Contribution ID: **41** Type: **not specified**

Concurrency and Scientific Programming

Thursday, August 31, 2017 10:40 AM (40 minutes)

Modern CPU architectures for scientific computing are characterised by having multiple CPU cores and wide vector registers. Effective use of these features requires different design patterns from those that ran on the single core systems of the past. Instead of serial processing, parallel processing needs to be at the heart of the new model, with many operations proceeding concurrently. In this talk I discuss what that challenge means for scientific code and which programming patterns are most likely to be useful to scientists. I will look at some common libraries used to safely implement concurrency and give some examples from various fields. I will also emphasise the role of memory layout in achieving good throughput in high performance scientific code.

Presenter: Dr STEWART, Graeme (University of Glasgow)

Session Classification: Plenary