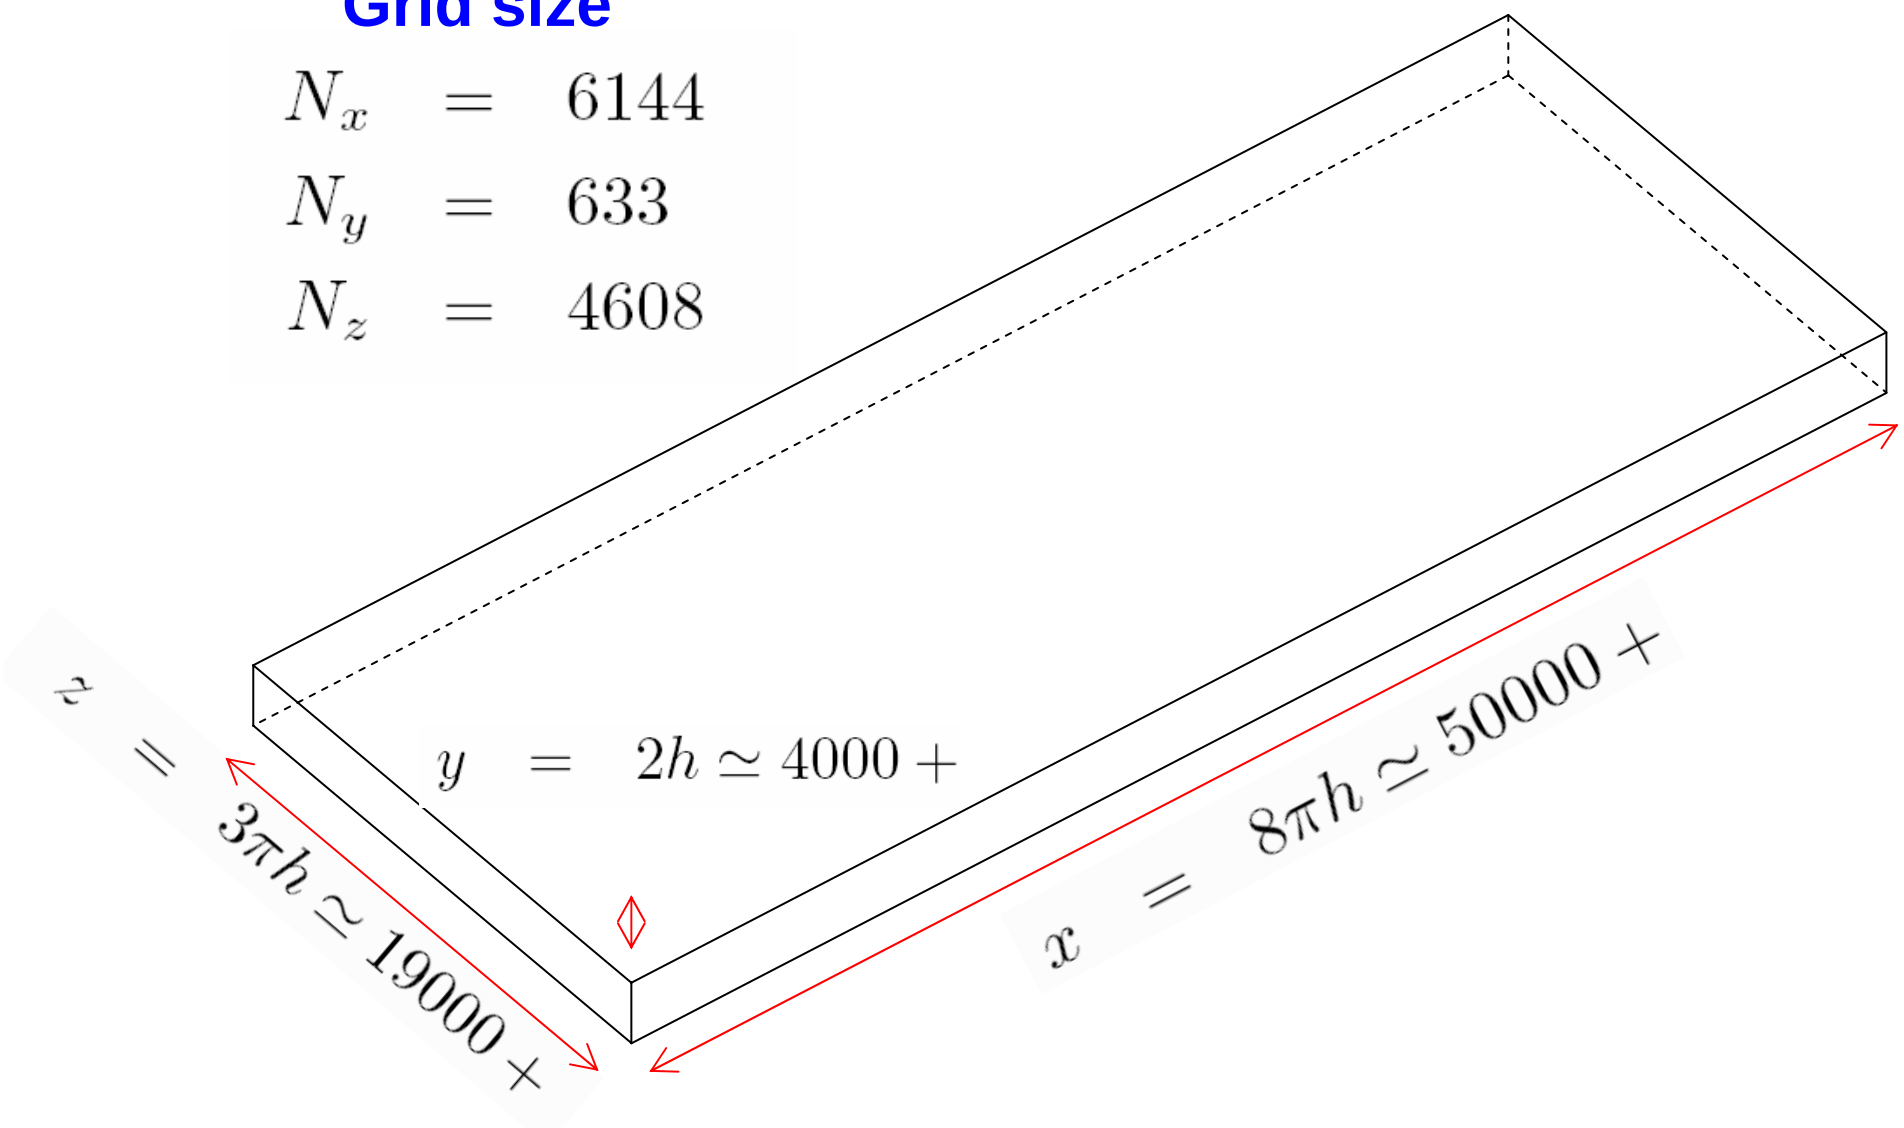

Log Layer

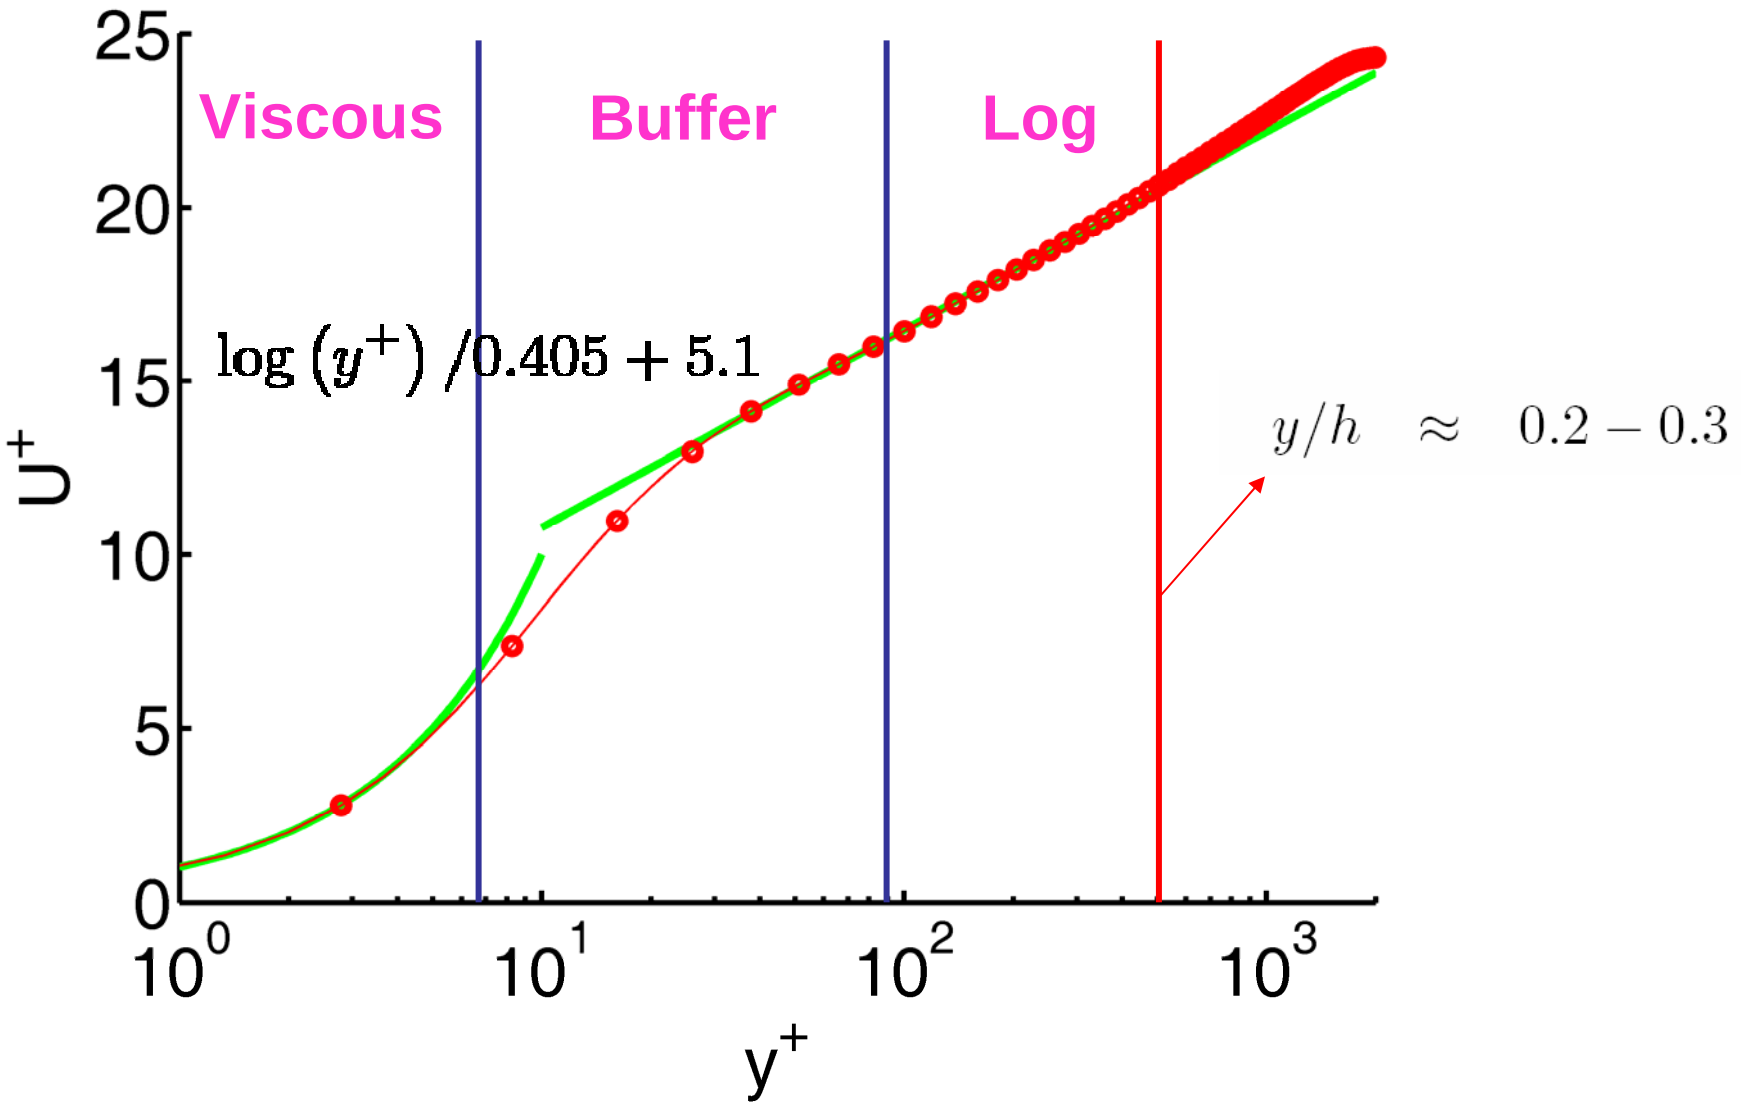
Viscous
Buffer
Log
$\log(y^+)/0.405 + 5.1$
$y/h \approx 0.2 - 0.3$
U^+
y^+
25
20
15
10
5
0
10^0
10^1
10^2
10^3

Previous simulations versus channel 2000

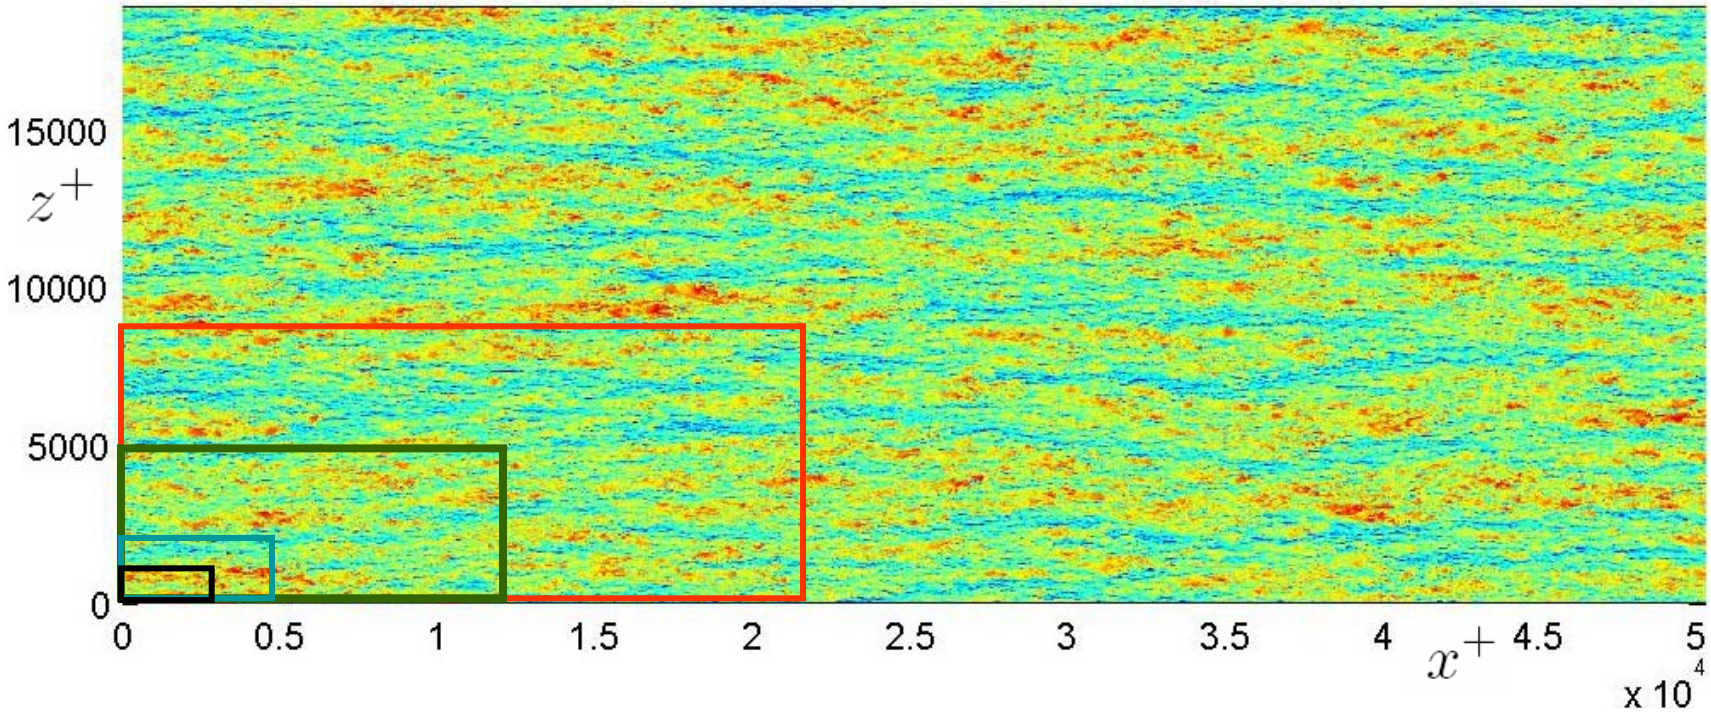

Kim, Moin and Moser, 1987, 180 (Cray XMP, NASA Ames)

Del Álamo and Jiménez, 2003 (Kadesh, CEPBA) 180

Del Álamo and Jiménez, 2003 (Kadesh, CEPBA) 550

Del Álamo, Moser, Jiménez and Zandonade, 2004 (Blue Horizon) 950

Navier-Stokes Equations

$$\frac{\partial u_i}{\partial t} = -\frac{\partial p}{\partial x_i} + H_i + \frac{1}{Re}\nabla^2 u_i,$$

$$\frac{\partial u_i}{\partial x_i} = 0,$$

$$\vec{H} = \vec{u} \times \vec{\omega} - \frac{1}{2}\nabla(\vec{v} \cdot \vec{v})$$

Velocity-Vorticity form

$$\frac{\partial}{\partial t}\phi = h_v + \frac{1}{Re}\nabla^2\phi \quad \Leftarrow \quad \phi = \nabla^2 v$$

$$\frac{\partial}{\partial t}\omega_y = h_g + \frac{1}{Re}\nabla^2\omega_y$$

70-80% of time

99% of communication

$$h_g = \frac{\partial H_1}{\partial z} - \frac{\partial H_3}{\partial x}$$

$$h_v = -\frac{\partial}{\partial y}\left(\frac{\partial H_1}{\partial x} + \frac{\partial H_3}{\partial z}\right) + \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial z^2}\right)H_2$$

Kim, Moin and Moser (1987)

