Developing a scalable and flexible code for high-resolution DNS of two-phase flows

Iain Bethune(1), Toni Collis(1), Mike Jackson(1), Lennon Ó Náraigh(2), Prashant Valluri(3), David Scott(1)

(1) EPCC, University of Edinburgh, Edinburgh, UK
(2) School of Mathematical Sciences and CASL, UCD, Dublin, Ireland
(3) Institute for Materials and Processes, University of Edinburgh, Edinburgh, UK

TPLS

TPLS (**Two-Phase Level Set**) is an open-source program for simulation of two-phase flows in 3D channel geometries using high resolution DNS. TPLS solves the incompressible Navier—Stokes equations for a two-phase flow. A regular grid finite-volume discretization is employed based on an idealized channel geometry with a range of different inlet conditions that can be prescribed by the user. The interface between phases is tracked with either a Level-Set method or the Diffuse Interface Method. The code evolves the physical variables (pressure, fluid velocities, and interface configuration) through discrete time steps. At each timestep, the key computational tasks performed amount to the solution of large systems of sparse linear equations with tens of millions of unknowns, for the key physical variables. In addition, regular I/O is required to save the system state for later analysis and visualization, or to restart in the case of hardware failure.

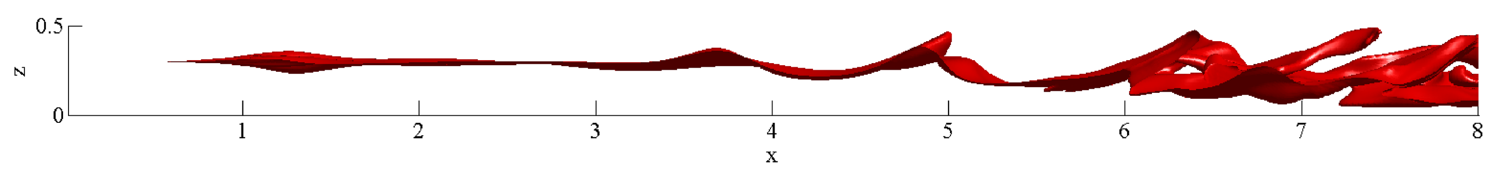

The code is implemented in Fortran90, initially with MPI parallelization using a 2D domain decomposition and bespoke Jacobi/SOR iterative solvers. Over the last two years, we have improved the TPLS code in several respects to give better performance, scalability and usability, moving from an in-house code specialized for use by the original developers, to an open-source flexible program which can easily be used by others, including academic and industrial users. The culmination of this work is TPLS version 2.0, presented herein.

http://sourceforge.net/projects/tpls/

Using linear solvers from PETSc

The most expensive part of the calculations in TPLS is the pressure solve, which requires the solution of Poisson's equation at each timestep. The original version of TPLS used successive over-relaxation (SOR). This has been replaced by calls to the PETSc library which allows experimentation with various preconditioners and Krylov subspace solvers. In addition, some optimization of MPI communications has been carried out. An improvement in performance of up to 54% has been achieved with respect to the SOR version, and the pressure calculation now exhibits excellent strong scaling to over 2000 cores.

Cores	Total Time	I/O	Total Computation	**Pressure**	Other
1024	18680	12497	6183	**3432**	2751
2048	16270	12560	3710	**1784**	1926
Efficiency	0.6	0.5	0.8	**1**	0.7

Strong scaling on a 512x144x152 computational grid using a benchmark simulation [1]

With the improvement in the pressure calculation comes an obvious realisation that a major bottleneck is the I/O, which is done in the original version by periodically gathering all data to a master node for writing to a single file. This is the motivation for the development of *parallel output*.

1. Performance enhancement and optimization of the TPLS and DIM two-phase solvers, David Scott *et al.* Technical Report, EPCC, 2013.

Parallel Output

The bottleneck caused by the serial output has been removed by parallelizing the I/O and writing the output in the highly compressed NetCDF format. With this change the code is very efficient: close-to-linear speedup is obtained for process counts up to p=1000.

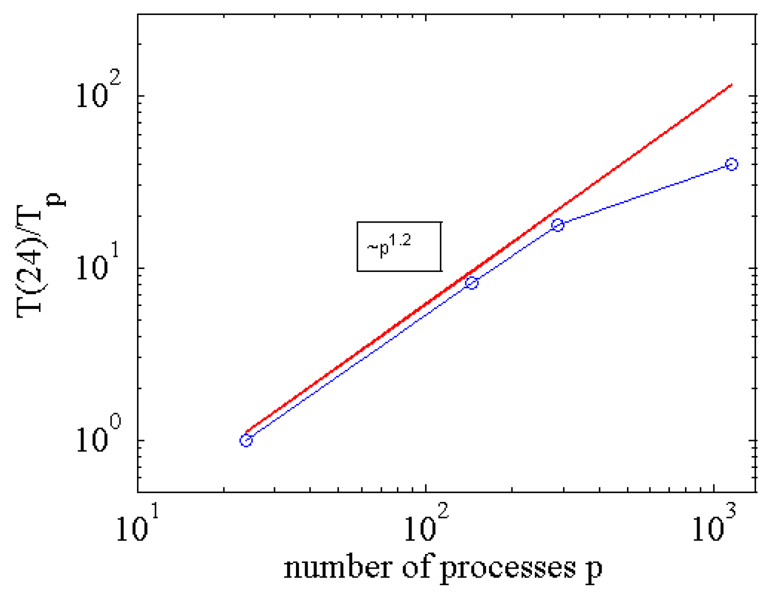

Strong scaling study for a 257x145x153 grid run for 5,000 timesteps using the parallel NetCDF I/O

Storage efficiency: The previous version of TPLS used ASCII formatted text files for the output data. For a simulation on a typical grid (such as the one above), this produces 1.5GB of data per output operation. In contrast, using the NetCDF format, the same simulation generates 225 MB per output operation - a compression factor of 6.67.

Some physical applications

TPLS has already been applied to a range of physical problems, including:

- Simulating the genesis and evolution of three-dimensional waves in two-phase laminar flows
- Understanding the flow-pattern map in a falling-film reactor

In these investigations, TPLS has revealed how three-dimensional waves form in systems where two-dimensional disturbances dominate, thereby contributing the resolution of a paradox in the two-phase flow literature [2].

TPLS is used by a number of European research groups, including:

- Dr Lennon Ó Náraigh (UCD, Dublin)
- Dr Prashant Valluri (IMP, Edinburgh)
- Professor Peter Spelt (Ecole Centrale, Lyon)
- Professor Luca Brandt (KTH, Stockholm)

Acknowledgements

This work was funded under the embedded CSE programme of the ARCHER UK National Supercomputing Service (http://www.archer.ac.uk).

2. Ligament Instability, nonlinear instability and ligament dynamics in three-dimensional laminar two-layer liquid-liquid flows, Ó Náraigh *et al.* Journal of Fluid Mechanics (2014), vol. 750, pp. 464-506