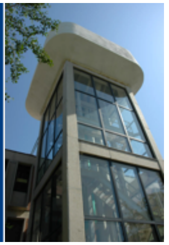

Large Eddy Simulation of Sediment Transport and Hydrodynamics at River Bifurcations using a Highly Scalable Spectral Element based CFD Solver

Som Dutta
Ven Te Chow Hydrosystems Lab
Dept. of Civil and Environmental Engineering
University of Illinois at Urbana-Champaign

PI: Prof. Marcelo H. Garcia
Ven Te Chow Hydrosystems Lab
Dept. of Civil and Environmental Engineering
University of Illinois at Urbana-Champaign

Co-PI: Prof. Paul Fischer
Dept. of Computer Science, and Dept. MechSE
University of Illinois at Urbana-Champaign

Bulle-Effect: the non-linear distribution of near bed sediment between the lateral and the main channel of a steam/river diversion.

Bifurcation: when a river/stream splits into two. e.g. the Pannerdensche Kop bifurcation on Rhine River, Netherlands.

(image courtsey :http://www.citg.tudelft.nl/uploads/RTEmagicC_Rivers_Rijn_by_Gelderlander.jpg.jpg)

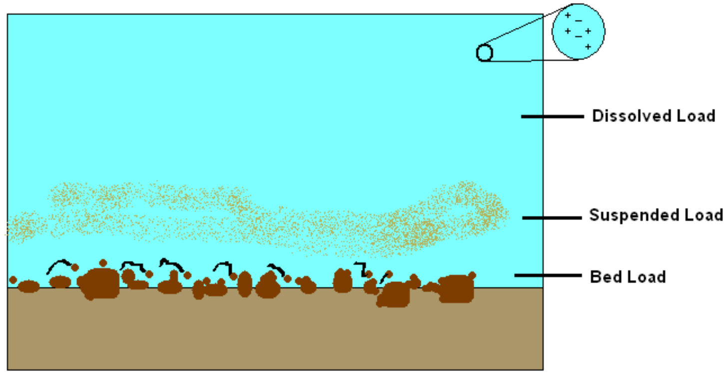

$$\frac{S_{side}}{S_{main}} = a\left(\frac{Q_{side}}{Q_{main}}\right)^{b}$$

Diversion: A special type of bifurcation where the main-channel continues along the original path. e.g. Mississippi River, West Bay diversion.

(image courtsey http://media.nola.com/hurricane_impact/photo/9034828-large.jpg)

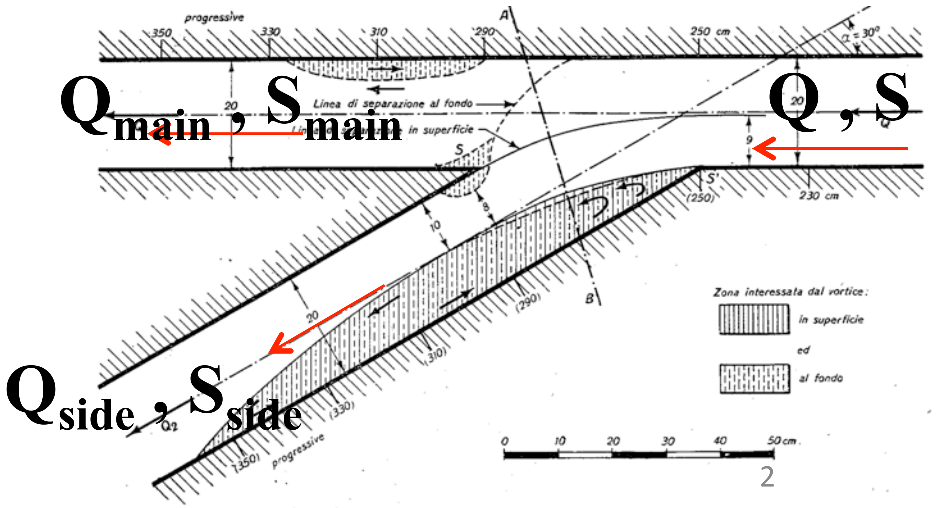

Bulle-Effect: what does the experiments say ?

- Data reproduced from Bulle's experiments from 1926.
- All the cases had a constant water discharge of Q = 5 l/sec = 0.005 m^3s^{-1}
- Experiments were done for five different diversion angles (30, 60, 90, 120 and 150 degree).
- Experiments were also done for different water discharge ratios.

$$\frac{S_{side}}{S_{main}} = a\left(\frac{Q_{side}}{Q_{main}}\right)^b$$

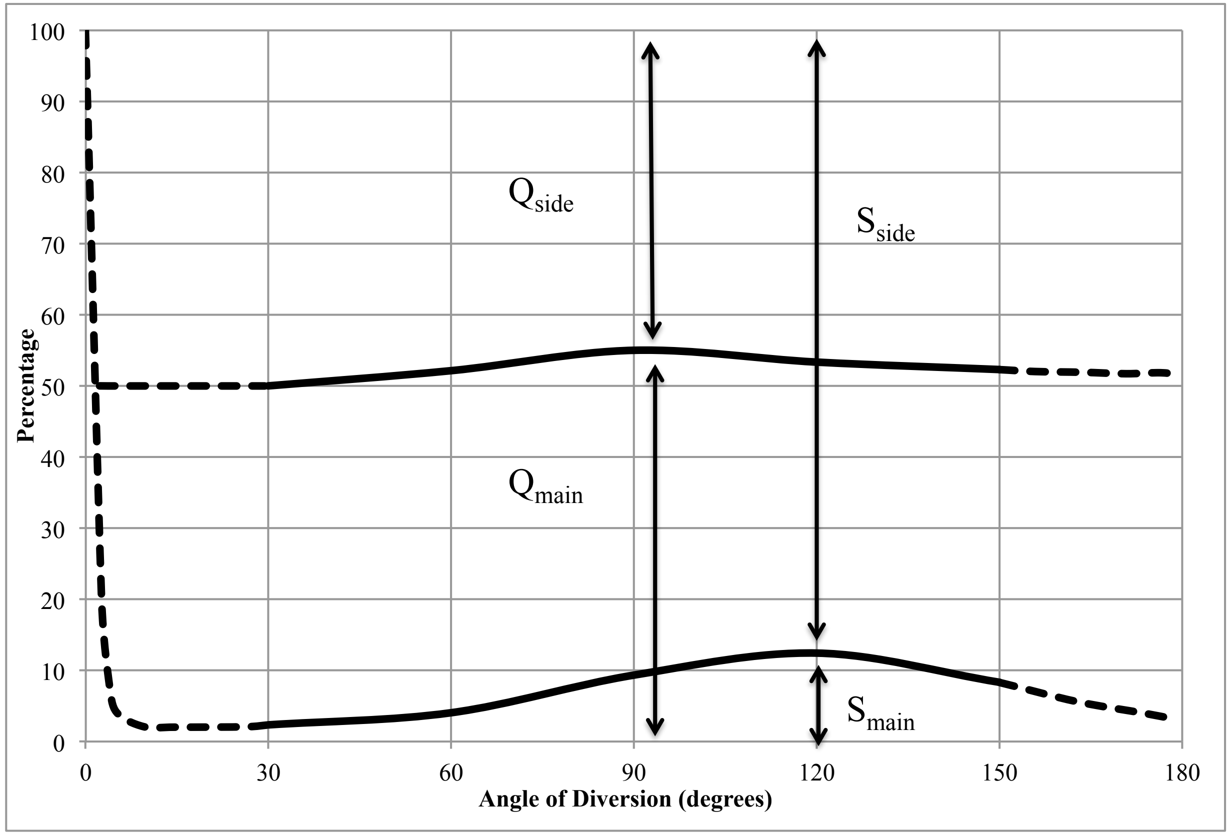

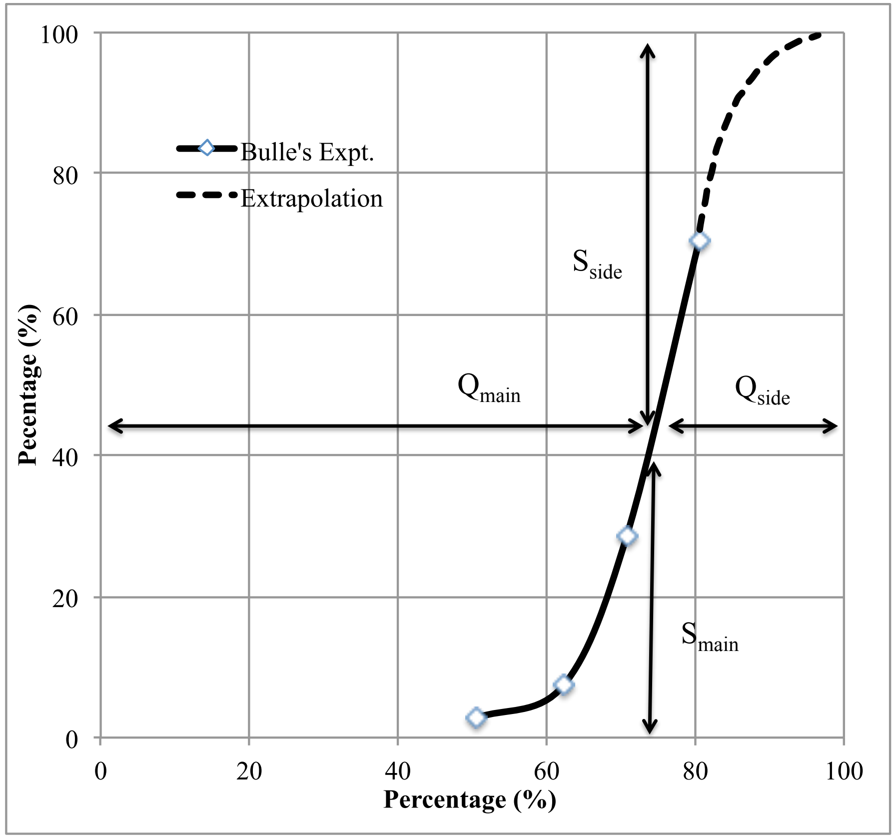

Bulle-Effect: what does the experiments say ?

- 20 years later Mario Cristani extended Bulle's experiments, and found that the data matched the trend predicted by Bulle.
- Albert Dancy (1946) found the dependence of the phenomenon on size of sediment (particle fall velocity).

- Experiments over the years, have further confirmed the phenomena, but a clear idea about the mechanism is still elusive.

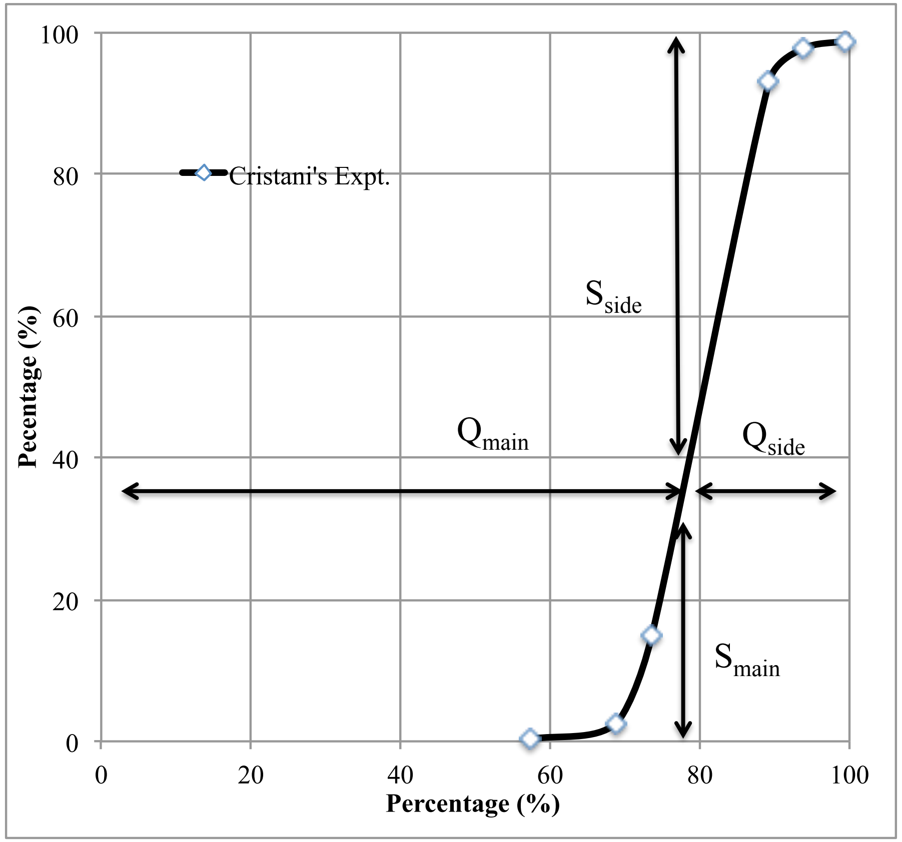

Dancy, 1946 >>

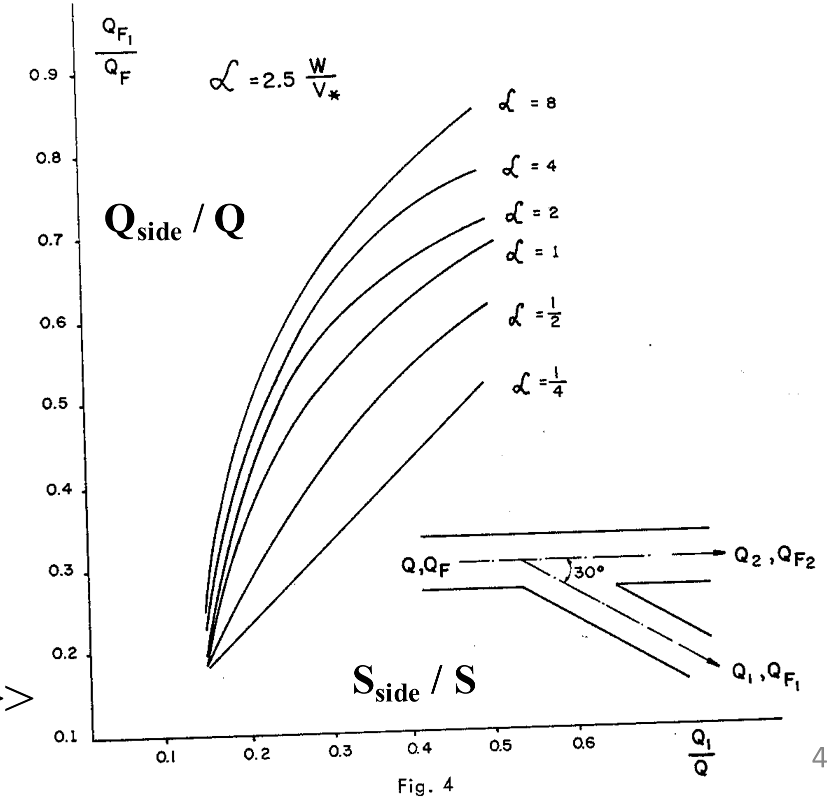

Why is Bulle-Effect important ?