Discretization

Time stepper: third order Runge-Kutta method (Spalart et al 1991)

$$\partial_t u = L(u) + N(u)$$

$$\begin{aligned} u' &= u_n + \Delta t \left[L\left(\alpha_1 u_n + \beta_1 u'\right) + \gamma_1 N_n \right], \\ u'' &= u' + \Delta t \left[L\left(\alpha_2 u' + \beta_2 u''\right) + \gamma_2 N' + \xi_1 N_n \right], \\ u_{n+1} &= u'' + \Delta t \left[L\left(\alpha_3 u'' + \beta_3 u_{n+1}\right) + \gamma_3 N'' + \xi_2 N' \right] \end{aligned}$$

Wall-normal: Compact Finite Differences (Lele, 1991)

$$D_j + \sum_{m=1}^{M} b_m (D_{j+m} + D_{j-m}) = \frac{1}{h} \sum_{n=1}^{N} a_n (u_{j+n} - u_{j-n})$$

N=7

$$D_j^{(2)} + \sum_{m=1}^{M} b_m (D_{j+m}^{(2)} + D_{j-m}^{(2)}) = \frac{a_0 u_j + \sum_{n=1}^{N} a_n (u_{j+n} + u_{j-n})}{h^2}$$

N=5

Parallelization

Planes code

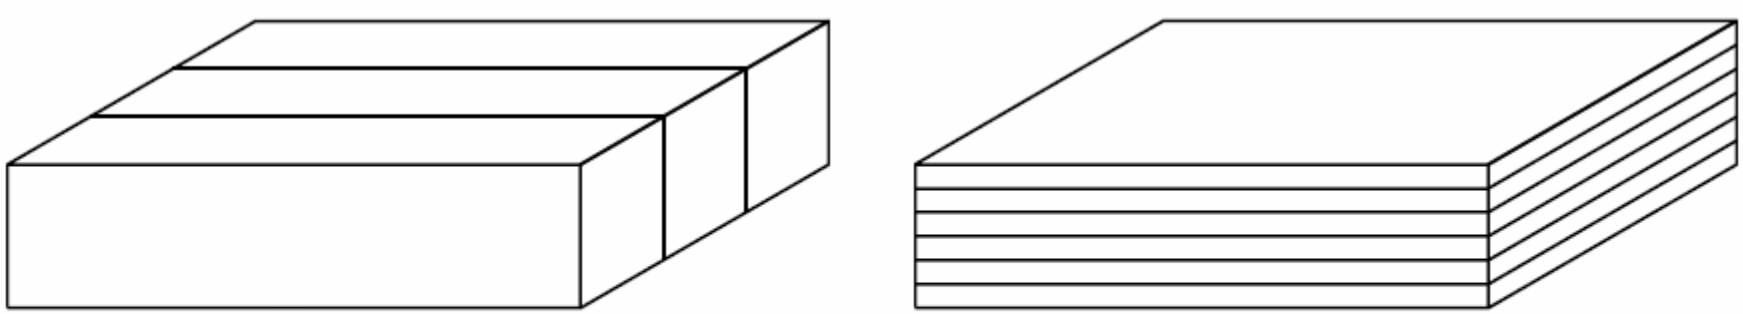

$$N_{proc} \leq N_y = 633$$

Line-plane code

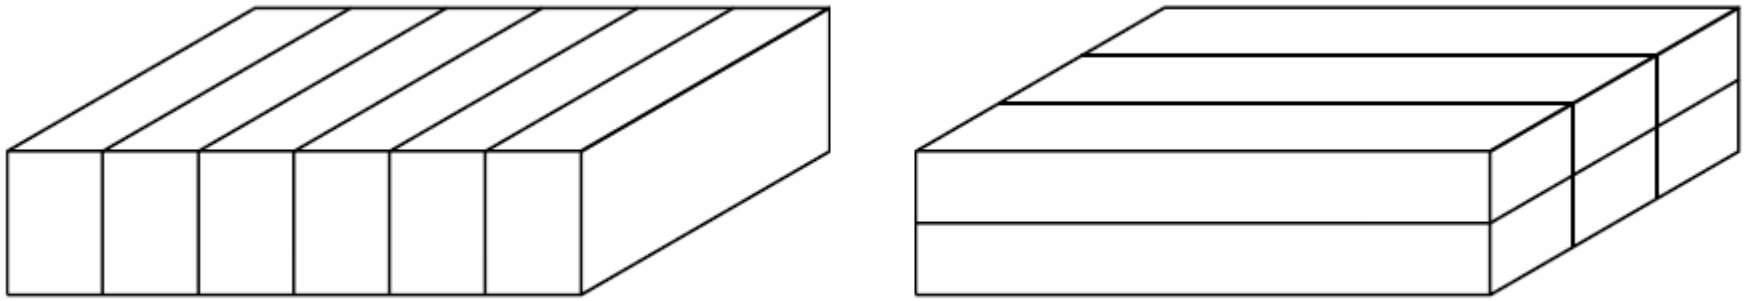

$$N_{proc} \leq N_x/3 = 2048$$

Parallelization

MareNostrum

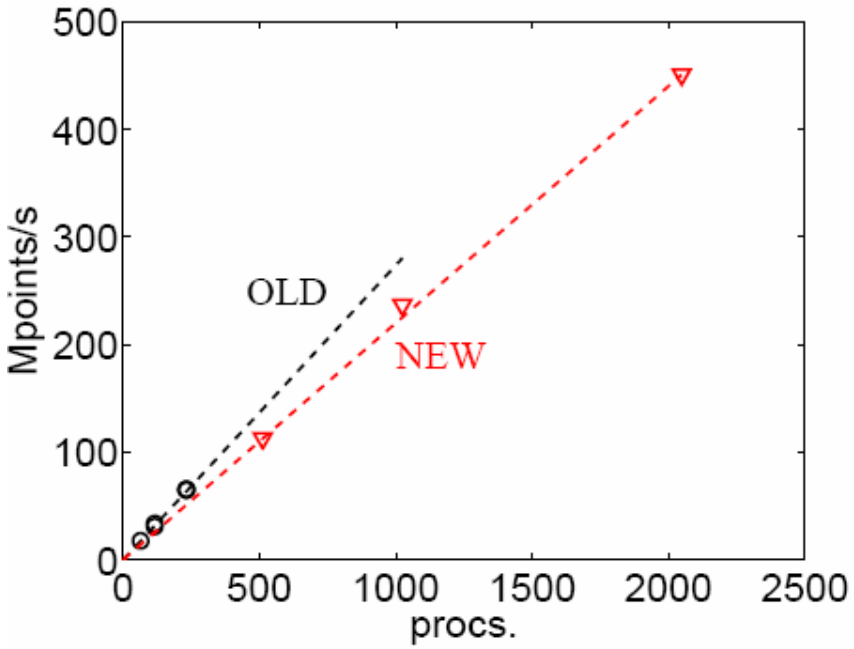

OLD CODE

9 Fourier buffers

8 buffer transposes

Comm. Time/Total = **28%**

$$N_{proc} \leq N_y$$

NEW CODE

12.5 Fourier buffers

13.5 buffer transposes

Comm. Time/Total = **58%**

$$N_{proc} \leq N_x/3$$

Dns are very expensive!

	Processor	Total
Memory	**0.2GB**	**400GB**
Time steps	***125.000***	***125.000***
Time for each Runge-Kutta step	**40s**	**40s**
Total CPU-hours	***2800h***	***6e6h (4e6h)***
Total human hours	**4 months**	**4months**
Data transferred between processors for each step	***500GB***	***1PB***
Total data transferred	**61PB**	**122EB**
Total amount of raw data obtained	***25TB***	***25TB***
Flops obtained	**50GF**	**3TF**
Total Flops	***18.3PF***	***3.6EF***

