

CNeRF Revisited + The Sim Theory Scheduler: A Case Study

Goal

In December 2024, Simulation Theory integrated the Sim Theory proprietary Scheduler technology into the CNeRF project¹. This integration resulted in approximately a **39% increase** in CPU occupancy.

Once the performance boost the Simulation Theory Scheduler achieved was clear, the team wanted to find out how the Scheduler compared to other solutions. All publicly available threading and scheduling solutions were assessed to find out which had a feature set closest to the Sim Theory Scheduler for the most fair comparison.

The requirements were:

- Multiple options for parallelization of work
- CPU agnostic
- Readily available
- Provides a C/C++ API

Intel oneAPI² and oneTBB³ were chosen because they meet the requirements in the following ways:

- Support a work stealing job system
- Support parallel for
- Support x64 and community support for aarch64
- Readily available and our team has experience using it
- C++ API

¹ <https://github.com/rafaelanderka/cNeRF>

²

<https://www.intel.com/content/www/us/en/docs/oneapi/programming-guide/2024-1/intel-oneapi-threading-building-blocks-onetbb.html>

³ <https://github.com/uxlfoundation/oneTBB>

Hypothesis

Simulation Theory's Scheduler technology can perform better than Intel's oneTBB when integrated with CNeRF at the top-level domain of the application. Neither Simulation Theory's Scheduler nor Intel's oneTBB should disrupt or modify the fork and join nature of OpenMP within PyTorch.

Out of the box, running on the CPU, CNeRF was achieving an effective CPU occupancy of 39.5% using about 25 of the 64 available logical CPU cores. This performance is due to PyTorch, its use of OpenMP, and the parallelization OpenMP provides.

Integration

Simulation Theory approached integration with CNeRF in the following ways:

Simulation Theory's Scheduler

- **Subdivide the work into equally-sized blocks.** The Scheduler was initialized with N threads. All work was then broken up inside the Scheduler into N evenly-sized blocks and processed one block per thread. The result of this is a small number of large tasks processing simultaneously on N threads.
- **Place all work in a work buffer using an atomic incrementor.** The Scheduler was initialized with N threads. An atomic incrementor was used to gate access by the N threads to the data buffer. This is the most efficient use of the Scheduler and results in fine-grained tasks processing simultaneously on N threads.

Intel's oneTBB

- **Use the oneTBB job system to create equally-sized blocks.** oneTBB was initialized with N threads. All work was divided into jobs and submitted to oneTBB for execution. The result of this is a small number of large tasks processing simultaneously on N threads.
- **Use the oneTBB parallel for solution to mimic the approach that was used with the Simulation Theory Scheduler and an atomic incrementor.** This effectively resulted in fine-grained tasks processing simultaneously on N threads.

Results

Building on top of Simulation Theory's prior integration of the Scheduler, integrating oneTBB into CNeRF only required a couple of hours of additional work. For all of the following testing, the Simulation Theory Scheduler and Intel's oneTBB were initialized with the same number of threads.

CPU Occupancy

The overall CPU occupancy was **39.5%** on 25.267 out of 64 logical CPU cores in the baseline implementation of CNeRF.