The knowledge gained from understanding the phenomena will help **Save the Deltas** around the world.

A

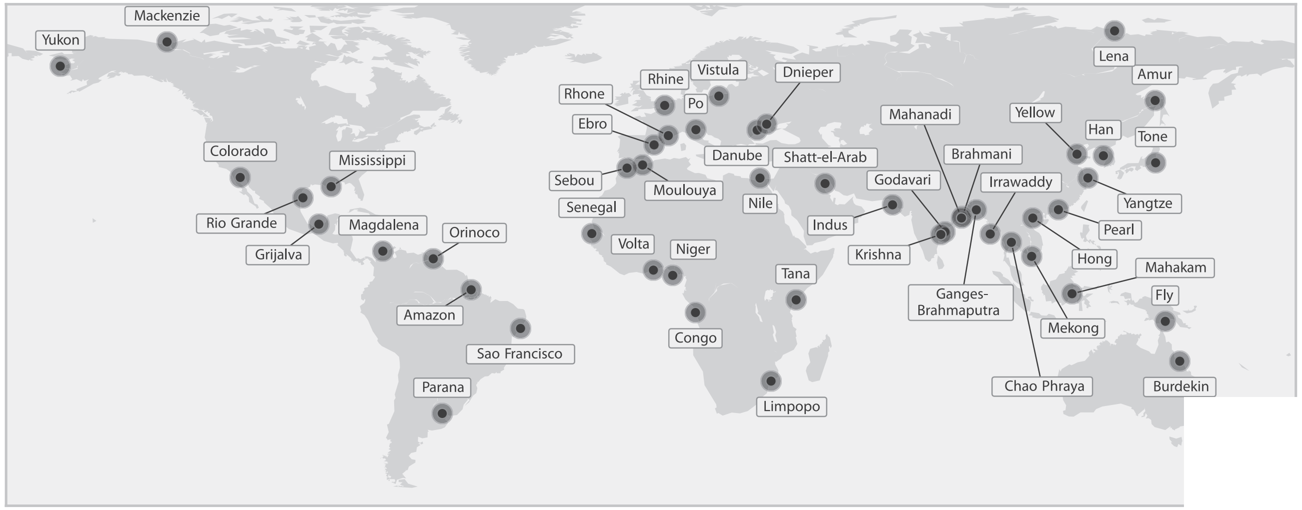

Profiling risk and sustainability in coastal deltas of the world
Z. D. Tessler *et al.*
Science **349**, 638 (2015);
DOI: 10.1126/science.aab3574

- Only way the vulnerable deltas can be saved is through creating **diversions** to build land.
- Also the structure of the river in a delta is highly dendritic and full of bifurcations/natural diversions.

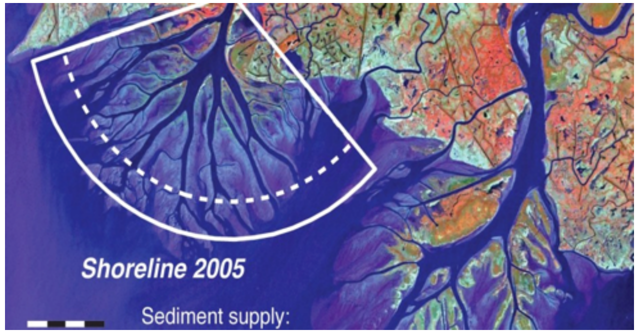

Image source: http://www.mississippiriverdelta.org/

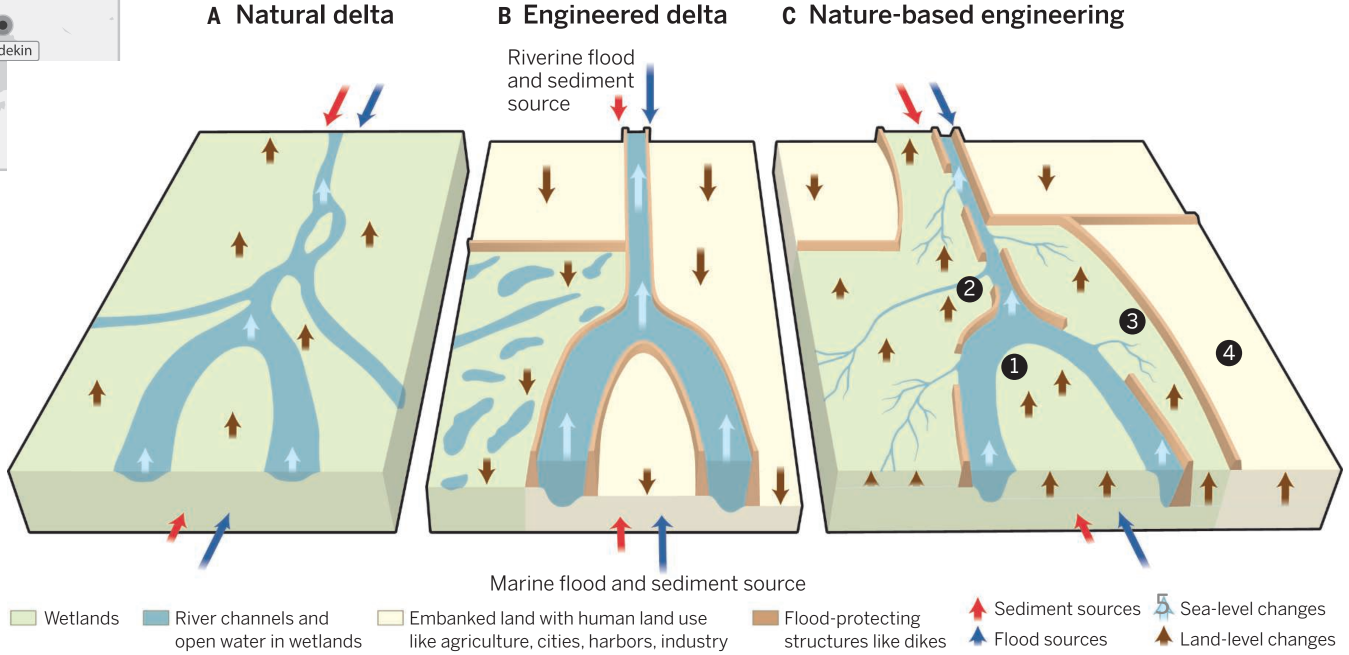

Building land with a rising sea
Stijn Temmerman and Matthew L. Kirwan
Science **349**, 588 (2015);
DOI: 10.1126/science.aac8312

An example where delta rebuilding is already on the drawing board

The Mississippi River Delta

The knowledge gained from understanding the mechanism behind *Bulle-Effect* will help design these diversions for efficiently.

Source:
http://www.mississippiriverdelta.org/

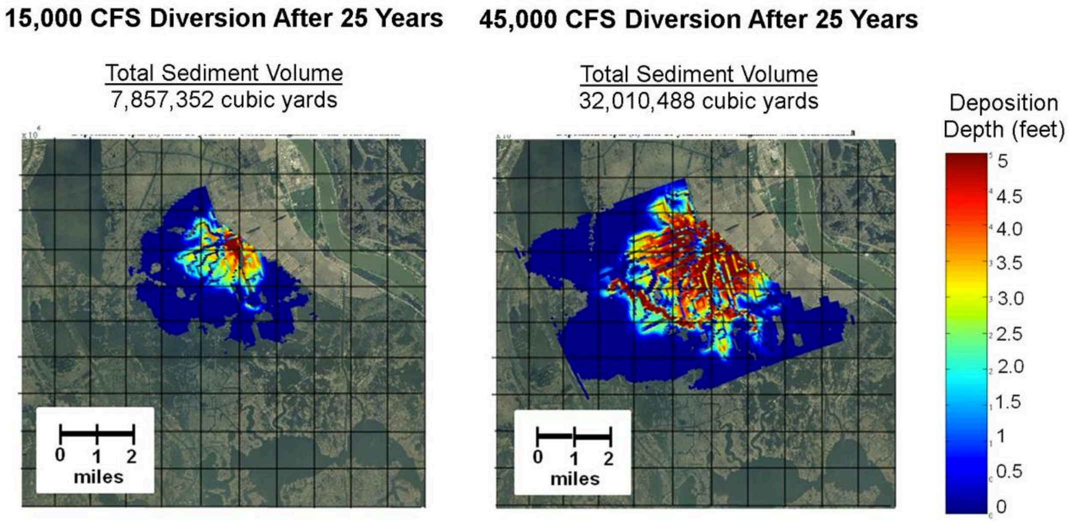

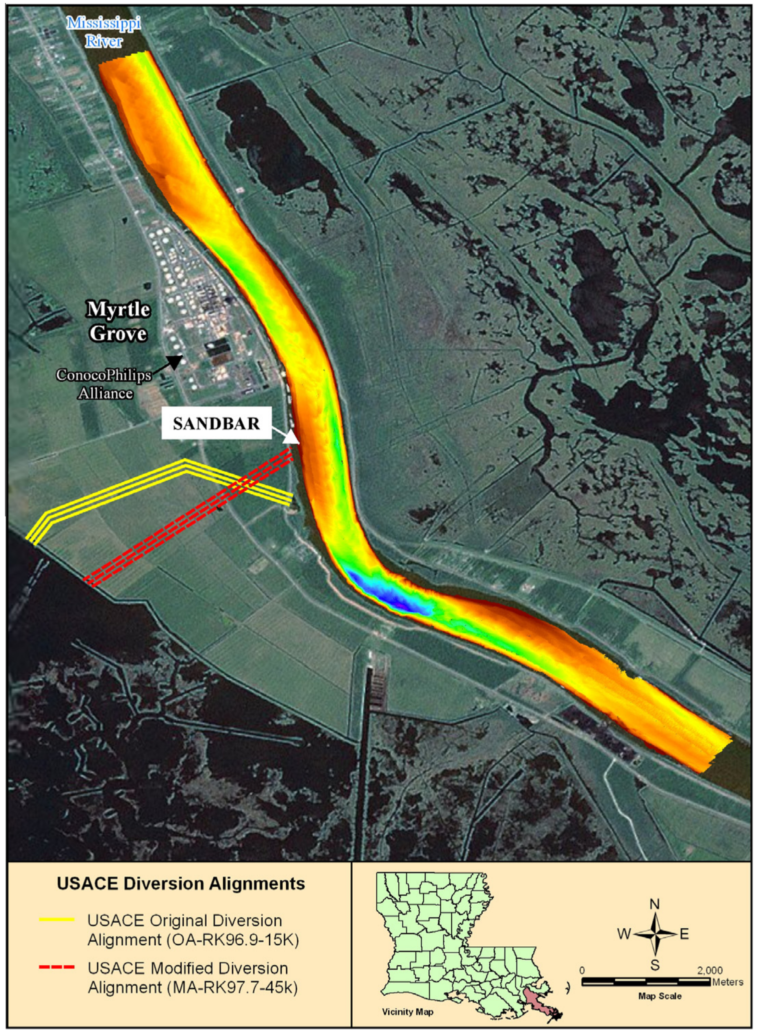

(Meselhe et al. 2012)

Objectives of the study

- Conduct high quality **Large Eddy Simulations** of the flow and sediment transport **at the scale of Bulle's experiments**. The simulations will **resolve all the important features of the flow**, thus simulate the hydrodynamics accurately. And the **Lagrangian particle model** will **help capture the response of the sediment to the complex flow field** accurately.

- In order to understand the effect of Reynolds number of the flow on Bulle-Effect, simulations were conducted for a range of bulk Reynolds numbers (Re), from 10 to 25,000, for diversion angle of 90 degrees.

- For Reynolds numbers in the range 10-7000, the resolution of the mesh is good-enough for Direct Numerical Simulation (DNS).

- Most of the simulations were conducted for the discharge ratio Q_{side}:Q_{main} of 50:50, though in order to study the effect of different discharge ratios on Bulle-Effect, simulations with five different discharge ratios (15:85, 35:65, 50:50, 65:35, 85:15) were conducted for Re = 7000, 25000.

- For Re = 25000, simulations were also conducted for diversion angles of 30, 60, 120 and 150 degrees.

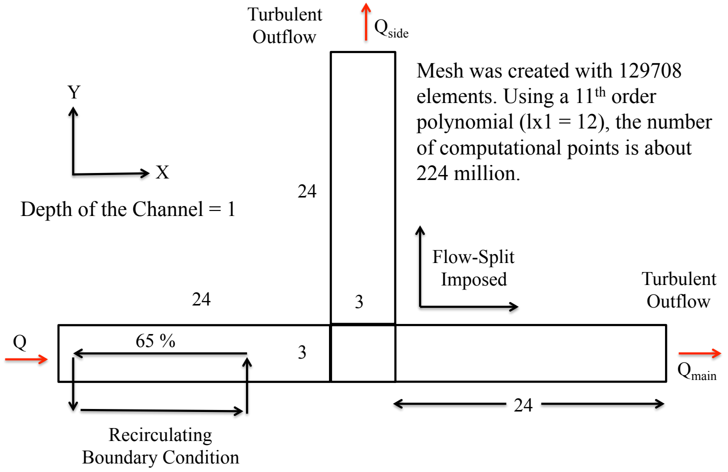

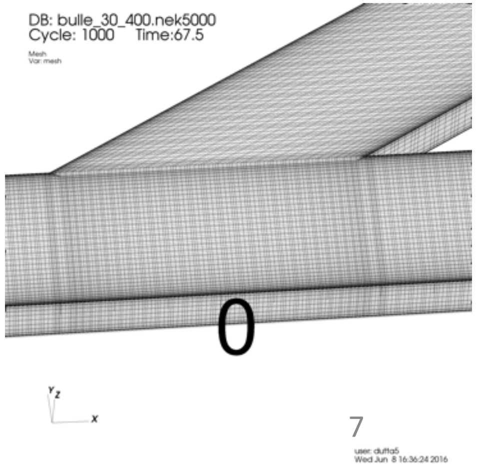

Nek5000: the highly scalable incompressible Navier-Stokes solver

- Simulations were conducted using Nek5000, which is an open-source spectral element based incompressible Navier-Srokes solver. (Fischer et al., https://nek5000.mcs.anl.gov/)

- The spectral element method (SEM) combines the accuracy of spectral methods with the flexibility of local approaches, like Finite Element Methods.

