

Direct Numerical Simulation of Fluid Turbulence at Extreme Scale with psOpen

Jens Henrik Göbbert^{1,2)}, Michael Gauding³⁾, Cedrick Ansoerge⁴⁾,
Bernd Hentschel^{1,2)}, Torsten Kuhlen^{2,5)}, Heinz Pitsch⁶⁾

¹⁾ Jülich Aachen Research Alliance, JARA-HPC, Germany

²⁾ Virtual Reality Group, RWTH Aachen University, Germany

³⁾ Chair for Numerical Thermo-Fluid Dynamics, Technische University Freiberg, Germany

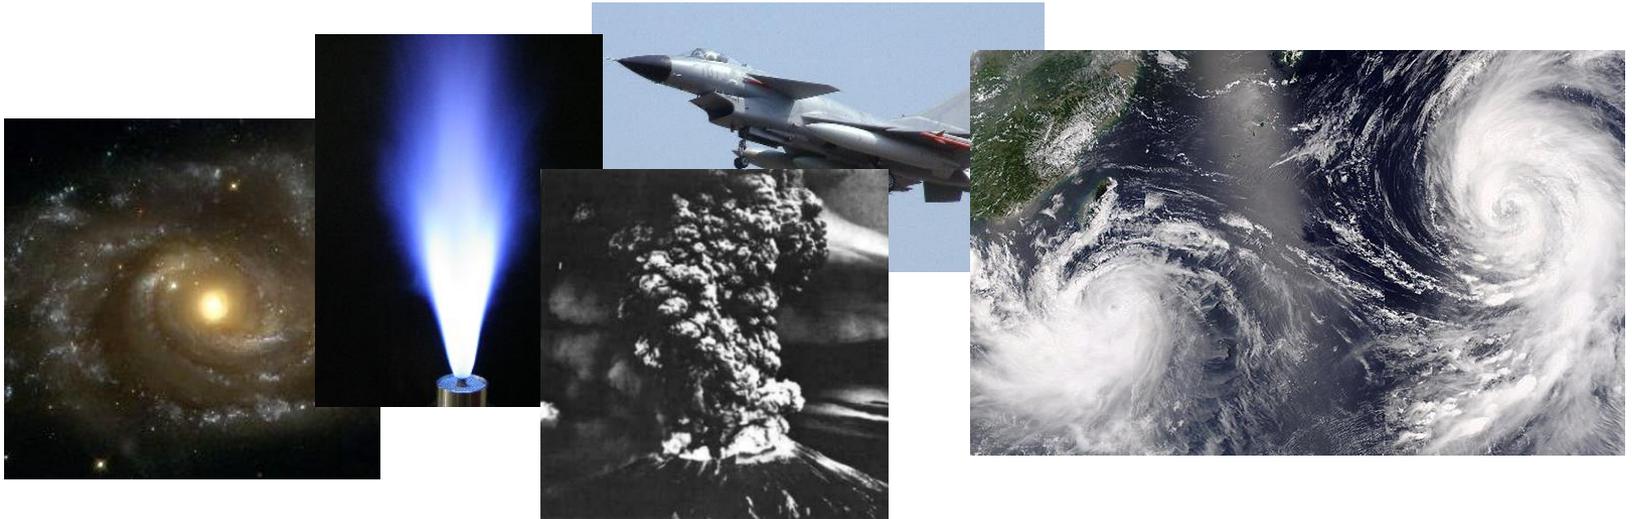
⁴⁾ Max Planck Institute for Meteorology, Hamburg, Germany

⁵⁾ Jülich Supercomputing Centre – Forschungszentrum Jülich GmbH, Germany

⁶⁾ Institut for Combustion Technology (ITV), RWTH Aachen University,, Germany

Turbulence

phenomenologically a fluid regime characterized by
chaotic and stochastic property changes



Motion of turbulent flows is
one of the unsolved problems of classical physics

- introduction to pseudo-spectral DNS
- 3D Fast Fourier Transformations
 - implicit filtering
 - overlapping communication and computation
- scaling
- conclusion

- exact solution of Navier-Stokes equations $\frac{\partial \mathbf{u}}{\partial t} + \boldsymbol{\omega} \times \mathbf{u} = -\nabla \left(p + \frac{1}{2} \mathbf{u}^2 \right) + \nu \nabla^2 \mathbf{u}$
 - partial differential equation
 - with non-linear term $\nabla \cdot \mathbf{u} = 0$,

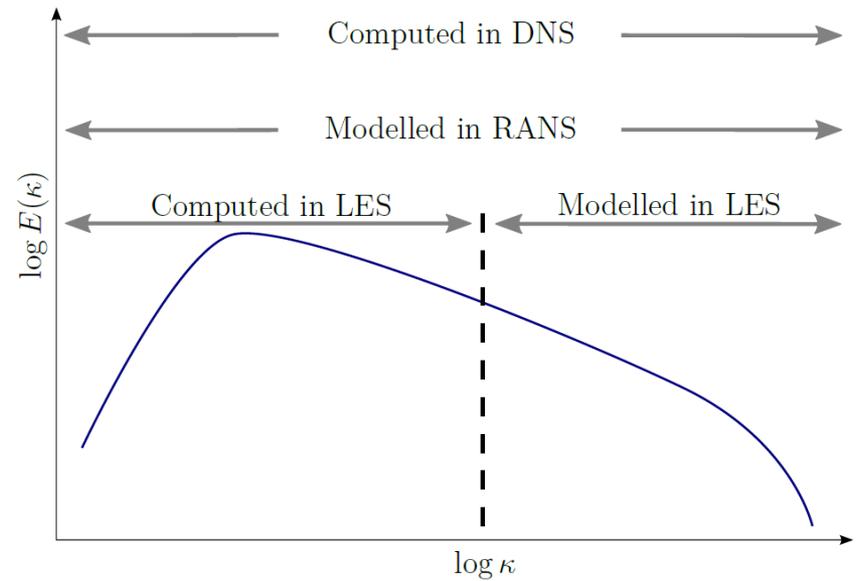
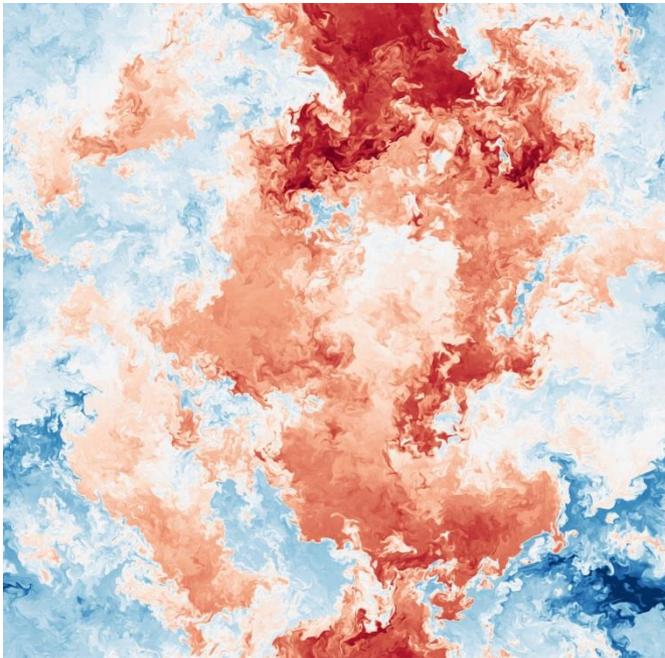
Computational challenges in DNS

- smallest structures of the flow have to be resolved by the computational grid
- ratio of largest to smallest structures increases with the Reynolds number
- the total number of grid points is $N^3 \propto Re^{9/4}$

