



Status, challenges and trends of data-intensive supercomputing

Jia Wei¹ · Mo Chen² · Longxiang Wang¹ · Pei Ren² · Yujia Lei¹ · Yuqi Qu¹ · Qiyu Jiang¹ · Xiaoshe Dong¹ · Weiguo Wu¹ · Qiang Wang¹ · Kaili Zhang² · Xingjun Zhang¹

Received: 28 July 2021 / Accepted: 20 May 2022 / Published online: 12 July 2022
© China Computer Federation (CCF) 2022

Abstract

Supercomputing technology has been supporting the solution of cutting-edge scientific and complex engineering problems since its inception—serving as a comprehensive representation of the most advanced computer hardware and software technologies over a period of time. Over the course of nearly 80 years of development, supercomputing has progressed from being oriented towards computationally intensive tasks, to being oriented towards a hybrid of computationally and data-intensive tasks. Driven by the continuous development of high performance data analytics (HPDA) applications—such as big data, deep learning, and other intelligent tasks—supercomputing storage systems are facing challenges such as a sudden increase in data volume for computational processing tasks, increased and diversified computing power of supercomputing systems, and higher reliability and availability requirements. Based on this, data-intensive supercomputing, which is deeply integrated with data centers and smart computing centers, aims to solve the problems of complex data type optimization, mixed-load optimization, multi-protocol support, and interoperability on the storage system—thereby becoming the main protagonist of research and development today and for some time in the future. This paper first introduces key concepts in HPDA and data-intensive computing, and then illustrates the extent to which existing platforms support data-intensive applications by analyzing the most representative supercomputing platforms today (Fugaku, Summit, Sunway TaihuLight, and Tianhe 2A). This is followed by an illustration of the actual demand for data-intensive applications in today’s mainstream scientific and industrial communities from the perspectives of both scientific and commercial applications. Next, we provide an outlook on future trends and potential challenges data-intensive supercomputing is facing. In a word, this paper provides researchers and practitioners with a quick overview of the key concepts and developments in supercomputing, and captures the current and future data-intensive supercomputing research hotspots and key issues that need to be addressed.

Keywords Data-intensive supercomputing · I/O intensive supercomputing · High performance data analytics · Parallel processing systems · Supercomputing storage

1 Introduction

Supercomputing, or High performance computing (HPC), has been an important engine of progress for human society since its inception in the middle of the last century. Its development is a comprehensive reflection of the highest-end computer hardware and software technologies, and is also the frontier track for countries around the world to compete in science and technology. From the birth of the first

electronic computer in 1946, mankind entered the era of supercomputing. HPC served as a high-end computing infrastructure for solving computationally intensive scientific and engineering problems. The core problem to be solved by supercomputing is how to efficiently utilize the evolving parallel processing units as well as parallel architecture to serve scientific computing tasks such as high energy physics, astronomy, bioinformatics, dynamics and fluid dynamics. The traditional supercomputing aims to increase the number of floating point operations per second, focusing on computational performance, programmability, scalability, energy efficiency and other issues, and is dedicated to handling computationally intensive tasks.

With the technological development of hardware and software for supercomputing and the change in demand for

✉ Longxiang Wang
wx419@xjtu.edu.cn

¹ School of Computer Science and Technology, Xi’an Jiaotong University, Xi’an 710062, China

² Huawei technologies Co. Ltd., ShenZhen 518129, China

supercomputing applications caused by social progress, the role played by supercomputing is also changing (Lu et al. 2018). The data generated by typical HPC applications such as earthquake prediction, oil exploration, and climate analysis are growing explosively, and are characterized by volume, velocity, and variability, which are typical big data scenarios. On the other hand, AI technology is being used in various HPC scenarios—such as deep learning-based climate analysis and earthquake simulations—and has even won a Gordon Bell Award. Joseph II, Executive Director of IDC HPC User Forum, pointed out that the advent of the Big Data era has made HPDA (High performance data analytics) applications the next hot spot for HPC, with 67% of HPC resources currently devoted to HPDA—of which AI is a typical application. Zheng Weimin, an academician of the Chinese Academy of Engineering, also pointed out that the future trend of supercomputing is AI + scientific computing at the 2019 HPC China Conference. Consequently, data-intensive tasks are becoming a mainstream application of supercomputing.

Figure 1 shows the industry application distribution from the 2018 China HPC TOP 100. The applications are distributed across eight major areas, with the largest number of applications used for big data processing at 27%; the rest are cloud computing (20%), scientific computing (17%), supercomputing centers (13%), government telecommunications (6%), security (4%), industrial manufacturing (4%), and energy and oil (2%). This indicates that HPC applications are becoming more data-intensive. The design goals of supercomputing have also shifted from simply pursuing computational power in terms of FLOPS, to pursuing both computational and data processing power. In recent years, the benchmark for assessing supercomputing capabilities has been supplemented by the I/O 500 and GRAPH 500 benchmarks, which have shifted from purely assessing compute-intensive capabilities to more data-intensive processing capabilities. GRAPH 500 focuses on a newly designed

benchmark based on graph search and shortest path problems, with a focus on the ability to handle data-intensive applications.

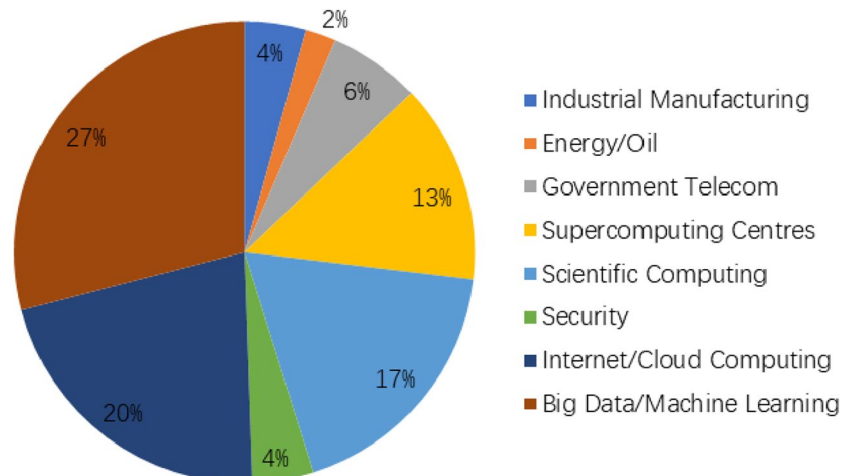
There are currently two main academic perspectives that distinguish between data-intensive and compute-intensive tasks. One considers the amount of data and computation required by the application itself; and the other considers the most significant challenging problems that exist for the application.

From the perspective of the amount of data and the computations required, AM Middleton considers data-intensive computing tasks as a class of parallel computing tasks that use data-parallel methods to process large amounts of data, which are typically terabytes or petabytes in size and are often referred to as big data. Computing tasks that spend most of their execution time on computational requirements are considered computationally intensive, while computing tasks that require large amounts of data and spend most of their time on I/O and data processing are considered data intensive (Middleton 2016).

