The Parallel Universe, Issue 30

Abstracts

Driving Performance with Intel® Advisor's Flow Graph Analyzer

Optimizing Performance for an Autonomous Driving Application

Vasanth Tovinkere, Architect, Intel® Flow Graph Analyzer; Pablo Reble, Software Engineer; Farshad Akhbari, Perceptual Computing Technical Lead; and Palanivel Guruvareddiar, Perceptual Computing Software Architect; Intel Corporation

The advanced driver assistance systems (ADAS) and autonomous driving technologies deployed in modern cars rely on environmental awareness through sensors such as radar, LIDAR, and cameras. The underlying machine learning or deep learning algorithms that provide this awareness are compute-intensive—and become even more demanding as the sensor resolution grows. In this article, we explain how to design and analyze the parallelism in such an application using Intel® Advisor's Flow Graph Analyzer (FGA), a tool from Intel that supports parallel applications using the Intel® Threading Building Blocks (Intel® TBB) flow graph interface.

Welcome to the Adult World, OpenMP*

After 20 Years, It's More Relevant than Ever

Barbara Chapman, Professor, Stony Brook University, and Director of Computer Science and Mathematics, Brookhaven National Laboratory

The single most amazing thing about OpenMP is that after 20 years, it's still relevant. Indeed, it's used more extensively than ever. And with implementations on all major platforms, it's broadly available. The work to evolve the specification is driven by an active team of vendor representatives, along with participants from research laboratories and universities who have committed themselves for the long haul. For a compiler-based programming interface, this is huge.

Enabling FPGAs for Software Developers

Boosting Efficiency and Performance for Automotive, Networking, and Cloud Computing

Bernhard Friebe, Senior Director of FPGA Software Solutions Marketing, Intel Corporation, and James Reinders, HPC Enthusiast

There couldn't be a better time to examine field programmable gate arrays (FPGAs). A new era in computing is emerging thanks to the new programmability of FPGAs. With the onslaught of data in the world, there's an incredible need for the power-efficient computing available with custom silicon designs—but with the flexibility available with the flexibility of FPGAs. In this article, we share our

thoughts on what it takes to truly enable FPGAs for software developers across all fields—including automotive, networking (e.g., 5G), and end-to-end cloud computing (e.g., the data center).

Modernize Your Code for Performance, Portability, and Scalability

What's New in Intel® Parallel Studio XE

Jackson Marusarz, Technical Consulting Engineer, Intel Corporation

Whether you're working on HPC clusters, remote clouds, local workstations, or anything in between, Intel® Parallel Studio XE is a workhorse—with compilers, libraries, and tools to help you improve your productivity and application performance. And Intel continues to innovate with the latest release of Intel Parallel Studio XE. Besides adding support for the newest hardware and programming language standards, several new features address growing technologies and new environments including Intel® Xeon Phi™ processors and the Intel® Xeon® Scalable processor family.

Dealing with Outliers

How to Find Fraudulent Transactions in a Real-World Dataset

Oleg Kremnyov, QA Engineer; Mikhail Averbukh, Software Engineer; and Ivan Kuzmin, Software Engineering Manager; Intel Corporation

One challenging—but also very important—task in data analytics is dealing with outliers. We generally define outliers as samples or events that are inconsistent with the rest of data population. An outlier often contains useful information about abnormal characteristics of the systems and entities that impact the data generation process. In this article, we focus on one of the most common applications of outlier detection—credit card fraud. With some simple outlier detection approaches, it's possible to find 75 to 85 percent of fraudulent transactions—with false alarms less than 1 percent of the time—on a real-world dataset.

Tuning for Success with the Latest SIMD Extensions and Intel® Advanced Vector Extensions 512

Best Practices for Taking Advantage of the Latest Architectural Features

Xinmin Tian, Senior Principal Engineer; Hideki Saito, Principal Engineer; Sergey Kozhukhov, Senior Staff Engineer; and Nikolay Panchenko, Staff Engineer; Intel Compiler and Language Lab, Intel Corporation

Intel® Advanced Vector Extensions 512 (Intel® AVX-512), the latest x86 vector instruction set, has up to two fused-multiply add units plus other optimizations. It can accelerate the performance of a wide range of workloads. In this article, we give a brief overview of the Intel AVX-512 Instruction Set Architecture (ISA) and describe what's new in the Intel 18.0 compilers for Intel® Xeon® Scalable processors. Next, we present several new simd language extensions for Intel AVX-512 support in Intel's latest compilers. Finally, we share our best practices in performance tuning to achieve optimal performance with the Intel AVX-512 ISA.

Effectively Using Your Whole Cluster

Optimizing SPECFEM3D_GLOBE Performance on Intel® Architecture*

Rama Kishan Malladi, Technical Marketing Engineer, Intel Corporation

Intel is a leading provider of both hardware and software for datacenter users—such as the latest Intel® Xeon® and Intel® Xeon Phi™ processors. However, many high-performance computing (HPC) applications don't make full use of the processors' advanced capabilities. In this article, we'll provide a step-by-step methodology to improve the performance of SPECFEM3D_GLOBE*, a software package that simulates three-dimensional global and regional seismic wave propagation and performs full waveform imaging (FWI) or adjoint tomography based on the spectral element method (SEM).

Is Your Cluster Healthy?

Must-Have Cluster Diagnostics in Intel® Cluster Checker

Brock A. Taylor, HPC Solution Architect, Intel Corporation

Intel® Cluster Checker is a powerful tool for quickly identifying and solving issues in high-performance computing (HPC) clusters. Subtle and sometimes simple issues on a system can impact cluster performance and blunt the efforts of fine-tuning and parallelizing an application. Often, the first signs of a system issue appear when applications run too slowly—or simply stop running altogether. Intel Cluster Checker provides a methodical way to help quickly determine if the underlying reason an application is experiencing problems is actually a problem with the cluster.

Optimizing HPC Clusters

Enabling On-Demand BIOS Configuration Changes in HPC Clusters

Michael Hebenstreit, Data Center Engineer, Intel Corporation

Since around 2000, most high-performance computing (HPC) systems have been set up as clusters based on commodity x86 hardware. These clusters consist of one- or two-socket servers to perform the actual computations, plus storage systems and administrative nodes. Modern Intel® Xeon® processor-based systems, as well as the Linux* kernel, provide many ways to optimize both hardware and operating system (OS) for a specific application. It's easy to do if your cluster is only used for a specific workflow—but for more complex usages, it's beyond the ability of most cluster managers. We outline a way to perform complex optimizations on a per-job basis. There's a price to pay in the form of added complexity and job startup times, but for Intel's HPC benchmarking cluster, Endeavor, this feature became a very important way to boost performance over the last year. Your gains might be even higher.