- Nek5000 uses high-order Legendre polynomials as the basis function, along with a Gauss- Lobatto-Legendre grid.

- Using high-order polynomial eliminates dispersion errors, which is very important for large-scale and long-term turbulence calculations. (Kreiss & Oliger 72, Gottlieb et al. 2007)

- Time-stepping is done using the combination of 3rd order Backward Differencing (BDF) and Extrapolation (for the non-linear terms).

- In case the dissipative scales are not resolved at higher Reynolds number (LES), a local element based explicit cutoff filter (a spectral filter) in the wave number space is used to remove energy from the highest wavenumbers (Fischer and Mullen, 2001).

Parallel scalability of Nek5000

- Nek5000 used MPI for parallelization.
- Nek5000 has a history of scaling efficiently on different HPC platforms throughout the world.
- Recently it has shown strong scaling up to a million MPI ranks on MIRA (at ALCF), with parallel efficiency of 60 percent. The tested problem had about 2 billion computational points, showing strong scaling for granularity of ~ 2000 points/processors.
- Nek's very efficient parallel scalability is due to, a scalable 'gather-scatter' kernel, that helps reduce the communication cost during the simulation, especially for the global communication needed in the pressure (Poisson) solver.
- Scalability will depend on the architecture and hard-ware of the machine.

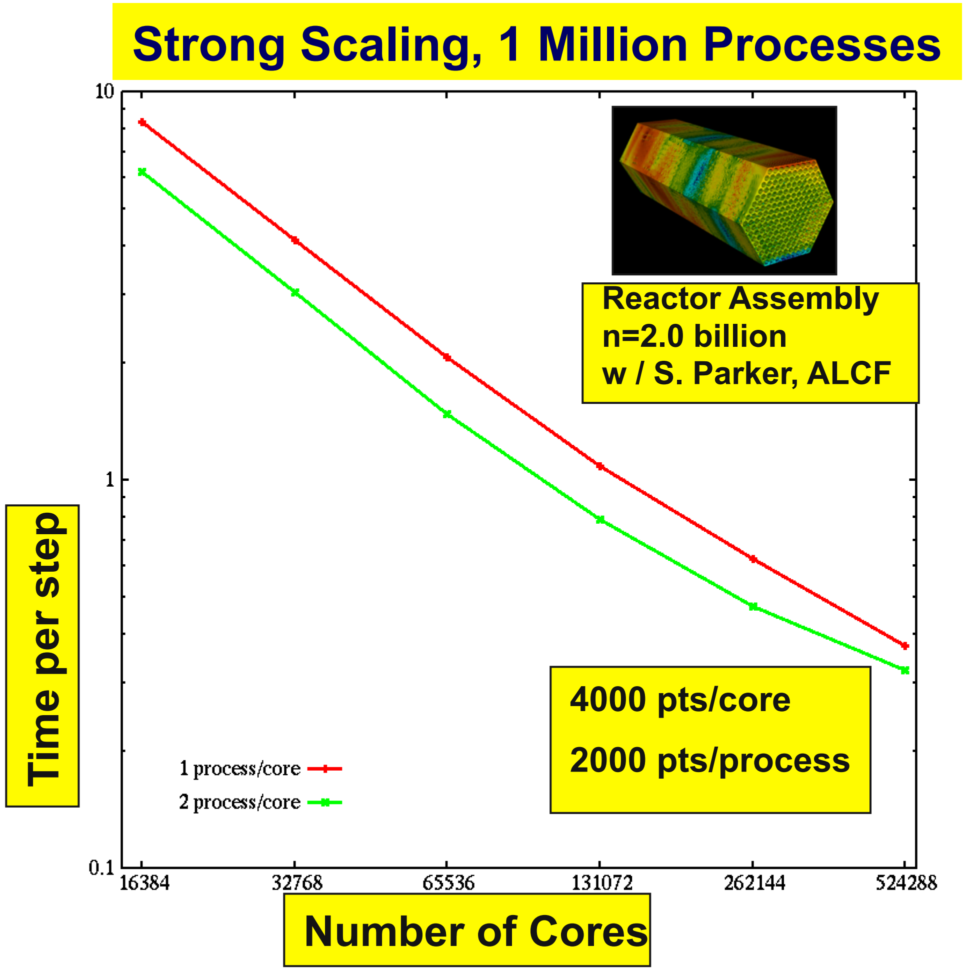

Configuration of Current Problem

- Dimensions of the simulations are same as Bulle's experiments.
- The 90-degree mesh has around 130,000 elements, and combined with the 12 collocation points for each direction at each element, in total about 224 million computational points. The mesh for the 30-degree case has around 242.74 million grid points.
- Part of the inlet channel is used for recirculating the flow, in order to have a fully-formed turbulent flow.
- The flow-split at the bifurcation was imposed using a using a fast implicit enforcement of the flow division. This helped in accurate yet faster convergence of the simulation.
- Sediment particles were modeled using a semi-implicit Lagrangian particle tracking algorithm developed for the current study. (the algorithm has been presented during the poster session here.)

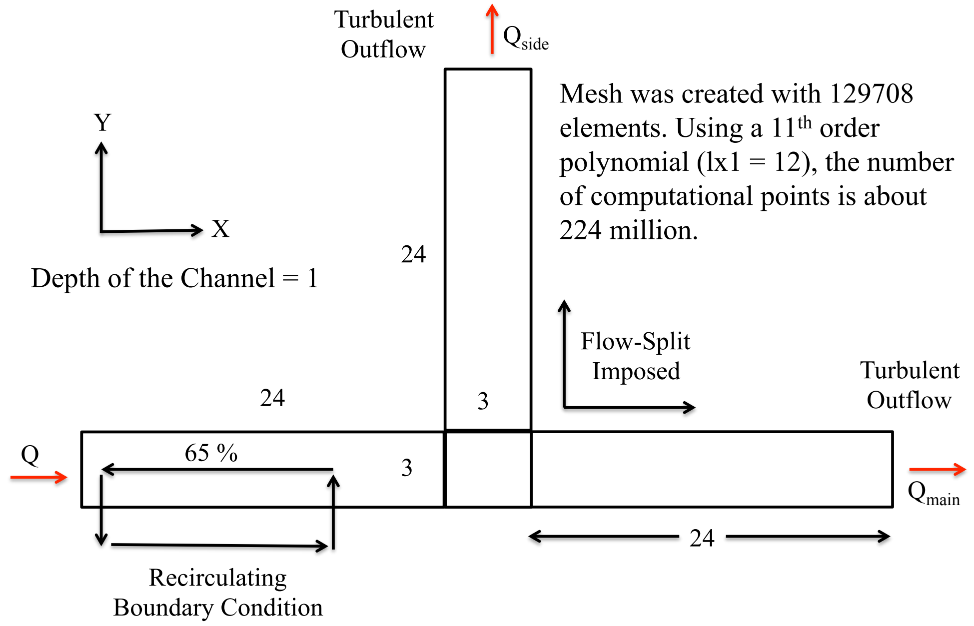

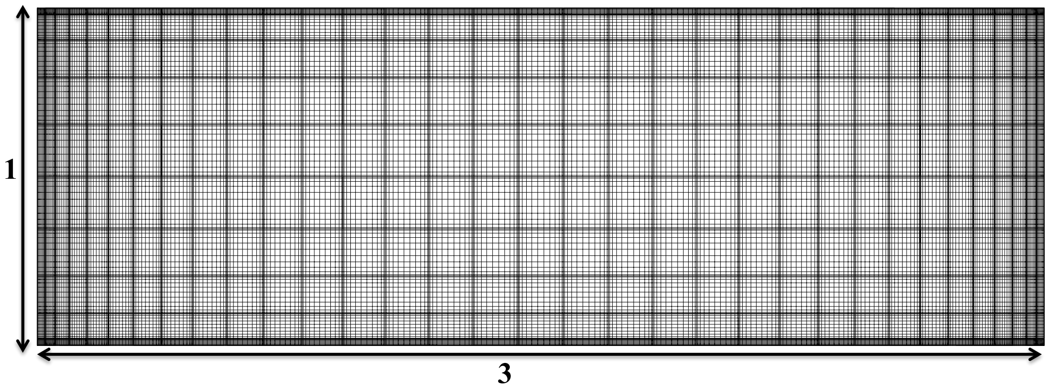

For the points closest to the walls, $z^+ = 0.058$ and $y^+ = 0.65$

Why Blue Waters ?

- Scale of the simulations are same as Bulle's experiments, which makes them one of the few simulations in the field of River Mechanics that has conducted high-quality LES at the scale of experiments.

- The number of **computational points is in the range of 224 million to 242.74 million**, along with 120,000 particles.

- Each simulations were run long enough to reach a statistical steady state, which can range from 90 to 150 convective time units depending on the flow-split. And then sediment particles were added to the domain, which took around 40-50 time units to move out from the main channel. **So 130 -200 convective time-units.**

- The Re = 25000 cases takes approximately 256 node hours for 1 convective time units, **which means it can take up to 51200 node hours for a complete simulation** … this **would not have been possible without a petascale system which can provide sustained performance = Blue Waters**.

Performance of Nek5000 on Blue Waters

- Parallel scalability on Blue Waters, was tested for the target problem using a case with 224.136 million gird points
- We always use XE nodes, with 32 processors on each node.
- Nek5000 was found to scale strongly up to 32768 mpi ranks, with linear speedup up to 4096 mpi ranks and relatively efficient scaling up to 16384 mpi ranks (40 %).
- Thus on Blue Waters, Nek5000 achieved strong scaling till n/P ~ 6840, though at n/P ~ 27360 the parallel efficiency has already reduced to 68.3 %. (n is number of grid points, and P is number of processors)
- Thus Nek5000, for the target problems scales efficiently, but not to the level achieved on MIRA.
- But Why ?

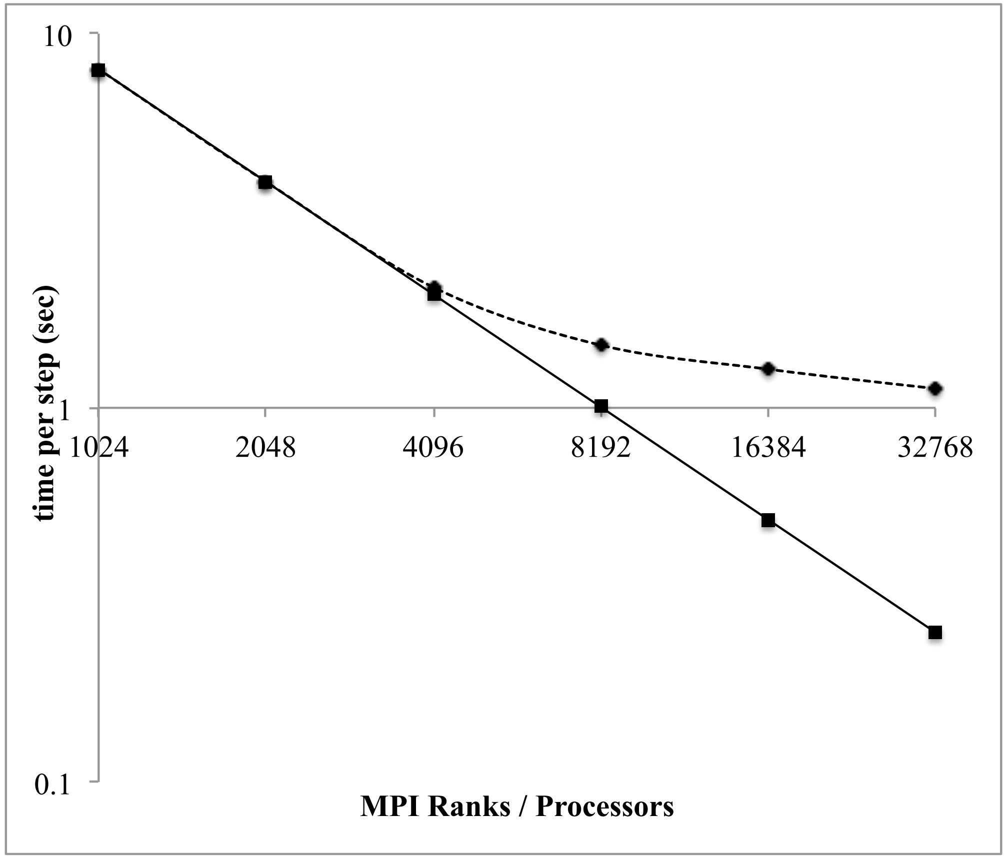

Performance of Nek5000 on Blue Waters compared to Mira (BGQ/ANL)

Nonoverlapping: $T(n, P) \;=\; T_a/P \,+\, T_c(n, P) \,+\, c_0$

Overlapping: $T(n, P) \;=\; \max[T_a/P, T_c(n, P)] \,+\, c_0$

n = number of computational points
P = number of processors
T = total time to compute
T_a = time to complete a parallelizable process on 1 processor
T_c = communication time
C_0 = other overheads and non-parallelizable work

For parallel efficiency of 100 %

$$T_a/P \quad \gg \quad T_c(n, P) \,+\, c_0,$$

For reasonable parallel efficiency, that is 50 to 70 %

Analysis inspired by Fischer et al., 2015, AIAA

$$T_a/P \quad \geq \quad T_c(n, P) \,+\, c_0.$$

So, we will try to come up with the theoretical values of n/P to which we can expect our problem to scale on Blue Waters.

Performance of Nek5000 on Blue Waters compared to Mira (BGQ/ANL)

Inter-processor communication costs

$$t_c(m) \;=\; (\alpha \;+\; \beta m)\; t_a$$

Number of 64-bit word messages passed

internode latency (non-dimensional)

Inverse bandwidth

inverse of the flop rate observed for the given algorithm on the computer in question, in the absence of communication

- t_a estimated through matrix-matrix product performance for sets of noncached matrices of order 10. (typical for Nek5000).