In terms of the most significant challenging problems of the application, Kleppmann argues that if data problems—including the amount of data, the complexity of the data, and the frequency of data exchange—are the main challenges for a task, then the task is considered as data intensive. Conversely, if computing is the main challenge, then the task is computationally intensive (Kleppmann 2017).

Both definitions are widely accepted by the industry. We argue that these two definitions are equivalent. When the time spent on processing data accounts for a large portion of the overall completion time in a task, the problems associated with data processing must become a bottleneck for the entire task, and the bottleneck device is a key limiting factor in improving the performance of the overall system (Harchol-Balter 2013). Therefore, this paper argues that data-intensive tasks are a class of applications where the size of data is massive, the variety of data is rich, the

Fig. 1 Application Distribution of HPC TOP100



data operations are frequent, and where the data processing becomes a critical bottleneck.

The evolution of applications into a data-driven paradigm requires supercomputing platforms to make adaptive developments. The future development trend of supercomputing platforms is focused on the explosive increasing volume of data for computational processing tasks, diversification of computing power, higher reliability and availability requirements, and deeper integration of supercomputing centers with data and intelligent computing centers.

First, the amount of data processed by supercomputing platforms increases significantly. Further, data types are becoming more diverse. In addition to the structured data commonly used in traditional supercomputers, unstructured data, such as image, video, and voice data, are increasingly being used. The ultimate goal of supercomputing is to better serve academia and industry (Hager and Wellein 2010). Since 1980, supercomputing applications have started to enter the commercial area, and typical applications such as gene sequencing (Schmidt and Hildebrandt 2017), earth simulation (Wang and Yuan 2020), engineering design (Henz et al. 2018), and drug development (Puertas-Martín et al. 2020) have an intense demand for accessing large-scale data. In particular, in recent years, the rapid development of deep learning, reinforcement learning, and other related AI technologies has greatly contributed to performance breakthroughs in image recognition, speech recognition, time series prediction, and other applications. AI technologies support the rapid development of smart cities, autonomous driving, precision medicine, and other applications (Zhong et al. 2019). In addition, traditional compute-intensive tasks on HPC, such as weather prediction, are gradually being transformed to incorporate new AI technologies, which are more data-intensive. AI techniques often benefit from fitting accurate models to large-scale data. For example, the ResNeXt WSL (Orhan 2019), the most powerful pre-training model for image classification to date, developed by Kaiming and he's team, has 800 million parameters and uses 940 million images for training. The development of AI technology has brought huge demands and challenges for storage capacity. According to the forecast data of the global HPC market growth rate in 2020 by Hyperion Research, an international authoritative analyst firm, the storage market is growing at the first rate, much higher than the overall market average growth rate. In addition, supercomputing systems are increasingly demanding the performance of metadata. The VPIC, a representative scientific computing application, needs to save the particle track into a file for each process. This requires storage systems to have the ability to handle millions of file creation, writes, and other operations simultaneously.

Second, the computing power of supercomputing platforms has increased significantly and has diversified.

Existing supercomputing platforms can often provide hundreds of Peta FLOPS of computing power; and in recent years, they will break through Exa FLOPS of computing power. Therefore, a few single tasks can consume all the cluster computing power, and most scenarios are multi-task concurrent, such as the Shanghai Jiaotong University Supercomputing Center with 100 PFLOPS of computing power. Further, the average number of concurrent tasks is approximately 50. These tasks have different storage requirements, and some have high bandwidth requirements, and some have high IOPS requirements. At the same time, the dramatic increase in computing power has also led to an explosion in the size of data output from scientific computing—requiring HPC storage systems to be able to handle a variety of different load capacities. With the rise of diverse computing power such as GPUs, FPGAs, and various AI chips, efficient data transfer from the storage side to the accelerated computing device side has become a serious bottleneck in traditional HPC storage architectures.

Third, ever-rich applications have become more demanding in terms of the reliability and availability of supercomputing platforms. Traditional supercomputers were mainly used for scientific research, which had a high tolerance for failures and could accept occasional machine failures as long as the correct results could be obtained. Currently, supercomputers are more often used in production systems, which have high requirements for the reliability and availability of results and processes, and have a very low tolerance for data loss.

Fourth, supercomputing centers, data centers, and AI computing centers are gradually integrating to provide more diversified services; such as AI computing, cloud computing, big data, virtualization, and disaster recovery. For example, Tianhe 2A is no longer limited to providing scientific computing services, but also supports big data analysis services (Orion big data analysis platform). The biggest problem they face is the flow of large amounts of data. According to a study by the Internet Data Center (IDC), the total amount of global data reached 44 ZB in 2020 (1 ZB = 1 billion TB = 1 trillion GB), while the amount of data generated by China alone reaches 8 ZB, accounting for approximately one-fifth of the total global data volume. China's total data volume is expected to grow to 48.6 ZB by 2025, accounting for 27.8% of the world's total data volume. These data include file storage for supercomputing applications, block storage for virtualization, object storage for AI, HDFS storage for big data, and archival storage of supercomputers, all of which are fragmented. Further, accessing the data flow is the biggest challenge today.

The change in supercomputing platforms presents a significant challenge for storage systems. The increasing volume and richness of data types, the balance of bandwidth and I/O capabilities, the guarantee of processing and

results reliability, and data flow are the challenges faced by the supercomputing industry. These four challenges place higher demands on storage systems. Therefore, storage systems have become a key bottleneck in data analysis. We argue that future supercomputing platforms need to balance computing-intensive tasks and data-intensive tasks, as the contradiction between growing data processing needs, and relatively poor storage capacity deepens.

Through analysis of the characteristics of data-intensive applications and the development of supercomputing, we argue that storage systems for supercomputers need to focus on the following key issues.

- Complex data-type supporting problems. HPDA applications have complex data types, including structured data, semi-structured data, and unstructured data. The datasets may consist of large files, or many small files. How should the platform handle different types and sizes of data to fully utilize the platform's computational potential is a key challenge.
- Mixed-load optimization problems. How to ensure that the storage system performs well under various HPDA application workloads.
- Multiprotocol support and interoperability issues. The storage system should simultaneously support multiple data access protocols that HPDA applications may use, such as NFS, HDF, POSIX, SMB protocol, and Amazon S3, etc.

This paper first briefly introduces the key concepts of data-intensive supercomputing. We then introduce the typical data-intensive application requirements through both commercial and scientific supercomputing applications, and reveal the problems of data silos in the next generation of supercomputing in the face of multiple computing power paradigms in complex application scenarios. Next, we introduce the most representative supercomputers in the TOP 500 and IO 500 lists. Finally, we present the challenges faced by data-intensive supercomputers, such as data transfer bottlenecks, data silos, resource seizure, and data management.

