Scalable coupling of Molecular Dynamics (MD) and Direct Numerical Simulation (DNS) of multi-scale flows

Lucian Anton[1], Edward Smith[2]
[1]*Numerical Algorithms Group Ltd, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR, UK*
[2]*Department of Mechanical Engineering, Imperial College London, Exhibition Road, South Kensington, London SW7 2AZ, United Kingdom*

May 31, 2012

Abstract

The goal of this dCSE project was to develop the software infrastructure to couple high preformance fluid and molecular dynamics solvers. This captures important details of localised fluid flow at the molecular level whilst preserving the efficiency of the continuum computation in the rest of the domain. The result is a verified and scalable library which links a continuum fluid dynamics (CFD) code to a molecular dynamics code.

Contents

1 Introduction

1.1 General aspects of coupling CFD and MD

The ability to influence macroscopic flow behaviour by manipulating interfaces holds promise for significant impact in future engineering technologies. For example, surface features of micron-size resolution and hydrodynamic coatings are sought for turbulent drag reduction, in much the same way as secretions of tono-filaments and lipid droplets reduce drag on a dolphin skin; morphing surface textures are envisaged for "fly- by-feel" technologies that mimic the ability of insects to optimise lift; and use of electric fields is explored for manipulating the interfacial behaviour of droplets in electro-sprays.

The key to achieving these challenging goals is to recognise that the dependency between the micro- and the macro-scale is bi-lateral: while the microscopic interface region can significantly alter the macroscopic flow, the outer state co-equally affects the microscale dynamics, thus establishing a feedback loop. A faithful representation of the dynamics must therefore span 7-8 orders of magnitude in spatial extent and 9-10 orders of magnitude in time. As a result, this level of multi-scale modelling has remained beyond the most advanced engineering simulation capabilities.

In this dCSE project we have developed the software infrastructure for coupled continuum-molecular systems. The Molecular Dynamics (MD) simulations track the motion of individual discrete atoms based on their mutual interactions inside a confined domain. On the other hand, the continuum software uses Direct Numerical Simulations (DNS) to solve the non-linear, partial differential equations for the fluid velocity at all points within a large computational domain.

Coupling between discrete and continuous simulations is a very active research field, dating back to a seminal paper by O'Connell and Thompson [1]. There have been a number of key papers over the last eighteen years, including, [2, 3, 4, 5, 6, 7, 8, 9]. As the number of researchers utilising coupled techniques continues to grow, so does the need for an efficient, scalable and robust ways to implement the coupled methodology.

To this end, the coupling of an MD algorithm ($\mathcal{S}treamMD$) and a DNS solver ($\mathcal{T}ransFlow$) within a single, coupled framework has been achieved in the first phase of the dCSE proposal. This will form the basis of a multi- purpose computational coupling utility which will be made available to the HECToR community.

1.2 Project goal and summary of the work done

The goal of this project was to build a software infrastructure to support simultaneous and coupled runs of large scale, parallel codes – in this case, a continuum fluid dynamics code and a discrete molecular dynamics code. The coupler is packaged as a library which is linked to both codes and manages all communication between them.

The coupling infrastructure was implemented by the following software components:

- An internal set of routines for data packing and communication. In order to ensure independence of the coupler from the specific data layout of the DNS and MD solvers, the coupler constructs internal data objects and communication maps at initialisation stage. In this way, coupling of any new DNS or MD solvers requires minimal changes.
- An interface module accessible by the DNS and MD applications. It includes all the necessary routines for synchronisation and data exchanges between the applications.

2 Coupler algorithm and implementation details

2.1 Coupler infrastructure

The coupler has the following main functions: (a) initialisation of all required communicators, (b) definition of the logical relations and communication maps between the MPI topologies of the DNS and MD applications and (c) data transfer operations. Below we provide a brief description of the algorithms implemented in each sector.

(a) The continuum solver and the molecular dynamics code run concurrently using the multiple program multiple data paradigm. The initialisation of the global communicator (`MPI_COMM_WORLD`)

is followed by a split into continuum (CFD_COMM) and molecular (MD_COMM) communicators corresponding to the DNS and MD applications. In addition, an inter-communicator is also defined, which includes only the processes from the continuum and molecular dynamics which are physically connected by the coupler.