- Internode latency and inverse bandwidth are parameters that tells us about the inter-node communication. These are estimated using Ping-Pong tests.

Performance of Nek5000 on Blue Waters compared to Mira (BGQ/ANL)

Ping-Pong test results, on Cray XK7 Titan, BGQ (ANL)

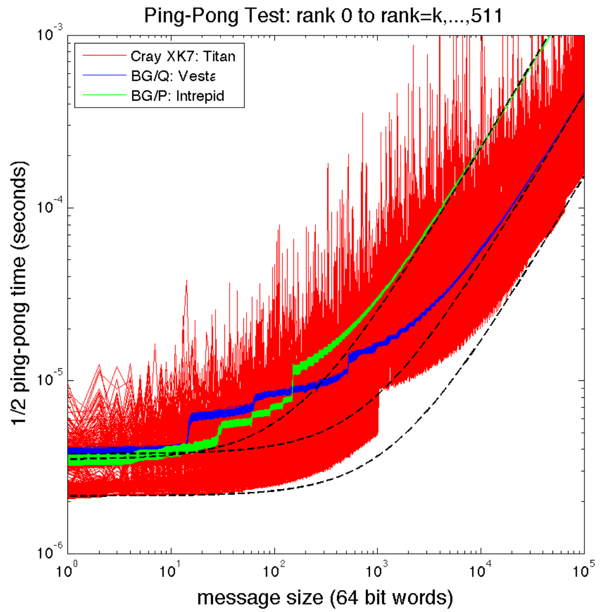

Year	t_a (μs)	αt_a (μs)	βt_a (μs/wd)	α	β	m_2	machine
1986	50	5960	64	119.2	1.28	93	Intel iPSC-1 (286)
1987	0.333	5960	64	17898	192	93	Intel iPSC-1/VX
1988	10	938	2.8	93.8	0.28	335	Intel iPSC-2 (386)
1989	0.25	938	2.8	3752	11.2	335	Intel iPSC-2/VX
1990	0.1	80	2.8	800	28	29	Intel iPSC-i860
1991	0.1	60	0.8	600	8	75	Intel Delta
1992	0.066	50	0.15	760	2.3	333	Intel Paragon
1995	0.02	60	0.27	3000	13.5	222	IBM SP2 (BU96)
1996	0.016	30	0.02	1875	1.25	1500	ASCI Red 333
1998	0.006	14	0.06	2333	10	233	SGI Origin 2000
1999	0.005	20	0.04	4000	8	500	Cray T3E/450
2005	0.002	4	0.026	2000	13	154	BGL/ANL
2008	0.0017	3.5	0.022	2060	13	160	BGP/ANL
2011	0.0007	2.5	0.002	3570	2.87	1250	Cray Xe6 (KTH)
2012	0.0007	3.8	0.0045	5430	6.43	845	BGQ/ANL
2015	0.0004	2.2	0.0015	5500	3.75	1467	Cray XK7

Ping-Pong test results, on Blue Waters

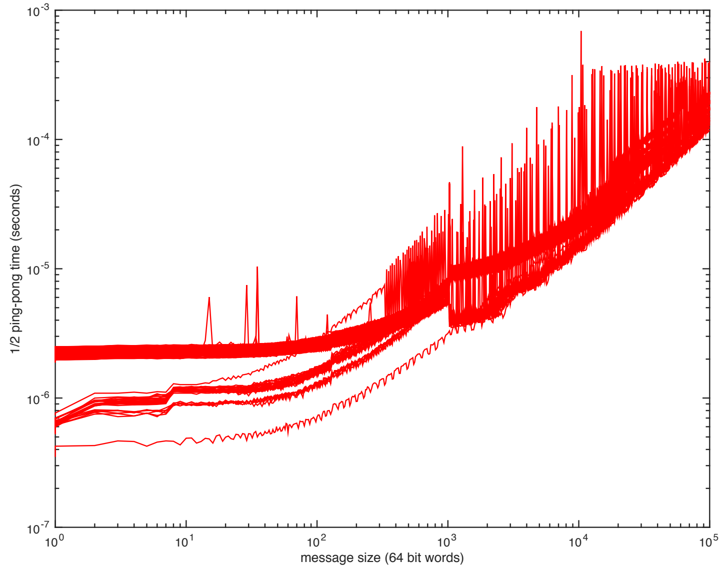

For Blue Waters: we found t_a = 0.00045 , averaged alpha = 4800, averaged betta = 4.91

For Blue Waters: we found t_a = 0.00045 , max alpha = 5530, max betta = 7.47

Performance of Nek5000 on Blue Waters compared to Mira (BGQ/ANL)

If the work the computer is doing primarily involves Jacobi Iterations: $$\frac{T_{cJ}}{T_{aJ}} = \frac{6(\alpha + \beta(n/P)^{\frac{2}{3}})}{14n/P} \leq 1.$$

MIRA > n/P ~ 1700 , Blue Waters (avg) > n/P ~ 2500 , Blue Waters (max) > n/P ~ 3000

If the work the computer is doing primarily involves Jacobi preconditioned Conjugate-Gradient iterations (e.g. Pressure-Poisson) : $$\frac{T_{cCG}}{T_{aCG}} = \frac{6(\alpha + \beta(n/P)^{\frac{2}{3}}) + 4\,\alpha \log_2 P}{27n/P} \leq 1.$$