→ DNS of high Reynolds number flows are computational very expensive

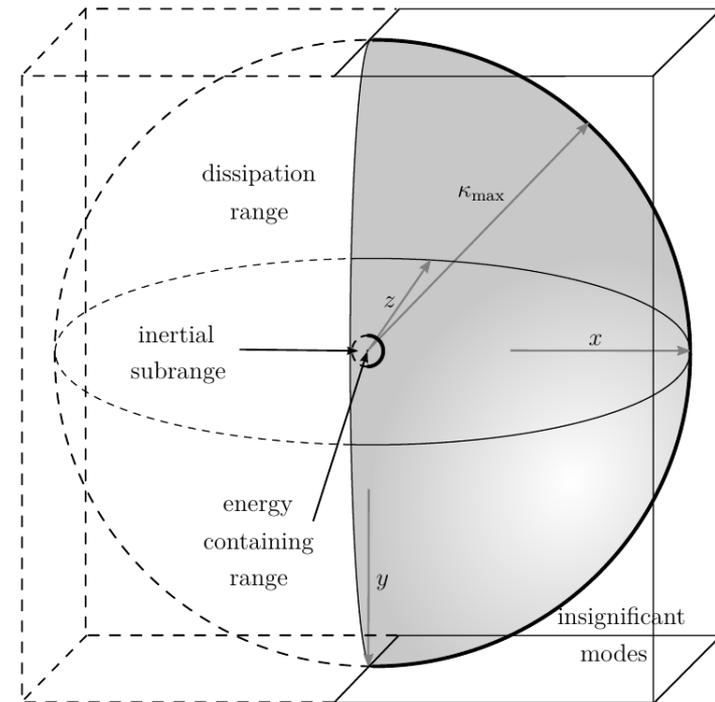
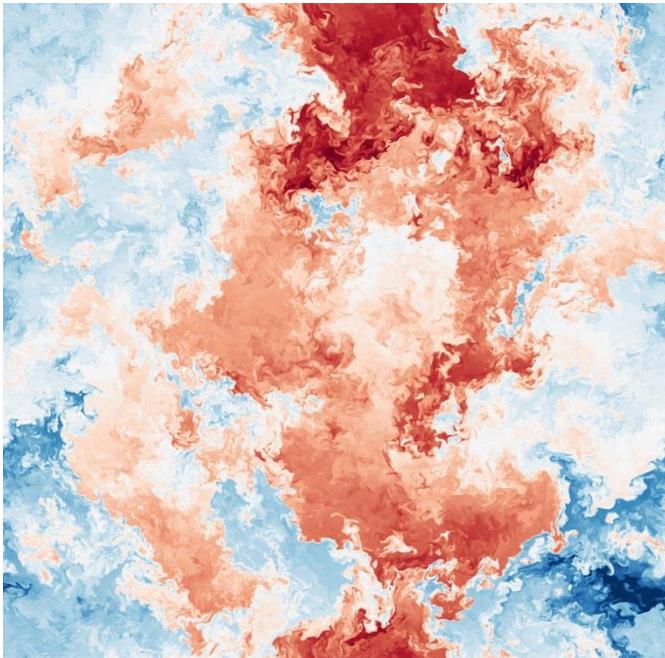
Turbulence Research

resolving smallest scales



Turbulence Research

resolving smallest scales



Navier-Stokes equation is efficiently solved by pseudo-spectral method

- spatial highly accurate due to utilization of Fourier transform
- derivatives are point operations in spectral space

$$\mathcal{F} \left[\frac{\partial}{\partial x_i} \phi(\mathbf{x}) \right] = i\kappa_i \hat{\phi}(\boldsymbol{\kappa}) \quad \mathcal{F} \left[\frac{\partial^2}{\partial x_i^2} \phi(\mathbf{x}) \right] = -\kappa_i^2 \hat{\phi}(\boldsymbol{\kappa})$$

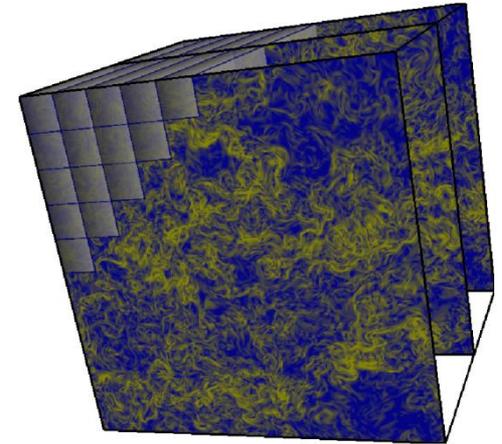
but

- multiplications within non-linear terms are computed in real space

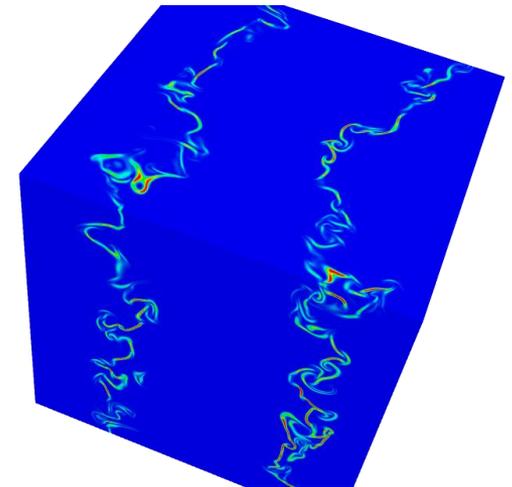
→ 3d Fast Fourier Transformations (3d-FFT) required

psOpen – DNS by a pseudo-spectral approach

- isotropic homogeneous decaying/forced turbulence
- turbulent channel flow
- turbulent premixed flames
- scalar mixing
- Fortran 95
- MPI+OpenMP , single/double precision
- I/O using HDF5
- developed with the help of JSC 2007-2015

