CRESTA

MASSIVELY PARALLEL LIQUID CRYSTAL SIMULATION

Liquid crystals (LCs) are widespread in technology (including displays and other optical devices) and also in nature, but much is yet to be understood about the range of possible LC configurations. Simulations are vital in paving the way to improved knowledge and exciting new applications.

The below simulation pictures show, in yellow, the defect lines that naturally occur within bulk blue-phase 1 LC. The addition of particles (shown in blue) can dramatically alter the properties of the substance. Small particles (left) will be attracted to the nearest defect line, which can have practical advantages in increasing LC stability. But particularly interesting is inclusion of relatively large particles (right): these are seen to have a more profound effect by completely disrupting the defect structure. We aim to tune the properties of the system such that the LC can act as a template to guide particles in self assembly, allowing for new substances with special optical properties. However we need to properly resolve the structure whilst having a large enough simulation to include enough particles: this is extremely computationally demanding.

The work within CRESTA has enabled the use of massively parallel GPU-accelerated supercomputers to solve this problem. We have successfully used 4,096 GPUs in parallel to simulate a system that combines large particles and the LC system (see image on back page). We are now in the process of tuning the system to properly explore physical properties, in collaboration with our experimental colleagues.

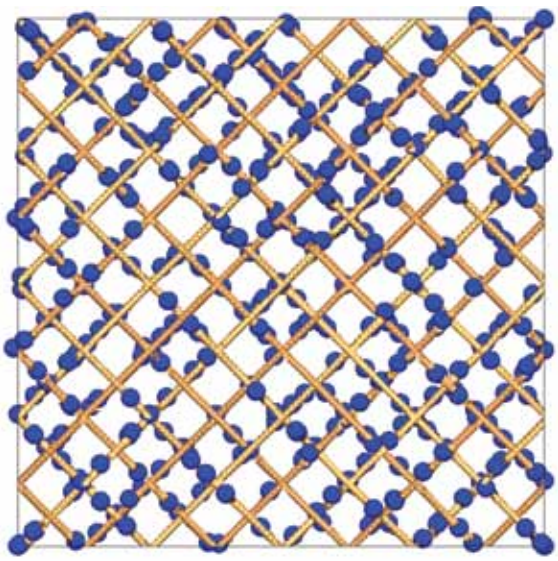

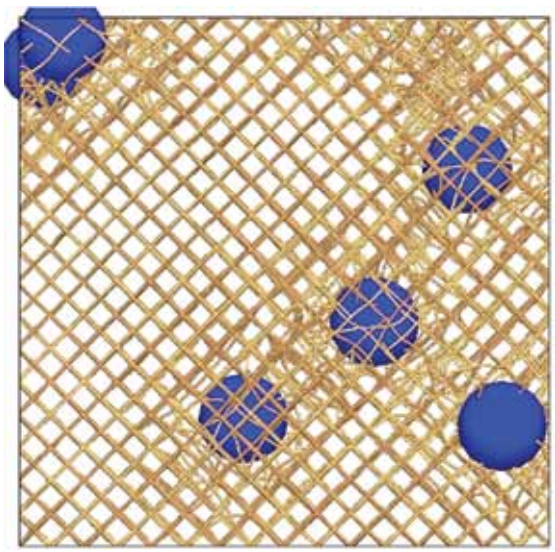

The CRESTA project receives funding from the EU's Seventh Framework Programme (FP7/2007-2013) under the grant agreements nr. 287703.

Optimising for thousands of GPUs in parallel

Ludwig is a versatile software package capable of simulating soft matter substances using lattice Boltzmann (LB) and finite-difference techniques. This lattice-based approach is particularly suitable for parallelisation: we can exploit large scale supercomputers and for many years Ludwig has been solving problems using multiple CPUs in parallel. However, complex problems such as combinations of liquid crystals and large particles require the very fastest of supercomputers, and today's top machines feature accelerators such as Graphics Processing Units (GPUs). The push towards the Exascale demands such power efficient many-core architectures, but mapping the scientific problem to the hierarchy of distinct computational components and memory spaces, in a way that achieves optimal performance, is a major challenge.

To **minimise data movement** we offload all computational components in each timestep in order to be able to keep data resident on GPU. Particles resident on the CPU interact with the fluid resident on the GPU through transfers limited to the small subset of affected fluid data. The **optimal decomposition** is dependent on the kernel: for most we use separate CUDA thread for each site, but for some a team of threads works on each site to aid caching. To **maximise memory performance** the data layout is tailored to ensure memory coalescing, and the constant cache and texture memory pipe exploited where possible for read-only data. We **minimise multi-GPU communication overheads** through use of multiple CUDA streams and non-blocking MPI to overlap different stages of the communication phase.