MIRA > $P = 10^9$ n/P ~ 12000 , $P = 10^9$ n/P ~ 17000

MIRA + hardware support for MPI all reduce > $P = 10^9$ n/P ~ 2200

Blue Waters (avg) > n/P ~ P = 32768 n/P ~ 12500

Blue Waters (max) > n/P ~ P = 32768 n/P ~ 14000

So, the difference maker was **“hardware supported MPI all reduce”**

Results: 90-degree, Re=20000, 50:50 flow split

Re = HU/v

H – channel height

U – Mean velocity

v - viscosity

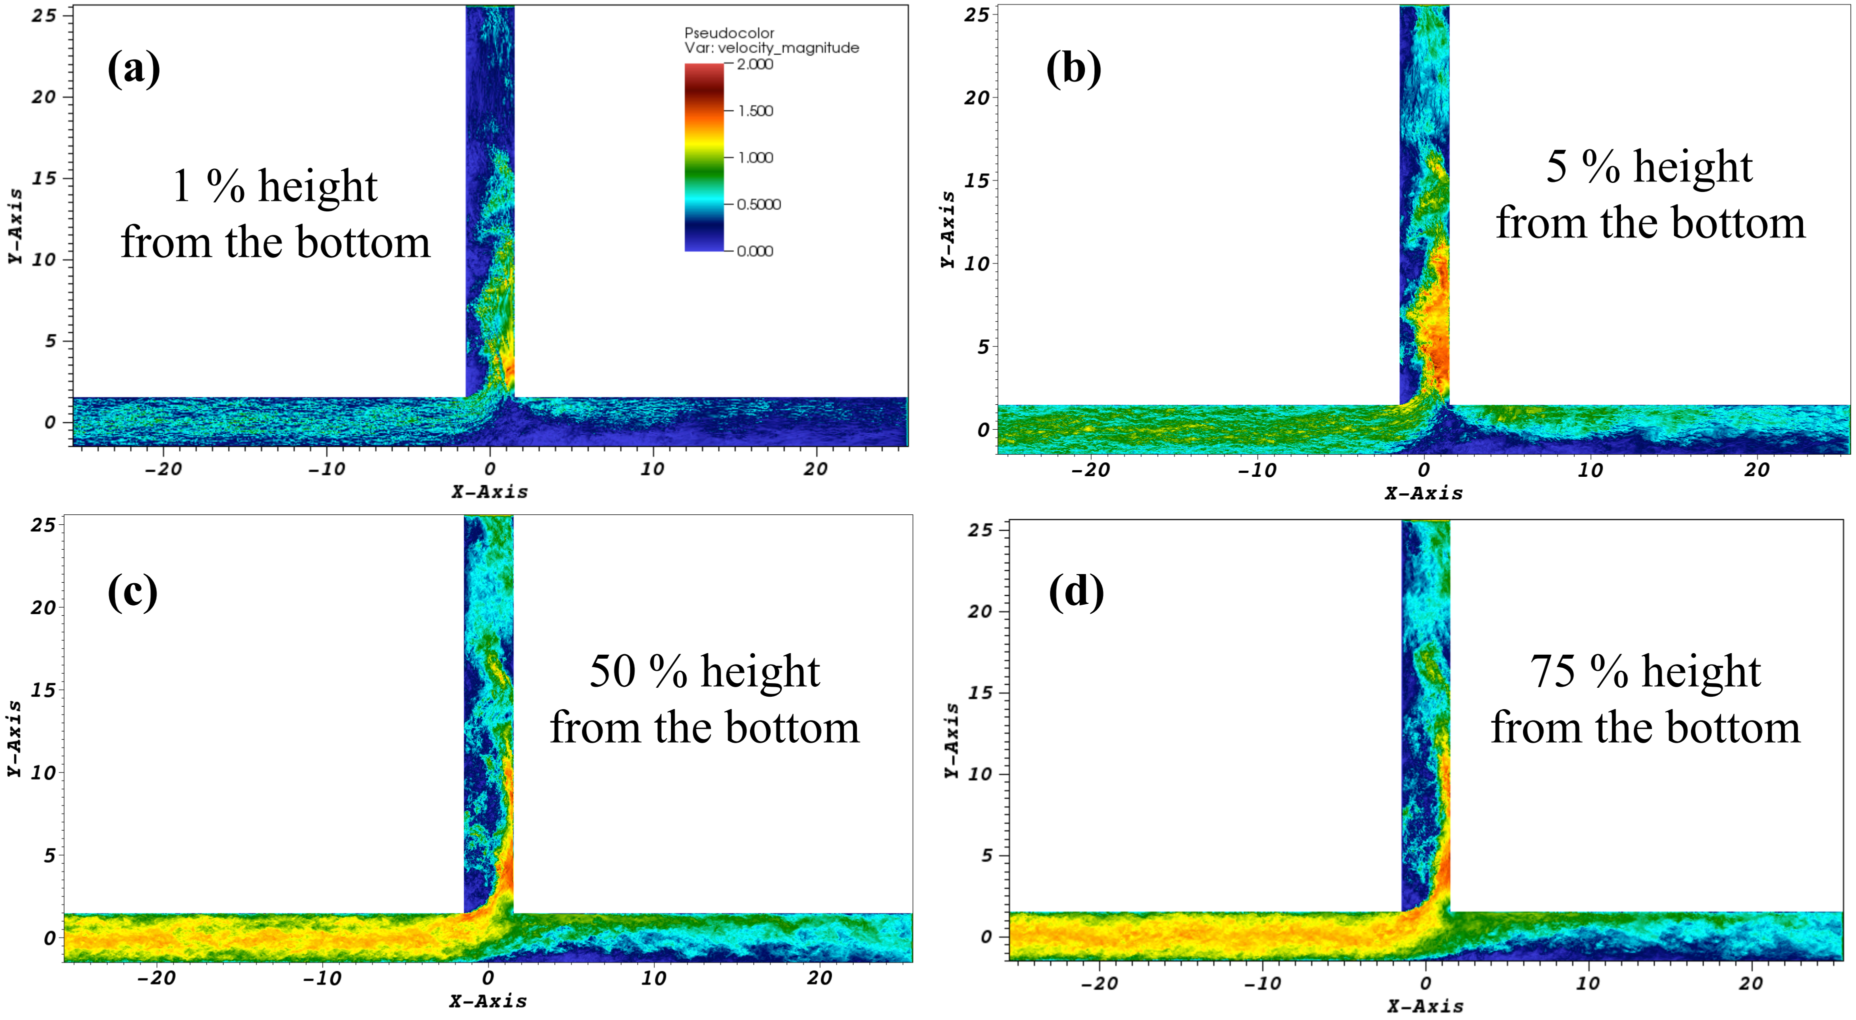

Instantaneous Velocity Magnitude at different levels

Results: 90-degree, Re=20000, 50:50 flow split

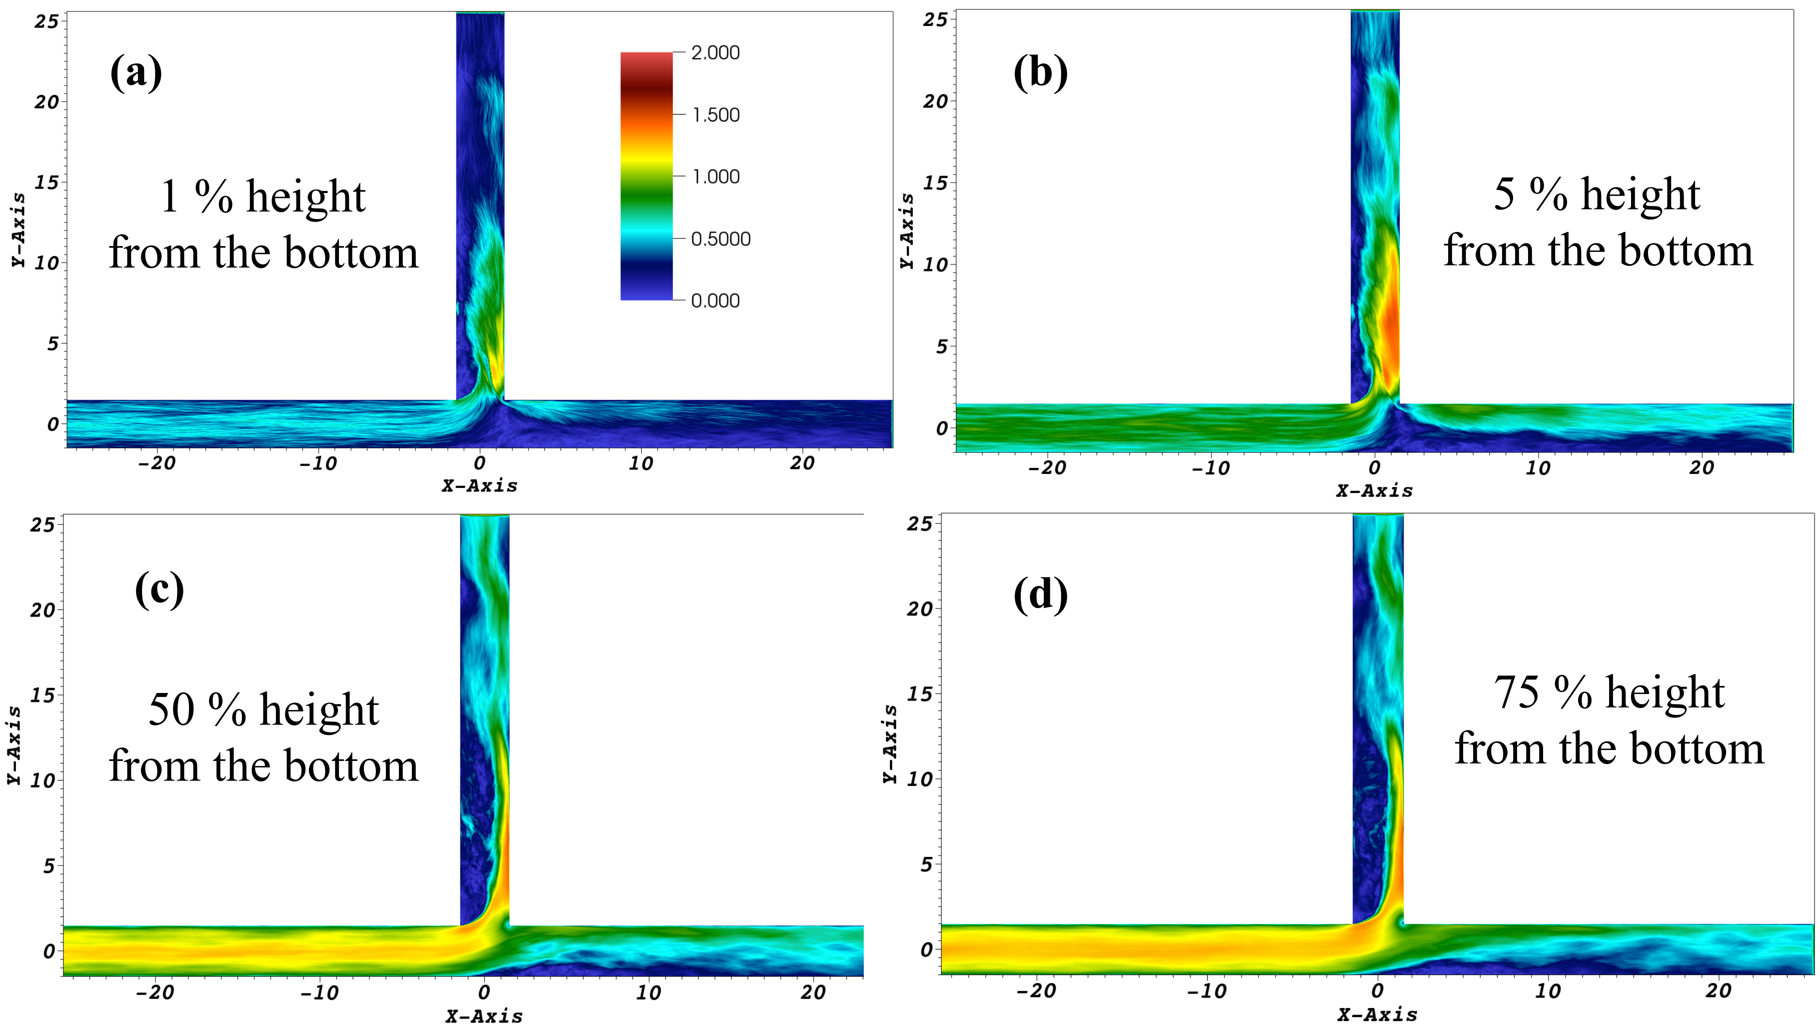

Time averaged Velocity Magnitude at different levels

Results: 90-degree, Re=20000, 50:50 flow split

- In the panels (a), (b), and (c), the general direction of the flow is out of the plane.
- For cross-section (d), it is into the plane

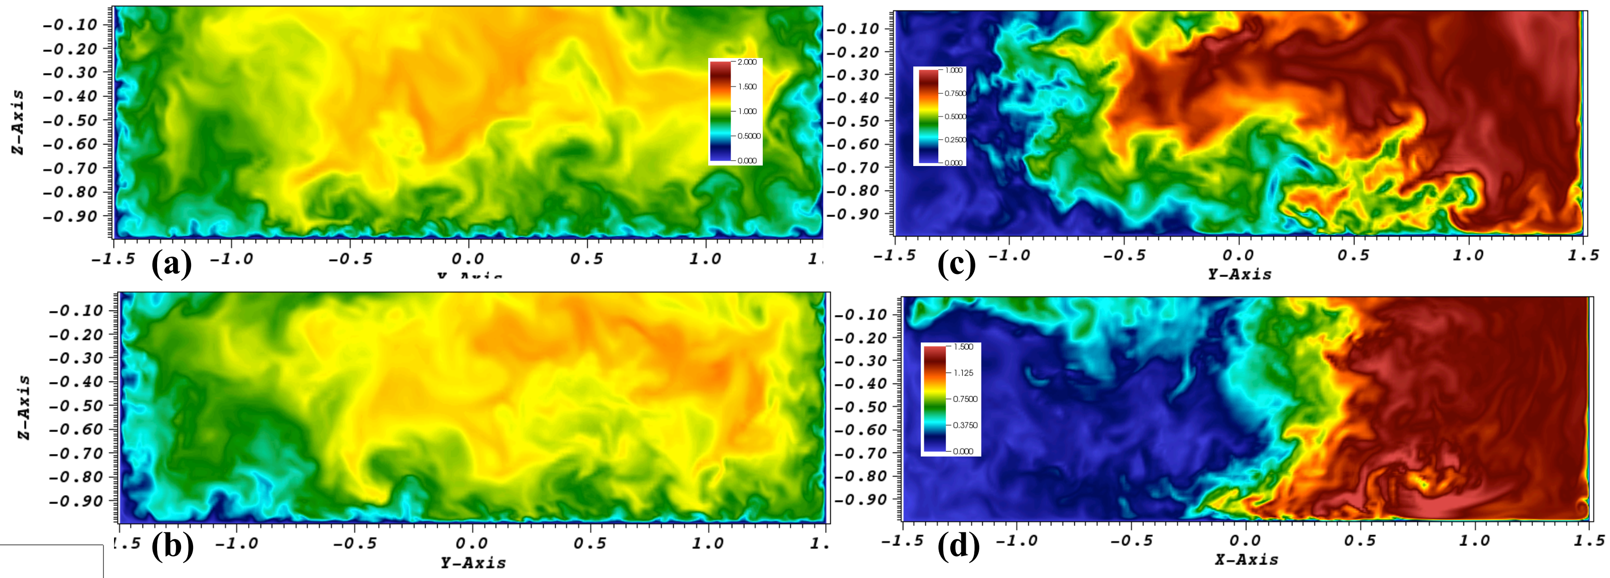

Instantaneous velocity magnitude at different cross-sections

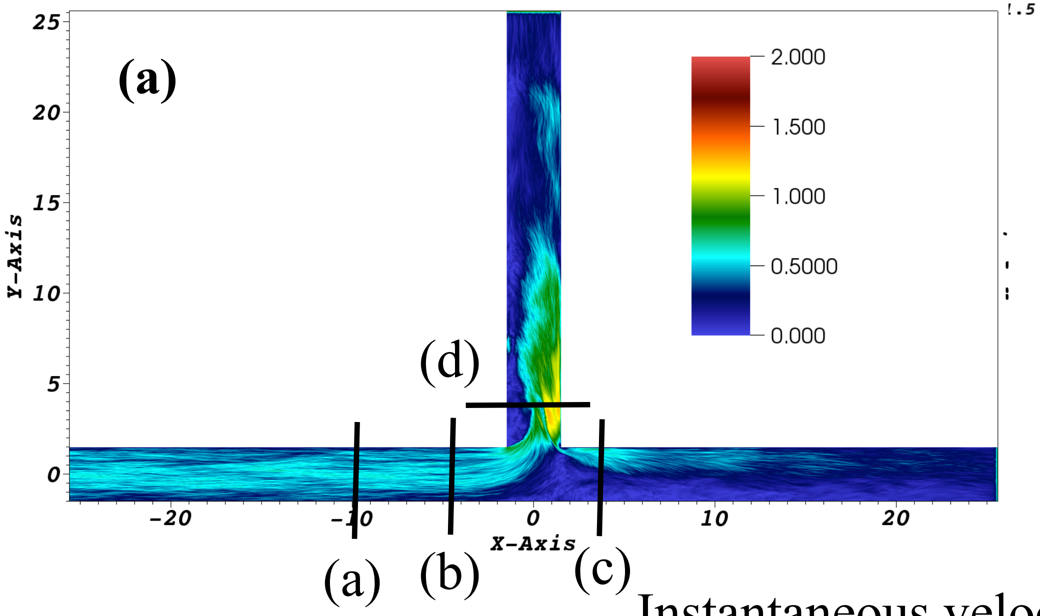

Instantaneous velocity in the z-direction >>

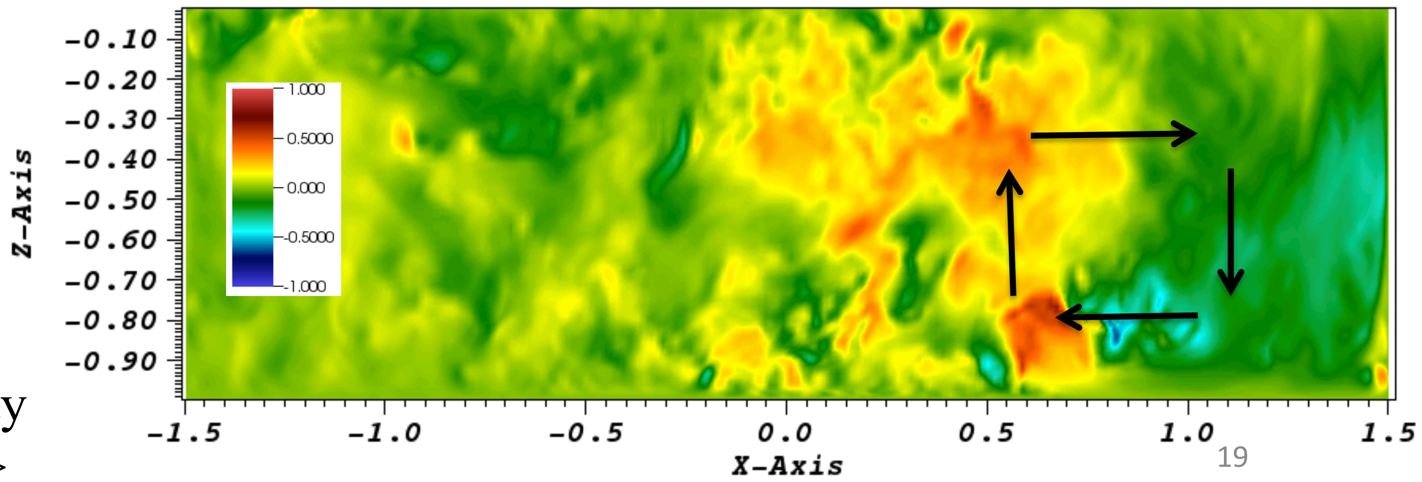

Results: 90-degree, Re=20000, 50:50 flow split

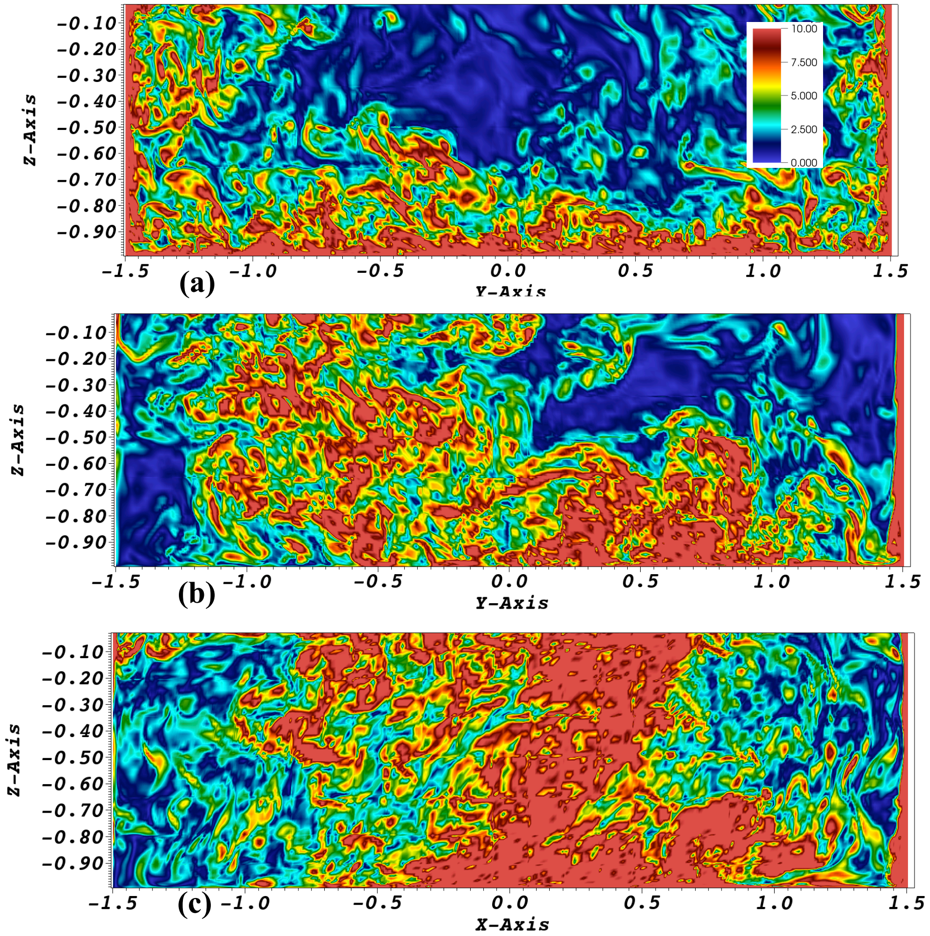

Vorticity at different cross-sections.
(a) main-channel before bifurcation. (b, c) are from the two channel after the bifurcation.

