


A decorative graphic consisting of a series of blue and light blue rectangular blocks arranged in a staircase pattern, ascending from left to right.

Interpreting and publishing SPEC Results


Mathew Colgrove, NVIDIA, Portland Group

Robert Henschel, Indiana University




spec

Contents



- Preparing a result for submission
- Comparing results

Contents



- Preparing a result for submission
- Comparing results

Output files of a SPEC run

- In the **result** subdirectory
- Text files, “.txt”,
 - Preview of the Result as it would look on the SPEC website
- Log files, “.log”, “.log.debug”
 - Verbose output of the benchmark run
- Raw files, “.rsf”,
 - Above the “line” are editable fields about the run such as system or software configuration
 - Below the “line” are the encoded results. Tampering with the results will corrupt the file.

Preparing a result for submission

- Flags and Platform files
 - XML files containing detailed descriptions of the compiler flags and platform settings.
 - Required for a valid result (TODO: Add slide on this)
- Rawformat
 - Script used to format a raw file into text, html, Postscript, or PDF
 - Also performs a submission check to determine result is valid.

```
$> rawformat outputfile.rsf  
$>  
$> rawformat -F path/to/flagsfile.xml
```

Runs offline verification of result (similar to submission), produces same output as online


Adds flags-file to the result

- Hands-on
 - Edit the enclosed “.rsf” and use “rawformat” and the flags files to format a valid result.

Reportable results and Estimates

- Reportable, compliant results fully conform to the run and reporting rules of the benchmark
 - Different benchmark suites may have different rules
 - Valid results just pass the syntax, compliant results are more stringent. Such things as:
 - Are the software and hardware components:
 - Specified using customer-recognizable names, generally available within certain time frames, documented,
 - Supported and be of production quality?
 - Do the compiler flags
 - Generate correct code and improve performance for a class of programs larger than the SPEC suites,
 - Are generally available, documented, and supported?
- Estimates
 - Results that do not meet all the run and reporting rules can still be used, if clearly marked as an Estimate
 - An estimate may be fully compliant except for one issue, such as an experimental component
 - But could be numbers pulled out of thin air.


Submitting a result to SPEC



- Submitting results to SPEC for publication is **not** required for most SPEC benchmarks, including the SPEC HPG benchmarks discussed today.
- Results submitted to SPEC are reviewed before publication and may require a publication fee for non SPEC members. See: https://www.spec.org/hpg/submitting_results.html
- Process your rsf-file through rawformat to check for anything missing or faulty
- Attach your rsf-file to an email to the appropriate benchmark submission email.
 - For a list of submission emails, please contact info@spec.org
- You will receive a reply with a .sub-file attached
- In case you need to update anything, modify the sub-file, and attach to an email to
 - resubaccel@spec.org


- Beyond creating compliant results, how the results can be used is governed by SPEC
- The source of the result must be clear
- The date the result must be clear and correct
- All SPEC trademarks must be referenced
- Metrics must be disclosed.
 - Derived metrics may be used provided the SPEC metric is given.
- Basis of comparison is disclosed (if applicable)
- Full fair use rules can be found at: <https://www.spec.org/fairuse.html>

Research and Academic Use




- While academics are encouraged to comply with SPEC run and reporting rules, it's understood that this may not be possible.
- Hence, academics are allowed to modify and more or else abuse the benchmarks as the needs of their research dictates. Provided that:
 - Results are marked as “estimates”
 - Deviations from the rules be clearly disclosed

Contents





- Preparing a result for submission
- Comparing results

Use Cases (at IU)




- Comparing performance and energy
- Comparing performance of hypervisors
- Comparing HPC systems at a site
- Compare compiler performance over time
- Compare performance of different compilers
- Scalability study for different interconnects
- System setup questions like to use HT or not, which OS to use
- Compare accelerator performance

SPEC ACCEL OpenCL





SPEC ACCEL OpenACC





SPEC MPI2007 Medium

Cray XE6 (AMD Opteron 6380)





SPEC MPI2007 Medium

IBM iDataPlex (Intel Xeon L5420)





SPEC MPI2007 Medium

Available HPC Systems at IU



SPEC OMP2012

Performance and Energy



SPEC ACCEL

The effect of ECC (Using result #21 and #22, NVIDIA K40c, base)