> "Whilst undertaking our CRESTA project work to enable Ludwig for massively-parallel GPU accelerated supercomputers, we have strived to ensure that the techniques can be transferred to other applications and scientific problems."
>
> – Alan Gray, EPCC, The University of Edinburgh

The simulations were performed on Titan at Oak Ridge National Laboratory, a Cray XK7 featuring 18,688 NVIDIA K20X GPUs

targetDP: portable performance for lattice-based applications

Modern computing systems feature a hierarchy of parallelism: the programming challenge is to expose algorithmic parallelism in a way that maps to hardware with optimal performance, intuitive expression of algorithmic content, and portability across different systems such as those with and without accelerators.

targetDP is a lightweight framework which targets the Data Parallelism inherent in lattice-based applications to the hierarchy of hardware parallelism. The same application source code can be compiled for GPUs or SIMD multi-core CPUs. We are now in process of integrating targetDP fully within the Ludwig complex fluid application.

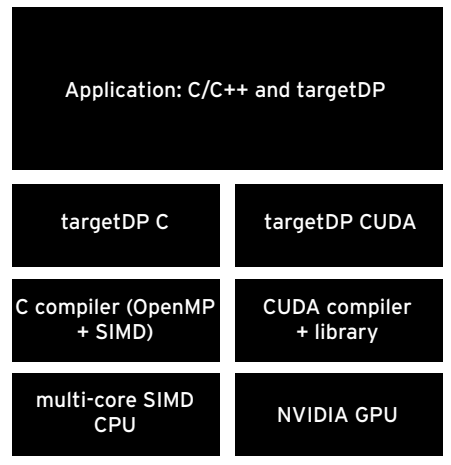

The lightweight targetDP framework consists of a set of C preprocessor macros, and a small C library interface for memory management.

> "Performance portability is vital to allow us to use the evolving range of leading machines effectively."
>
> Kevin Stratford (Ludwig lead developer), EPCC, The University of Edinburgh

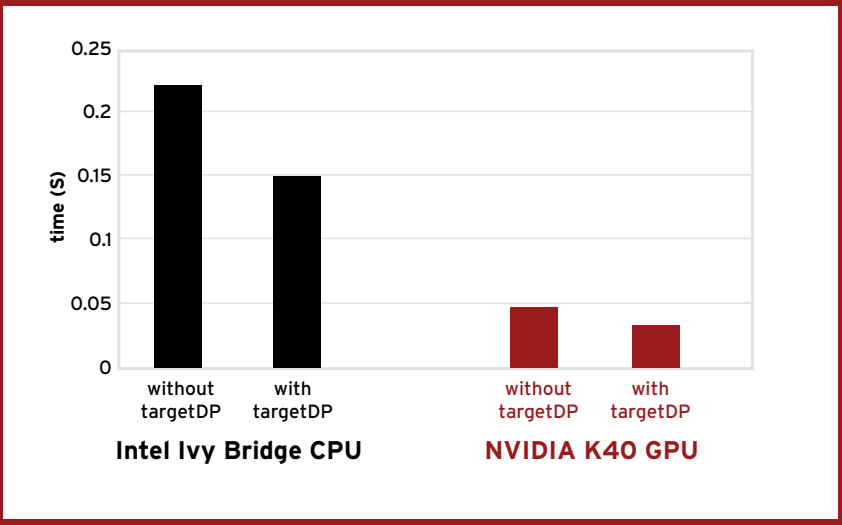

targetDP maps lattice-based parallelism to thread-level parallelism (TLP) and instruction-level parallelism (ILP). TLP maps to OpenMP threads on the CPU and CUDA threads on the GPU. The intelligent exposure of ILP has performance benefits for both CPUs and GPUs, as seen in the above graph which shows time for a binary collision LB computational kernel extracted from Ludwig.