2 Basic Concept

2.1 Supercomputers and Clusters

In most computing contexts, a supercomputer refers to a computer capable of performing high-speed operations that cannot be handled by ordinary personal computers, with specifications and performance much more powerful than those of personal computers, and in this context supercomputers and high-performance computing systems are different expressions of a unified concept. Supercomputer can

usually be divided into five categories: parallel vector processors (PVP), symmetric multiprocessors (SMP), massively parallel processors (MPP), distributed shared storage multiprocessors (DSM) and clusters (Rajak 2018). The only three massively parallel processors left are the "Fugaku" (Shimizu 2020) supercomputer developed by Japan's RIKEN and Fujitsu, the "Sunway TaihuLight" (Gao et al. 2021) supercomputer developed by China's Wuxi National Supercomputing Centre and the "Perlmutter" supercomputer developed by American National Energy Research Scientific Computing Centre.

In summary, clusters are one type of supercomputer and have taken over the mainstream of supercomputers. A cluster is an inexpensive, easy to build, and highly scalable parallel computer system. It consists of multiple homogeneous or heterogeneous independent computers interconnected by a high-performance network or local area network to collaborate on a specific parallel computing task. Nowadays, many enterprises' data centers or cloud computing centers are implemented in the form of computer clusters, such as Huawei Cloud Computing Centre, Baidu Cloud Computing Centre, Ali Cloud Computing Centre, Google Data Centre, Amazon Cloud Computing Centre, etc.

2.2 High Performance Data Analysis

In recent years, high-resolution simulations, observations from advanced instruments, and a vast network of sensors and edge components have provided a vast and constant flow of data—driving a variety of new scientific discoveries and technological advances. In such an evolving environment, where the scale, velocity, and complexity of data continues to expand, scientific advances rely on the practical ability of analytical tools to effectively extract and process valuable information from the sea of data. For example, real-time, high-frequency stock trading, as well as highly complex analytical problems found in scientific research, all require the dual needs of both high-volume and high-speed processing. Therefore, there is an urgent need to develop new methods and strategies to deal with such extreme scale data-centric scenarios.

HPDA has been created to solve the above problems. High-performance data analytics combines supercomputing with data analysis. It uses the parallel processing of supercomputers to run powerful analysis software at speeds in excess of a trillion floating point operations, which allows large datasets to be examined quickly, and valuable information to be found.

The development of high-performance data analytics has been attributed to technological advances in supercomputing, but also to the development of data management techniques and tools designed to address the challenges of big data applications. To leverage high-performance data

analytics for scientific development and commercial applications, both supercomputing and data-intensive analytics research are considered to be two of the most fundamental and complementary aspects of high-performance data analytics (Usman et al. 2020).

3 Support of Mainstream Supercomputer Systems for Data Intensive Applications

3.1 Fugaku

The Japanese ARM-based Fugaku (Shimizu 2020; Kodama et al. 2020) supercomputers are recognized as the most powerful machines in the world. As of June 2020, the system held the 1st rank in the TOP 500 list, which lists the world’s 500 most powerful supercomputers. Fugaku emphasizes multiple capabilities, such as computation and storage, and also performs well on other benchmarks—such as Graph500 (Nakao et al. 2020) and HPL-AI (Kudo et al. 2020).

Fugaku uses a A64FX CPU based on the Armv8.2-A architecture (Odajima et al. 2020; Zhang et al. 2020). The A64FX is a general-purpose processor with a high computational performance and memory bandwidth. Figure 2 illustrates the architecture of A64FX with 48 compute cores, and two or four auxiliary cores. The auxiliary core handle interrupts from the operating system and communications. The users can select either 2.0 or 2.2 GHz as the CPU frequency for a given job. The A64FX has four core memory groups (CMGs) consisting of 12+1 cores and HBM2 (8 GiB, 256 GB/s). Because the four CMGs are connected by a network-on-chip (NOC), the number of processes per node should be a divisor of four for general applications.

Fugaku uses Tofu Interconnect (Schmidt and Hildebrandt 2017) for network interconnection, whose topology is a six-dimensional grid/circle. The locations of the nodes can be specified using the XYZ abc axes. The ac axes are grid topologies consisting of only two nodes, and the other axes are circular topologies. The dimensions of the abc axes are fixed at (a, b, c) = (2, 3, 2), and the dimensions of the XYZ axes depend on the system. Because (X, Y, Z) = (24, 23, 24) for Fugaku, the total number of nodes is 158, 976. Fig. 3 also




Fig. 2 The Architecture of Fugaku




Fig. 3 The Architecture of A64FX

shows that the XY Zb axes use two ports, and the ac axes use one port. A64FX has six tofu network interfaces (TNIs) and is capable of communicating at 6.8 GB/s in six directions simultaneously.

Fugaku’s storage system is divided into three tiers, as shown in Fig. 4. The first tier serves as a dedicated high-performance area for job execution. The second tier provides a large-capacity shared area for users and jobs. The third tier provides a commercial cloud storage.

The main problems of Fugaku are insufficient multi-protocol support, poor metadata performance, and the inability to support million-scale file operations simultaneously. Moreover, it cannot support multiple computing power.

3.2 Summit

The Summit (Kahle et al. 2019; Wang et al. 2020; Luo et al. 2020; Hernández et al. 2020; Womble et al. 2019) consists of 256 racks with 18 IBM Power System AC922 computing nodes. Figure 5 shows a node architecture consisting of two IBM POWER9 CPUs and six NVIDIA V100 GPUs connected using the NVLink 2.0. The CPUs and GPUs are

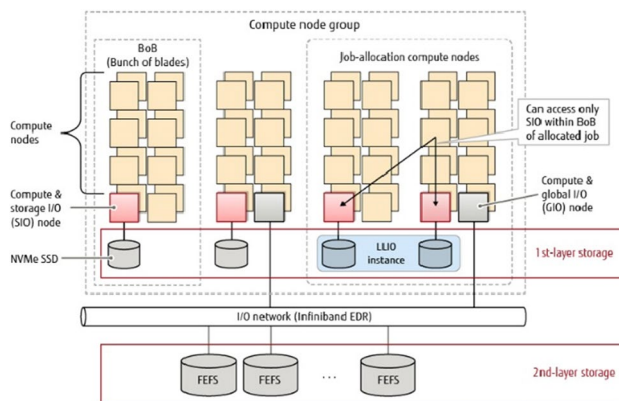


Fig. 4 The Storage Architecture of Fugaku

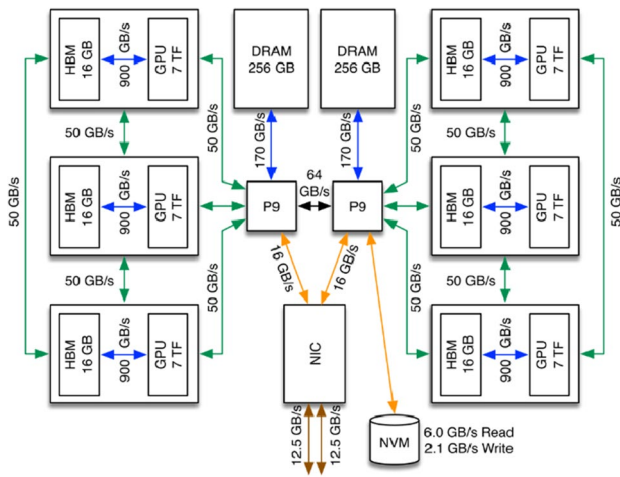


Fig. 5 The Computing Node Architecture of Summit

connected to over 5 TB of addressable memory through a high-bandwidth NVLink. The Merano enhanced data rate (EDR) and infinite bandwidth (IB) link the nodes in a non-blocking fat tree topology. The intra-node CPU bandwidth is up to 64 GB/s and the GPU bandwidth is up to 50 GB/s. Each node has 1.6 TB of local high-speed temporary storage (NVMe) with 6.0 GB/s read bandwidth, and 2.1 GB/s write bandwidth. A large, high-performance, global temporary file system is provided as a fully centric GPFS.