(b) Using the problem parameters (e.g. DNS grid specifications, physical domain extents) the algorithm identifies the required coupling between the DNS and MD simulations and stores this as a mapping between adjacent MPI ranks in both codes. This ensures localised MPI communication – preserving the individual scaling of the coupled codes.

(c) Using the logical layout of processes and the internal maps, the coupler implements all data transfers between the applications during the computation. As the spatial relationship between MPI ranks is established during the initialisation, coupled data transfer subroutines require only simple MPI type send-receive calls. This guarantees the generality of the developed modules and easy integration of new continuum or MD solvers, or new coupling schemes.

The names of the main interface subroutines of the coupler with a brief description of their functionality are listed below. The coupler flowchart is presented in Figure 1 .

```
! coupler_create_comm (cfd+md) splits MPI_COMM_WORLD, creates
!                               intercommunicator between CFD and MD
! coupler_create_map  (cfd+md) creates correspondence maps between
!                               the CFD grid and MD domains
! coupler_cfd_init    (cfd)    initialises coupler and CFD parameters
!                               using data from MD or COUPLER.in
! coupler_md_init     (cfd)    initialises coupler and MD parameters
!                               using data from CFD or COUPLER.in
! coupler_send_data   (cfd+md) sends grid data exchanged between
!                               realms ( generic interface)
! coupler_recv_data   (cfd+md) receives data exchanged between
!                               realms ( generic interface)
! coupler_cfd_get     (cfd)    returns coupler internal parameters
!                               for CFD realm
! coupler_md_get      (md)     returns coupler internal parameters
!                               for MD realm
```