Simulations of channel kinematics

NEAR-WALL

$$Re_\tau \gg 60$$

LOG LAYER

$$0.2Re_\tau \gg 60$$

$$Re_\tau \gtrsim 1000$$

ASYMPTOTIC

Largest $\gg$ Streaks

$$10Re_\tau \gg 1000$$

$$Re_\tau \gtrsim 1000$$

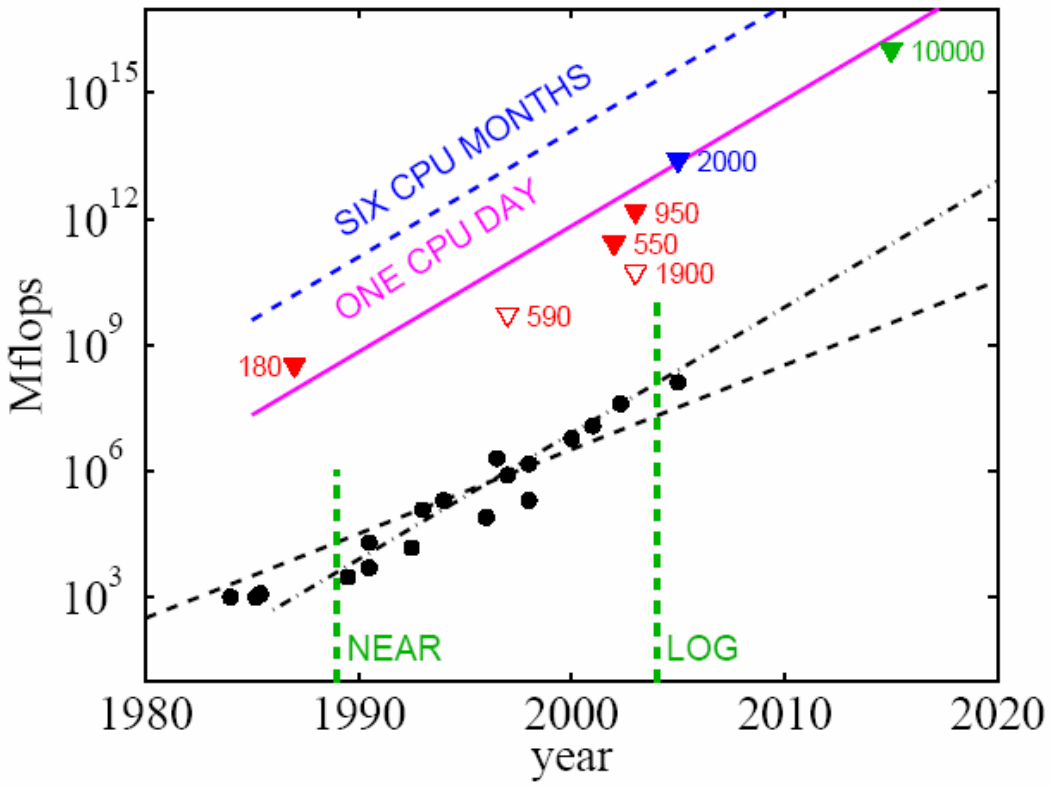