For more information see: www2.epcc.ed.ac.uk/~alang/publications

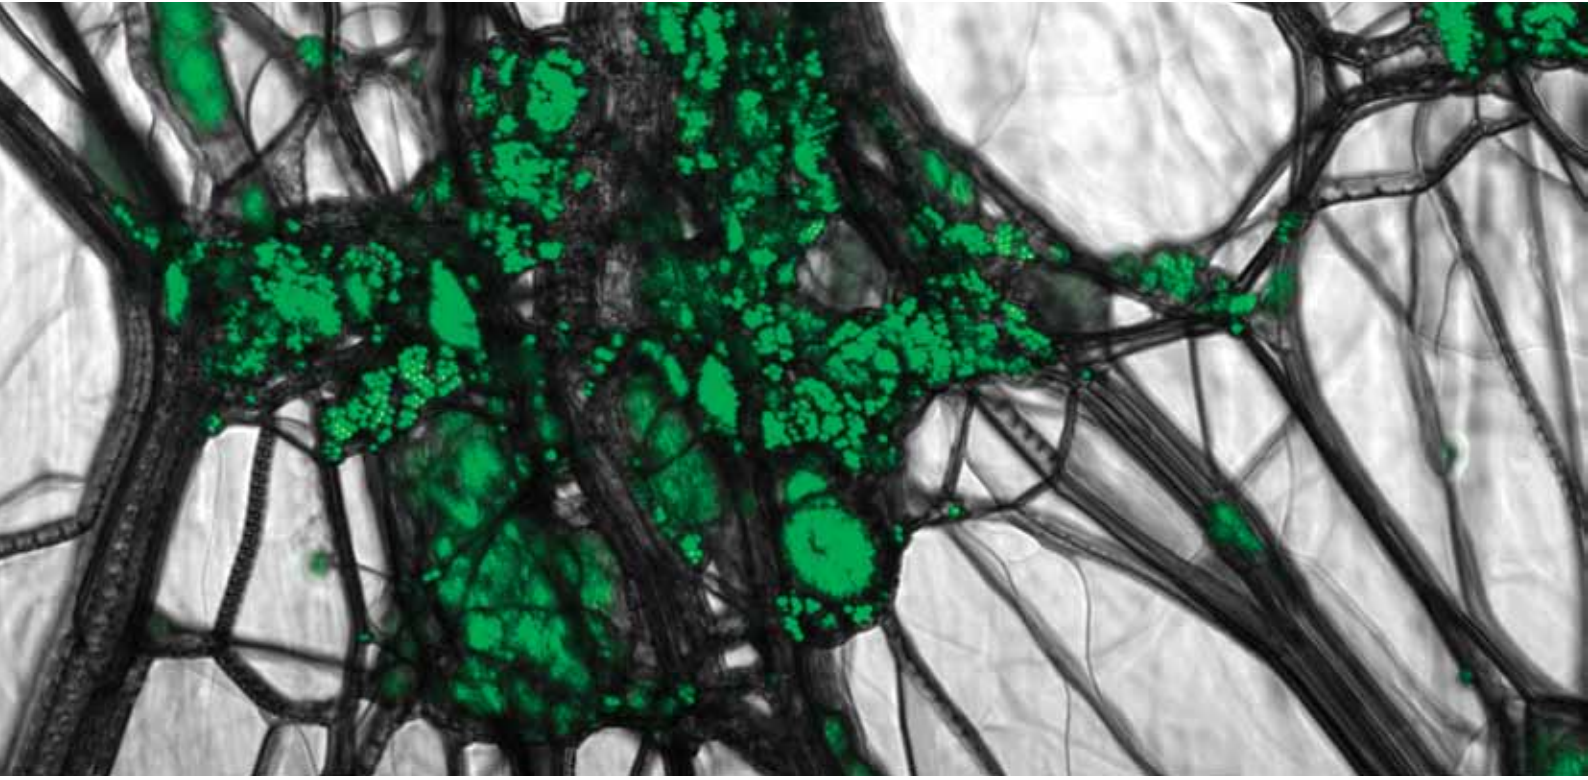

The Collaboraitve Research into Exascale Systemware, Tools and Applications (CRESTA) project form's part of Europe's exascale research program. This is focused on developing applications with exascale potential and through a co-design process, developing the software required to support these applications at the exascale.

"This work demonstrates the potential of running highly tuned parallel applications on the fastest of parallel supercomputers."

– Alistair Hart, Cray

"Strong scaling is THE hard problem in parallel computing, so it is extremely exciting to see projects that successfully harness thousands of massively parallel GPUs in unison to solve important scientific problems."

– David Luebke, Senior Director of Research, NVIDIA

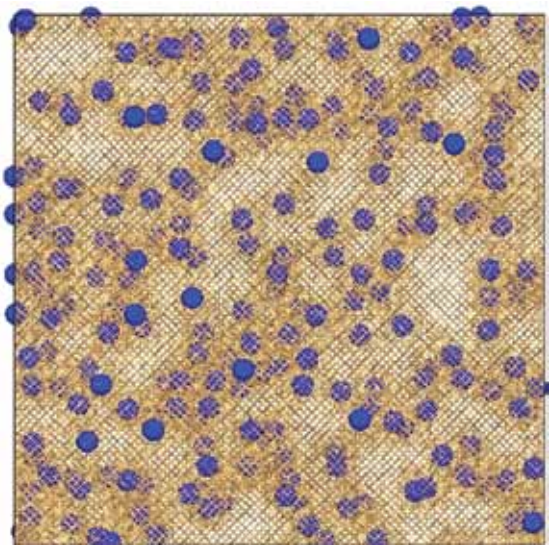

LEFT: A visualisation from the first large-scale simulation of bulk blue-phase 1 liquid crystal with large colloidal particles. This 1024^3 lattice was partitioned across 4,096 GPUs on the Titan supercomputer at Oak Ridge National Laboratory. RIGHT: A polarising optical microscopy image from an experiment showing 3 colloidal particles interacting with cholesteric LC.