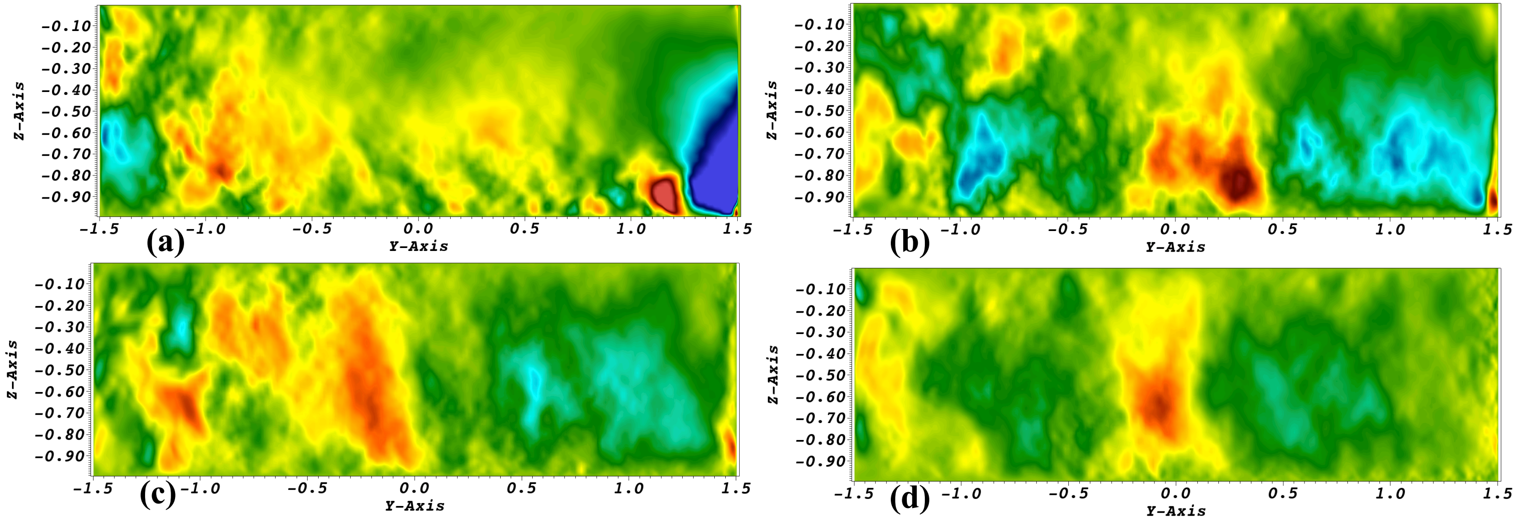

Time-averaged velocity in the z-direction, in the two channels after the bifurcation. Top panel is for the main-channel.

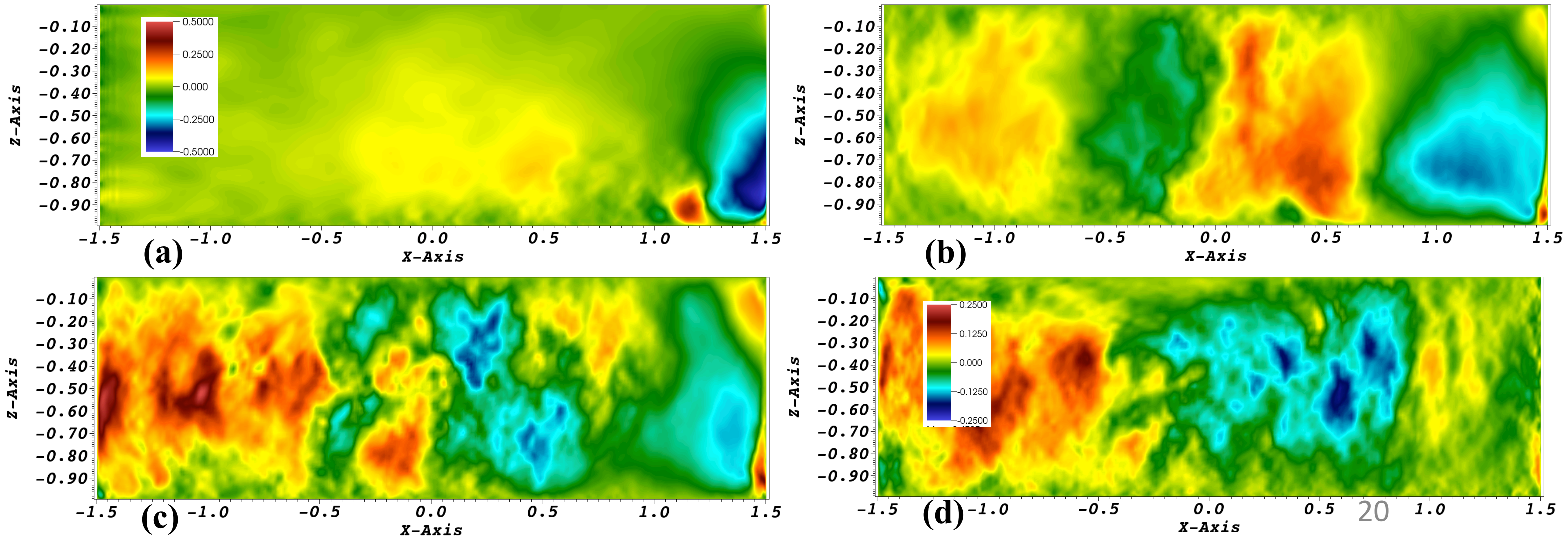

Results: 90-degree, Re=20000, 50:50 flow split

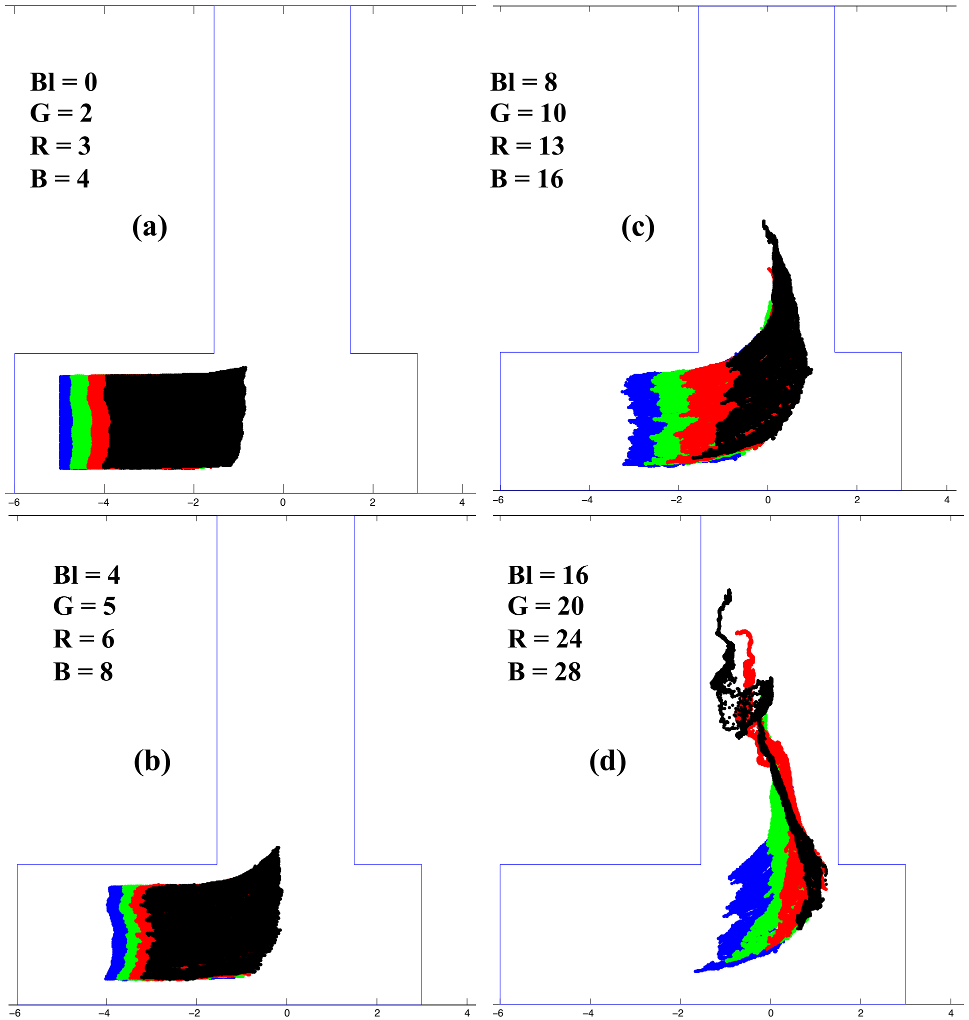

- 11250 sediment particles of size 0.015 (actual size 1.05 mm, density 2.65 kgm^{-3}) were released at the same time, upstream of the bifurcation.

- In agreement with the experiments, a very small quantity of the total sediment (4.29 %) entered the main-channel after the bifurcation.

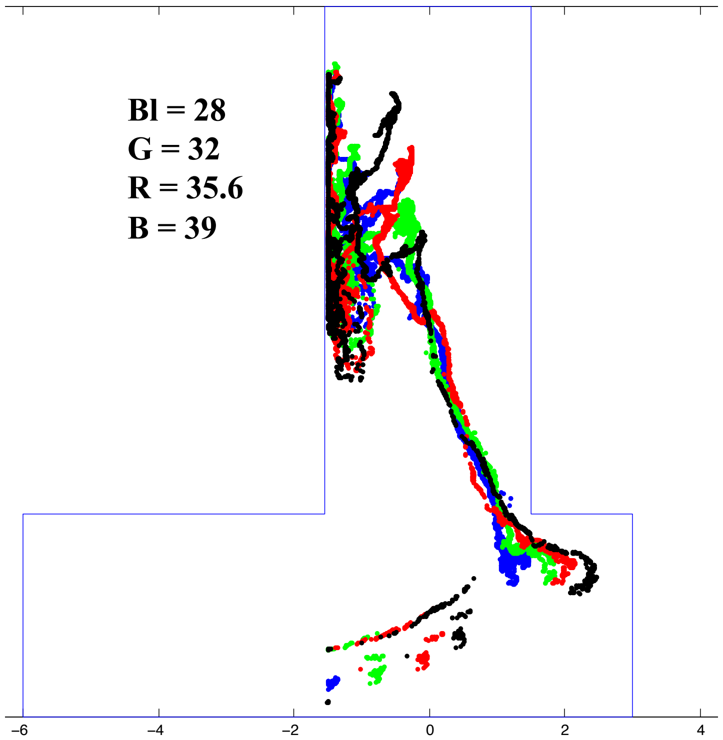

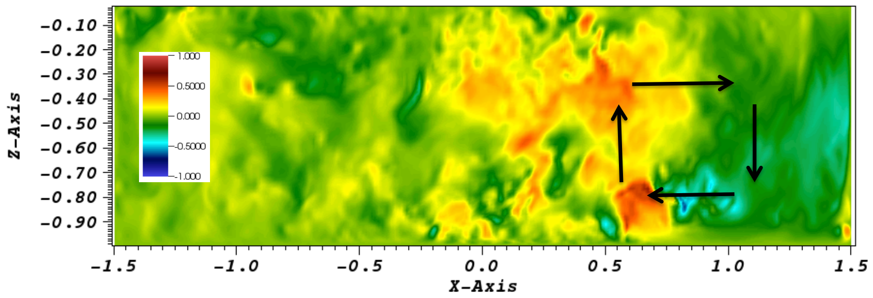

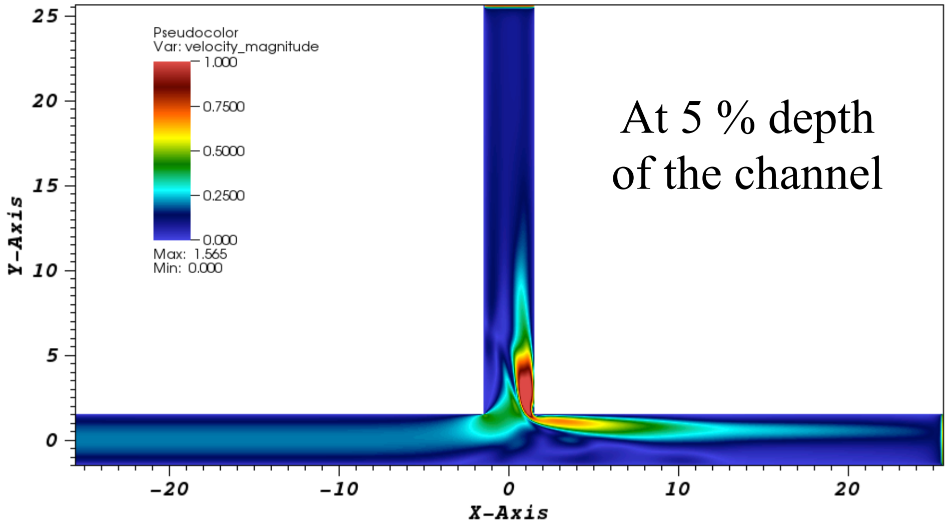

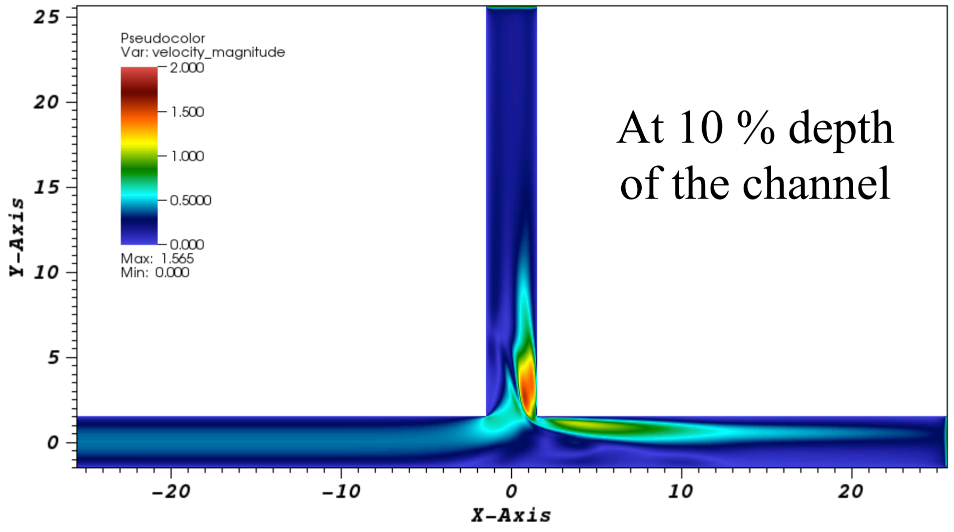

Simulation with Bulk Reynolds number 300 (and 50:50 split)