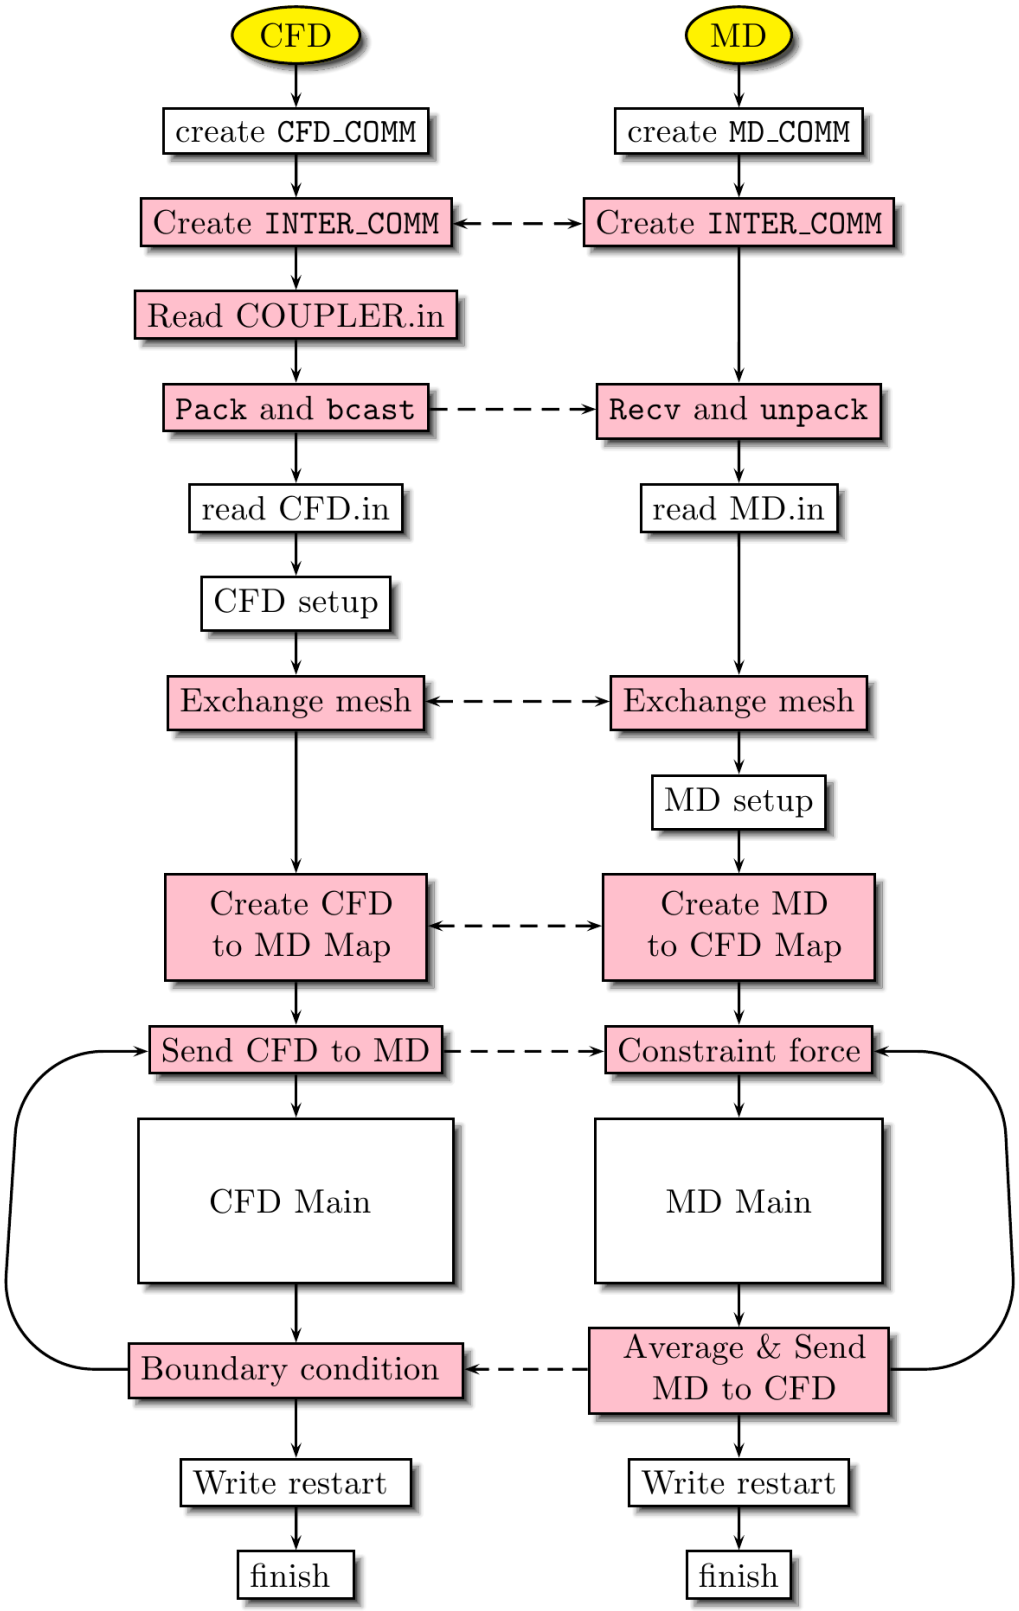

Figure 1: Flowchart representation of the DNS–MD coupler, the coloured boxes represent coupler sectors.

2.2 Benchmarking and scalability of the coupled algorithm

Before benchmarking the coupled algorithms, the scaling of the individual codes is presented in figure 2. An efficient coupler should ensure the scaling of the individual codes is maintained Performance testing for $\mathcal{T}ransFlow$ was carried out using a curvilinear mesh, for turbulent flow around a compressor blade. The flow was simulated using $370,000$ grid points per MPI-task, on an SGI Altix 4700, Intel Itanium 2 (HLRB, Super-Computing Centre, Germany). Parallel efficiency is defined with respect to the performance on 16 cores. As the number of cores is increased, the number of nodes per core is held fixed. At 768 cores, the algorithm performed at 94.4% parallel efficiency, shown in figure 2a.

As part of the development under work package one, $\mathcal{S}treamMD$ has been benchmarked using 1024 cores on HECToR. The profile shown in figure 2b is for 5000 iterations of a Lennard Jones system with $3,317,760$ molecules. The number of processors is increased for this system size and the parallel efficiency is defined as the simulation time relative to the ideal scaling of a single core. The results compare favourably with those obtained from profiling of LAMMPS. The higher efficiency of $\mathcal{S}treamMD$ is in part due to the numerical method, but it should also be noted that LAMMPS was tested with a smaller system size ($32,000$ molecules) and hence efficiency is quickly reduced at larger numbers of cores.