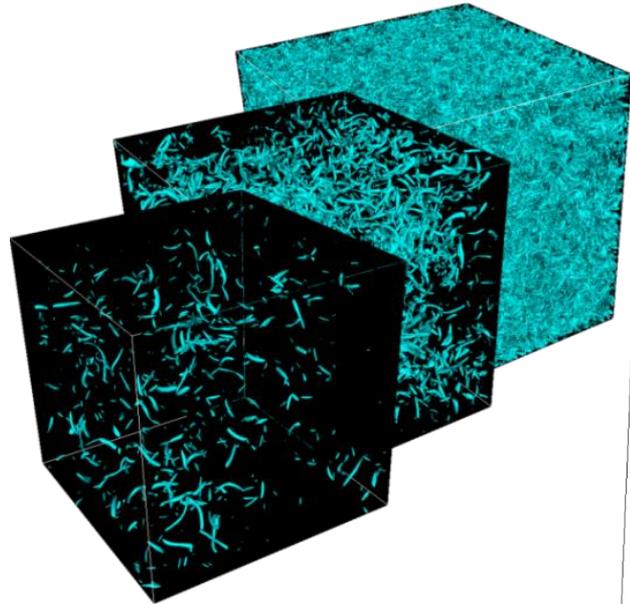


domain decomposition

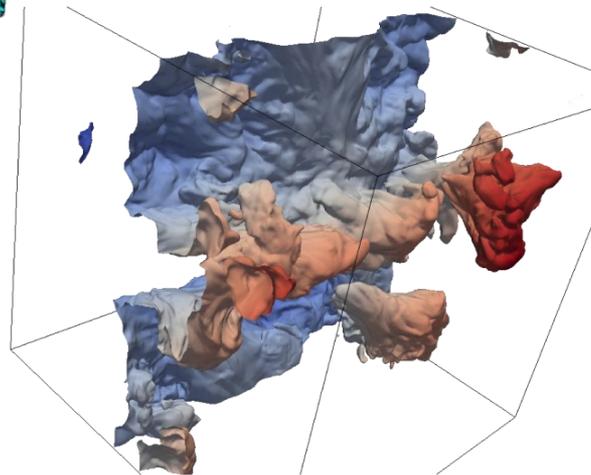


scalar mixing

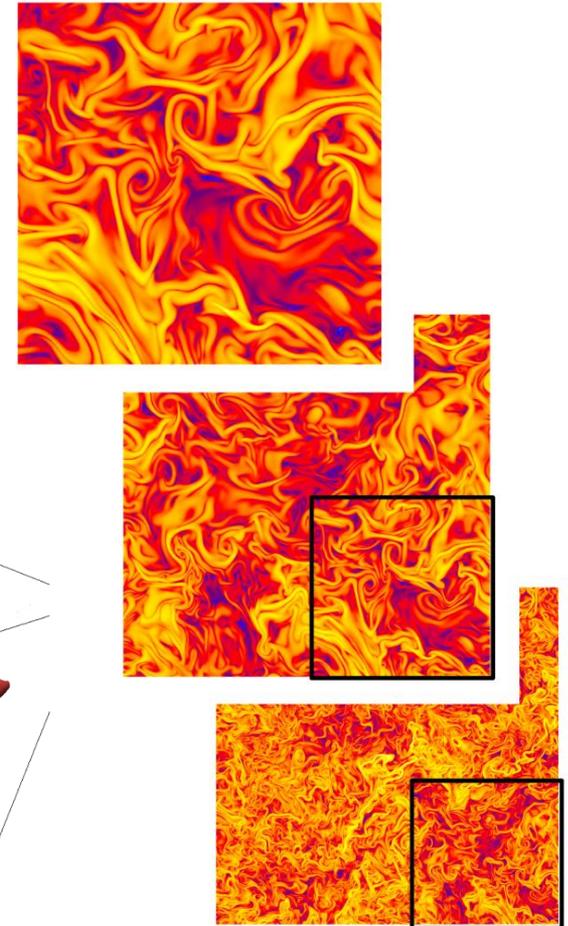
psOpen – DNS by a pseudo-spectral approach



vortex structure in
decaying turbulence



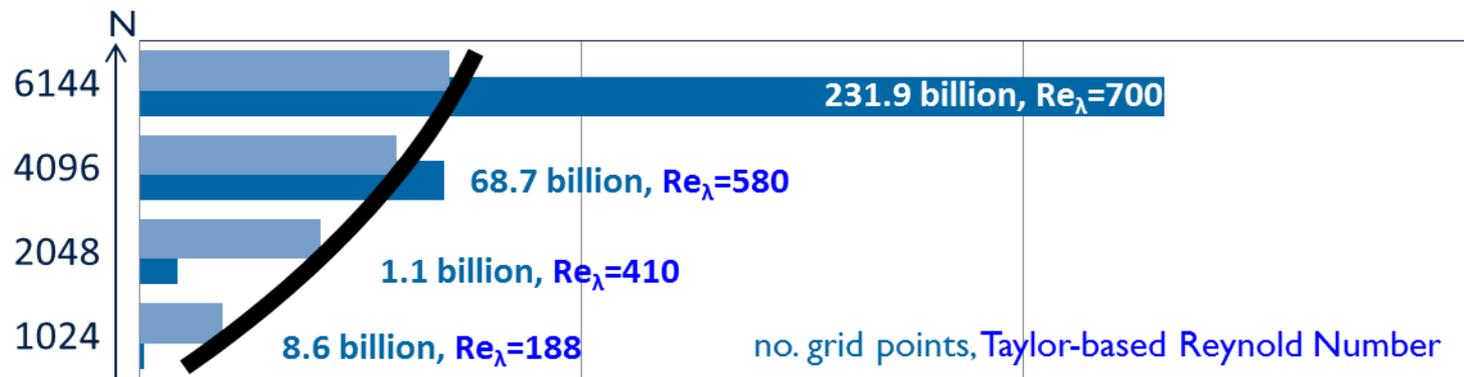
highly resolved structure of turbulent premixed flame



scalar dissipation

current largest simulation – isotropic homogeneous forced turbulence

- $N^3 = 6144^3 = 231.9$ billion grid cells in physical space
- production runs on 8 192 - 16 384 nodes
- 262 144 – 524 288 MPI threads
- 32 MPI processes x 2 OpenMP threads per node
- input/output file up to 4084 GB



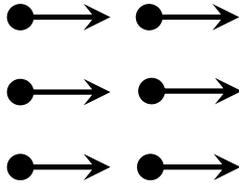
comparing setups of different computations on JUQUEEN

	A0	A1	A2	A3	A4	A5	A6
N^3	512 ³	1024 ³	1024 ³	2048 ³	2048 ³	4096 ³	4096 ³
Re_λ	88	119	184	215	331	529	754
file size (GB)	8	64	64	512	512	4096	4096
M	180	60	60	10	10	10	10
data size (TB)	1.44	3.81	3.81	5	5	22	22

	B0	B1	B2	B3	B4	B5	B6
N^3	720 ³	1440 ³	1440 ³	2816 ³	2816 ³	5632 ³	6144 ³
Re_λ	84	115	173	207	297	529	770
file size (GB)	22	177	177	1331	1331	5324	6912
file size compressed (GB)	6.6	52.6	52.6	393.2	393.2	1572.8	2041.9
M	40	20	20	10	10	5	5
data size (TB)	0.88	3.54	3.54	13.3	13.3	22.4	34.5
data size compressed (TB)	0.26	1.05	1.05	3.9	3.9	6.6	10.3

Data Analysis

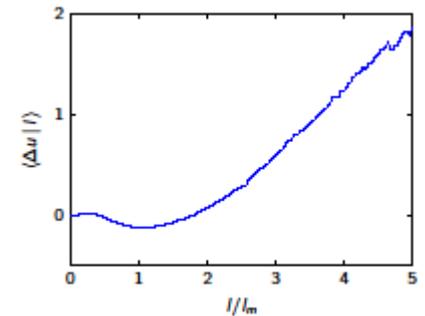
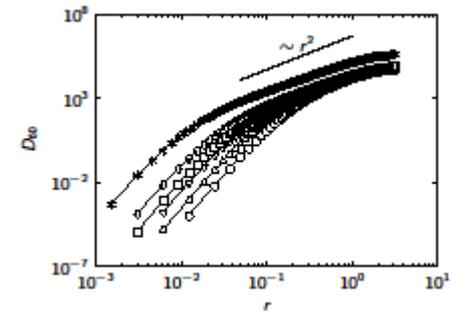
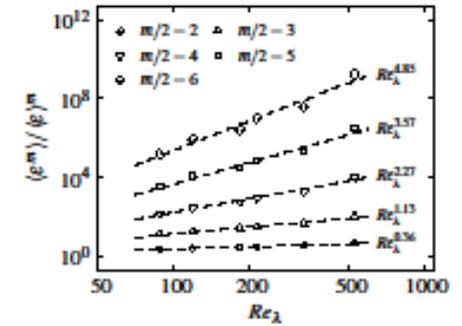
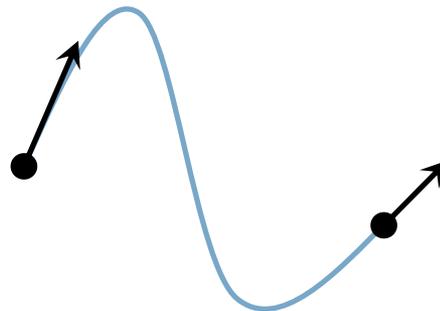
1-point statistics