Summit uses the IBM Spectrum Scale platform to provide storage capacity. Spectrum Scale offers excellent scalability, flash based acceleration performance, and automatic policy-based storage tiering (from flash, disk to tape). The Spectrum Scale provides a single file namespace for all data, Swift/S3 object services, and interfaces to storage services such as Openstack and Big Data. It is a high-performance, highly scalable storage software for file, object, block, and big data analytics, with the ability to store, run, and access data anywhere. Spectrum Scale has added Hadoop big data and objects to GPFS, and enhanced storage capabilities with rich features such as cache acceleration, lifecycle management, unified namespace, and multi-site. It also supports features such as encryption, hardware, and software decoupled deployment—and provides rich access interfaces such as file interfaces (POSIX, NFS, CIFS),

3.3 Sunway TaihuLight

The Sunway (Gao et al. 2021; Fu et al. 2016) is a converged architecture. One part is oriented toward high-speed computing systems for traditional supercomputing, and this part favors computational power. The other part is oriented toward auxiliary computing systems for emerging

applications such as big data, and this part takes storage into account. However, in general, Sunway still focuses on computing power, and the rest of it exists only as an auxiliary computing system. The system architecture is shown in Fig. 6.

The high-speed computing part of the system is composed of the SW26010 processor with an autonomous Sunway instruction set, which uses a heterogeneous multicore architecture, a cluster of on-chip computing arrays, and a combination of parallel shared storage. The architecture diagram is shown in Fig. 7.

Sunway’s storage system (Chen et al. 2020) consists of online storage system and near-line storage system. The online storage system consists of storage service nodes with SSDs, high-performance dual-controller fiber serial SCSI disk arrays, and metadata service nodes, which are responsible for providing users with high-speed and reliable online data storage access services with an I/O aggregation bandwidth of 220 GB/s. The near-line storage system consists of metadata service nodes, storage service nodes, and high-capacity fiber storage area networks, which are responsible for providing cloud-oriented and user-service storage services, as shown in Fig. 8.

The storage scale and bandwidth provided by Sunway is much smaller than other systems such as Summit and Fugaku. It only supports POSIX interfaces and does not optimize the storage performance for multiple computing power scenarios. Therefore, it is difficult to adapt to data-intensive application requirements.

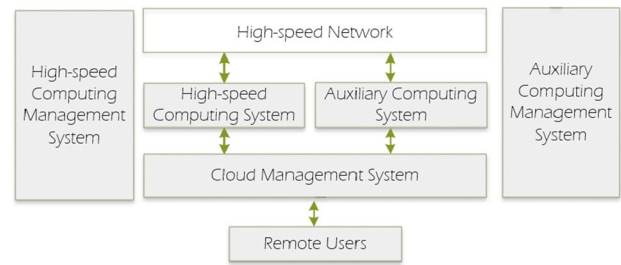


Fig. 6 Architecture of Sunway




Fig. 7 Architecture of SW26010

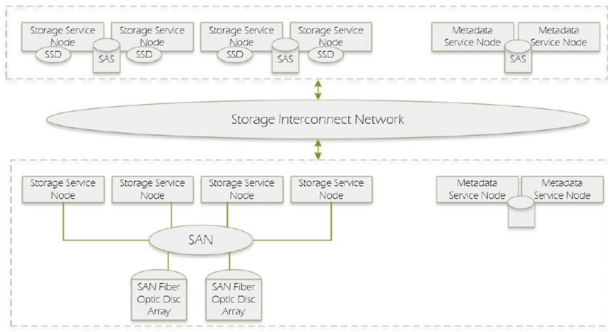


Fig. 8 Storage Architecture of Sunway

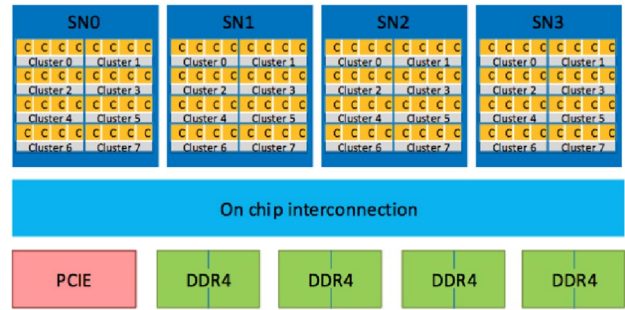


Fig. 10 The Architecture of Matrix-2000 accelerator




Fig. 9 The Node Architecture of Tianhe-2A

3.4 Tianhe-2A

The Tianhe-2A (Lu et al. 2020) consists of two parts: a computational processing module (CPM) and an accelerated processor unit (APU) module. The CPM module has four CPUs, and the APU is equipped with four Matrix-2000 accelerators. The logical architecture of Tianhe-2A is shown in Fig. 9.

The two CPUs are connected by an Intel QPI, each with four memory channels and eight memory slots. They are connected to dedicated NICs and Matrix-2000 accelerators with eight memory channels each by x16 PCI Express 3.0. In a compute node, the CPU is equipped with 64 Gb DDR3 memory, and the accelerator is equipped with 128 Gb DDR4 memory.

Matrix-2000 (Fig. 10) is configured with four super nodes (SNs) that are connected by a scalable on-chip

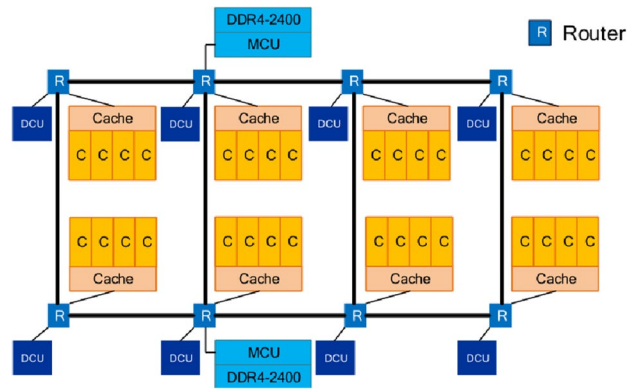


Fig. 11 On-chip network topology

communication network. Each SN has 32 compute cores, meets cache coherence requirements, supports eight DDR4-2400 channels, and integrates x16 PCI Express 3.0 endpoint ports. The compute core is an 8 to 12 level reduced instruction set computer (RISC) pipeline that extends the 256-bit vector instruction set architecture (ISA).

The network-on-chip (NoC) topology within an SN is a 4 x 2 grid (Fig. 11). The NoC has a total of eight routers, each connected to a cluster consisting of four compute cores, and a direct control unit (DCU). Each SN has three FIT link ports, and each port is connected to other SNs. In Matrix-2000, four SNs were connected by a point-to-point interconnect network.