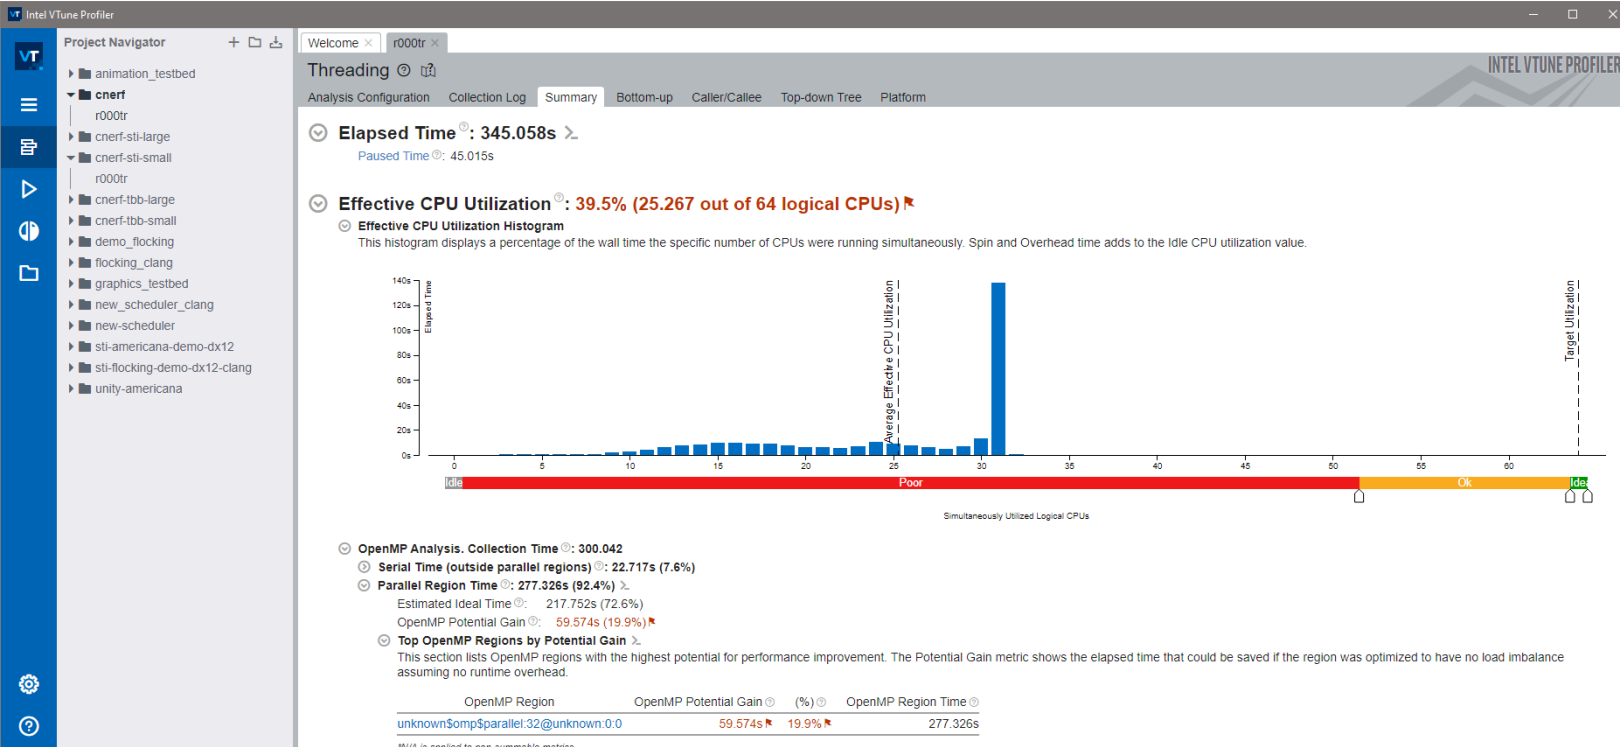
Large-block Work Distribution

- Using the large-block work distribution model of Intel oneTBB, CPU occupancy **improved** to **56.8%** on 36.361 out of 64 logical CPU cores. This was an **improvement of 30.5%**.
- Using the large-block work distribution model of the Simulation Theory Scheduler implementation, CPU occupancy **improved** to **57.4%** on 36.735 out of 64 logical CPU cores. This was an **improvement of 31.2%**.
- Simulation Theory's Scheduler technology **beat** Intel's oneTBB performance **by 1%** using the large-block work distribution method.