One point statistics

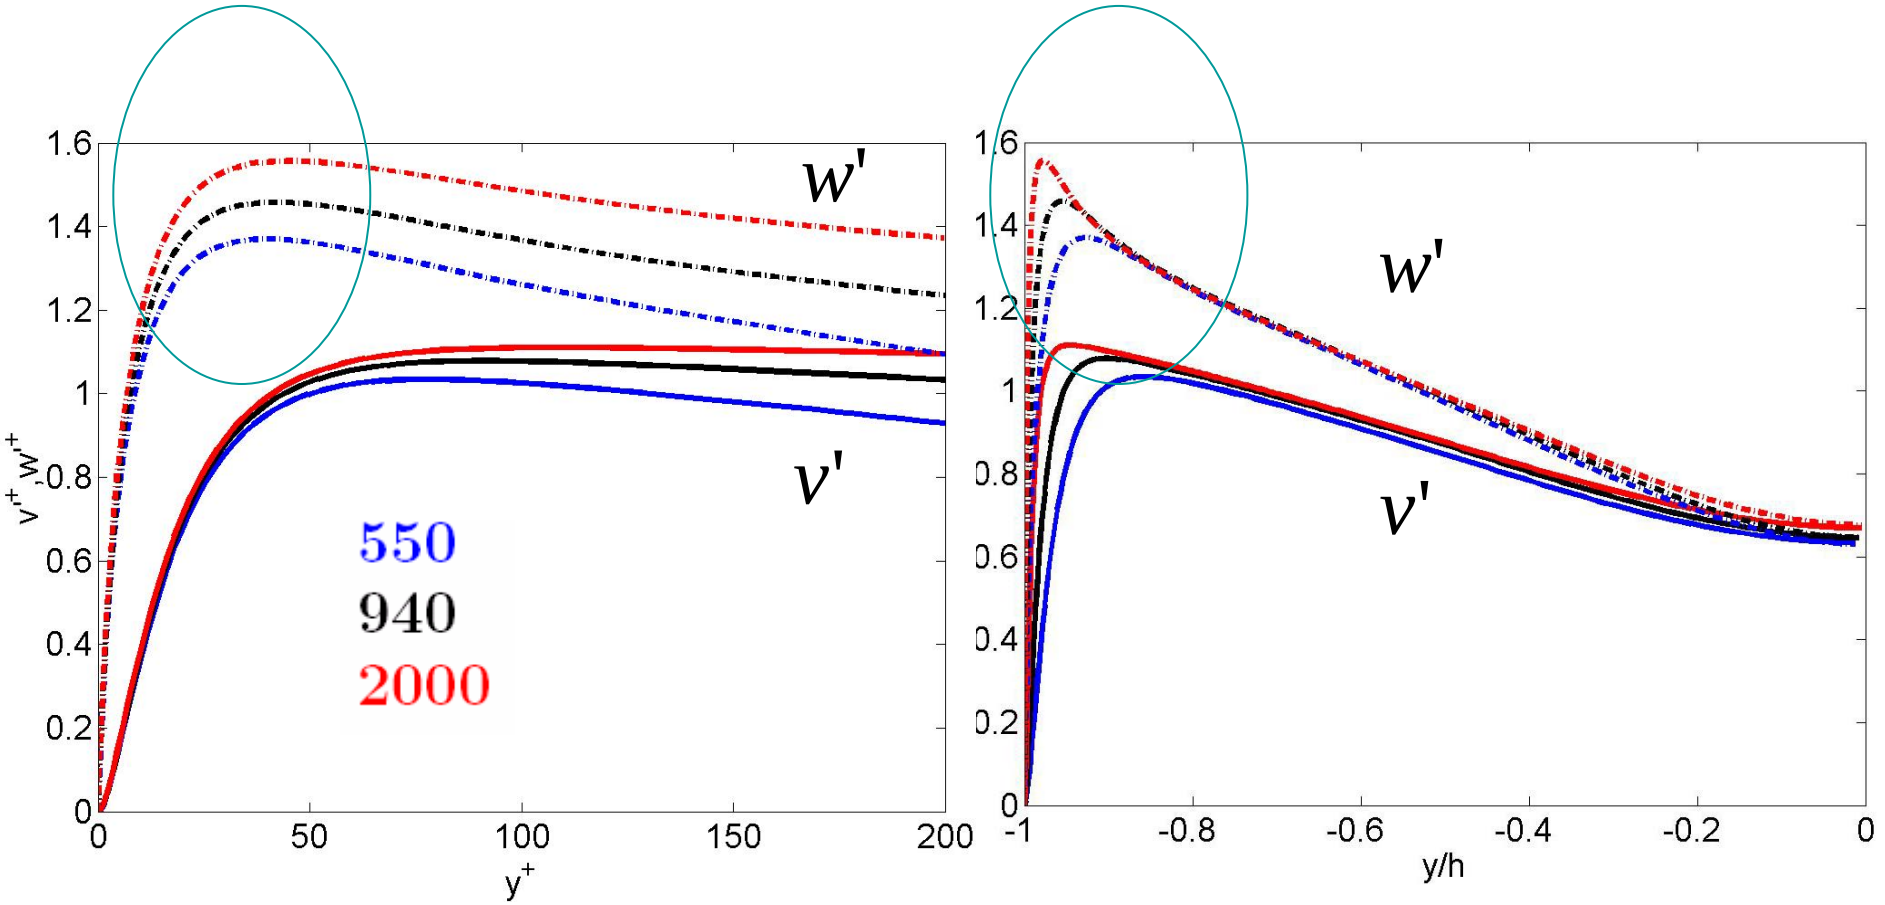

One point statistics

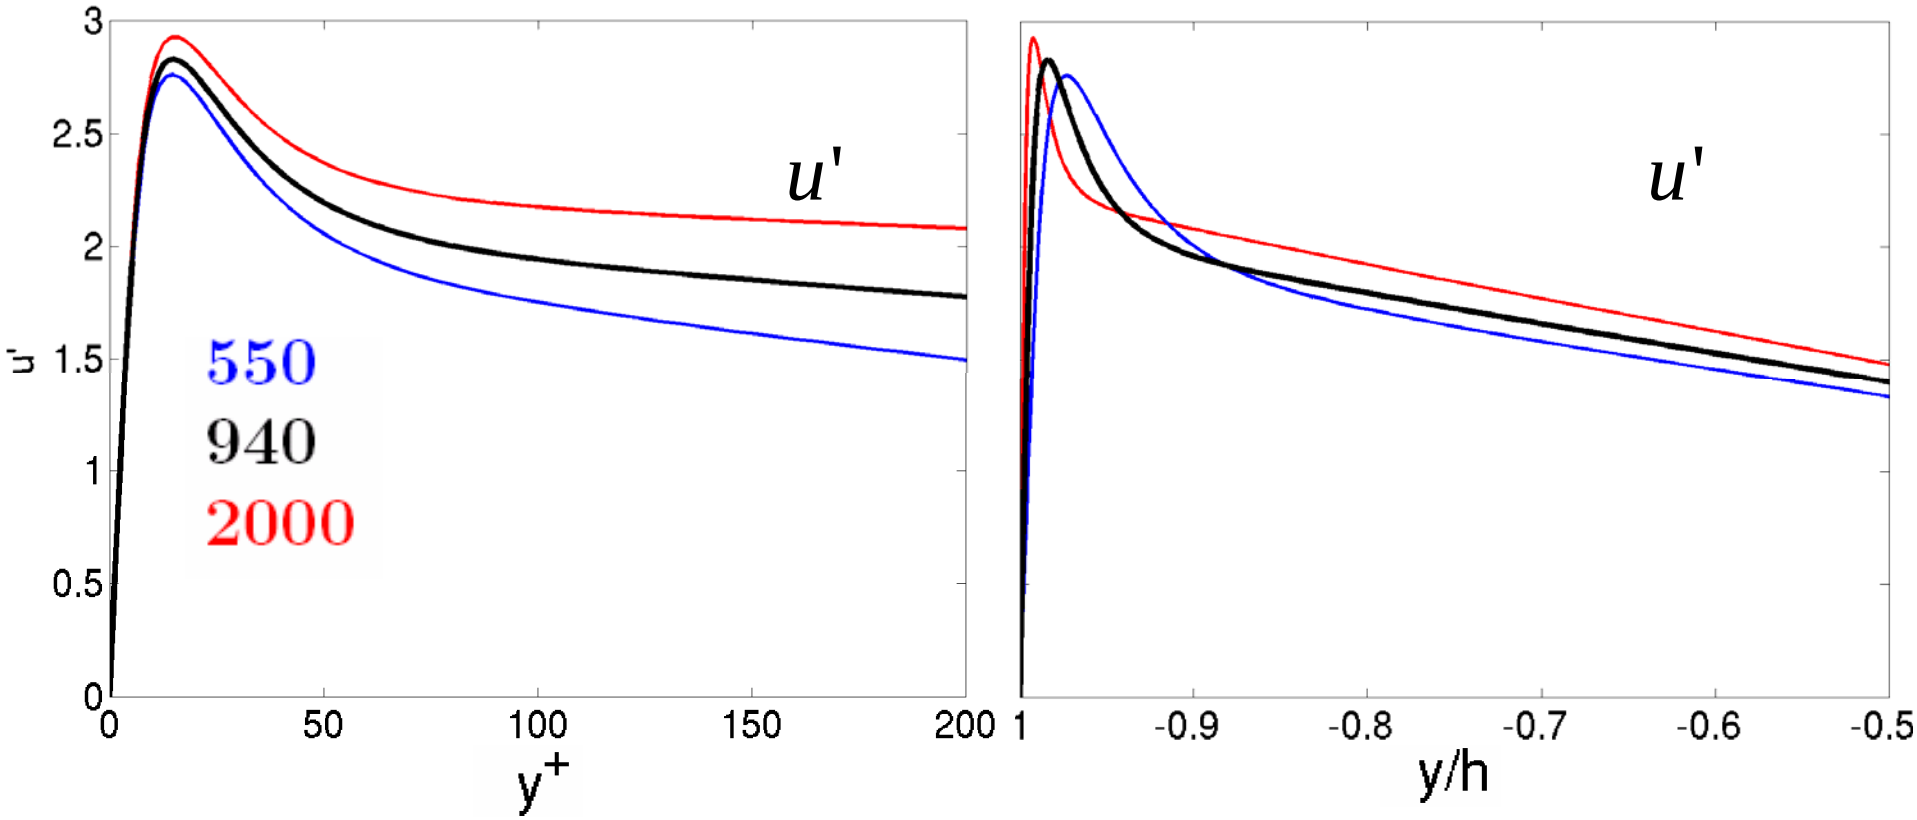

Scaling failures

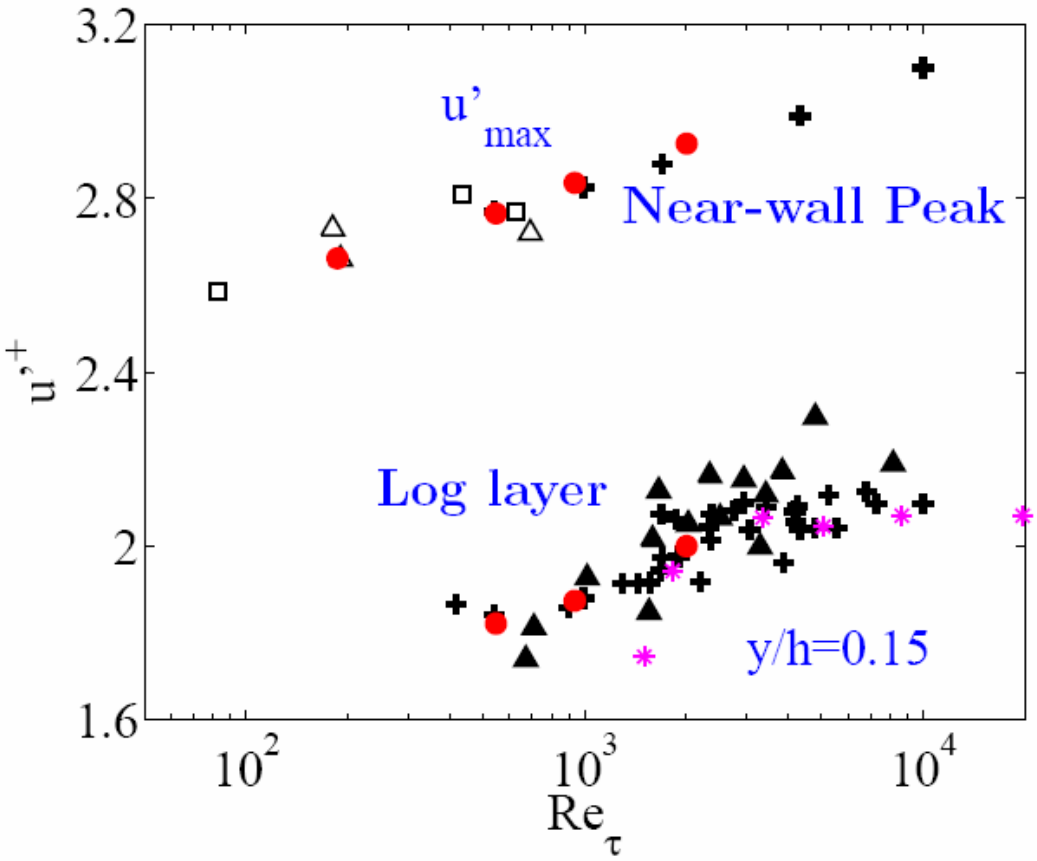

Simulations

Experiments

DeGraaf & Eaton (2000)