2-point statistics

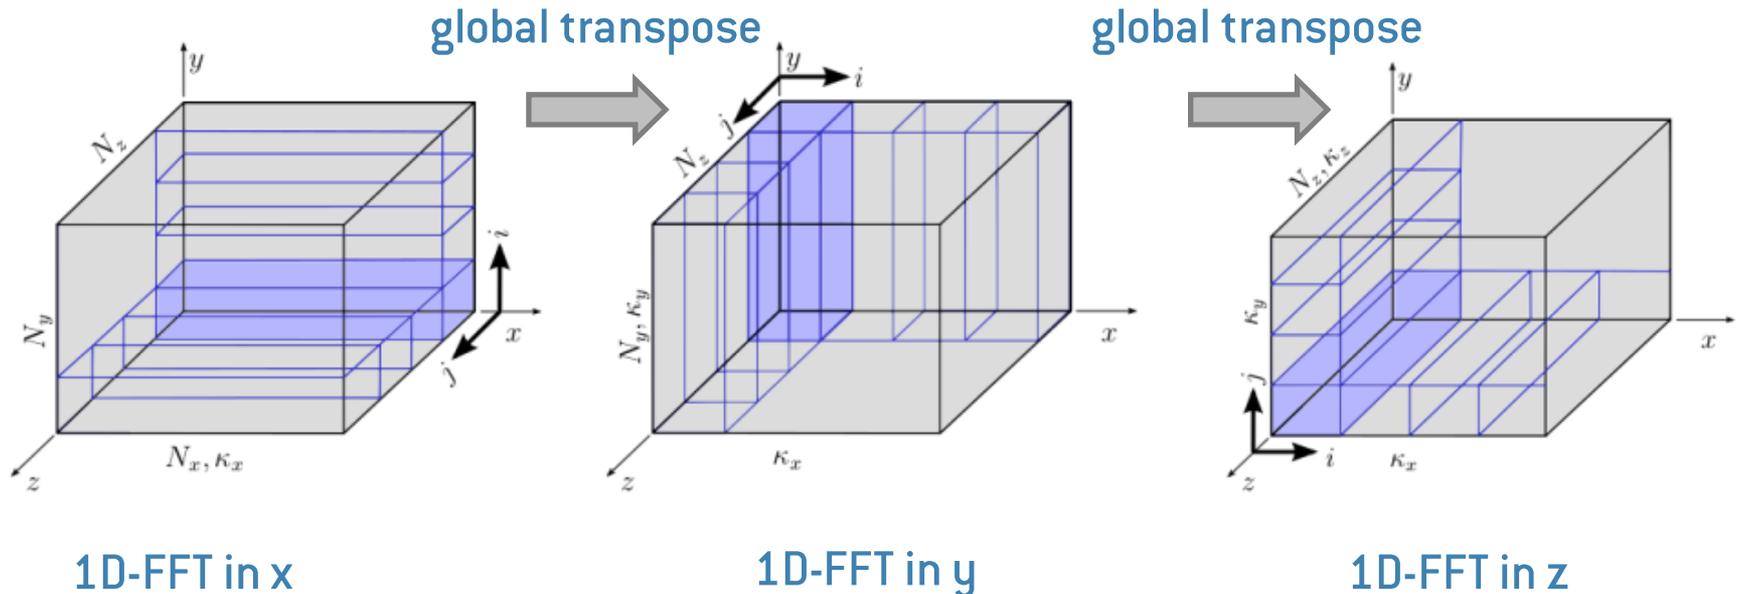


statistics along path lines



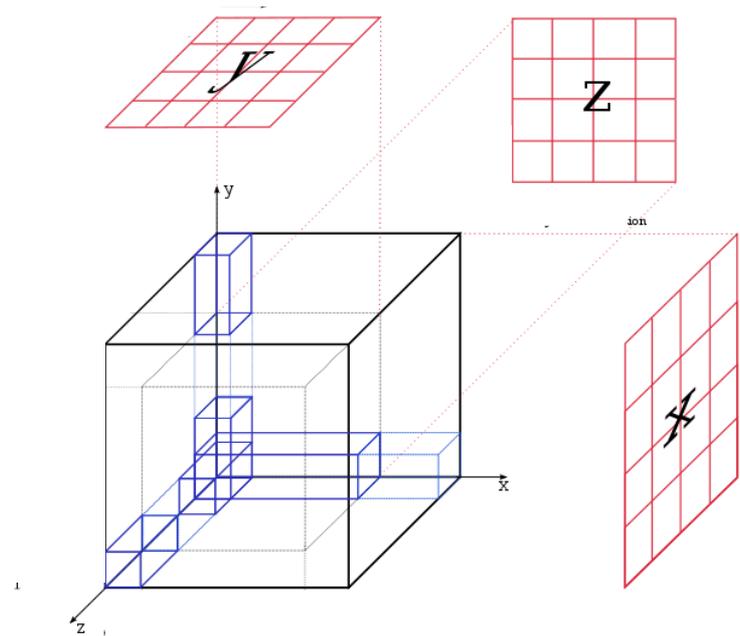
3D-FFTs on distributed memory systems

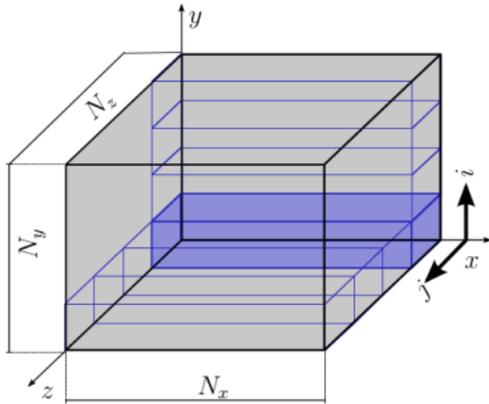
- distribute computational domain over all processors (1D, 2D)
- global transpose between computation of 1D-FFTs



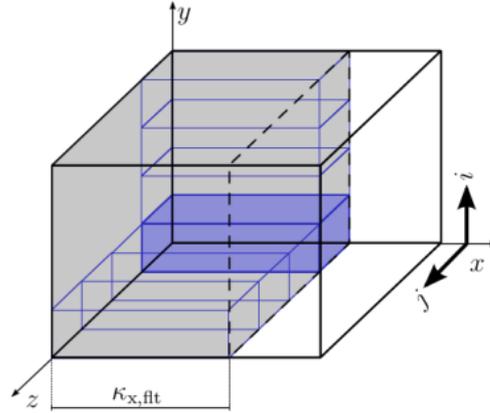
3D-FFTs on distributed memory systems

- 80% of the compute time required by 3d-FFTs
 - all-to-all communication
 - 2x data redistribution among processes
- development of an improved 3d-FFT library 'nb3dffft'
 - data reduction including the filter process
 - overlapping communication and computation

