
## Why ViON?

### Experience, Expertise, Solutions

- Veteran-Owned business with over 40 years' experience serving government and commercial customers
- 100% dedicated focus on delivering optimum solutions to support our customers and their missions
- Industry pioneer in private cloud solutions
- Certified Storage, System and Network Engineers with 24/7/365 Support Center, Cleared and U.S. Based

### HPC Solutions from ViON

- Complete HPC portfolio, including Hybrid HPC, ViON Forever Data Cloud, Private Cloud and HPC Solution Services
- HPC solution partners include: Fujitsu, Supermicro, Arista, DDN, Mellanox, Pavilion I/O and more

HPC FEATURE OVERVIEW

# Fujitsu PRIMEHPC FX700 Compiler and Profiling Tools

Fujitsu has developed unique HPC compiler capabilities for the FX700 platform. These capabilities enable the FX700 to achieve even greater performance and scale when compared to other HPC compiler code also running on an FX700. Fujitsu HPC application development has evolved across many generations of products, (hardware and software). Fujitsu's software development studio has been specifically augmented to take advantage of hardware features designed into the A64FX CPU and includes the SVE extensions, the high bandwidth on chip memory and on chip networking.

The compiler runtime libraries support the HPC extensions of ARM SVE. High-speed execution of applications using SVE has been greatly increased through enhanced compiler optimization.

Fujitsu's Development Assistance Tool provides graphical display of large memory page allocation functions which are used to tune applications to achieve highly parallel execution. The development tools maximize the overall HPC performance by speeding up processing at each level (i.e., cores within a node and between nodes). FX700 utilizes HPC-X with SHARP for MPI and Data Streaming. The development studio supports the standard languages, C, C++, Fortran, OpenMP and MPI.

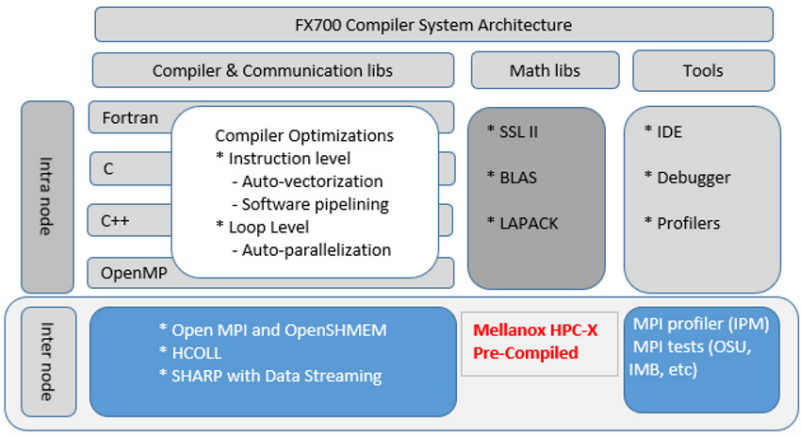

## Core Speed-up Features

The Compilers have vectorizing functions that utilize the SVE features of the A64FX. These compiler directives can be used at compile time to programmatically vectorize complex loops that contain IF statements and math functions. This is in addition to vectorization using SIMD instructions with SVE masking, providing  significant benefits with shorter loops.

**Another important feature of the compilers is their instruction-scheduling function called pipelining. The advantage of this technique is that cycles in a loop can overlap each other, resulting in faster execution for parallel operation.**

## Software pipelining function

Another important feature of the compilers is their instruction-scheduling function called pipelining. The advantage of this technique is that cycles in a loop can overlap each other, resulting in faster execution for parallel operation.

## Prefetch Function

The A64FX has 2 type of prefetch functions, a hardware driven prefetch which handles simple memory access patterns and a software derived prefetch that handles more complex memory access patterns.

## Sector Cache Instruction

The A64FX has a sector cache function which allows a portion of L2 cache to be used as high-speed memory. The compilers have specific sector cache instruction that specify data to be stored in the sector cache area. When applied to data that is used repeatedly, this enhances performance by avoiding the latency to access the HBM2 memory.

## FP16 Support

Native support for Floating Point 16 (FP16) data types is critical for many AI type applications. The A64FX C and C++ compiler support FP16 data types. The performance of SIMD calculations are two times better with FP16 than FP32 (single precision floating-point type). Using FP16 can improve conventional scientific calculation applications as well as AI applications.