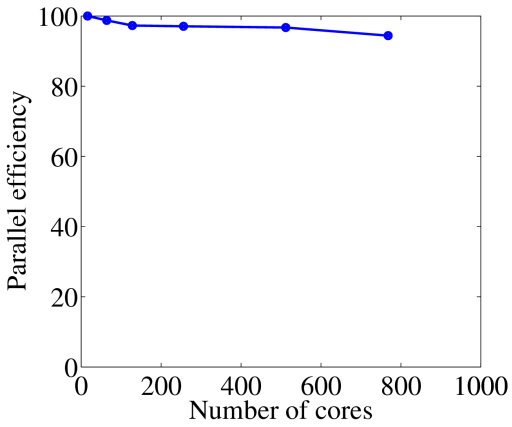

(a) Parallel performance of $\mathcal{T}ransFlow$.

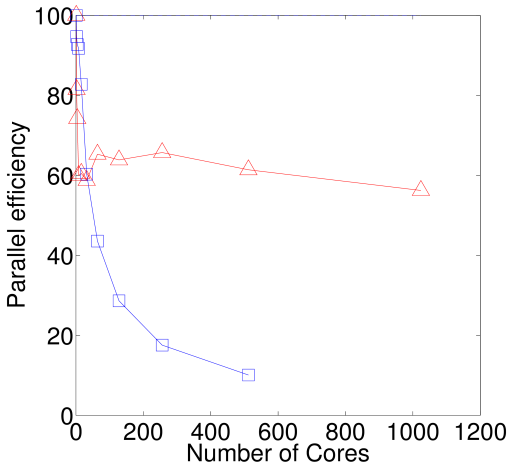

(b) Parallel performance of $\mathcal{S}treamMD$ (△) vs. LAMMPS on a Cray X5 (□)

Figure 2: Parallel performance of the two codes profiled individually.

The coupling methodology follows the work of [3] and the accuracy of the coupled algorithm has been verified by recreating results from that paper. The simulated problem is the canonical sheared Couette flow, and therefore it is possible to compare the numerical result to an analytical solution. Figure 3a shows the recreation of the results presented in [3] along with the analytical results. The coupled continuum–MD code accurately reproduces the analytical solution.

Scalability of the coupled algorithm was also evaluated. For the case of laminar Couette flow, the computational requirements of the continuum solver are almost negligible. The speedup of the code therefore depends almost entirely on the scaling of $\mathcal{S}treamMD$ and the coupler. If the coupler is performing efficiently, this combined speedup can be expected to be similar to the scaling of $\mathcal{S}treamMD$ alone. The scaling of the coupler is compared to $\mathcal{S}treamMD$ in figure 3b, up to 1024 processes. Figure 3b demonstrates that the coupler performance only slightly deteriorates speedup.

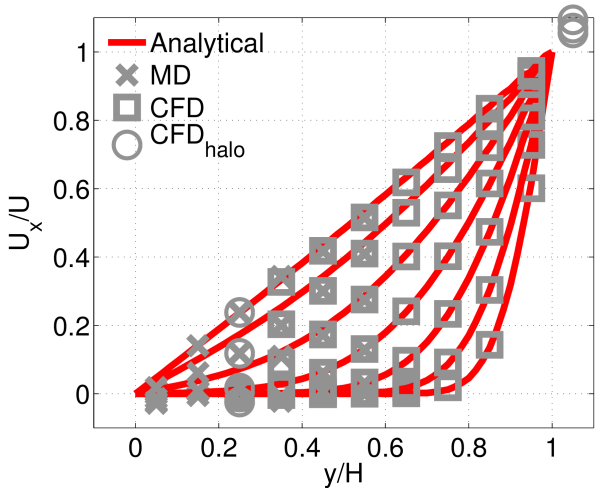

(a) Verification of the coupled codes against analytical solution for Couette flow.

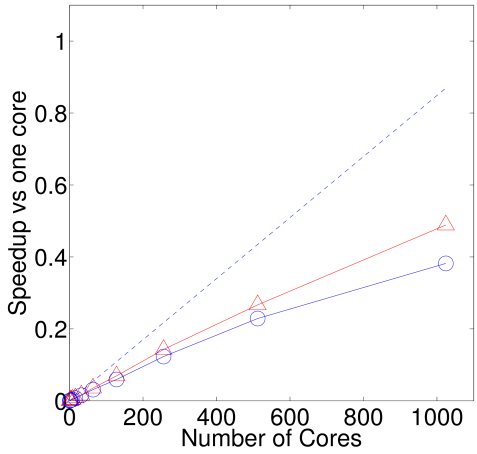

(b) Parallel speedup of $\mathcal{S}treamMD$ only ($\triangle$), coupled code ($\circ$) against the ideal speedup ($--$).