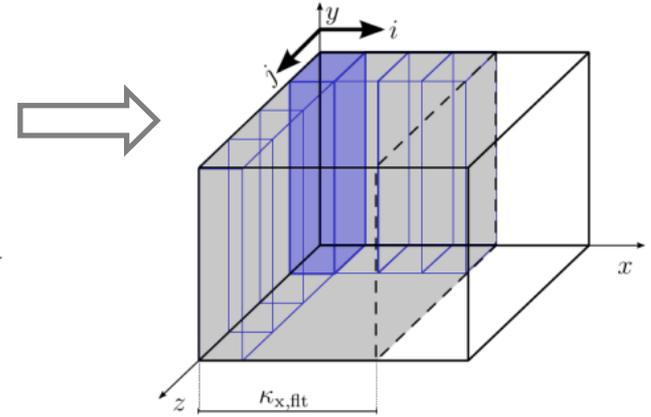




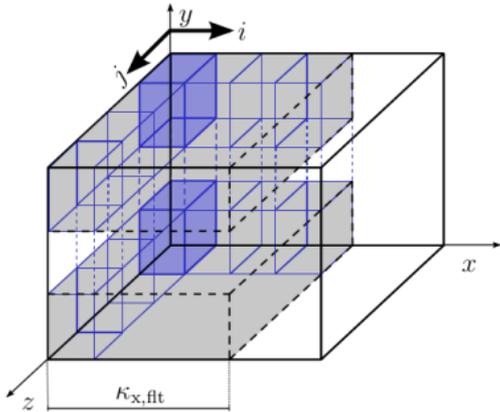
(a) 1d-FFTs applied in x direction



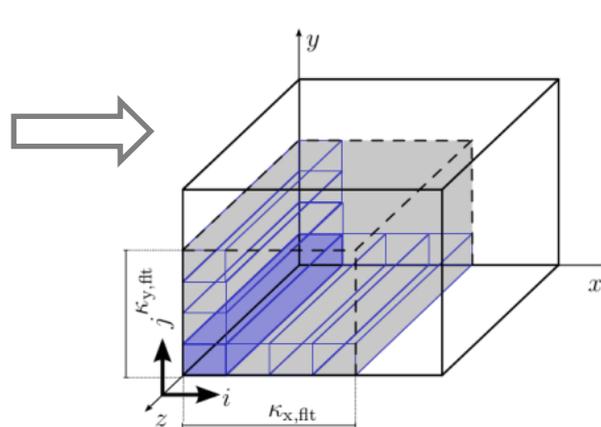
(b) cut-off filtering in x while packing for the transpose



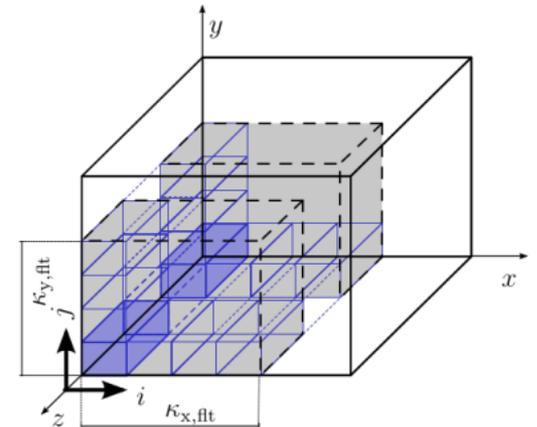
(c) unpacking and applying 1d-FFT in y direction



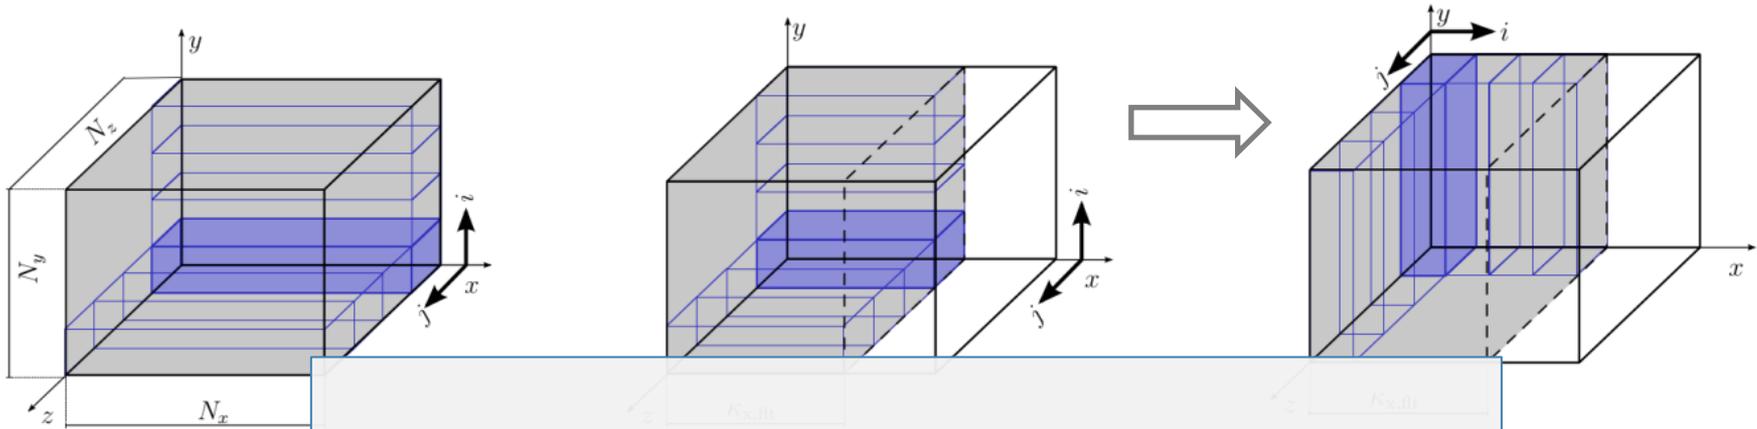
(d) cut-off filtering in y while packing for the transpose



(e) unpacking and applying 1d-FFT in z direction



(f) cut-off filtering in z considering access patterns



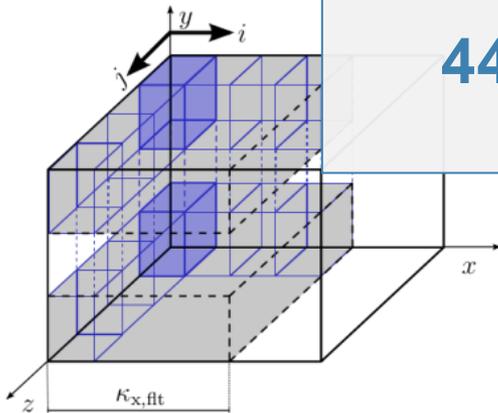
(a) 1d-FFTs applied in x direction

70.4% data reduction

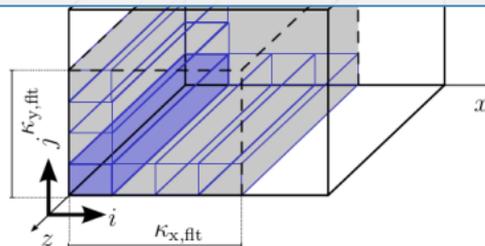
44.4% less data to send

44.4% less 1d-FFTs to compute

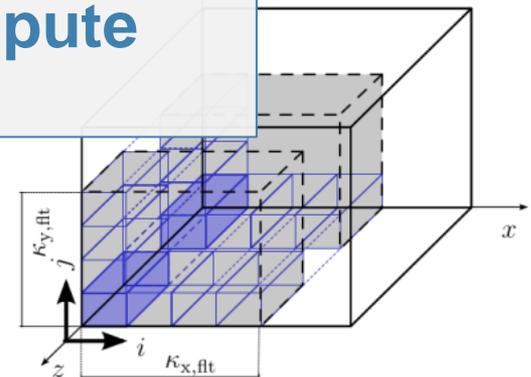
(c) unpacking and applying 1d-FFT in y direction



