

数学与系统科学研究院

计算数学所学术报告

报告人: **Prof. Guo Xiaohu**

*( Scientific Computing Department, Science and Technology  
Facilities Council, Daresbury Laboratory, Daresbury Science and  
Innovation Campus, Warrington, Cheshire )*

报告题目:

**Exploring the thread-level parallelisms  
for the next generation geophysical fluid  
modelling framework Fluidity-ICOM**

邀请人: 崔涛 博士

报告时间: **2013 年 5 月 14 日 (周二)**

**下午 16:00-17:00**

报告地点: **科技综合楼三层 301**

**计算数学所小报告厅**

## **Abstract:**

The major challenges caused by the increasing scale and complexity of the current petascale and the future exascale systems are cross-cutting concerns of the whole software ecosystem. The trend for compute nodes is towards greater numbers of lower power cores, with a decreasing memory to core ratio. This is imposing a strong evolutionary pressure on numerical algorithms and software to efficiently utilise the available memory and network bandwidth.

Unstructured finite elements codes have been effectively parallelised using domain decomposition methods, implemented using libraries such as the Message Passing Interface (MPI) for a long time. However, there are many algorithmic and implementation optimisation opportunities when threading is used for intra-node parallelisation for the latest multi-core/many-core platforms. The benefits include reduced memory requirements, cache sharing, reduced number of partitions and less MPI communication. While OpenMP is promoted as being easy to use and allows incremental parallelisation of codes, naïve implementations frequently yield poor performance. In practice, as with MPI, the same care and attention should be exercised over algorithm and hardware details when programming with OpenMP.

In this talk, we highlight our progress in implementing a hybrid OpenMP-MPI version of the unstructured finite element application Fluidity-ICOM. We demonstrate that utilising non-blocking algorithms and libraries are critical to mixed-mode application so that it can achieve better parallel performance than the pure MPI version. In the matrix assembly kernels, the OpenMP parallel algorithm utilises graph colouring to identify independent sets of elements that can be assembled simultaneously with no race conditions. The TCMalloc are used here to tackle performance issues arising from automatic arrays memory allocations. The sparse linear systems defined by various equations are solved by using threaded PETSc and HYPRE is utilised as a threaded preconditioner through the PETSc interface. With explicit communication overlap using task-based parallelism, a significant speedup over the pure-MPI mode and efficient strong scaling for PETSc sparse matrixvector multiplication kernels has been achieved. Since unstructured finite element codes are well known to be memory bound, particular attention has to be paid to ccNUMA architectures where data locality is particularly important to achieve good intra-node scaling characteristics. With mixed mode MPI/OpenMP, Fluidity-ICOM can now run well above 32K cores job.

**欢迎大家参加!**