Parallelising Tracer: Quantitative Vortex Identification in the High Reynolds Number Regime

Researcher: Marius Koch (PhD Candidate)
Duration: 6 months (from June 2020 to November 2020)

January 24, 2020

1 Background and Motivation

With the emergence of unsteady peta-scale turbulent flow simulations comes the generation of a significant amount of high-fidelity data. Insight from this data could help address a range of open problems in the field of turbulent flow physics. There are numerous unanswered questions regarding the dynamics of turbulent flow features. E.g. quantitative statistics on the abundance of attached eddies that reach into the log-layer can benefit the development of improved wall models [2]. Currently the size of a data set that is fully resolved in both time and space is too large to interrogate using the classical approach of writing the data to disk and analysing it in a post-processing step. We have addressed this *big data* challenge by developing Tracer, a framework which applies a range of so called *in-situ* analysis technologies that can run concurrently with the simulation to extract insight while the data is still in the device memory.

In particular vortices are identified and classified over time in order to analyse their evolution and interaction. This is achieved by a topology based domain segmentation of the Q-criterion field of high-fidelity unsteady turbulent flow simulations. Isosurfaces of Q are considered to identify vortices. Obtaining the spatial extent of a vortex, i.e. the vertex cloud that makes up the vortex, requires the definition of a Q threshold, as is also the case for the visualisation of vortices. This threshold can vary drastically between vortices in highly vortical regions and those in low vortical regions. Identifying vortices with a single user-defined threshold is equivalent to current standard visualisation techniques, which render isosurfaces of a single Q criterion throughout the whole domain. Such an approach results in vortices being missed in low vortical regions and multiple vortices being merged into a single structure in high vortical regions. This issue is illustrated in Figure 1, which shows two isosurface sets of Q in a section of a Taylor-Green vortex and for comparison the topology based segmentation.

Topology based domain segmentation allows automatically defining an individual threshold for each vortex, which can be used for visualisation and, especially relevant in the context of this project, for identifying vortices and obtaining their individual point clouds in the whole range of the vortical regime. The knowledge about the individual point clouds allows one to gain further information about each vortex, e.g. its internal pressure distribution, or its shape.

The flow simulations are conducted by the highly parallel high-order computational fluid dynamics solver PyFR [6]. PyFR is well suited for producing the flow field, since its high-order capability enables efficient resolution of structures down to the smallest relevant scales and its highly parallel nature allows simulations of large problems. To be able to run concurrently with PyFR, Tracer is designed from the ground up to exploit massively parallel multi-GPU systems. To date Tracer is able to run alongside PyFR on a single GPU. The complexity of flow fields increases with its Reynolds number (Re). An increase in Re therefore requires a higher resolution of the domain, which is occupying more memory. A Nvidia V-100 GPU can currently fit a channel flow simulation by PyFR with Tracer up to $Re_\tau = 180$. However, investigating e.g. attached eddies that reach into the log-layer, requires channel flows in the regime of $Re_\tau \geq \mathcal{O}(10^3)$. To allow the analysis of such flows and of engineering problems, Tracer's capabilities need to be extended to being able to run on multiple GPUs.

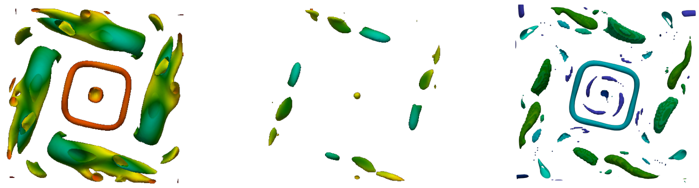

Q = 0.01   Q = 0.035   topology based segmentation

Figure 1: For Q = 0.01 a vortex ring is surrounded by four large structures, the vorticity of which is stronger than the ring's. If Q is chosen to be Q = 0.035 the four large structures break up into three smaller structures each and the ring disappears. Via topology based segmentation both the *weaker* ring and the surrounding structures are resolved, additionally other structures that have been missed with the Q isosurfaces appear.

2 Tracer

Tracer can identify vortices by accessing PyFR simulation data in-situ or by loading a vtu file from disk. We have shown that via Tracer significantly more turbulent structures can be identfied in a turbulent flow field compared to setting a global Q threshold. Additionally, while the identification of vortices is highly dependent on the user-defined Q value, the number of identified structures with topology-based segmentation is scarcely sensitive to user-defined parameters.