$k_x k_z E^+,\ y^+ = 15$

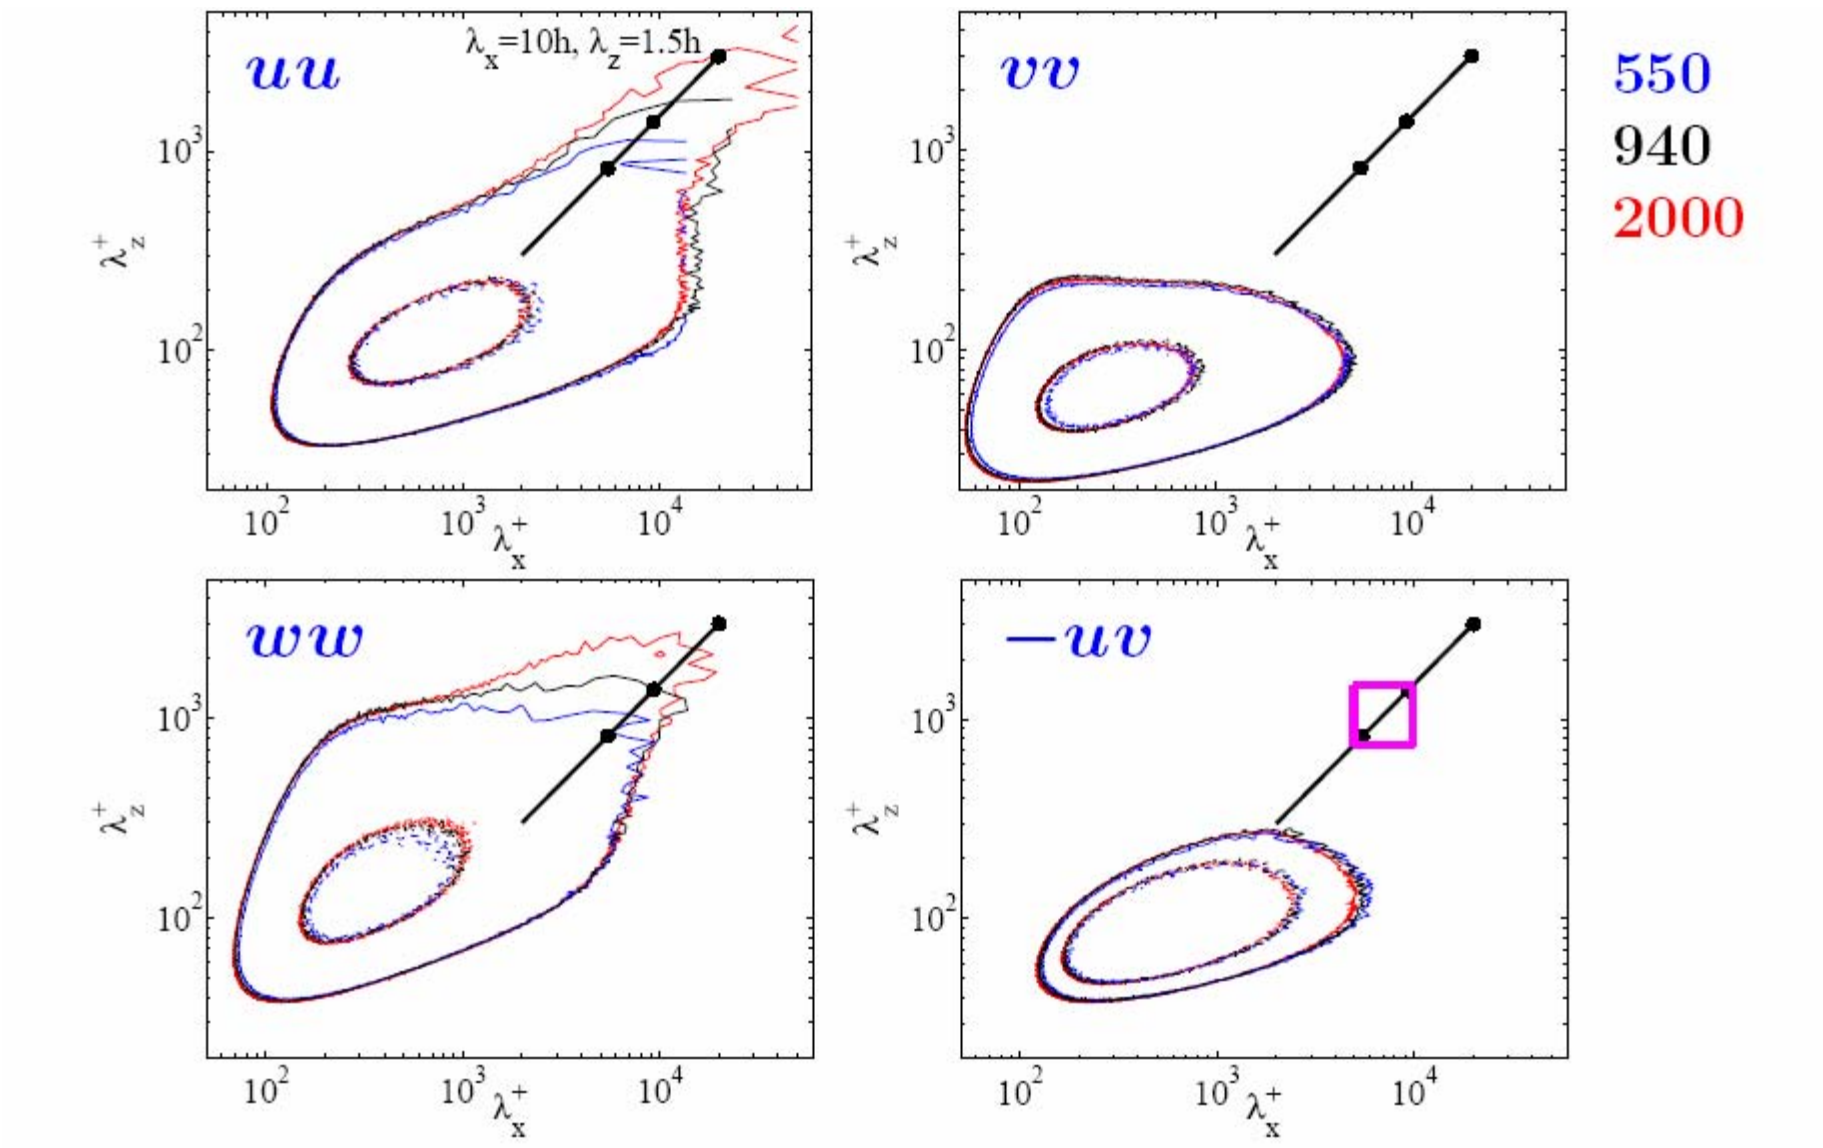

$Re_\tau = 2000$, E_{uu} and $-E_{uv}$

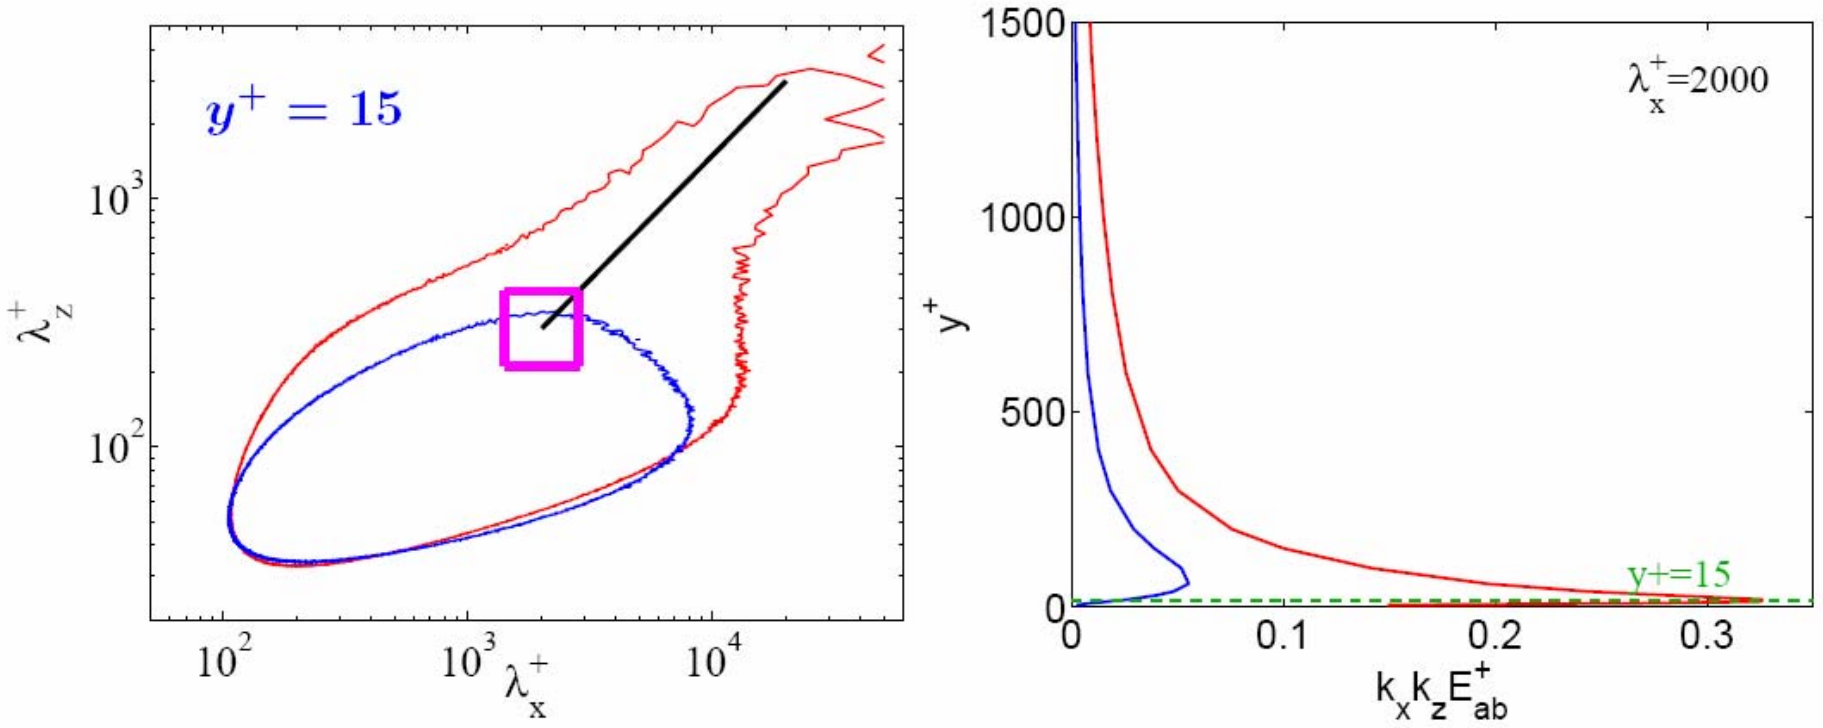

$Re_\tau = 2000$, E_{uu} and $-E_{uv}$

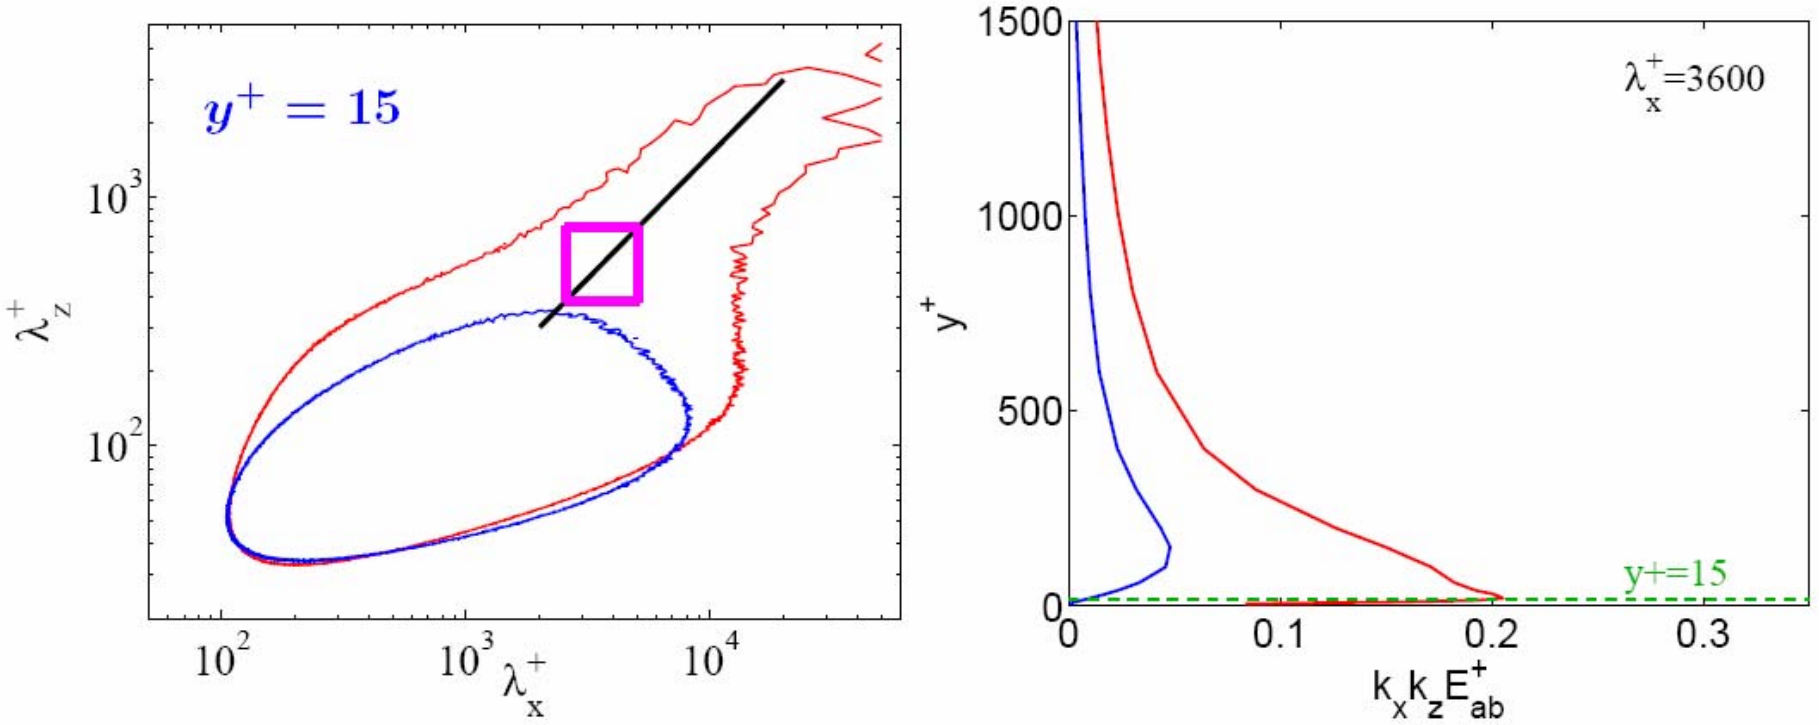

$Re_\tau = 2000$, E_{uu} and $-E_{uv}$

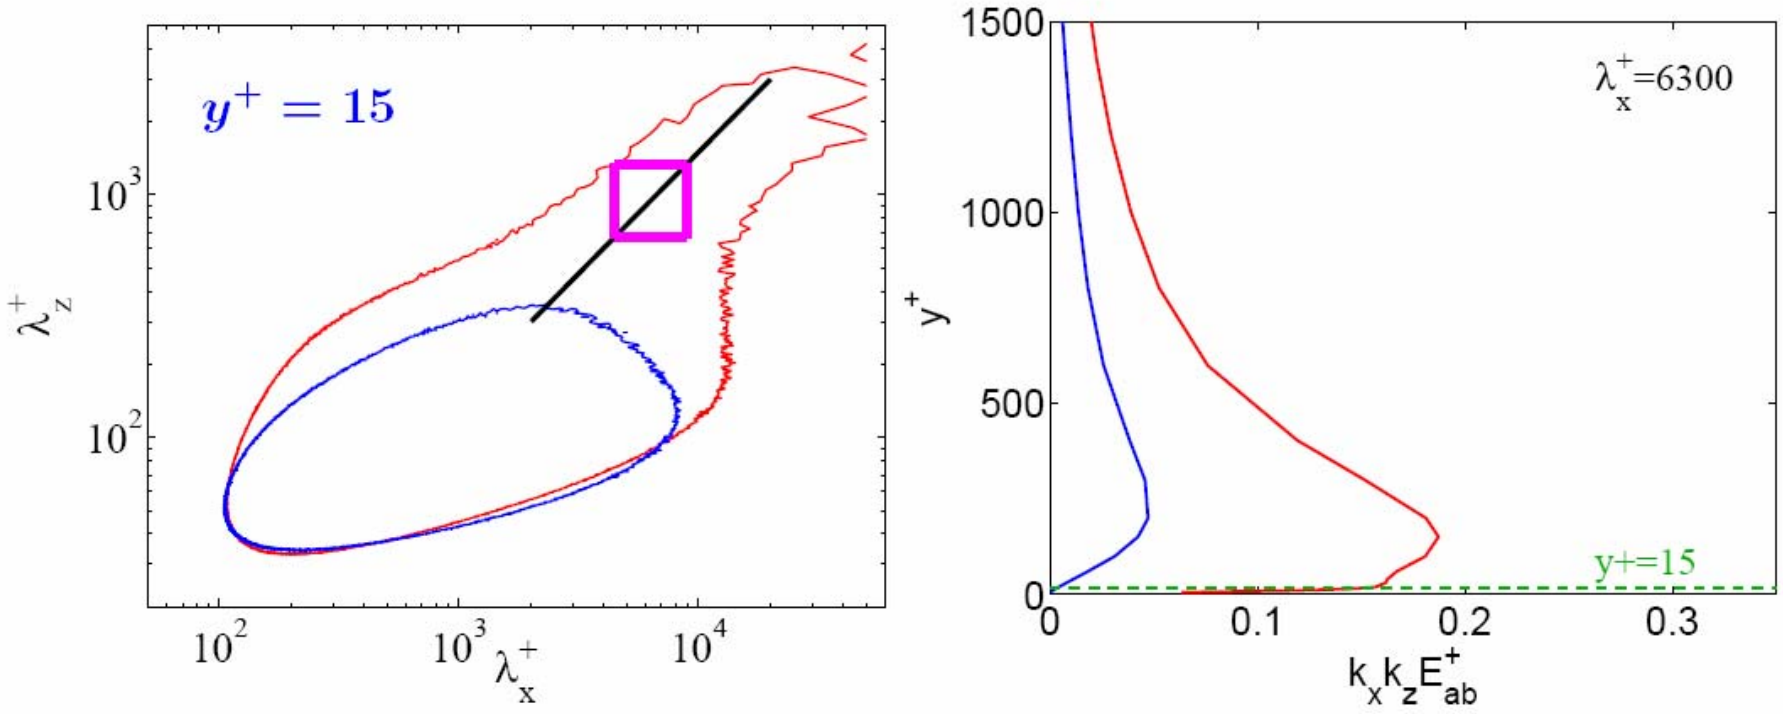

$Re_\tau = 2000$, E_{uu} and $-E_{uv}$

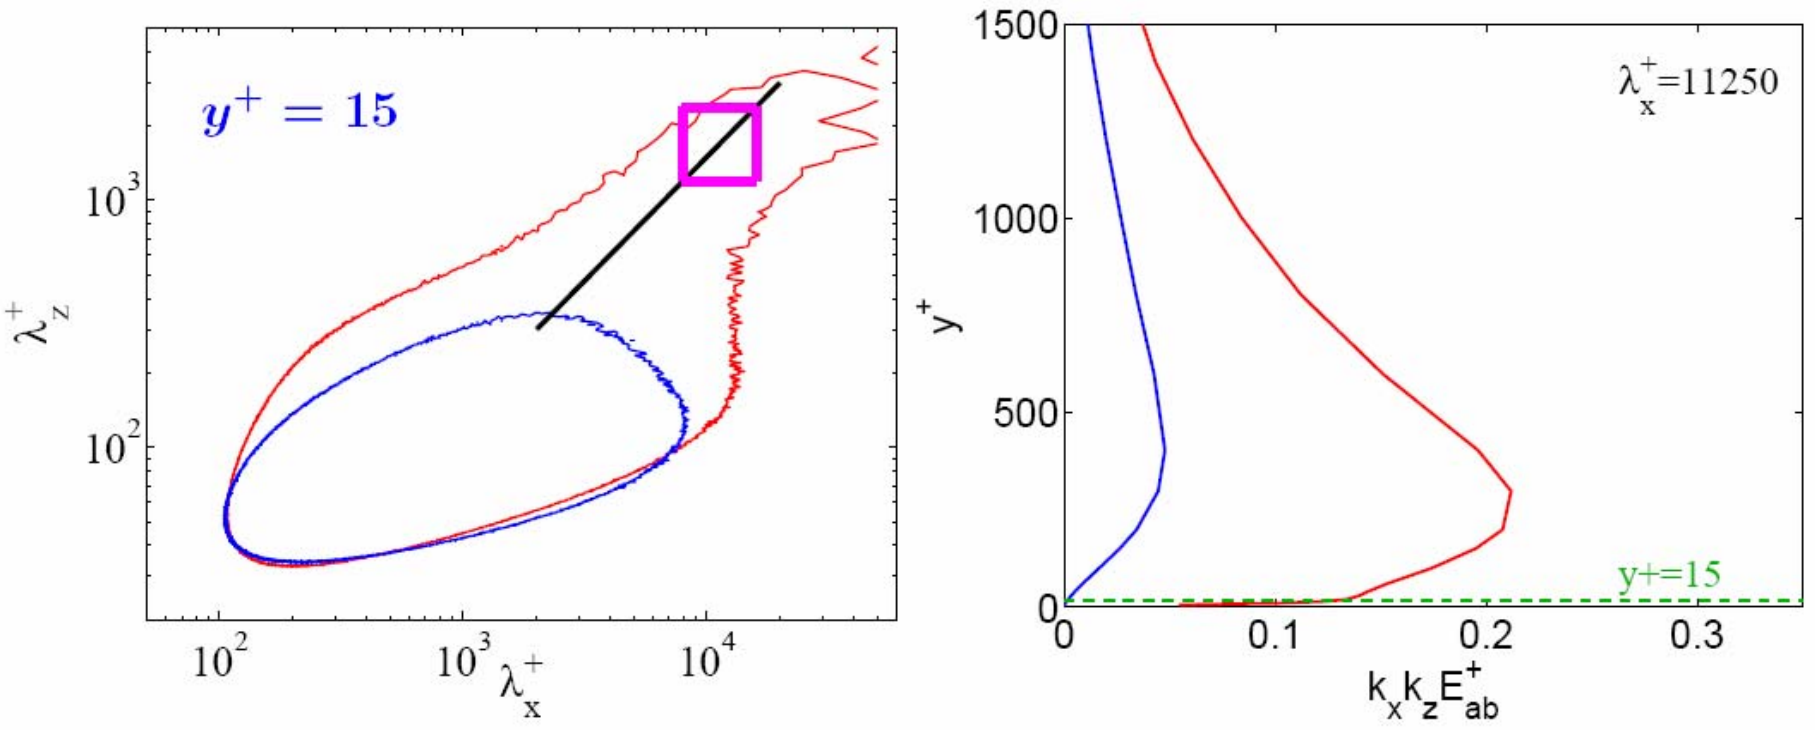

Cascades in wall turbulence

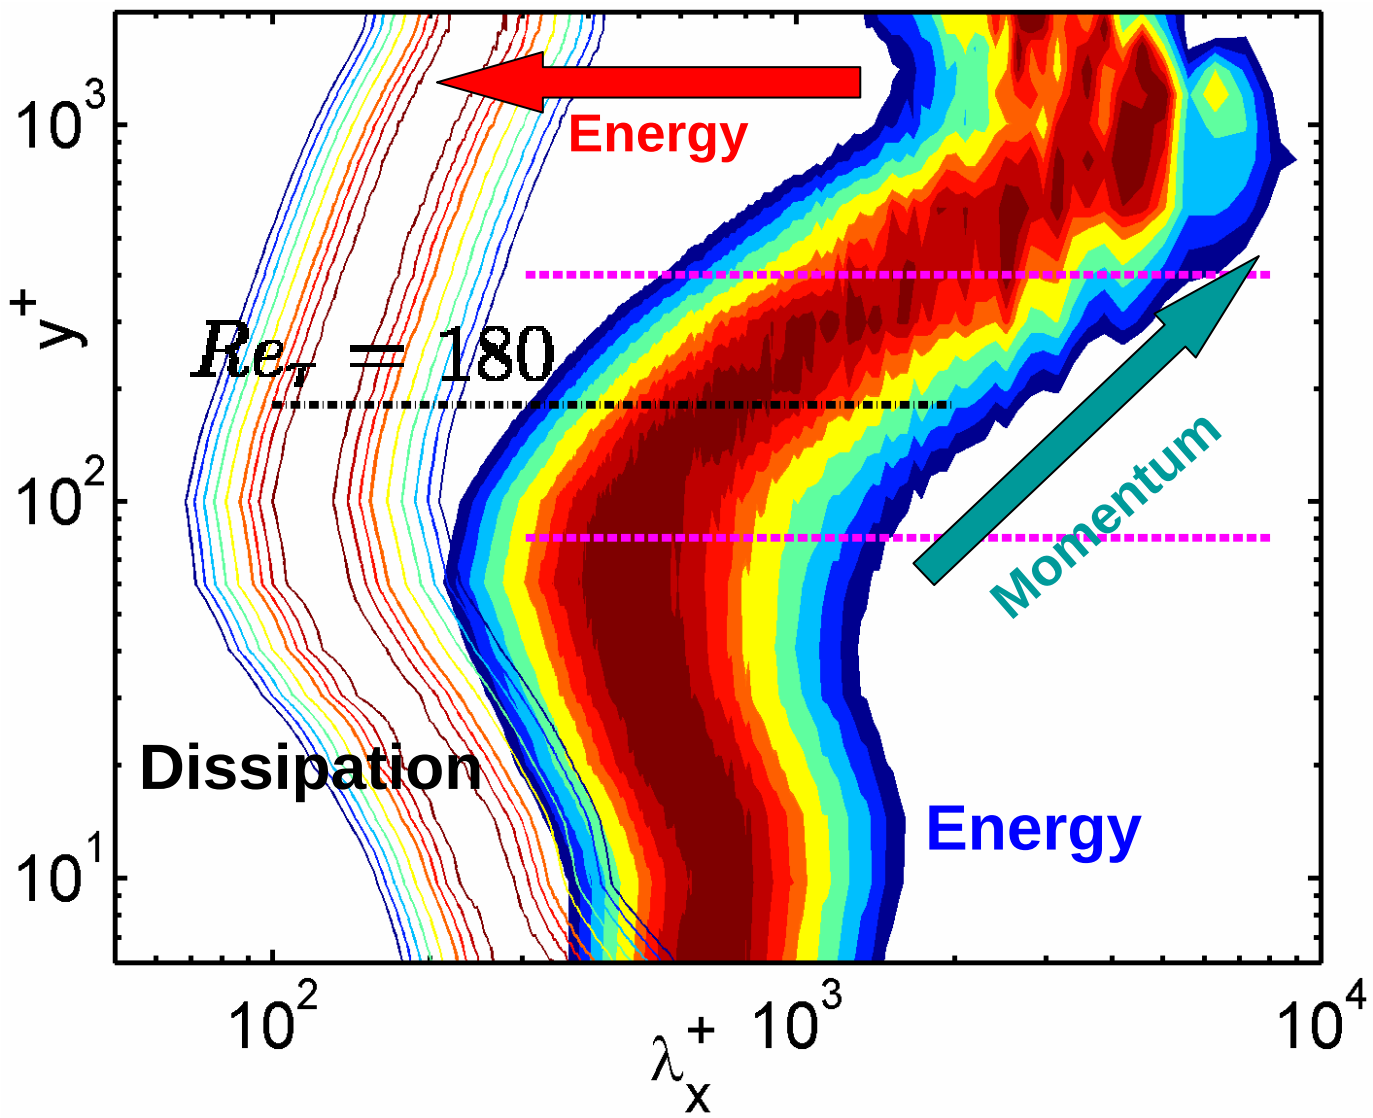

Outer

Log