Tianhe-2A has improved bandwidth and multi-protocol interface support to better face data-intensive application requirements. However, there are still shortcomings in the optimization of multiple computing power and metadata performance under massive file operations.

3.5 Perlmutter

Perlmutter is a next-generation supercomputer deployed by the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory. The compute

nodes use AMD EPYC 7763 processors and NVIDIA A100 SXM4 GPUs, with a total of 706,304 processor cores. The storage system uses the Lustre file system with 35 petabytes of available capacity and an all-NVMe connection, which improves read and write bandwidth by up to 15 times and 9 times respectively compared with the previous generation Cori system.

In its first participation in the TOP500 ranking in June 2021, Perlmutter reached the top 10 of both the TOP500, GREEN500 and HPCG rankings, achieving nearly 15% of its peak processing power using no more than 10% of the processor cores of the Fugaku supercomputer.

Perlmutter features extremely high bandwidth and IOPS due to its 35PB all-NVMe file system based on HPE Cray E1000(as shown in Fig.12), which supports Lustre, Grid-RAID, Idiskfs, Slingshot and LNet multi-track transport. Meanwhile, Perlmutter uses the most advanced NVIDIA A100 GPU available as an accelerator with a custom tensor processing unit to shine in data-intensive tasks such as AI and HPDA.

3.6 Exascale supercomputers

Frontier is a new generation of exascale supercomputers with peak FLOPS exceeding 1.5 EFLOPS, which was developed by Cray and AMD with the permission of the U.S. Department of Energy. The storage system of the frontier is similar to Summit, still using Luster as the parallel file system, but has 2-4 times the storage performance and capacity than Summit.

The current Chinese prototypes of exascale supercomputers mainly include the Sunway exascale prototype, the Dawning exascale, and the Tianhe-3 exascale prototype. All three completed their development and deployment in 2018.

Among them, the Tianhe-3 exascale prototype was ranked first in the latest GRAPH500 SSSP ranking in June 2021, demonstrating its strong ability to support data-intensive applications.

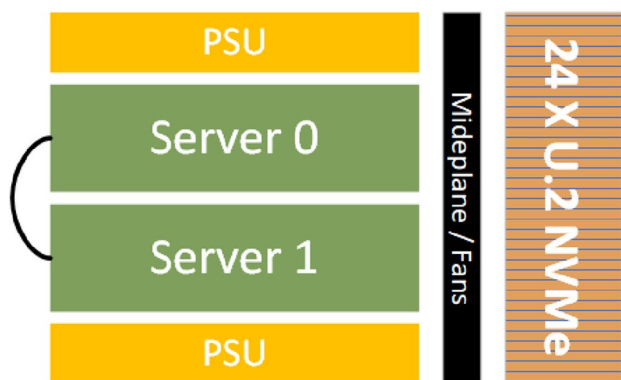


Fig. 12 Architecture of E1000

Similar to the Tianhe-3 exascale prototype, the main changes in the Sunway exascale prototype and the Dawning exascale supercomputers are the increase in computing power, with little change in the storage architecture.

4 Representative Parallel File Systems

Parallel file systems are widely used for supercomputing storage. It is able to support multiple simultaneous accesses by multiple threads, mapping multiple files and directories to many devices at the same time. We present the current development of parallel file systems using the representative BeeGFS(Abramson et al. 2020), CephFS(Zhang et al. 2019b), IBM Spectrum Scale(Iannone et al. 2019), and Lustre(Braam and Zahir 2001) parallel file systems, which are widely used in supercomputing, as examples.

4.1 BeeGFS

BeeGFS is a parallel cluster file system with POSIX interface designed for all performance-oriented application scenarios, including HPC, AI and deep learning, life sciences, and oil.

As shown in Fig. 13, the BeeGFS architecture consists of metadata servers, storage servers, and management servers. BeeGFS implements the separation of object data and meta-data. Object data is the data that the users want to store and

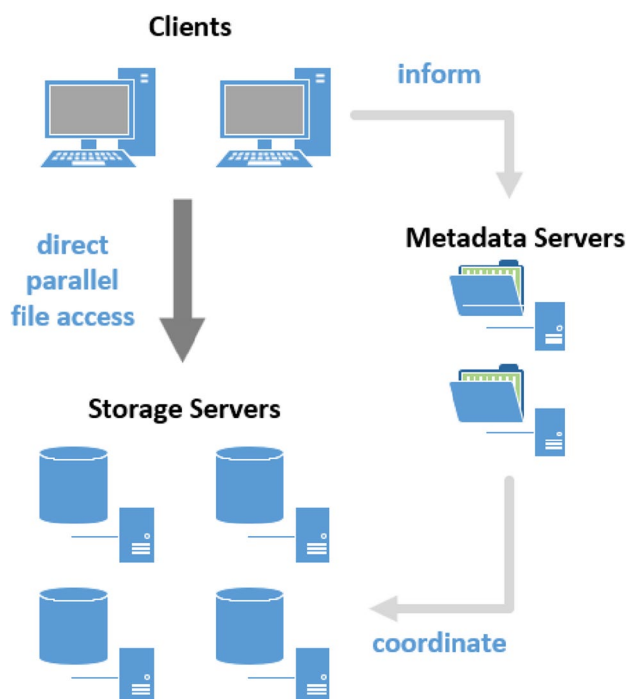


Fig. 13 Architecture of BeeGFS

is maintained by the storage servers. Metadata is the data about data, including access permission, file size and storage location, managed by metadata servers. The most important part of metadata is how to find the specific file from multiple file servers. In this way, the clients can directly talk to the storage servers to retrieve information after obtaining the metadata of the specific file or directory. Compared to Lustre, BeeGFS is easier to use and requires less maintenance, and the number of storage and metadata servers can be flexible and scalable.

4.2 CephFS

The Ceph File System (CephFS) is a POSIX-compliant distributed file system that provides access to files on the Ceph storage cluster and uses the Ceph storage cluster to store user data. CephFS stores data and metadata separately, providing excellent performance, reliability and scalability for upper-level applications.