The vehicle for finding these individual Q thresholds is the join tree (JT). A JT is a graph representation of a scalar field that records appearances and joins of contours, which in the 3D case are isosurfaces, while the isovalue sweeps from plus infinity to minus infinity. Constructing the JT is part of the contour tree construction. Tracer employs the necessary steps of a state-of-the-art contour tree construction algorithm [1] to construct the JT. On the basis of the individual Q thresholds point clouds, that make up vortices, can be identified. Additionally the JT allows to efficiently compute size, pressure distribution, and various other properties of the vortices. Using a classification algorithm these vortices can be categorized allowing one to obtain statistics on a significant amount of samples, which can help answering open questions regarding the dynamics of vortices. Problems that can be addressed are, amongs others, the height of hairpin vortices as they detach from the wall in a flat plate turbulent boundary layer, and the ratio of the number of disc shaped vortices to the number of rod shaped vortices in a homogeneous isotopic turbulence. The developed framework can furthermore be applied to other scalar fields, isosurfaces of which are also considered to identify vortices, e.g. λ_2. This can pave the way to a comparison of their quality. [5].

3 Project Objectives

Currently Tracer is able to analyse flow fields that fit in the memory of a single GPU. However, the vast majority of simulations conducted with PyFR run on multiple GPUs. This project aims to parallelise Tracer to enable it to analyse all of PyFR's turbulent flow simulations. This shall be achieved by each process building its JT locally and preparing it for combination with JTs of neighbour ranks before they are merged on the host. The milestones for this are as follows:

- **Communicating MPI boundaries** from PyFR to Tracer.
- **Combining JTs** which are spread over multiple ranks.
- **Applications** of Tracer to turbulent flow fields spread over multiple GPUs.

4 Work Plan

I. Communicating MPI Boundaries: 1 Month

PyFR passes pointers to the simulation information to Tracer. This data needs to be extented with information about MPI boundaries and its neighbouring ranks. In order for two JTs of two regions sharing an interface to be stitched together they both need to be augmented with all maxima of their interface. Tracer will include all maxima in the context of boundary components in the locally built JTs.

II. Combining Join Trees: 2 Months

Since the JTs are augmented with interface maxima we can choose from a range of subtree combination algorithms, e.g. [3, 4], to construct trees that span multiple partitions. For doing so it is enough to communicate the *pure* JT, the size of which is roughly two orders magnitude smaller than the size of the domain or the size of the *fully augmented* JT. The combined JT can then be communicated back to the associated ranks on which it is augmented with the rank-local vertices. In that way a parallel analysis of the domain can be conducted.

III. Applications: 3 Months

There remain a lot of unanswered questions on the relation between turbulent structures and general properties in wall bounded turbulent flows [5]. The domain sizes for high Re flows are too large to fit on a single GPU. The implemented parllelisation would enable Tracer to produce statistics on significant amount of certain types of vortices. Specifically we would be able to get statistics of abundance and hight of attached vortices which can help improve wall models [2].

5 Impact

To the best of our knowledge there is no comparable tool to Tracer that conducts a topological analysis in-situ with a high-performance CFD solver on an unstructured mesh.

Parallelising Tracer would enable the analysis of flows at higher Re than currently possible. Specifically this would allow studies on turbulent channel flows to get more insight in vortices of the log-layer, which are hypothesised to be linked to skin friction. Another unanswered question is whether hairpin vortices persist in a fully turbulent regime. Such knowledge could help to better understand wall-bounded turbulence and will eventually lead to better turbulence wall models.

The application of Tracer to real world industrial flows can furthermore directly benefit aerodynamic design processes. Our group is cooperating with Arup to investigate whether extreme low-pressure values on facades are caused by vortices.

Moreover Tracer reduces the burden of post-processing a vast amount of data. Diminishing this *big data* problem will help enable use of peta/exascale high-fidelity CFD in the context of industrial design frameworks.

6 Alignment with PRISM Strategy

The work plan is subdivided into three stages, and each section will be summarized in an individual report as well as conference presentations and journal articles.
Retention of key staff: Marius Koch is a member of Dr. Peter Vincent's group (PyFR developer) and has developed Tracer.
Collaboration within PRISM: This project is highly interdisciplinary spanning across topics in mathematics, computer sciences, and fluid dynamics, as such it will require close collaboration with and between Dr. Peter Vincent in the Department of Aeronautics and Prof. Paul Kelly in the Department of Computing.
Long-term research: This project aims to establish Tracer as a quantitative vortex identification and classification tool for massively parallel simulations of turbulent flow fields. Henceforth results produced by Tracer can amongst others benefit the construction of new wall models for LES.

References

[1] H. Carr et al., *2016 IEEE 6th Symp. Large Data Analysis and Visualization (LDAV)*, *pp.* 75-84, 2016.

[2] I. Marusic and J. Monty, *Ann. Rev. Fluid Mech.*, *vol.* 51, *pp.* 49-74, 2019.

[3] D. Morozov and G. Weber, *ACM Sigplan Notices*, *vol.* 48, *No.* 8, *pp.* 93-102, 2013.

[4] V. Pascucci and K. Cole-McLaughlin, *Algorithmica*, *vol.* 38, *pp.* 249-268, 2004.

[5] J. M. Wallace, *J. Turbulence*, *vol.* 13, 2012.

[6] F. D. Witherden et al., *Computer Physics Communications*, *vol.* 185, *No.* 11, 2014.