(d) cut-off filtering in y while packing for the transpose



(e) unpacking and applying 1d-FFT in z direction

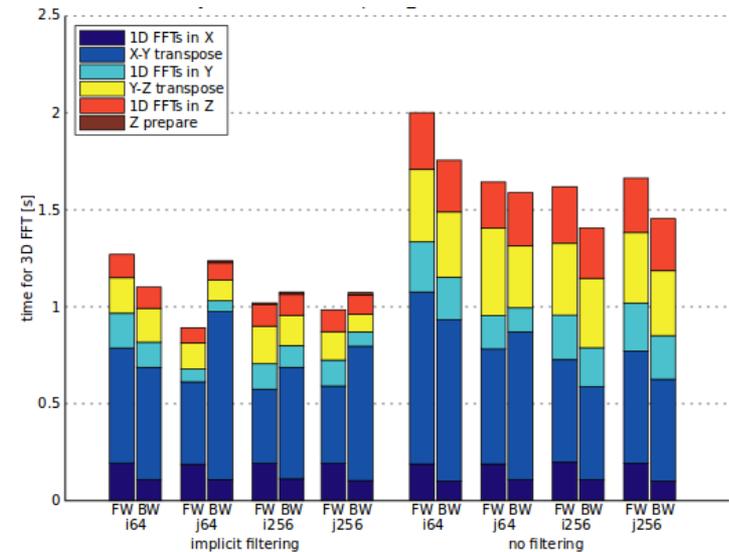
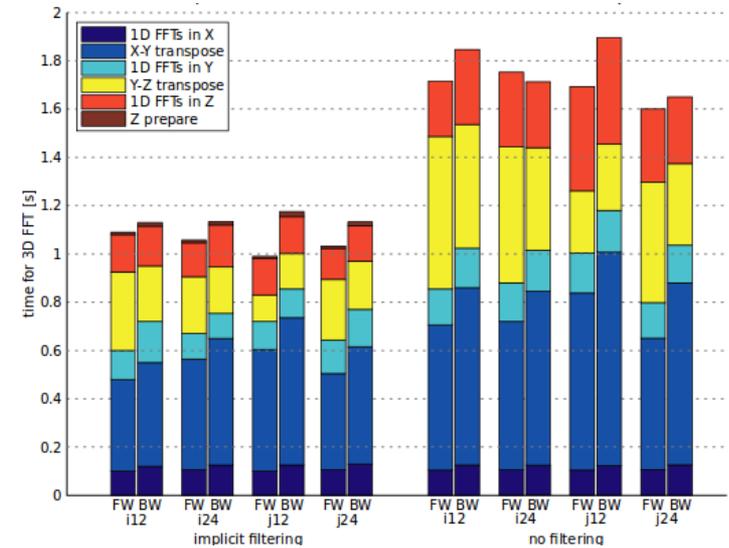


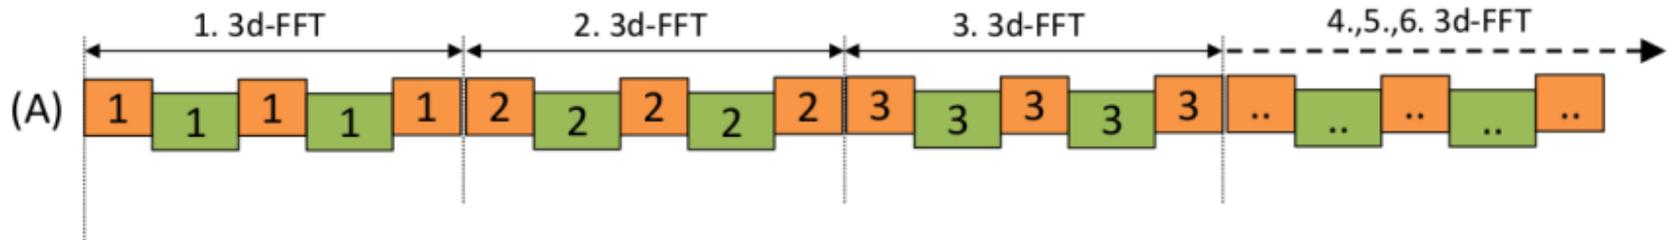
(f) cut-off filtering in z considering access patterns

