

# SPhot

## Summary Version

1.0

### Purpose of Benchmark

Confirm correct hybrid MPI + OpenMP performance, single CPU performance, parallel scaling.

### Characteristics of Benchmark

Parallel execution requires no communication except to distribute the problem parameters and collect results. Loads are very close to being perfectly balanced. OS “noise” can lead to apparent load imbalance.

The instruction mix is roughly 34% load/store, 34% fixed point, 24% floating point (of which slightly over half are done in FMA instructions), and 9% branch. L1 cache hit rate is challenging.

To more closely represent LLNL production codes, there are three OpenMP critical regions. The work in these critical sections is not all important for the correct operation of the benchmark; however, they must not be optimized out of the compiled code.

### Limitations of Benchmark

Although SPhot’s calculation phase scales well to very large numbers of processors, there are two output options that do not scale well. The input variable `print_flag` controls the output options, as follows.

- 0 – Two output files are created. `out_setup.txt` reflects the input read. `out_answer.txt` contains the “answer”(escape probability) used to confirm correctness, the figure of merit, and a few other performance details.
- 1 – An additional output file is created called `out_full.txt`, containing a one line summary of each sample “run” performed.
- 2 – The file `out_full.txt` will contain both the one line summary described in option 1, plus a complete report (~50 lines) for each sample “run”.

It is recommended that after checking correctness using the small test cases, `print_flag` should be set to 0.

### Mechanics of Building Benchmark

See instructions in the tar file.

### Mechanics of Running Benchmark

See the Sequoia-specific test problem set instructions on the Sequoia Benchmarks Web page.

**Note:** Because an additional variable has been added to the input deck, the starting point for any input deck changes should be the sample provided in the latest tar file on the Sequoia Benchmarks Web page.

## **Verification of Results**

Output for specific problem sizes is provided in the tar file. Because this benchmark makes use of random numbers (from a self-contained random number generator), exact comparisons can only be made at specific problem sizes. For very large problems, the average escape probability should be very close to 0.262904, and the standard deviation should be very close to 0.00051246.