DNS OF TURBULENT PIPE FLOW: HOW FAR CAN WE REACH TODAY?

M.Quadrio[1] & P. Luchini[2]

[1]Politecnico di Milano, Dip. Ing. Aerospaziale
[2]Universitá di Salerno, Dip. Ing. Meccanica

Bertinoro, Oct 13–15, 2007

Outline

The well-designed DNS

Pipe or channel?

The (far) future of DNS

What can we do today?

Conclusions

In a DNS no modelling of Reynolds stresses is required.
However design choices are required, concerning:

- The numerical method
- The largest / smallest resolved temporal scales
- The largest / smallest resolved spatial scales
- The boundary conditions

The numerical method

- Formulation of the NS equations (primitive variables vs alternative formulations)
- Discretization: spectral methods vs. finite differences (compact: better resolution in wavenumber space)

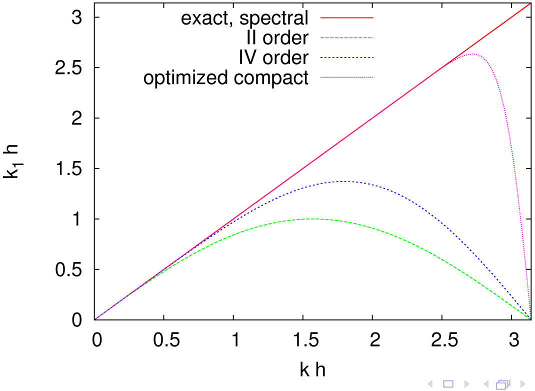

Smallest temporal scale

Choosing the timestep size

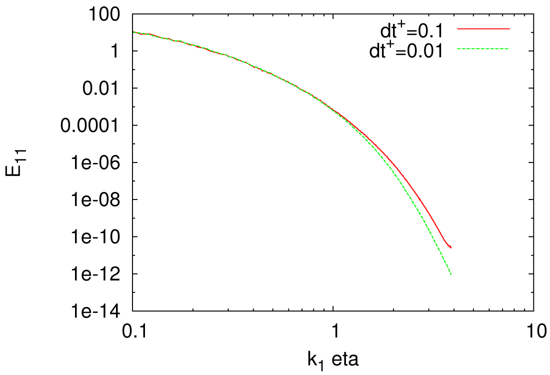

Common timestep values guarantee stability
Careful design is required to control accuracy

Largest temporal scale

Choosing the averaging time

- Issue in common with experiments
- Averaging time increases with order of statistical moments
- About 10 wash-out times are sufficient for low-order statistics

The smallest spatial scales

Choosing the spatial resolution

May be important for certain statistics: dissipation of a passive scalar, $Re_\tau = 160$ and $Sc = 1$

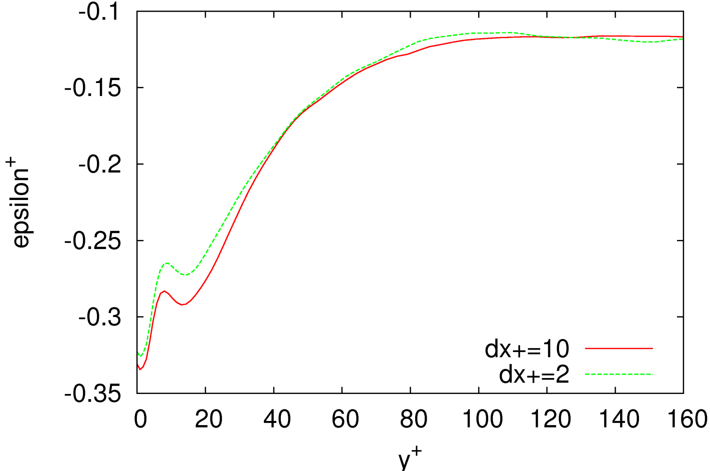

The smallest spatial scales

Choosing the spatial resolution

May be important for certain statistics: dissipation of a passive scalar, $Re_\tau = 160$ and $Sc = 1$