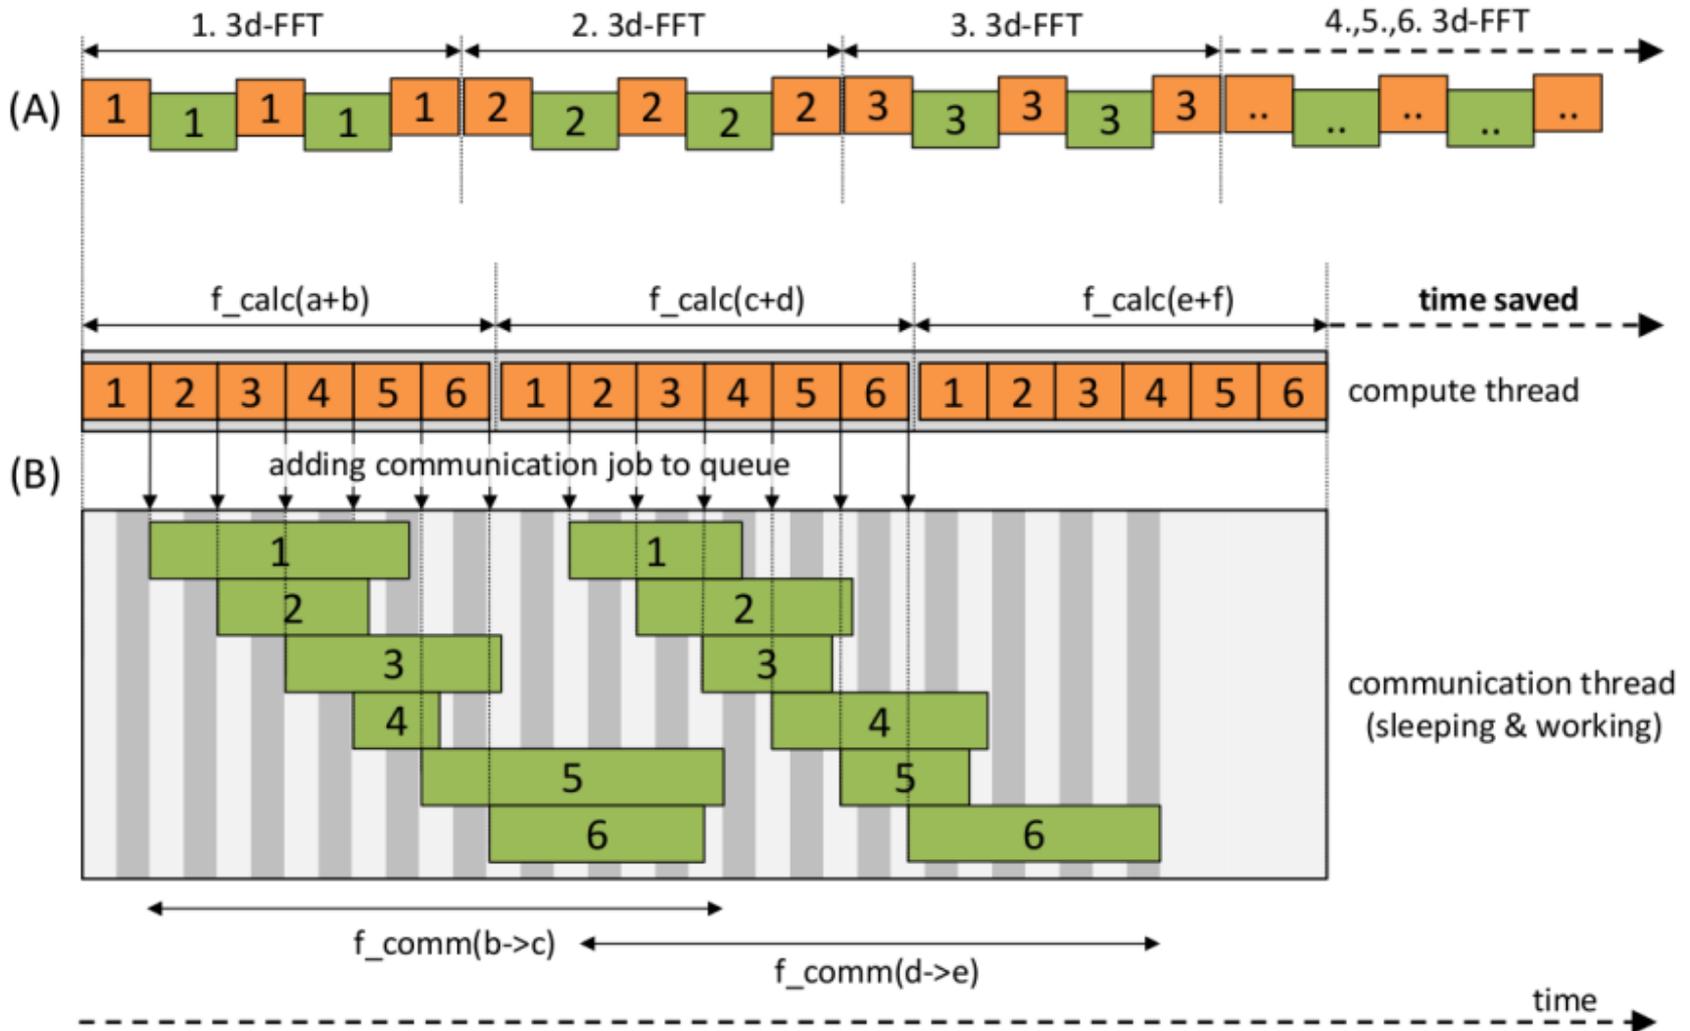
Performance Measurements

- RWTH Compute Cluster
 - 672 MPI processes on 56 nodes
 - domain size 2048^3

- JUQUEEN
 - 16.384 MPI processes on 8192 cores
 - domain size 4096^3







```

1: procedure ITERATE VELOCITY ( $n \rightarrow n+1$ )
2:   input:  $\hat{u}_{1|2|3}^n, \hat{G}_{4|5|6}^{n-1}, \hat{f}_u^{n-1}$ 
3:   for all wave numbers ( $\kappa$ ) do:                                ▷ compute vorticity
4:      $\hat{\omega}_{x,10}^n \leftarrow \kappa_z \hat{u}_{y,2}^n - \kappa_y \hat{u}_{z,3}^n$ 
5:      $\hat{\omega}_{y,11}^n \leftarrow \kappa_x \hat{u}_{z,3}^n - \kappa_z \hat{u}_{x,1}^n$ 
6:      $\hat{\omega}_{z,12}^n \leftarrow \kappa_y \hat{u}_{x,1}^n - \kappa_x \hat{u}_{y,2}^n$ 
7:      $\omega_{10|11|12}^n \xleftarrow{\mathcal{F}} \hat{\omega}_{10|11|12}^n$ 
8:     ▷ 3x in-place backward 3d-FFT
9:      $u_{7|8|9}^n \xleftarrow{\mathcal{F}} \hat{u}_{1|2|3}^n$ 
10:    ▷ 3x out-of-place backward 3d-FFT
11:    for all grid points ( $N$ ) do:                                ▷ compute  $G^n = \omega^n \times u^n$ 
12:       $G_{x,13}^n \leftarrow u_{y,8}^n \cdot \omega_{z,12}^n - u_{z,9}^n \cdot \omega_{y,11}^n$ 
13:       $G_{y,14}^n \leftarrow u_{z,9}^n \cdot \omega_{x,10}^n - u_{x,7}^n \cdot \omega_{z,12}^n$ 
14:       $G_{z,15}^n \leftarrow u_{x,7}^n \cdot \omega_{y,11}^n - u_{y,8}^n \cdot \omega_{x,10}^n$ 
15:       $\hat{G}_{13|14|15}^n \xleftarrow{\mathcal{F}} G_{13|14|15}^n$ 
16:      ▷ 3x in-place forward 3d-FFT
17:       $\hat{G}_{13|14|15}^n \xleftarrow{\text{filter}} \hat{G}_{13|14|15}^n$                                 ▷ dealiasing
18:       $\hat{G}_{13|14|15}^n \xleftarrow{\text{zero copy}} \hat{G}_{4|5|6}^{n-1}$                                 ▷ swap memory
19:       $\hat{u}_{16}^n \xleftarrow{\text{zero copy}} \hat{u}_2^n$                                 ▷ swap memory
20:       $\hat{f}_u^n \leftarrow F(\hat{f}_u^{n-1})$                                 ▷ compute forcing energy
21:      for wave numbers ( $\kappa_{ijk} \mid 0 \leq i, j, k \leq 2$ ) do:    ▷ add forcing energy
22:         $\hat{u}_{1|2|3}^{\text{force}} \leftarrow F_{\text{force}}(\hat{u}_{1|16|3}^n, \hat{f}_u^n)$ 
23:        for all wave numbers ( $\kappa$ ) do:                                ▷ advance velocity
24:           $\hat{u}_{1|2|3}^{\text{step1}} \leftarrow F_{\text{step1}}(\hat{u}_{1|2|3}^{\text{force}}, \hat{G}_{4|5|6}^{n-1}, \hat{G}_{13|14|15}^{n-1})$ 
25:          for all wave numbers ( $\kappa$ ) do:                                ▷ apply projection tensor
26:             $\hat{u}_{1|2|3}^{n+1} \leftarrow F_{\text{step2}}(\hat{u}_{1|2|3}^{\text{step1}})$ 
27:             $\langle \hat{u}_{1|2|3}^{n+1} \rangle = 0$                                 ▷ set mean velocity to zero
28:          output:  $\hat{u}_{1|2|3}^{n+1}, \hat{G}_{4|5|6}^n, u_{7|8|9}^n, \hat{u}_{16}^n$ 

```



```

1: procedure ITERATE PASSIVE SCALAR ( $n \rightarrow n+1$ )
2:   input:  $\hat{\phi}_{17}^n, \hat{J}_{18}^{n-1}, u_{7|8|9}^n, \hat{u}_{16}^n$ 
3:    $\nabla \hat{\phi}_{10|11|12}^n \leftarrow \hat{\phi}_{17}^n$                                 ▷ 3x compute derivatives
4:    $\nabla \hat{\phi}_{10|11|12}^n \xleftarrow{\mathcal{F}} \nabla \hat{\phi}_{10|11|12}^n$ 
5:   ▷ 3x in-place backward 3d-FFT
6:    $J_{10}^n \leftarrow u_{7|8|9}^n \cdot \nabla \hat{\phi}_{10|11|12}^n$                                 ▷ compute convective term
7:    $\hat{J}_{10}^n \xleftarrow{\mathcal{F}} J_{10}^n$ 
8:   ▷ 1x in-place forward 3d-FFT
9:    $\hat{J}_{10}^n \xleftarrow{\text{filter}} \hat{J}_{10}^n$                                 ▷ dealiasing
10:  for all wave numbers ( $\kappa$ ) do:                                ▷ advance passive scalar
11:     $\hat{\phi}_{17}^{n+1} \leftarrow F_{\text{adv}}(\hat{\phi}_{17}^n, \hat{J}_{18}^{n-1}, \hat{J}_{10}^n, \hat{u}_{16}^n)$ 
12:     $\hat{J}_{18}^{n+1} \xleftarrow{\text{zero copy}} \hat{J}_{18}^{n-1}$                                 ▷ swap memory
13:  output:  $\hat{\phi}_{17}^{n+1}, \hat{J}_{18}^{n+1}$ 