Viscous and buffer

Instantaneous velocity and vorticity

Colored on u. Shaded on ω_z

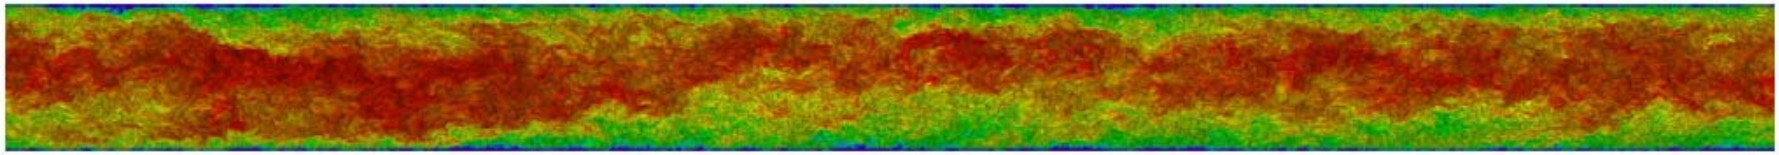

Streamwise: u & ω_z

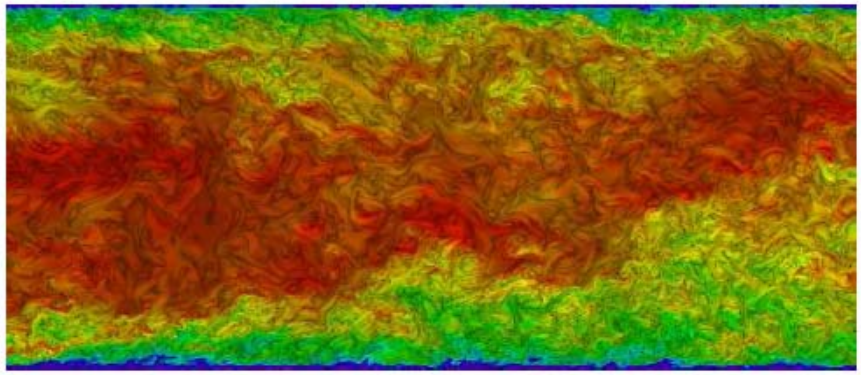

Spanwise: u & v

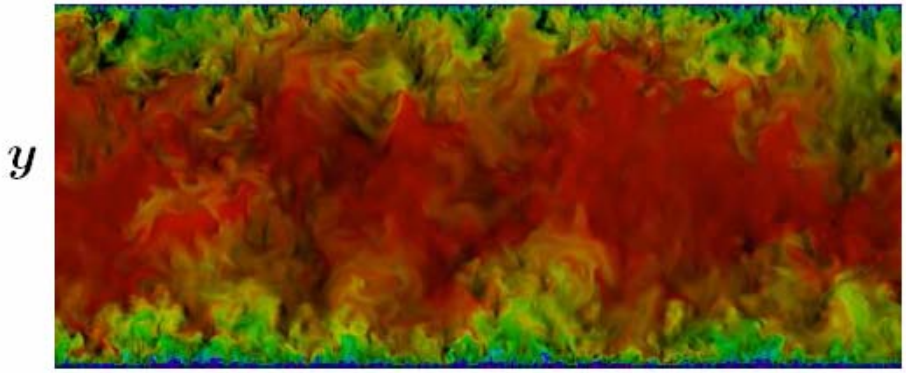

Hoyas & Jiménez, 2005

Conclusions

- First DNS on Turbulent Channel Flows with an appreciable log layer and inverse momentum cascade
- This DNS has been a very expensive simulation, but not more that one experiment of the same magnitude, and we can compute almost any imaginable quantity.
- We have obtained 25TB of data that we are analyzing: Pressure, Energy balances…
- We have confirmed some trends but we also have found new questions.

Future work?

When a channel 4000 ?

- Grid size

(12288,901,9216)

- Estimated time per step:

140s on 4096 processors, 280 on 2048

- Number of steps needed

250.000

- Total time

20 million CPU-Hours, between 800 and 1400 days

- Do you have a new MareNostrum?