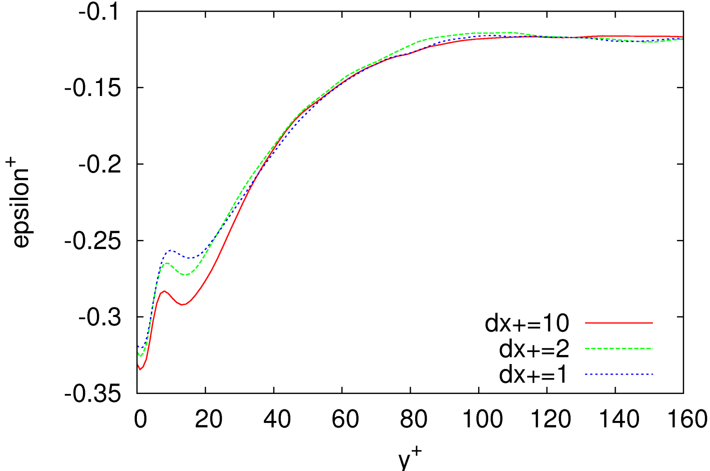

The largest spatial scales

Choosing the domain size: The issue

- Finite axial length of the computational domain
- Periodic boundary conditions create additional trouble
- Rather unexplored issue for pipe flow
- Recent results from several groups highlight the importance of representing large-scale structures (at least $20-30h$)

Space-time correlations

QL, PoF 2003

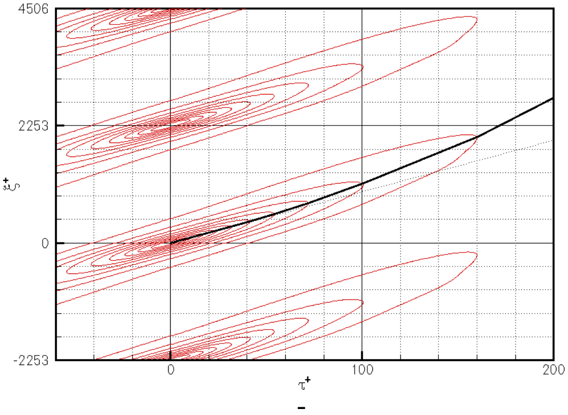

LITERATURE REVIEW

PIPE

1994 Eggels et al JFM:
$Re_\tau = 180$, $L/R = 10$

1996 Orlandi Fatica JFM:
$Re_\tau = 180$, $L/R = 15$

2000 Quadrio Sibilla JFM:
$Re_\tau = 200$, $L/R = 20$

2001 Satake et al:
$Re_\tau = 1050$, $L/R = 15$

- only a few others!!

CHANNEL

1987 Kim Moin Moser JFM:
$Re_\tau = 180$ and $L/h = 12$

1999 Moser Kim Mansour PoF:
$Re_\tau = 590$ and $L/h = 6$

2005 Iwamoto et al:
$Re_\tau = 2320$ and $L/h = 19$

2006 Hoyas Jiménez PoF:
$Re_\tau = 2003$ and $L/h = 25$

- many others

Pipe or channel flow?

Pipe

- Easier for experiments
- One periodic b.c.
- One wall
- Axis singularity
- Problem of azimuthal resolution
- Domain size?

Channel

- Easier for DNS
- Two periodic b.c.
- Two walls
- No axis
- No problem of azimuthal resolution
- Domain size?

The azimuthal resolution

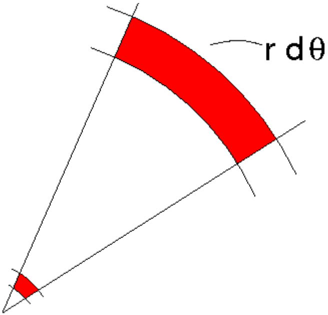

- Azimuthal resolution decreases with r
- Physical considerations set minimal resolution at $r = R$
- Waste of computational resources near the axis
- Unacceptable constraint on time step size due to stability problems