## Faster Math Libraries

The Fujitsu Compiler Suite provides faster math libraries. Internal core speed-up function is applied to libraries listed below. BLAS and LAPACK are well-known for linear algebra work. SSL II is a Fujitsu library used for many fields and its fast quadruple-precision basic-arithmetic library delivers high performance by handling quadruple-precision values in the double-double format. Application will speed up significantly by calling these math libraries.

| Sequential Libraries (thread-safe) | |
|---|---|
| **BLAS, LAPACK** | These linear math libraries are de facto standard. Some BLAS routines support FP16. |
| **SSL II** | Fortran library routines covering a wide range of fields. |
| **C-SSL II** | C interface for SSL II |
| **Fast quadruple-precision basic-arithmetic library** | Represents and calculates quadruple-precision values in the double-double format. There are some routines that support thread parallelization. |

**The performance graph below demonstrates the performance gains for BLAS matrix product routine using 16FP, 32FP and 64FP.**

Figure 1: Performance Chart

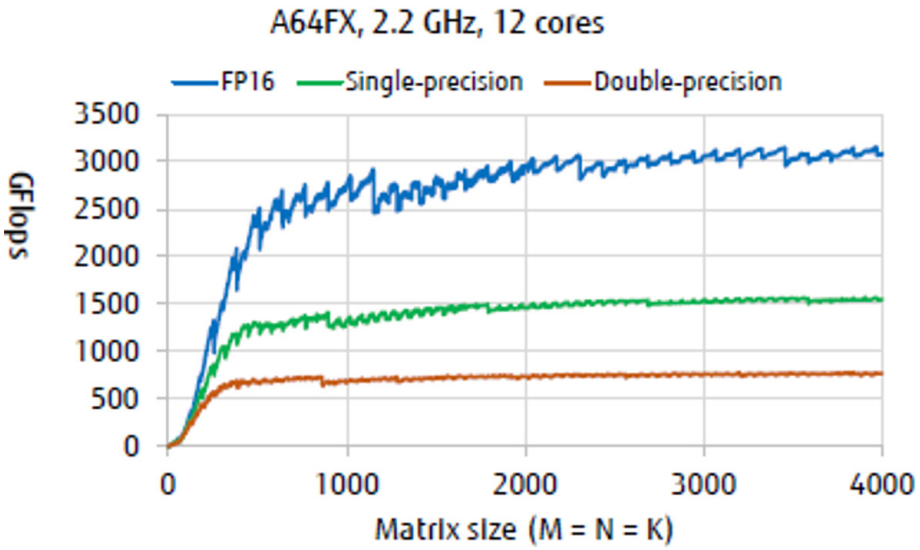

Figure 2: Performance Analysis Procedure

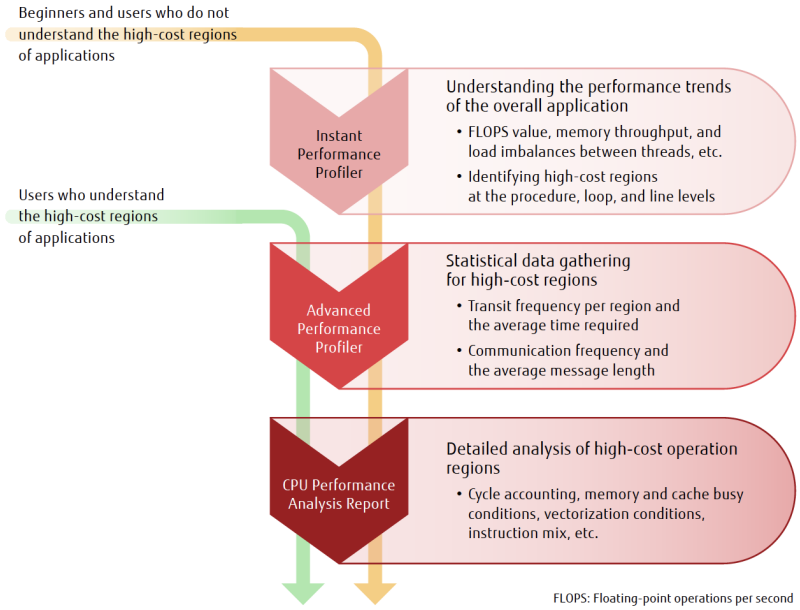


**The runtime library takes advantage of the A64FX hardware barrier mechanism for thread parallelization within a node. This speeds up the execution of parallel threads.**

## Thread Parallelization

The runtime library takes advantage of the A64FX hardware barrier mechanism for thread parallelization within a node. This speeds up the execution of parallel threads.

