Vectorization of CMSSW offline software

Dr Patrick Gartung[1] on behalf of the CMS Collaboration

[1] Fermi National Accelerator Laboratory

FERMILAB-POSTER-23-029-CMS-CSAID

Profiling with Intel's OneAPI Toolset

CMSSW uses Intel's Thread Building Blocks (TBB) library. Intel Vtune is used to profile the CMSSW application as it is the only profiler that can stitch together TBB thread stack frames. Intel Advisor is used to identify scalar loops with floating point operations that can potentially be replaced by SIMD instructions – auto-vectorization. Intel's legacy compiler ICC and small vector math library (SVML) can produce more vectorization with the SSE3 instruction set but are not validated for use in CMSSW.

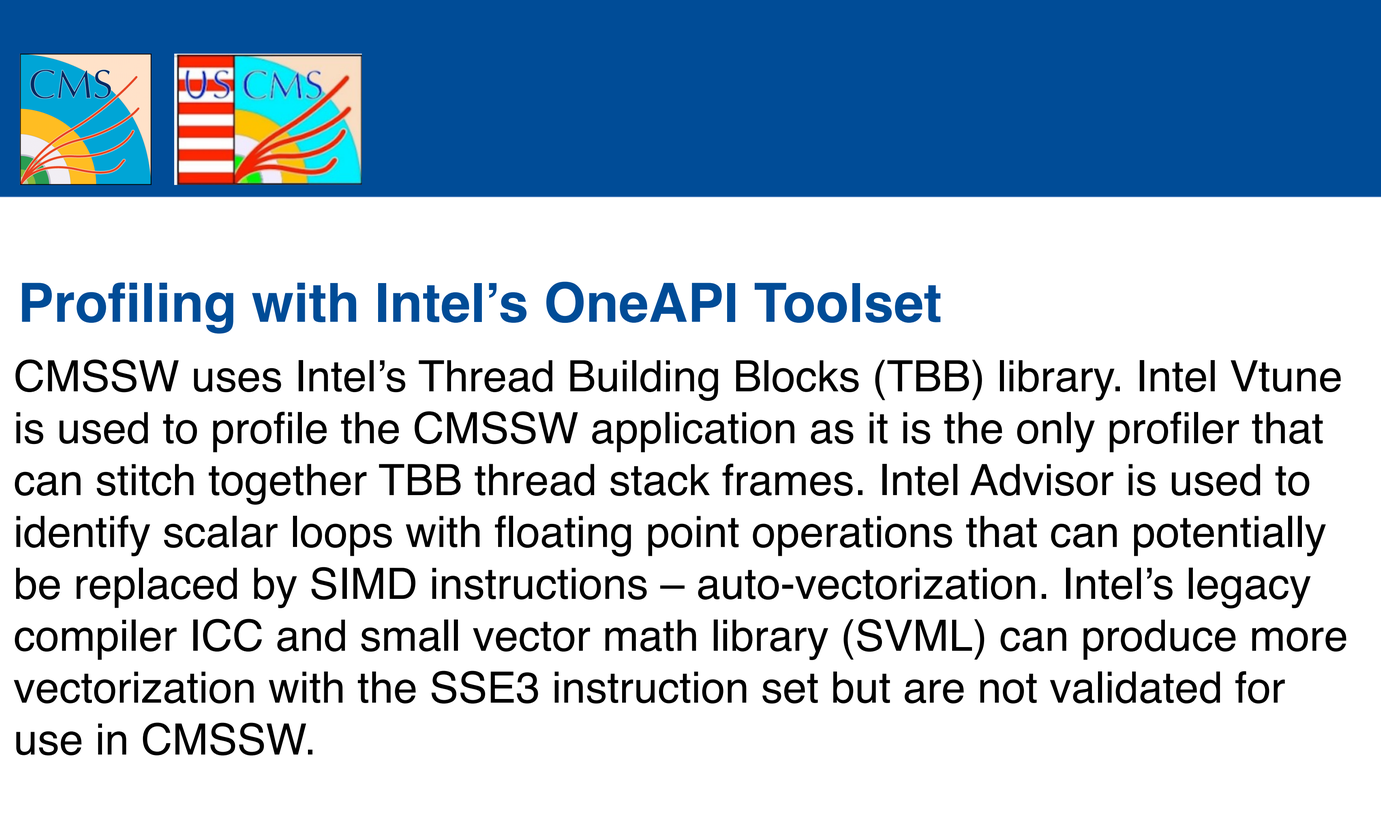

Screen caps of intel Advisor showing Vectorization Advisor pane, Threading Advisor pane, Survey pane with Source view of loop with timing, and Survey pane with Assembly view of loop with timing.

Example of changes to vectorize

An example of changes made to auto-vectorize a loop with GCC and the SSE3 instruction set

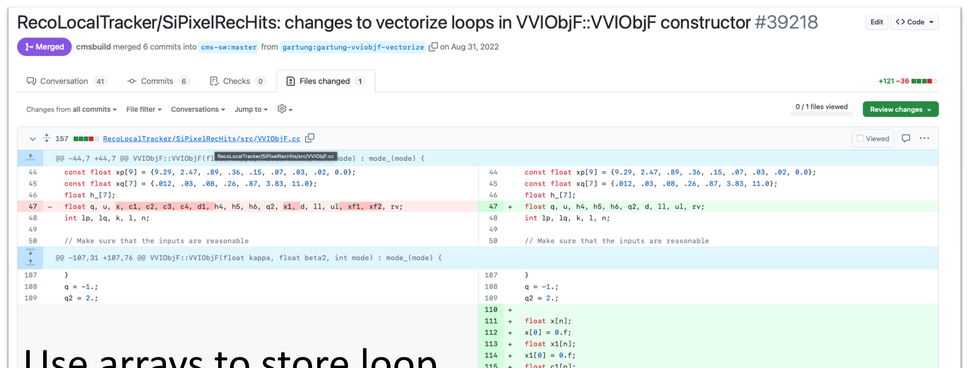

Screencap of partial diff from Github pull request for CMSSW https://github.com/cms-sw/cmssw/pull/39218/files

CMSSW multi microarchitecture releases and performance

For production, CMSSW is optimized for the SSE3 instruction set, the lowest common instruction set. Optimizing for higher instruction sets and vector widths enables more loops to be vectorized potentially increasing performance. The CMSSW build system was enhanced to produce three sets of libraries. These are compiled with the GCC march flags sse3, haswell and skylake-avx512. The library set is chosen automatically unless overridden by an environment variable.