As shown in Fig. 14, Ceph’s system architecture is divided into three parts: Clients client, Metadata server cluster (MDS) and RADOS distributed object storage. MDS caches and synchronizes distributed metadata, which allows clients to securely store metadata about the file system and maintain consistency. RADOS stores data and metadata as object, performing other critical functions. CephFS is built on RADOS and inherits the fault tolerance and scalability of RADOS, supporting redundant copies and high data reliability. Clients query and update the corresponding metadata of MDS through RPC, and then get the file data directly from RADOS based on object information. Compared with

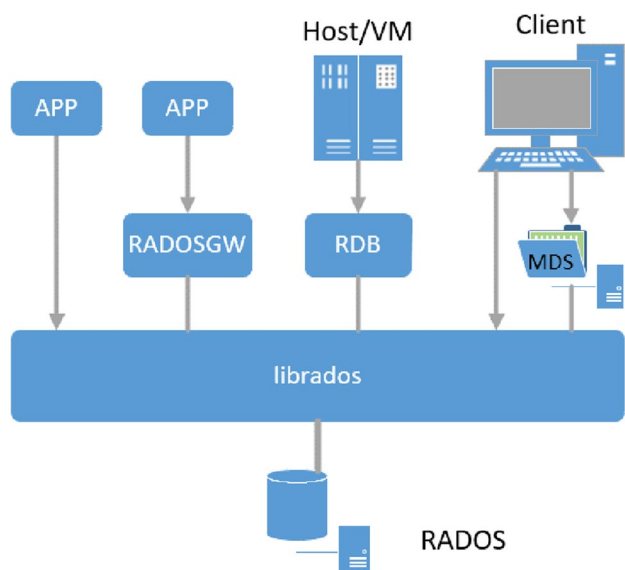


Fig. 14 Architecture of CephFS

Lustre, which is suitable for reading and writing large files, CephFS is more suitable for unified small file storage.

4.3 IBM Spectrum Scale

Spectrum Scale is a distributed, shared, parallel cluster file system developed by IBM. It ensures that all nodes within a resource group can access the entire file system in parallel. The file system operation provided by Spectrum Scale supports both parallel and serial applications, allowing parallel applications on any node to access the same file or different files at the same time, and providing a unified naming interface.

As shown in Fig. 15, the Spectrum Scale architecture is divided into three layers: storage layer, server layer, and client layer. The storage layer divides storage to server nodes, which configure and format these storage disks to form the Spectrum Scale parallel file system format. Finally, the server node communicates with the client nodes through Ethernet. NSD is a storage device visible to Spectrum Scale cluster for creating file systems, and NSD server is capable of providing I/O servers for specific NSD access. Spectrum Scale can span across all hosts within a cluster, providing better system performance and data consistency, and is a highly available file system.

4.4 Lustre

Lustre is a cluster-oriented storage system, which is an open source, global single namespace, POSIX-compliant distributed parallel file system. Lustre is highly scalable and high-performance, supporting tens of thousands of client systems, petabytes of storage capacity, and hundreds of GB of I/O throughput.

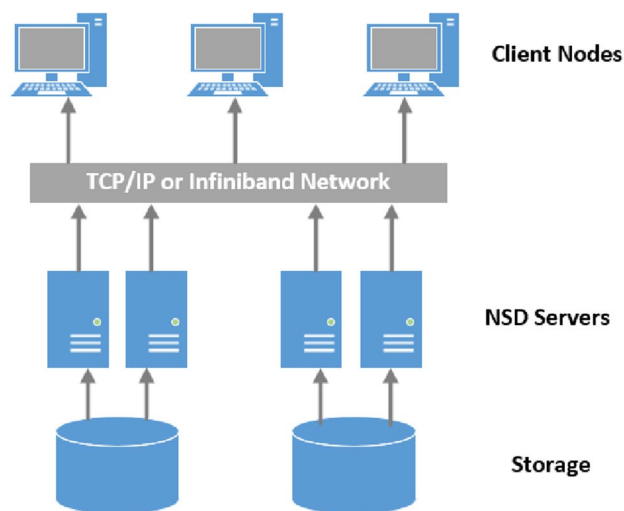


Fig. 15 Architecture of IBM Spectrum Scale

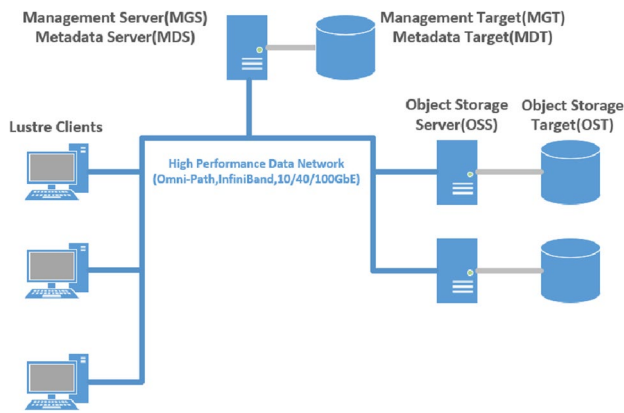


Fig. 16 Architecture of Lustre

As shown in Fig. 16, the Lustre file system architecture is a distributed, object-based scalable storage platform built on a computer network. The architecture consists of the following components: The Metadata Server (MDS) and the Object Storage Server (OSS), which provide namespace operations and I/O services with large amounts of data respectively. The Management Server (MGS) is a global registry of configuration information and is functionally independent of any one Lustre instance. The Client provides an interface for Lustre interaction between the application and the Lustre service. The software on the client provides a consistent POSIX interface for the end-user applications. The network protocol LNet provides a communication framework for binding services together.

5 Analysis of Typical Data-Intensive Application Requirements

With the rapid development of IoT technology, the amount of data that needs to be processed by supercomputing systems has exploded. Benefiting from the increasing capabilities of scientific observation instruments, millions of sensors around the world are observing and recording cosmic, meteorological, biological, physical, and chemical processes in real time—and are generating large amounts of observational data. In addition, large amounts of computing equipment around the world are running various scientific simulations that generate huge amounts of scientific data. The U.S. Department of Energy (DOE) studies on scientific data management indicate that “we are entering an information-led era. The ability to ride the information wave will be a marker that distinguishes whether it is the most suitable for science, business, and national security.”

Data-intensive computing heralds not only the evolution of informatics, but also a revolution in the way researchers collect and process information. Data-intensive computing

has become a common research problem for industry and academia at a time when large amounts of data and associated applications are emerging for individuals, enterprises, and industries. Researchers are driving the focus of data-intensive computing by focusing only on large datasets for broader problem domains (Lin et al. 2017; Dong et al. 2018), and are more involved in all aspects of real life (Seal et al. 2020; Hohman et al. 2019). Typical data-intensive computing applications are summarized in the following subsections.

5.1 Scientific and Engineering Applications

Scientific computing is superior in modern scientific research and engineering technology, where a large number of complex computation problems are encountered. Scientific computing is the process of applying computers to mathematical computation problems, which has gradually evolved into an important research area in the current data explosion era. It is undeniable that computing ability has become one of the most important symbols for measuring the level of national scientific and technological development. Data-intensive applications have become a necessary tool for scientists to discover new problems in frontier research areas.

Traditional applications, such as earthquake prediction, petroleum exploration, climate and environment prediction, and astrophysics are often simulated by dividing them into grids. With the continuous breakthrough in data collection and modeling technology, the grid is becoming increasingly more refined, which puts higher requirements on the calculation of numerical resolution and the processing ability of huge data in the process of scientific research. The scale of artificial intelligence and big data applications is so large that traditional supercomputing can no longer meet the demand for data storage and calculation (Table 1).