```

```

1: procedure ITERATE VELOCITY ( $n \rightarrow [n+1]$ ), PASSIVE SCALAR ( $[n-1] \rightarrow n$ )
2:   input (velocity):  $\hat{u}_{1|2|3}^n, \hat{G}_{4|5|6}^{n-1}, \hat{f}_u^{n-1}$ 
3:   input (passive scalar):  $\hat{\phi}_{17}^{n-1}, \hat{J}_{19}^{n-1}, \hat{J}_{18}^{n-2}, \hat{u}_{16}^{n-1}$ 
4:   for all wave numbers ( $\kappa$ ) do:                                ▷ compute vorticity
5:      $\hat{\omega}_{x,10}^n \leftarrow \kappa_z \hat{u}_{y,2}^n - \kappa_y \hat{u}_{z,3}^n$ 
6:      $\hat{\omega}_{y,11}^n \leftarrow \kappa_x \hat{u}_{z,3}^n - \kappa_z \hat{u}_{x,1}^n$ 
7:      $\hat{\omega}_{z,12}^n \leftarrow \kappa_y \hat{u}_{x,1}^n - \kappa_x \hat{u}_{y,2}^n$ 
8:      $\omega_{10|11|12}^n \xleftarrow{\mathcal{F}+3/2\text{filter}} \hat{\omega}_{10|11|12}^n$                                 ▷ 3x in-place backward 3d-FFT
9:      $u_{7|8|9}^n \xleftarrow{\mathcal{F}+3/2\text{filter}} \hat{u}_{1|2|3}^n$                                 ▷ 3x out-of-place backward 3d-FFT
10:     $\hat{J}_{19}^{n-1} \xleftarrow{\mathcal{F}+3/2\text{filter}} \hat{J}_{19}^{n-1}$                                 ▷ 1x in-place forward 3d-FFT
11:    for all wave numbers ( $\kappa$ ) do:                                ▷ advance passive scalar
12:       $\hat{\phi}_{17}^n \leftarrow F_{\text{adv}}(\hat{\phi}_{17}^{n-1}, \hat{J}_{19}^{n-1}, \hat{J}_{18}^{n-2}, \hat{u}_{16}^{n-1})$ 
13:       $\hat{J}_{19}^{n-1} \xleftarrow{\text{zero copy}} \hat{J}_{18}^{n-2}$                                 ▷ swap memory
14:      for all grid points ( $N$ ) do:                                ▷ compute  $G^n = \omega^n \times u^n$ 
15:         $G_{x,13}^n \leftarrow u_{y,8}^n \cdot \omega_{z,12}^n - u_{z,9}^n \cdot \omega_{y,11}^n$ 
16:         $G_{y,14}^n \leftarrow u_{z,9}^n \cdot \omega_{x,10}^n - u_{x,7}^n \cdot \omega_{z,12}^n$ 
17:         $G_{z,15}^n \leftarrow u_{x,7}^n \cdot \omega_{y,11}^n - u_{y,8}^n \cdot \omega_{x,10}^n$ 
18:         $\nabla \hat{\phi}_{10|11|12}^n \leftarrow \hat{\phi}_{17}^n$                                 ▷ 3x compute derivatives
19:         $\hat{G}_{13|14|15}^n \xleftarrow{\mathcal{F}+3/2\text{filter}} G_{13|14|15}^n$                                 ▷ 3x in-place forward 3d-FFT
20:         $\nabla \hat{\phi}_{20|21|22}^n \xleftarrow{\mathcal{F}+3/2\text{filter}} \nabla \hat{\phi}_{20|21|22}^n$                                 ▷ 3x in-place backward 3d-FFT
21:         $J_{19}^n \leftarrow u_{7|8|9}^n \cdot \nabla \hat{\phi}_{20|21|22}^n$                                 ▷ compute convective term
22:         $\hat{G}_{13|14|15}^n \xleftarrow{\text{zero copy}} \hat{G}_{4|5|6}^{n-1}$                                 ▷ swap memory
23:         $\hat{u}_{16}^n \xleftarrow{\text{zero copy}} \hat{u}_2^n$                                 ▷ swap memory
24:         $\hat{f}_u^n \leftarrow F(\hat{f}_u^{n-1})$                                 ▷ compute forcing energy
25:        for wave numbers ( $\kappa_{ijk} \mid 0 \leq i, j, k \leq 2$ ) do:    ▷ add forcing energy
26:           $\hat{u}_{1|2|3}^{\text{force}} \leftarrow F_{\text{force}}(\hat{u}_{1|16|3}^n, \hat{f}_u^n)$ 
27:          for all wave numbers ( $\kappa$ ) do:                                ▷ advance velocity
28:             $\hat{u}_{1|2|3}^{\text{step1}} \leftarrow F_{\text{step1}}(\hat{u}_{1|2|3}^{\text{force}}, \hat{G}_{4|5|6}^{n-1}, \hat{G}_{13|14|15}^{n-1})$ 
29:            for all wave numbers ( $\kappa$ ) do:                                ▷ apply projection tensor
30:               $\hat{u}_{1|2|3}^{n+1} \leftarrow F_{\text{step2}}(\hat{u}_{1|2|3}^{\text{step1}})$ 

```

- message sizes stay large
- number of messages do not change
- 2 synchronization points instead of 2x13

Performance increase by overlapping comm & comp

	1024 ³	2048 ³	4096 ³	6144 ³
1024 nodes	15.50%	21.60%	-	-
2048 nodes	-	26.63%	29.63%	-
4096 nodes	-	17.70%	19.78%	-
8192 nodes	-	11.64%	22.69%	17.53%

Results from Extreme Scaling Workshop 2015

Memory requirements:

- 20 datasets

float	6144 ³	8192 ³	12288 ³
4	34.6 TB	81.9 TB	276.5 TB ?
8	69.1 TB	163.8 TB	---

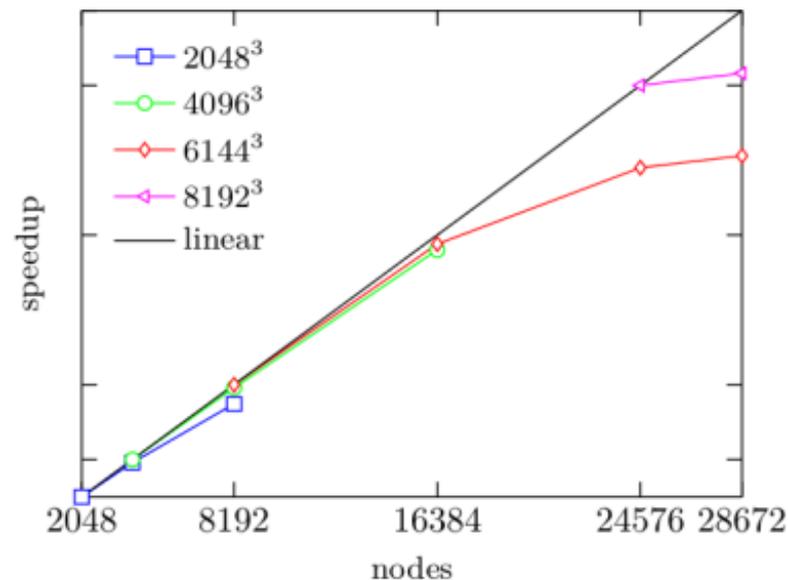


Table 1.2: Strong scaling tests of psOpen for various grid sizes.

bg_size	rpn	ranks	tpp	threads	time per iteration in seconds			
					2048 ³	4096 ³	6144 ³	8192 ³
2048	32	65536	2	131072	0.6295	3.0785	-	-
4096	32	131072	2	262144	0.3281	1.8815	-	-
8192	32	262144	2	524288	0.1805	0.9606	2.3276	-
16384	32	524288	2	1835008	-	0.4950	1.2132	-
24576	32	786432	2	1572864	-	-	0.9530	3.7811
28672	32	917504	2	1835008	-	-	0.9278	3.6801

- making pseudo-spectral codes for DNS scale on large systems is challenging
 - 3d-FFTs require lots of MPI_alltoall() calls, which cannot be avoided
- for best performance key functions have to be rewritten
- 3d-FFT can be optimized, if it is not thought as a single function call, but as part of the algorithm
 - Moving the required filter operation of the main algorithm into the 3d-FFT reduces the data to be sent by 44.4% and the 1d-FFTs by 44.4%
 - Calling multiple 3d-FFT for different fields at the same time allows overlapping of communication and computation -> ~20% performance gain

Thank you for your attention.