Micro-architecture	GEN-SIM	DIGI-HLT	RECO
default	0.079 ev/sec per thread	0.087 ev/sec per thread	0.082 ev/sec per thread
haswell	0.090 ev/sec per thread	0.099 ev/sec per thread	0.077 ev/sec per thread
skylake-avx512	0.080 ev/sec per thread	0.087 ev/sec per thread	0.078 ev/sec per thread

Throughput results on a 50% loaded node for a Run 3 workflow for micro-architectures 11th Gen Intel Core i7 i11700 @ 2.50Ghz

Micro-architecture	GEN-SIM	DIGI-HLT	RECO
default	0.056 ev/sec per thread	0.053 ev/sec per thread	0.045 ev/sec per thread
haswell	0.051 ev/sec per thread	0.053 ev/sec per thread	0.053 ev/sec per thread
skylake-avx512	0.052 ev/sec per thread	0.052 ev/sec per thread	0.052 ev/sec per thread

Throughput results on a 100% loaded node for a Run 3 workflow for micro-architectures 11th Gen Intel Core i7 i11700 @ 2.50Ghz

The use of higher vector widths can cause the boost clock to be lowered for thermal management[1]. For the 50% loaded RECO process, the throughput is lower for higher vector widths. This indicates that the boost clock was used but lowered. For the 100% loaded RECO process, the throughput is higher for higher vector widths. This indicates that the boost clock was not used and the improvement from higher vector widths can be seen. [1] https://cvw.cac.cornell.edu/vector/performance_turbo

This document was prepared by the CMS Collaboration using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359.