The flow is laminar and steady

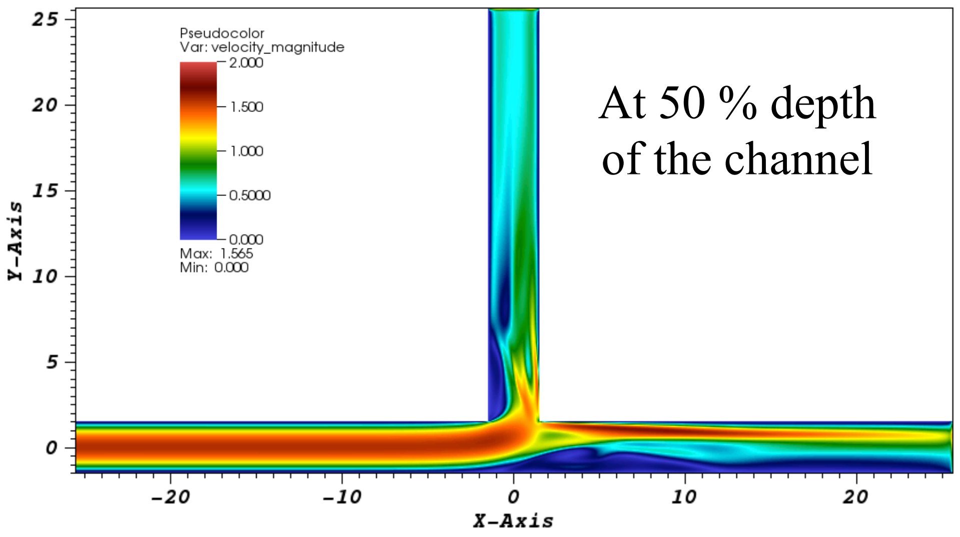

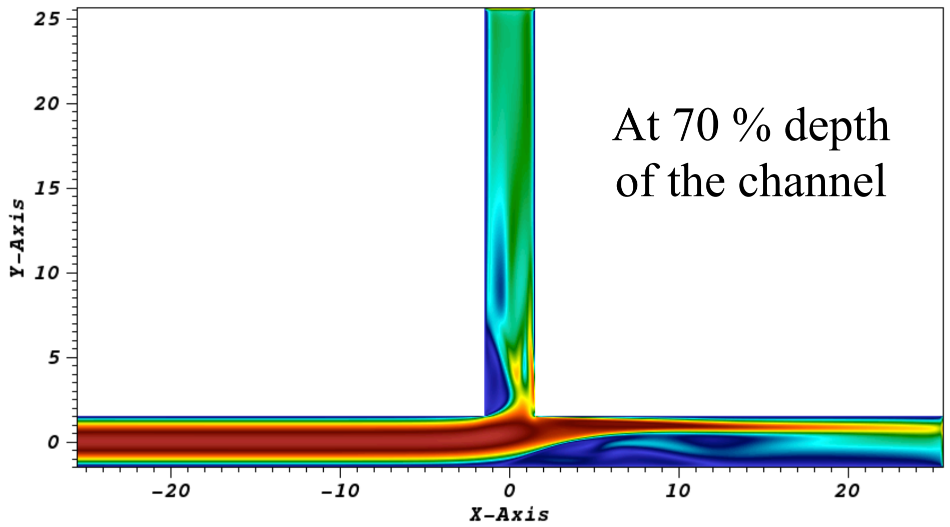

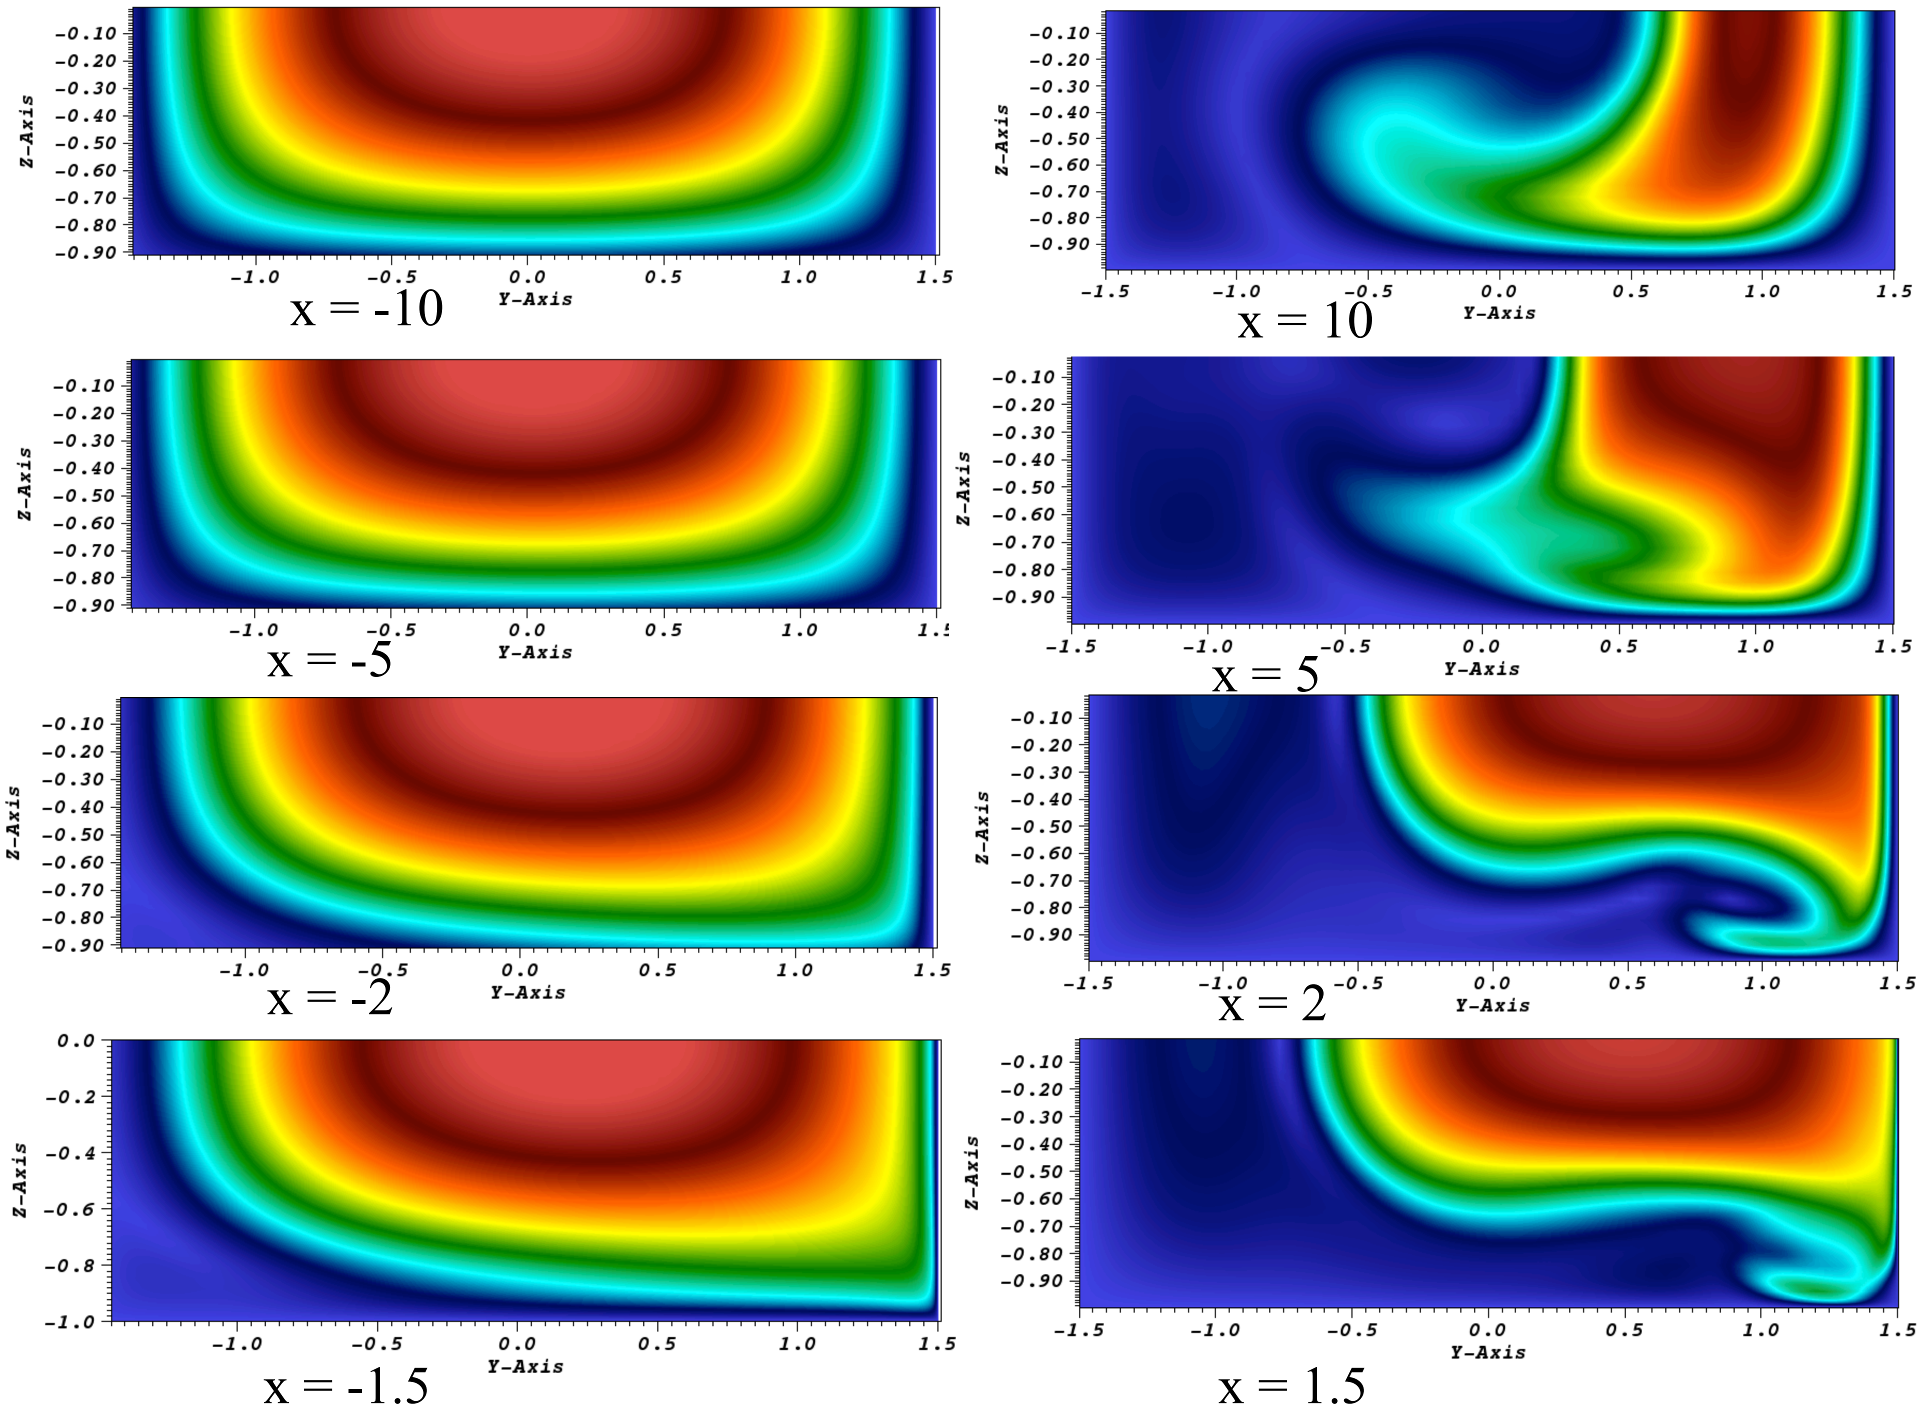
Z-Axis
Y-Axis
x = -10
x = 10
x = -5
x = 5
x = -2
x = 2
x = -1.5
x = 1.5