In the climate and environment field, several research areas are involved; such as meteorology, environmental science, and physical chemistry (Zhang et al. 2019a; Torbicki 2018; Shuman et al. 2001; Anh Khoa et al. 2021; Guo et al. 2020; Kappe et al. 2019; Belair et al. 2019; Massonnet et al. 2016; Kurth et al. 2018). Many meteorological satellites, weather sensors, and weather stations have been built in various countries to collect climate data, which has significantly increased the speed of collection, and also generated a large amount of information—leading to high storage capacity requirements. At the same time, the search time for useful data is inevitably prolonged, making it difficult to meet the standards for the access speed of weather data in storage systems. Berkeley National Laboratory’s research, which applies machine learning methods to extreme weather events, is gaining international attention. In 2018, a team led by Prabhat proposed using AI to predict how extreme

Table 1 Analysis of typical scientific application requirements(Lin et al. 2017; Dong et al. 2018; Seal et al. 2020; Hohman et al. 2019)

Application name	Exascale deep learning climate analysis	Practical evolutionary earthquake urban problem solver	Genomics-based analysis of drug addiction	Quantum dynamics simulation
Storage Scale	3.5TB	100TB to 1PB	10TB	TB
IOPS Requirement	1500 images per node	132665 degrees of freedom	2.36 billion	20 peta
AI Library support	Required	Required	Required	Required
Big Data Interface Support	Not required	Not required	Required	Not required
Storage Bandwidth Requirement	11.98 GB/s	10.2 GB/s	900 GB/s	975 GB/s
Data Reliability Requirement	High	High	High	High
Data Security Requirement	Medium	Medium	High	Medium
Real-time Requirement	High	Low	Low	Low
Fundamental Algorithms	DeepLabv3+	Artificial Neural Networks	Markov Clustering	Lattice Gauge Theory

weather will change in the future; that is, using variants of Tiramisu and DeepLabv3+ neural networks with supercomputers to identify extreme weather patterns from high-resolution climate simulations. Based on 3.5 TB of weather data, by making large-scale use of the dedicated tensor kernel built into Summit's NVIDIA GPUs, the team achieved a peak performance of 1.13 Exa FLOPS. It is the fastest deep-learning algorithm reported to date (Kurth et al. 2018).

In the field of earthquake prediction, seismic data are stored in large blocks of data and are developing toward refinement (Allen 2012; Minson et al. 2018; Mapar et al. 2012; Luo and Paal 2021; Yu et al. 2021; Amin and Ahn 2021; Hu et al. 2019; Ichimura et al. 2018). Scientists are creating more accurate finite element grids to represent the geological properties of the earth's crust, and to simulate the effects of geological disturbances and the probability of earthquakes. Tsumura Ichimura's team at the University of Tokyo proposed the application of Artificial Neural Networks and hybrid precision algorithms to accelerate the simulation of seismic physics in urban environments. A peak performance of 14.7% was achieved using 4096 nodes in Summit, and coupled ground and urban structure vibrations during large earthquakes into a unified simulation—showing a high scalability. As historical seismic data and newly observed data continue to increase and new models are being expanded, traditional supercomputing is facing new challenges. By integrating analytical methods of data-intensive computing and large-scale storage systems, we can dig deeper into historical data and explore new research methods—so as to accurately predict the locations of possible geological disasters, respond in time, and reduce economic losses.

In the field of life science, with the expansion of biological data, researchers need a unified platform for data storage and processing. Existing systems such as Dryad can provide efficient data retrieval, but there is no efficient data management service (Mazzucco et al. 2017; Chien et al. 2015;

Kumar et al. 2008; Lv et al. 2020; Sejdic et al. 2018; Joubert et al. 2018). With the help of supercomputers, storage systems can be developed for the integration and analysis of complex data. Many classical bioinformatics algorithms and processes can be targeted and optimized in combination with machine learning and other techniques. In 2018, the ORNL team developed a new "CoMet" algorithm (Joubert et al. 2018). They use Custom Correlation Coefficient network topologies to make dense vector similarity calculation and use Markov clustering and principal component analysis of the Laplacian transform of a network adjacency matrix to create SNP set. This algorithm allowed supercomputers to process large amounts of genetic data—four to five times larger than state-of-the-art technology—and to identify genes and effective treatments for greater vulnerability to pain and opioid addiction, achieving a peak throughput of 2.36 Exa FLOPS. In addition to processing genetic information, CoMet is currently being used in bioenergy and clinical genomics. Furthermore, the team designed an artificial intelligence system based on this algorithm to utilize 25,200 of Summit's NVIDIA Volta GPUs. This algorithm automatically designs an optimal deep learning network to extract structural information from raw atomic resolution microscopy data, enabling a major step forward in the field of material and genetic research. These applications are scalable and, to some extent, pose new challenges to the reliability and security of supercomputers.

In the field of quantum physics, accurate testing of the Standard Model in the nuclear environment requires a quantitative understanding of quantum chromodynamics, to which machine learning methods are being extensively applied (Hush 2017; Musser 2020; Gao et al. 2018, 2020; Li et al. 2019; Liu et al. 2019; Zhang et al. 2016; Duan et al. 2018; Chang et al. 2018). The importance of genetic algorithms makes them a benchmark for theoretical determination. The LLNL and LBNL teams used the Department of Energy's newest supercomputers—LLNL's Sierra and Oak

Ridge's Summit—to develop an improved algorithm that could determine neutron lifetimes more accurately. They use Lattice Gauge Theory and calculate nuclear observables such as neutron lifetime and develop an improved algorithm (utilize GPU Direct RDMA for inter-node transfers, Communication Autotuning, Management and Backfilling of Tasks) that exponentially decreases the time-to-solution. The team was able to achieve a sustained performance of approximately 20 peta FLOPS, which accounts for approximately 15% of Sierra's peak performance. Researchers say that just one percent more accuracy is what is driving physics forward. As a result, the high demand for arithmetic power and data storage management is driving the development of supercomputers to be more data-intensive.

In the field of oil exploration, the increasing difficulty and cost of extraction has become a serious challenge, and we have to seek innovations and breakthroughs in directional and geosteering drilling, deepwater drilling, enhanced recovery and the booming development of unconventional oil and gas. Seismic exploration is fundamental to oil and gas exploration. Shipilova et al. (2020) proposed the Matrioshka OMP algorithm, which uses homogenous seismic acquisition to construct subsurface images without actually penetrating the earth's crust. In addition, optical remote sensing has become a key discipline for analysing large-scale exploration areas. Gianinetto et al. (2016) proposed that remote sensing can be used to detect direct (hydrocarbon) and indirect (variable) seepage signals, and that the large microseepage signals identified are considered spectral anomalies that can verify the presence of hydrocarbon accumulation and sedimentation centres. Due to the high mixture and complexity of ground source data, how to efficiently analyse noise effects and accurately image them, and how to maximise the effectiveness of real-time monitoring data, requires data-intensive supercomputing to provide efficient arithmetic power and stability support.

5.2 Commercial Applications

High-performance computing applications are gradually being widely used in the commercial field. HPDA is a typical data-intensive application, which mainly involves smart cities, financial analysis, intelligent transportation, and many other fields. These applications require the collection and maintenance of massive datasets and intensive data analysis (Table 2). As the technology for collecting and storing data becomes more widespread and economical, an increasing number of data-intensive computing problems are emerging, and the need for HPDA applications is emerging in various industries.