Perspectives for supercomputing

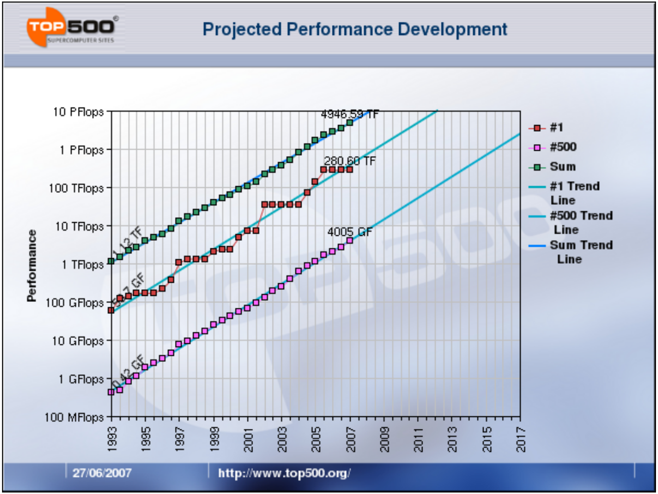

Perspectives for supercomputing

A not-too-bright outlook?

- Recent supercomputer growth due to uniprocessor improvements only
- Moore's law is slowing down
- More difficult to "ride on the coattail" of the Moore's law:
    - by 2020 800 loads and 0.09 MFlops for single memory access
    - by 2020 global latency equivalent to 1 MFlops
    - MTBF is decreasing (heat, size)

Perspectives for DNS of turbulence

- Wait for next-generation supercomputers

OR

- Explore single-precision computing
- Exploit GPUs
- Exploit new-generation coprocessors (e.g. Clearspeed)
- Envisage new simulation strategies

WHAT CAN BE DONE TODAY?

Q. Would it be useful to replicate the largest pipe flow DNS?

A. Yes, for example:

- to determine the scale-dependence of convection velocity (Taylor's hypothesis)
- to build a complete 5d correlation tensor
-
- *long list of useful aims*
-

Q. Would it be possible for us to replicate the largest pipe flow DNS?

A. Yes, with reasonable effort!

Designing the simulation

Design choices after Satake et al.

- $Re_\tau = 1050$ based on pipe radius
- $L/R = 15$
- Spatial resolution: $1024 \times 512 \times 768$ (nearly one billion d.o.f.)
- Required averaging time (extrapolated): 15 wash-out times

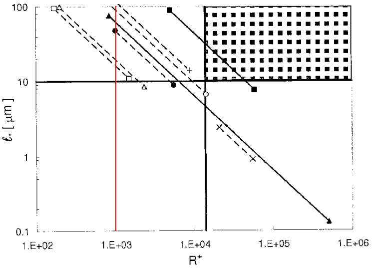

Introducing our DNS code

- Purposedly designed for high-Re turbulent pipe flow DNS
- Efficient parallel computing on commodity hardware
- Formulation in primitive variables
- Mixed spatial discretization (Fourier in homogenepous directions, compact explicit finite differences in radial direction)
- Smooth removal of azimuthal modes as the axis is approached
- Minimal bandwith requirements
- Minimal memory footprint

Introducing our Personal Supercomputer

A dedicated computing system

- tailored to the parallel algorithm
- 268 dual-core Opteron processors
- 2 1Gb interconnects
- 280 GB main memory
- 40 TB storage space
- 2.6 TFlops peak power

Comparing with classical supercomputers

ALTERNATIVE METRICS FOR COMPARISON?

	Earth Simulator	Our System
Energetic efficiency	3 Flops / Watt	70 Flops / Watt
Economic efficiency	1 MFlops / €	10 MFlops / €

Simulation data

- Wall-clock time: about 1 month with 16 machines
- RAM requirement: 7 GB (500MB when distributed)
- One full velocity field: about 3 GB
- Database size: about 300 GB
- Room for further optimizations

CONCLUSION

- Non-trivial to set up a well-designed simulation
- Resonably high Re_τ can be reached today with a commodity system
- Much higher *Re* are possible if other aspects (domain size, spatial resolution, etc) are sacrified