Fine-Grained Work Distribution

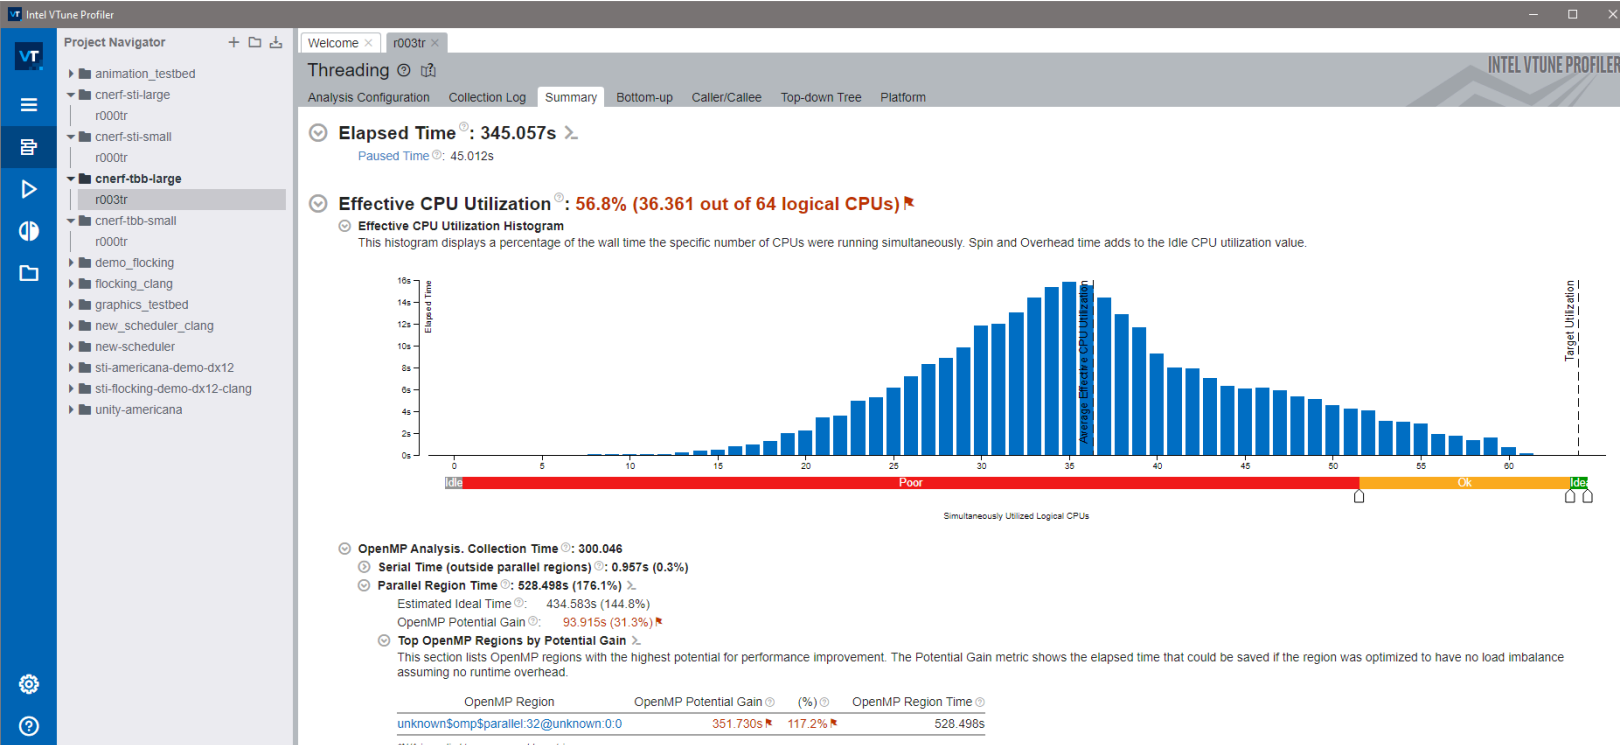
- Using the fine-grained work distribution model of Intel oneTBB, CPU occupancy **improved** to **57.4%** of 36.760 out of 64 logical CPU cores. This was an **improvement of 31.3%**.
- Using the fine-grained work distribution model of the Simulation Theory Scheduler implementation, CPU occupancy **improved** to **65.2%** on 41.743 out of 64 logical CPU cores. This was an **improvement of 39.47%**.
- Simulation Theory's Scheduler technology **beat** Intel's oneTBB performance **by 11.9%** using the fine-grained work distribution method.