## Faster MPI for Intranode Communication

For the MPI "Alltoall" primitive, the A64FX MPI algorithm is optimized to take advantage of the On-Chip Ring Bus, ensuring data moves in a one-directional ring thereby avoiding communication path conflict, which would adversely affect performance.

## Performance Analysis Procedure

As depicted in Figure 2, Performance Analysis Procedure can be carried out using the basic profiler or an advanced detailed profiler, along with the CPU analysis report.

Figure 3: Example of a CPU Performance Analysis Report

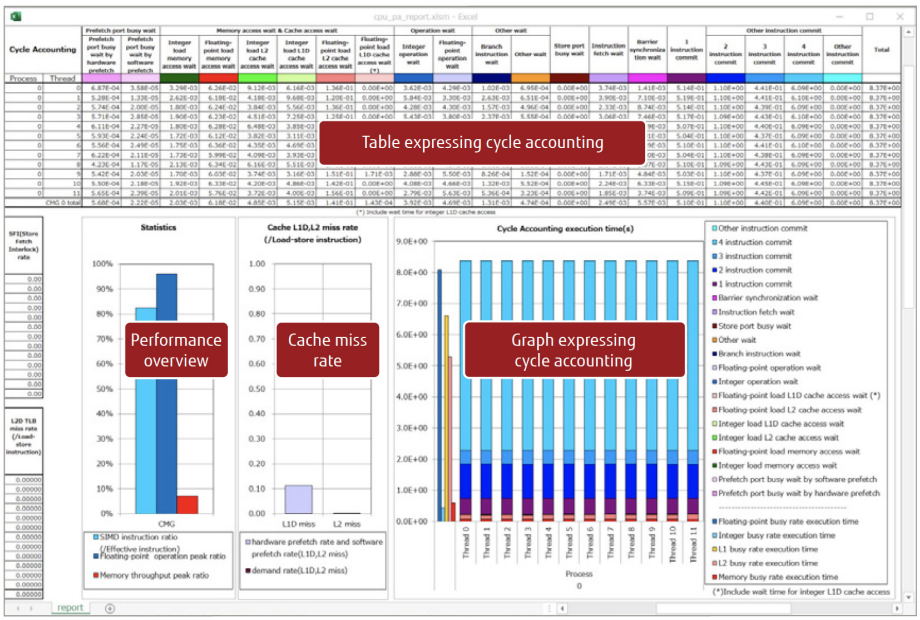


**With the detailed profiler, users can get a grasp of performance conditions and the statistical communication status of the MPI in the specified section.**

## Step-By-Step Performance Analysis With The Profilers

The basic profiler takes samples of the overall performance trend and cost distribution information for the application. The collected cost distribution information is displayed for individual procedures, loops and lines. This provides a high-level view of the overall application performance.

The detailed profiler is used to obtain more detailed performance information on hotspots identified by the basic profiler. With the detailed profiler, users can get a grasp of performance conditions and the statistical communication status of the MPI in the specified section.

The CPU performance analysis report uses graphs and tables to systematically and clearly display the PMU (performance monitoring unit) counters contained in the CPU. With the CPU performance analysis report, users can get a grasp of application bottlenecks in detail.

ViON®
Building your enterprise solutions

### About ViON Corporation

ViON Corporation is a cloud service provider with over 40 years' experience designing and delivering enterprise data center solutions for government agencies and commercial businesses. The company provides a large portfolio of IT as-a-Service, including infrastructure, multi-cloud and artificial intelligence (AI) solutions. Focused on supporting the customer's IT modernization requirements, ViON's Enterprise Cloud is changing cloud management for the market, providing a streamlined platform to audit and control technology in an evolving multi-cloud world. The ViON Marketplace® allows customers to research, compare, procure and manage a full range of everything as-a-Service solutions from leading manufacturers via a single portal. ViON delivers an outstanding customer experience at every step with professional and managed services, backed by highly-trained, cleared resources. A veteran-owned company based in Herndon, Virginia, the company has field offices throughout the U.S. (vion.com).

To learn more, go to vion.com/cloud or email us at info@vion.com