



OpenMP in the Real World: Experiences in porting a large legacy application with irregular data access

**Dr. Rob Thacker, Adjunct Assistant Professor, Physics
Queens University**

Parallelizing legacy codes can be daunting and time consuming. A lack of knowledge about the code, unclear coding practices, and opaque data dependencies can make it difficult to both extract and apply parallelism. Fortunately, the OpenMP API supports loop-level parallelism by the simple addition of pragmas to source code, and can often produce dramatic speed-ups for comparatively little effort. I'll review personal experiences from parallelizing a 7,000 line application that has irregular data access and a variable work pattern. I'll also present a Top-Ten list of tips to achieve good performance using OpenMP, and show how the tips improved the performance of the legacy application.