CNeRF Original Performance Graph⁴

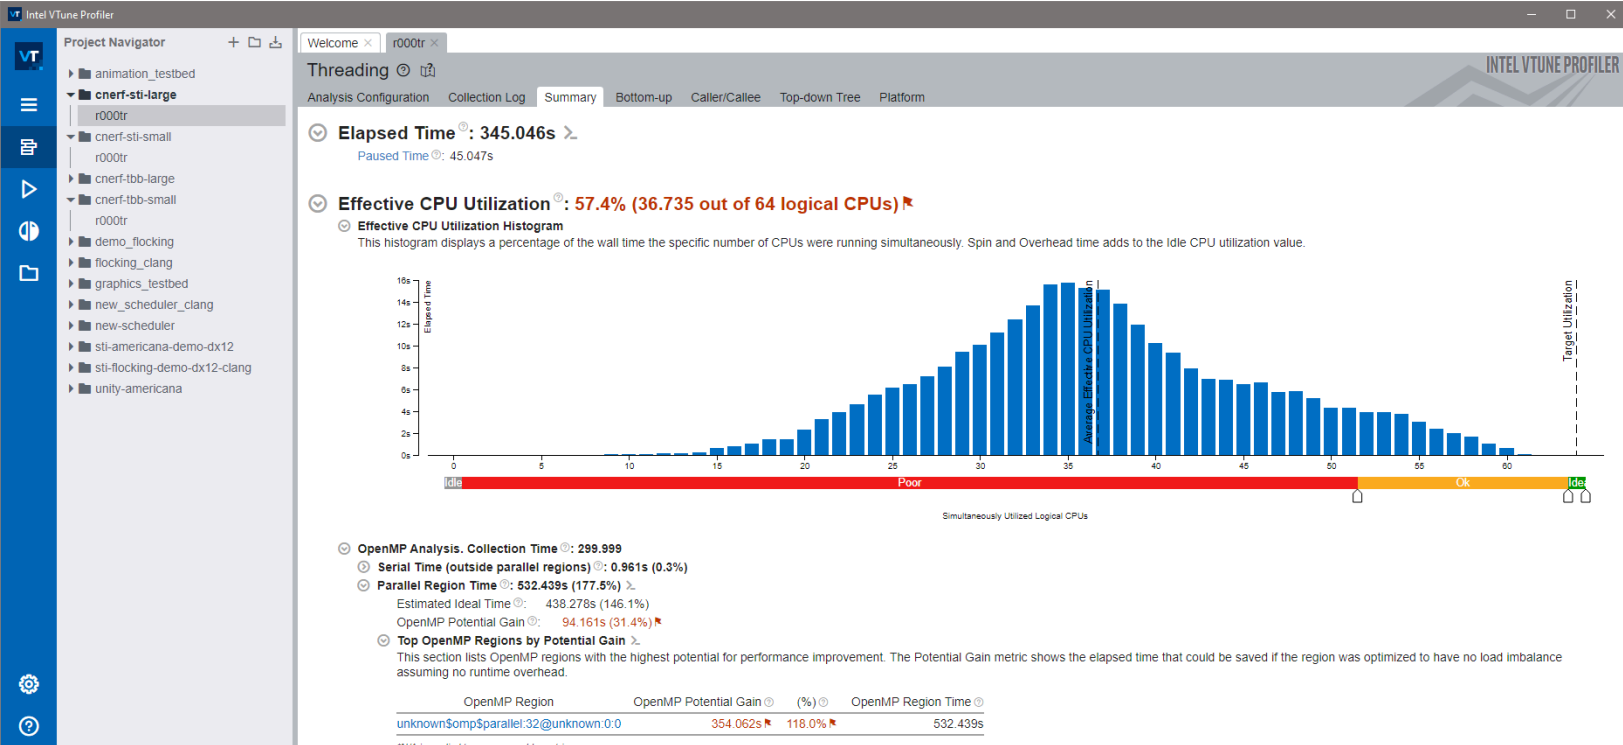


⁴ Simulation Theory used Intel V-Tune to capture these and the following performance metrics from the testing.

