



e-Seminar #29

Release the power of exascale using OpenMP

18 January 2023 2pm CET / 1pm GMT (1h duration)

Register for free here: <https://register.gotowebinar.com/register/4235623781431049741>

Modern supercomputer nodes now contain more than 100 cores, and often several GPUs. This means that the "MPI only" approach to parallel programming is coming under increasing strain and can no longer deliver the maximum performance from the hardware. OpenMP is the de facto standard for programming shared memory nodes, and now supports offloading to accelerator devices such as GPUs.

In this seminar, we will look at some ways that OpenMP can be used in the context of very large scale computing: combining OpenMP and MPI in a single application, portable offloading to GPUs, and using tasks for irregular parallelism.

This is the 29th in a series of online e-Seminars organised by CompBioMed.

Watch the full series at www.compbioimed.eu/training!

Dr Mark Bull is a Senior Research Fellow at EPCC in the University of Edinburgh. His research interests include benchmarking, parallel algorithms, parallel programming APIs, and novel uses of HPC. Mark acts as EPCC's representative on the OpenMP Architecture Review Board, and is a former chair of the OpenMP Language Committee, having had overall responsibility for Versions 2.5 and 3.0 of the OpenMP specification. He has been teaching OpenMP to scientists and students for more than 20 years.



Moderated by Tim Weaving, UCL



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 823712.

Organised in collaboration with:

