Contribution ID: **59** Type: **Poster Presentation**

Automatic dynamic stack management in large scientific applications: A case study using a global spectral model

Tuesday, 15 March 2016 17:00 (30 minutes)

Compute and data intensive scientific applications demand compilers to allocate more temporaries on the stack. For example, the change resolution component of a global spectral model changes the resolution of the input files using Nearest Neighbor Interpolation which requires large temporaries on stack. Temporaries include sub-arrays, automatic arrays and, sub-sections corresponding to actual arguments of a subroutine. If the infrastructure cannot provide adequate stack space at runtime relative to the total size of the temporaries, then the application program runs out of stack and aborts. Using Heap memory allocation as a solution to this introduced around 5% additional runtime to allocate and free the memory. This is observed in various components of a global spectral model.

We propose an automatic dynamic stack management framework which uses application profile information and the information of required stack memory at compile time. It does not mandate any hardware configuration changes. This technique manages stack frames on RAM by compiler-inserted code into the application binary. Our experiments with a global spectral model show that the average runtime savings of 21% along with a compile time overhead of 4%. The actual gain depends on the size of the temporaries in an application and the size of the RAM. Currently, it supports sequential and OpenMP applications, we further enhance our framework to deal with the complex MPI and GPU programming paradigms.

Primary author: Mr LAVETI, Ramesh Naidu (C-DAC)

Presenter: Mr LAVETI, Ramesh Naidu (C-DAC)

Session Classification: High Throughput & Supercomputing Systems and their Integration Session

Track Classification: High Throughput & Supercomputing Systems and their Integration