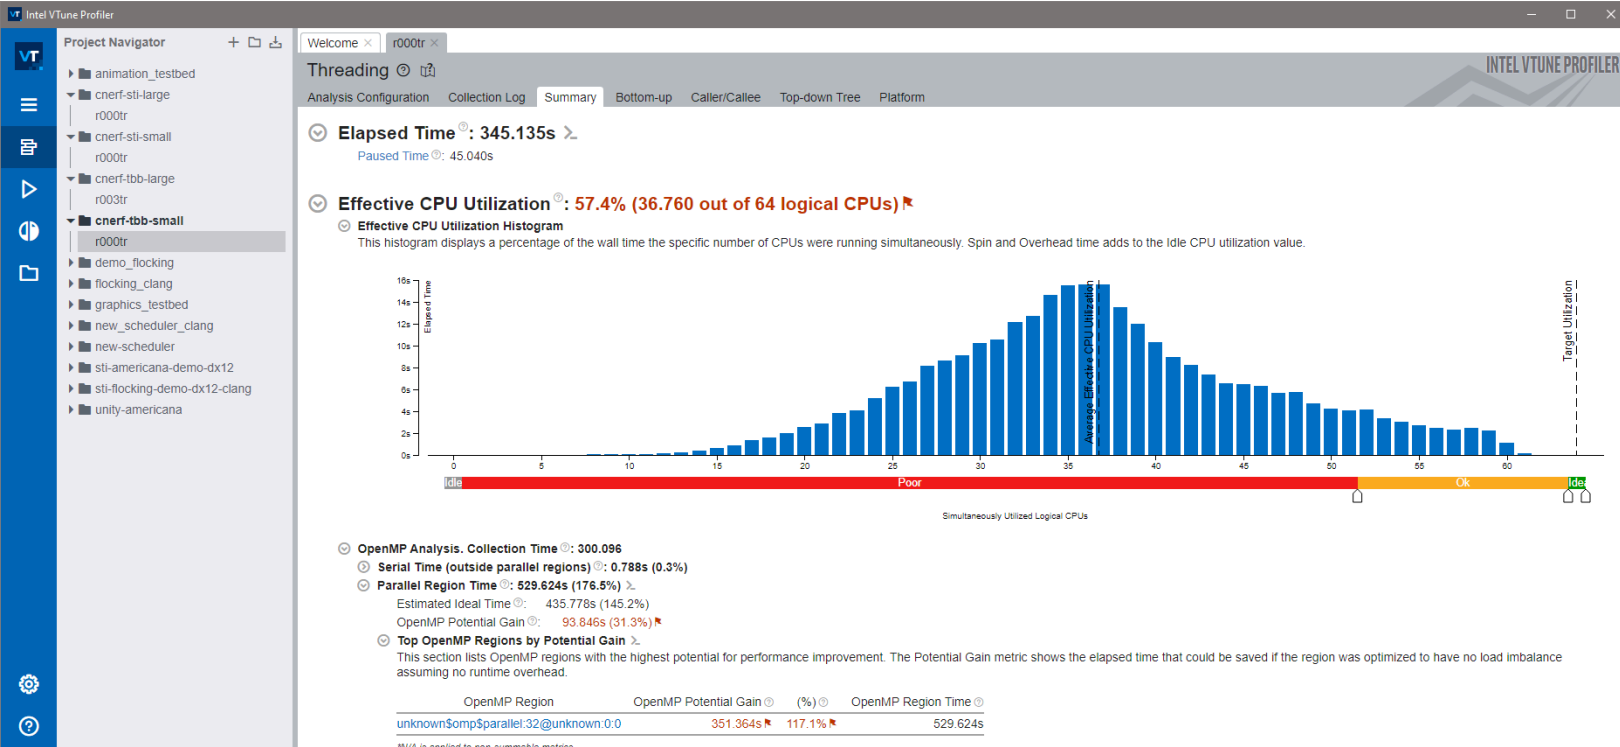
CNeRF Intel oneTBB Large, Equally-Sized Tasks Performance Graph



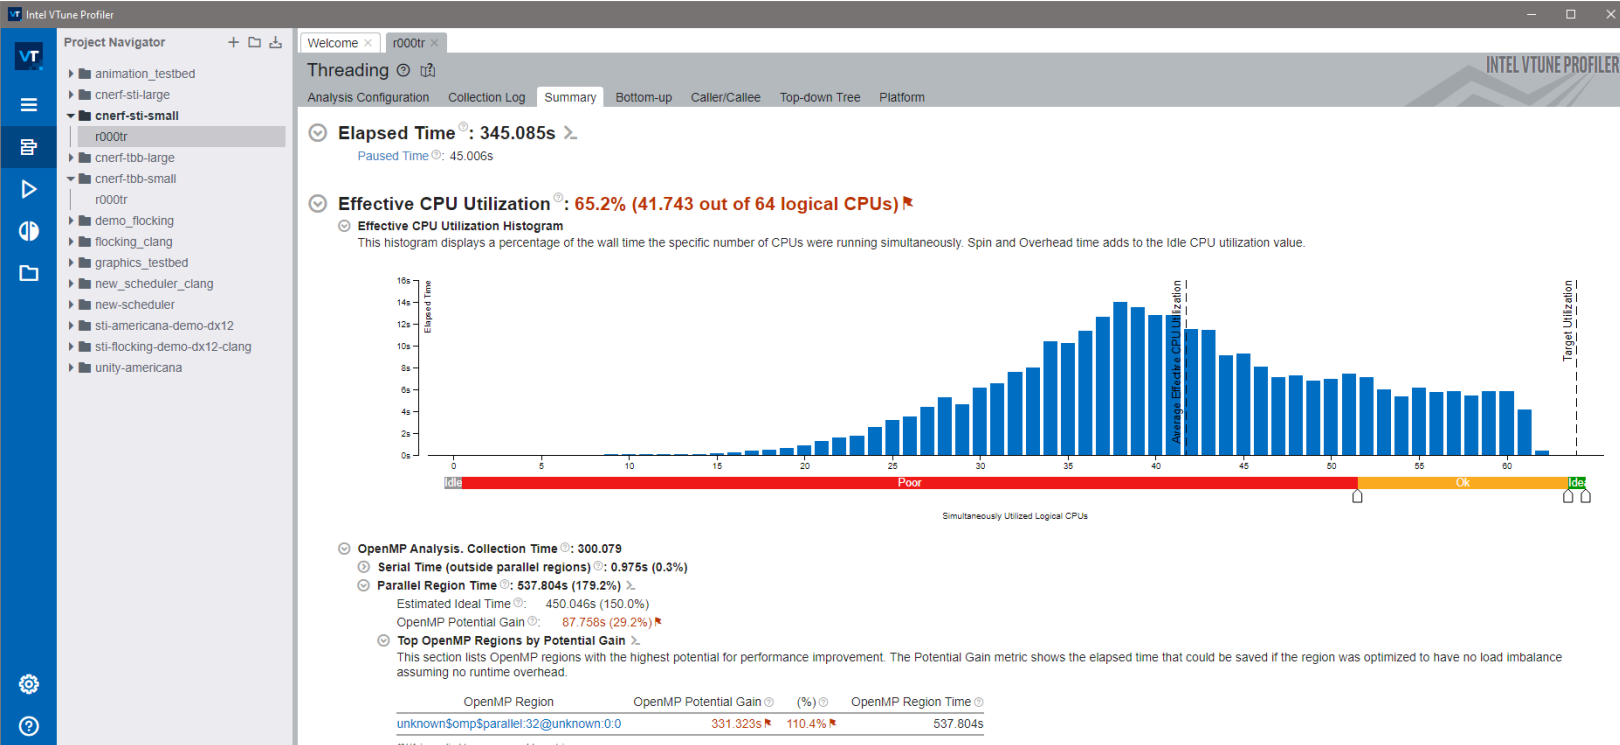
CNeRF Simulation Theory Large, Equally-Sized Tasks Performance Graph



CNeRF Intel oneTBB Small Tasks Performance Graph



CNeRF Simulation Theory Small Tasks Performance Graph



Author: Randy Culley, CTO
sales@simtheoryinc.com
simtheoryinc.com

Copyright © 2025. Simulation Theory, Inc. All Rights Reserved.