

# IRS

## Summary Version

1.0

### Purpose of Benchmark

The name IRS stands for implicit radiation solver. IRS solves a diffusion equation on a three-dimensional, block structured mesh. The code is written in C.

IRS is important for LLNL because the style of for loops and array indexing used in IRS is very representative of the coding style that is used widely in production codes at LLNL. Both parallel scaling efficiency and single processor performance of benchmark results in RFP submissions will be studied closely.

### Characteristics of Benchmark

This benchmark does use both MPI and OpenMP. This “hybrid” parallelization model is very important to the LLNL ASC program and key to the success of the Sequoia procurement.

On LLNL’s current computers, the principal challenge for parallel scaling efficiency is MPI collective performance.

The single CPU instruction mix on LLNL’s production computers is roughly 40% load/store, 18% floating point, 31% fixed-point, and 11% branch.

### Mechanics of Building Benchmark

Detailed instructions may be found in the file called “irs.readme.html” which is located in the tar file.

### Mechanics of Running Benchmark

Instructions for running the Sequoia-specific benchmarks can be found via a link on the main Sequoia benchmark page.

### Verification of Results

At the end of the run, both an official (raw) figure of merit (FOM) and a pass/fail verification message are printed.