Wall normal velocity at different cross-sections of the two branches

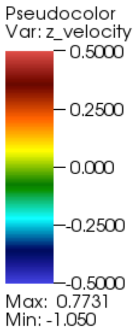

Main Branch

Diverted Branch

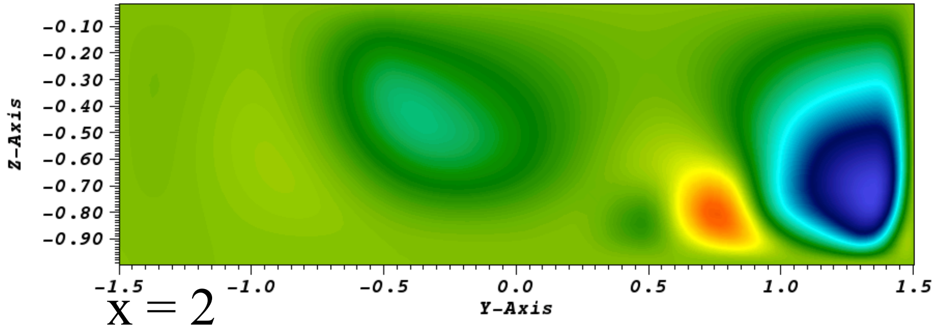

x = 2

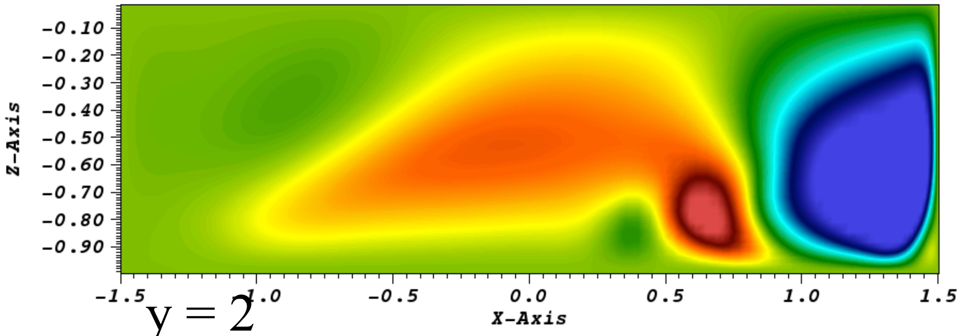

y = 2

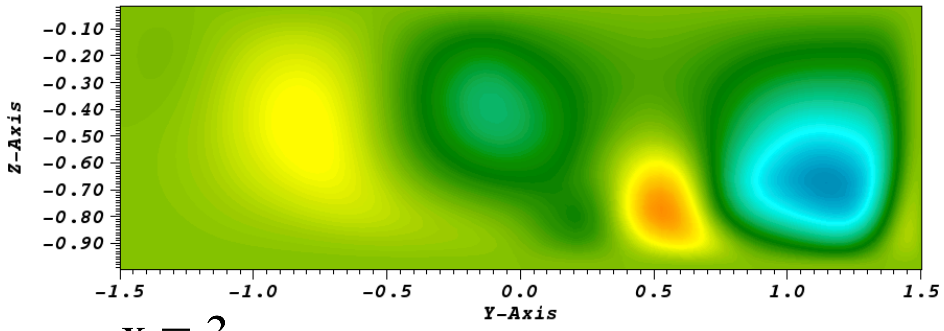

x = 3

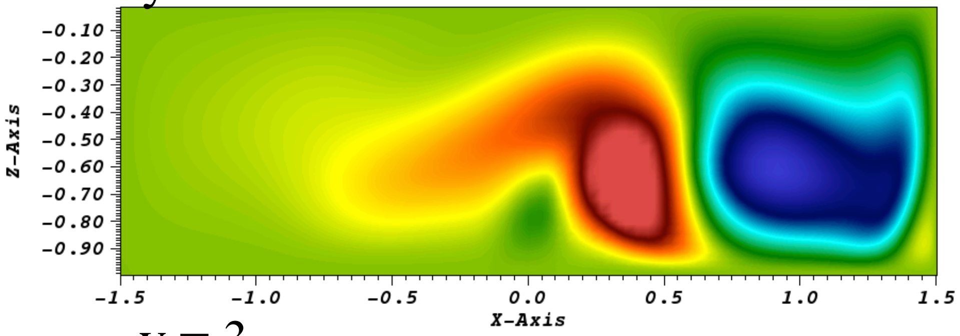

y = 3

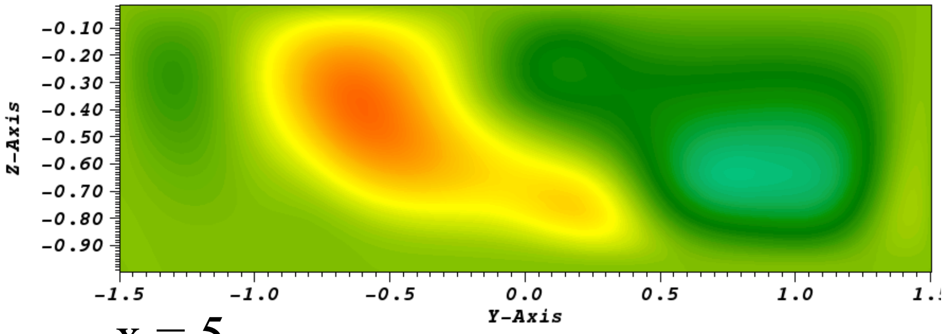

x = 5

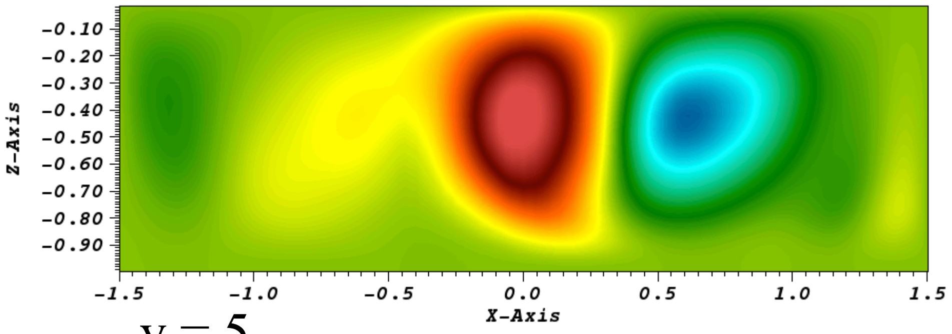

y = 5

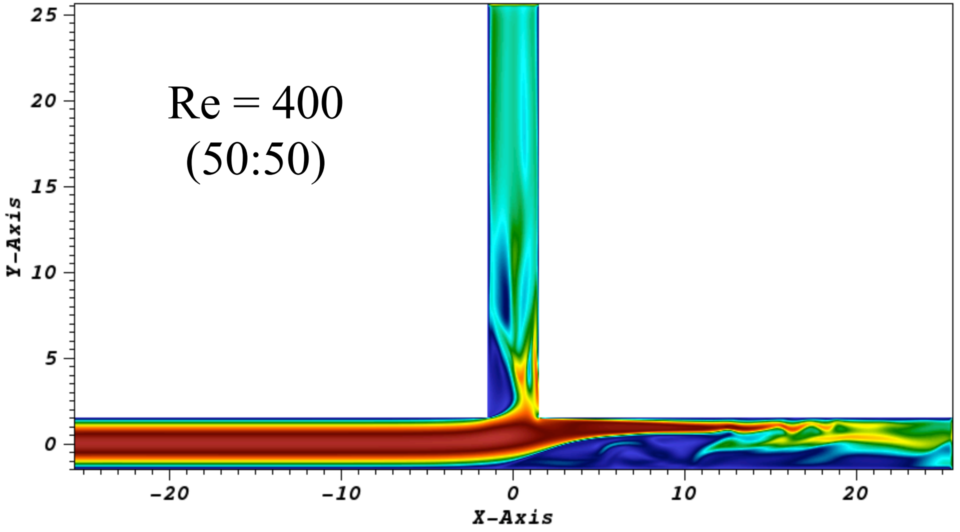

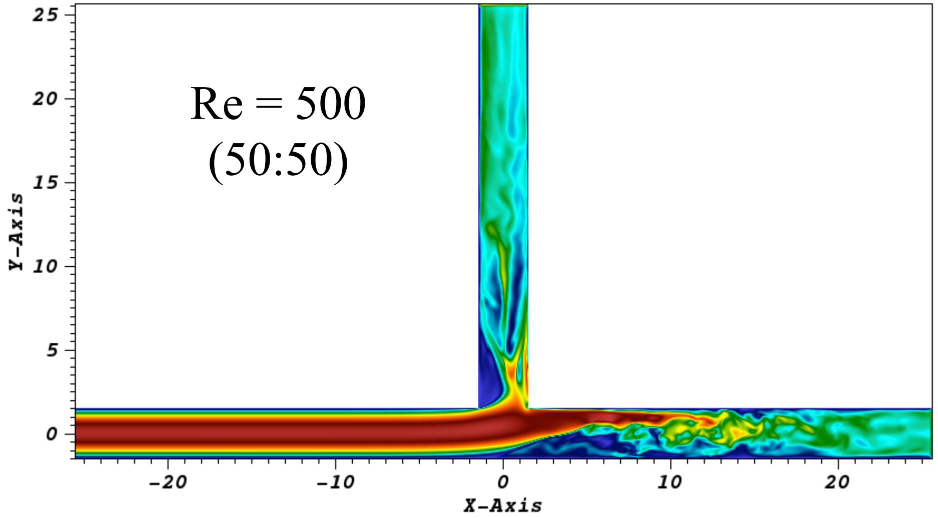

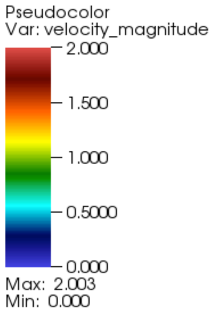

Regimes of the flow, changing with bulk Reynolds of the flow (Re)

The slice is taken at 70 % depth of the flow

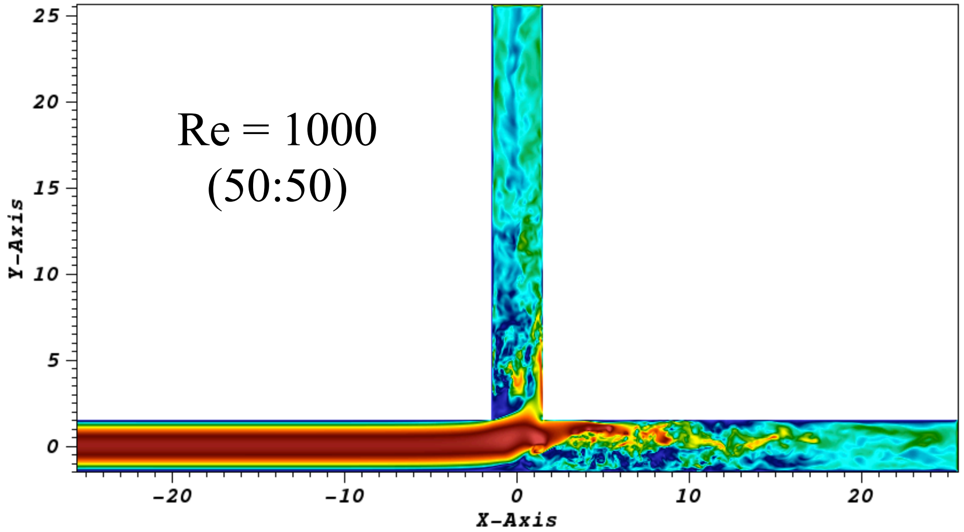

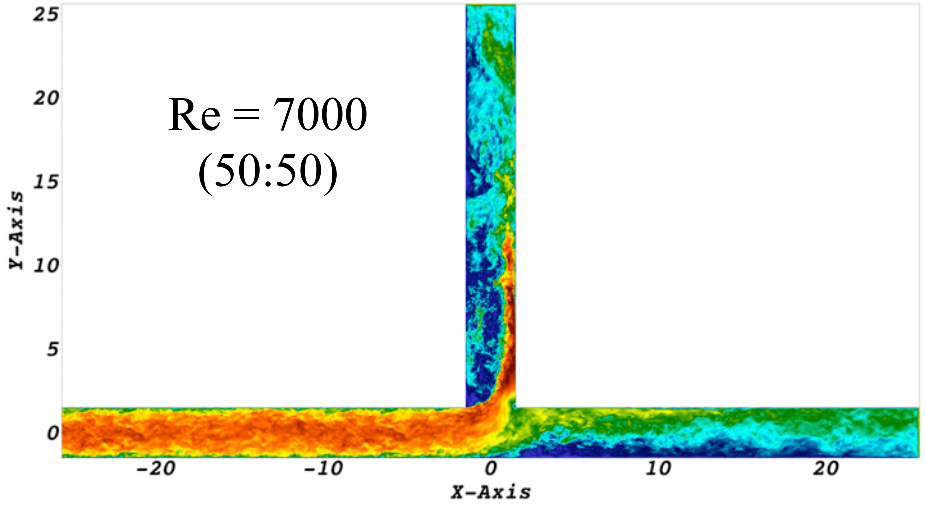

Conclusions and next-steps

- Flow and sediment transport was successfully modeled at an idealized 90-degree diversion for Re = 20000, and were able to capture the phenomena of Bulle-Effect.
- The driving mechanism for this highly non-linear phenomena has been identified, that is most of the flow near the bottom of the channel enters the lateral channel, taking along with it the near bed sediment.
- The flow patterns at other Re numbers were found to be similar, so it is expected that the phenomena will show up irrespective of Re.
- Currently simulations are being done to complete the sediment transport portion of the model for the range of cases mentioned before.
- Nek5000 was found to strongly scale up to 32768 mpi ranks on Blue Waters, though the efficiency reduces below 68.5 % after 8192 mpi ranks.
- Even though Nek5000's parallel scalability performance on Blue Waters is relatively good, it seems compared to MIRA the issue is the lack of "a hardware supported MPI all reduce". We need to find a way around ….

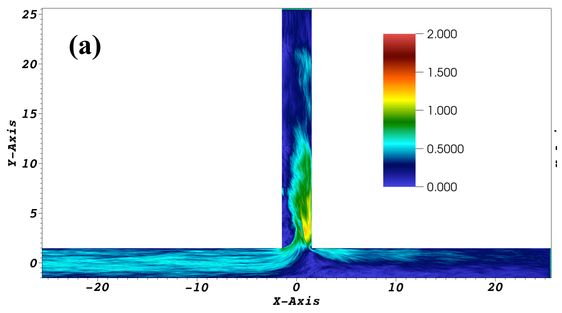

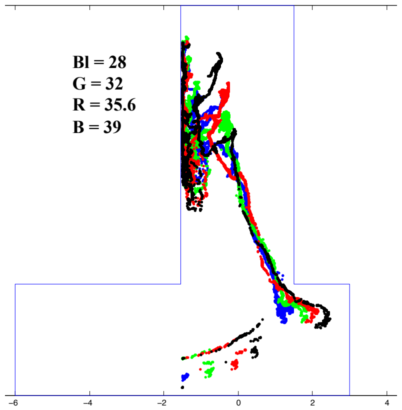

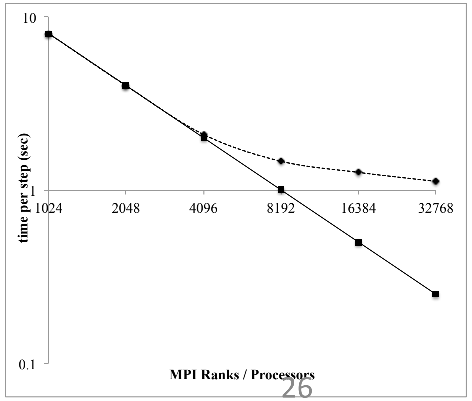

Acknowledgements and Collaboration with Blue Waters team

- Currently we use the PGI compiler, as we had been unable to compile with CRAY, so we are in communication with Dr. Tom Cortese to look in to the matter.

- We are also in communication with Dr. Rob Sineros to create an animation for one of the cases.