Figure 3: Accuracy and scalability of the coupled application.

3 Impact of the Work

The coupler's main purpose is to reduce computational requirements while still capturing essential molecular effects. For many cases of research interest, this will reduce the HECToR AU requirement by many orders of magnitude. In other cases of engineering interest, the use of coupling enables simulations which are impossible using either MD or DNS code alone. This has the potential to open the field of microscale simulation to a wide range of researchers and engineers.

As an example of the benefit of a coupled run, consider an entirely atomistic simulation of a micron sized channel. The number of atoms required to simulate even a $1\mu m \times 1\mu m \times 10\mu m$ channel would be of order 2×10^{11}. Assuming around $200,000$ molecules per core, this would require one million cores to simulate the entire channel. As most important molecular effects occur within a small ($\sim O(nm)$) region adjacent to the wall, Coupling could be used to reduce the MD region to $\sim O(10^9)$ molecules. This is a problem which could be easily run on $4,096$ cores.

This example demonstrates the potential of a coupling scheme in general and the benefit to HECToR users in particular.

4 Conclusion and future work

The coupling software was written and tested for correctness and scalability, as demonstrated in section 2.2.

Building on the success of the initial project, a subsequent dCSE project (in progress at time of writing) is continuing this work. The aim is to further develop the coupler so it will be useful to both coupling researchers and, in time, the wider fluids and nano-scale simulation community. The coupler will be expanded to pass a range of different variables (e.g. vectors and tensors) necessary for use in different coupling schemes. In addition, load balancing between the DNS and MD solvers and scalability in realistic problems will be considered. The development includes enhanced modularity of the various routines based largely on the widely successful MPI standard. The aim is to provide a minimal and robust set of coupler interface calls which

manipulate a protected set of variables accessible only through the coupler interface. This prevents side effects and make the coupler easily portable to other DNS/MD codes in various languages (e.g. LAMMPS in C++). In addition, an exhaustive parameter study will fully test the code and provide detailed error catching which, together with extensive documentation, should make the code attractive to HECToR users.

Acknowledgements

This project was funded under the HECToR Distributed Computational Science and Engineering (CSE) Service operated by NAG Ltd. HECToR A Research Councils UK High End Computing Service - is the UK's national supercomputing service, managed by EPSRC on behalf of the participating Research Councils. Its mission is to support capability science and engineering in UK academia. The HECToR supercomputers are managed by UoE HPCx Ltd and the CSE Support Service is provided by NAG Ltd. `http://www.hector.ac.uk`.

References

[1] S. T. O'Connell and P. A. Thompson. Molecular dynamics-continuum hybrid computations: A tool for studying complex fluid flow. *Phys. Rev.*, E 52:R5792, 1995.

[2] R. Delgado-Buscalioni and P.V. Coveney. Continuum-particle hybrid coupling for mass, momentum, and energy transfers in unsteady fluid flow. *Phys. Rev. E*, 67:046704, 2003.

[3] X. B. Nie, S.Y. Chen, W. N. E, and M.O. Robbins. A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow. *J. of Fluid Mech.*, 500:55, 2004.

[4] T. Werder, J. H. Walther, and P. Koumoutsakos. Hybrid atomistic-continuum method for the simulation of dense fluid flows. *J. of Comp. Phys.*, 205:373, 2005.

[5] N. G. Hadjiconstantinou. Hybrid atomistic-continuum formulations and the moving contact-line problem. *J. Comp. Phys.*, 154:245, 1999.

[6] M. K. Borg, G. B. Macpherson, and J. M. Reese. Controllers for imposing continuum-to-molecular boundary conditions in arbitrary fluid flow geometries. *Molec. Sims.*, 36:745, 2010.

[7] J. Li, D. Liao, and S. Yip. Coupling continuum to molecular-dynamics simulation: Reflecting particle method and the field estimator. *Phys. Rev. E*, 57:7259, 1997.

[8] W. Ren. Analytical and numerical study of coupled atomistic-continuum methods for fluids. *J. of Comp. Phys.*, 227:1353, 2007.

[9] E. G. Flekkøy, G. Wagner, and J. Feder. Hybrid model for combined particle and continuum dynamics. *Europhys. Lett.*, 52:271, 2000.