In the field of credit card fraud detection (Anbuvizhi and Balakumar 2016; Dai et al. 2016; Li et al. 2020; Koppers et al. 2017), credit card fraud seriously affects normal transactions. Therefore, it is necessary to use credit card fraud detection technology to detect fraudulent transactions and guarantee the safety of the credit card business. With the development of e-commerce, credit card transactions have been growing explosively, showing the characteristics of massive volume and high frequency. Its real-time transaction records are up to millions per second, and the size of historical transaction data stored in the system reaches the PB or even the EB level. The massive data storage requirements and high-frequency real-time detection bring unprecedented performance challenges to credit card fraud detection. For unbalanced and asymmetric credit data, the use of HPDA technology can provide a more objective and fair judgment and prediction ability than human subjective judgment, greatly reducing the occurrence of fraud and breach of contract in transaction activities. Forough and Momtazi (2021) use an ensemble model based on sequential modeling of data using deep recurrent neural networks and a novel voting mechanism based on artificial neural network to detect fraudulent actions.

Table 2 Analysis of typical commercial application requirements Shipilova et al. (2020); Gianinetto et al. (2016)

Application name	Credit Card Fraud Detection	Animation rendering	Smart City	Web search and knowledge capture
Storage Scale	PB to EB	GB to TB	PB	ZB
IOPS Requirement	Millions of times	30 frames	Millions of times	Tens of millions
Storage Bandwidth Requirement	100 MB/s	100 GB/s	100 MB/s	40 MB/s
AI library support	Required	Not required	Required	Required
Big Data Interface Support	Required	Not required	Required	Required
Data reliability requirement	High	High	High	High
Data Security Requirement	High	Medium	High	Medium
Real-time requirement	High	Low	Low	Low
Fundamental Algorithms	DRNN	RF, GBT, SVM, DL	SVM, Clustering, Ant Colony	DM, FM

In the field of animation production (Hong et al. 2016; Zhang et al. 2017), animation rendering is the core part of animation production, which is a typical HPDA application that transforms animation design into specific images mainly through the combination of model, light, material, shadow, and other elements. The core idea is to adopt and fine-tune multiple algorithms based on machine learning approaches, such as Random Forest (RF), Gradient Boosting Tree (GBT), Support Vector Machine (SVM), Deep Learning (DL). The rendering data generated by the 3D production software have high data-intensive characteristics. A highly fault-tolerant distributed cloud file system is typically used to achieve efficient storage and retrieval operations for large image frame data. It uses a cluster of thousands of machines to manage data and supports large datasets of GB to TB level for a single file—thus achieving efficient storage operations for large image frame data and producing richer and more realistic animation. The cloud file system contains name nodes that can manage the metadata of the entire file system, and data nodes that can store the actual application data files, thus providing a large data throughput.

In the field of smart city construction (Wang et al. 2021; de Assuncao et al. 2018; Mehmood et al. 2017; Tang et al. 2017), the data center integrated with IoT, data communication, and other technologies is the core of building smart cities, involving petabyte data volumes, while requiring maximum security of data and business. Methods of machine learning are widely adopted among this group of studies, such as support vector machines (SVMs), C-means fuzzy clustering, K-means clustering, hierarchical clustering, ant colony algorithm, etc. The storage and analysis of massive data information can effectively promote the construction of smart cities. Through HPDA technology, valuable information can be mined from massive data, which in turn can provide help and guidance for smart city construction and solve the existing problems. By building a big data center to strengthen data processing and interaction effects, it can promote the sharing of resources among different smart platforms and improve the city management effect. For building traffic models, predicting disease outbreaks, and reasonably allocating education resources, etc.; massive information can be integrated together, and can be fully refined and mined by HPDA to help build a smart city.

In the field of internet applications (Jun et al. 2018; Sapountzi and Psannis 2018; Stavrinides and Karatza 2017), the data handled by the internet are characterized by being massive, fast updating, diverse, and heterogenous. As the largest database in the world, it provides various types of data and contains diverse applications. The existing mainstream application models of the internet are web applications and SaaS applications; which provide services including e-government, e-commerce, web information retrieval applications (search engines, knowledge sharing, internet

news, etc.), web communication and interaction applications (instant messages, social networks, etc.), and web entertainment applications (videos, online games, etc.). They always use map-reduced based distributed parallel data mining algorithms. With the increase in data volume and the change in data characteristics, traditional data storage and indexing technologies can no longer handle these massive, dynamic, and distributed data. Therefore, HPDA becomes a key technology to drive the further development of Internet applications.

5.3 Summary

Data-intensive supercomputing opens up new possibilities for both scientific and commercial applications. With the development of information technology, the cost of hardware has gradually decreased. At the same time, with the increase in data volume, the applications put higher demands on the storage capacity, arithmetic power, and network bandwidth of supercomputing. More efficient data exchange between servers, higher reliability, and security requirements are required. Storage and computing uncoupling technology use distributed storage and supercomputing, flexible configuration of computing and storage units in the hardware architecture to meet the expansion needs, and has a wide prospect of use—becoming the key to solving the aforementioned problems.

6 Trends in Data-Intensive Supercomputing

Data-intensive supercomputing uses multiple computing power and EB-level storage to enable the processing of massive amounts of data. Data-intensive supercomputers are mainly oriented toward data-intensive tasks, aiming to solve data I/O, load balancing, data flow, data reliability, and other data storage-related problems during the execution of data-intensive tasks, with the goal of maximizing user productivity (rather than maximizing computational resource utilization). We argue that the data-intensive supercomputers should have the following characteristics.

6.1 Multiple Computing Power

According to a report by the IDC, the AI server market in China alone will be \$3.196 billion by 2020, which reflects the importance of computing power for all industries. And with nearly 20 ZB of new data currently being generated each year, the demand for AI computing power doubles every 3.5 months. Similarly, OpenAI published research shows that the demand for computing power has increased by an average of 10 times per year between 2012 and 2018, respectively. In the future, with the continuous development

of 5G and big data technology, the volume of data will explode, and the application scenarios will become more diverse—with the scale of data rising to tens of millions. This development trend will make it difficult for traditional CPU-based computing power to meet development requirements. In addition, the linear development model of a traditional supercomputer is difficult to adapt to new environments. Computing architecture and computing power will be diversified in the future.

To achieve the convergence of multiple computing powers, it is necessary to break away from traditional architecture design, and solve the compatibility and efficiency problems caused by multiple computing devices with different architectures. It should be able to achieve efficient aggregation and on-demand definition of multiple computing powers, implement storage and computing decoupling technology to improve flexibility, realize the combination of computing and data to break the upper limit of storage, and to achieve scalability through open and interconnected network technology. Supercomputers should adopt an open network chip with an open architecture so that it can support current mainstream network chips. It should have an open network operating system to improve the flexibility of computing power and to meet the needs of cloud computing in terms of interfaces, software definitions, and rapid iterations of the network.

6.2 EB Level Storage

Traditional supercomputers treat increasing computing power as their fundamental goal. From 1976 to the present, the computing power of the most powerful supercomputer has increased by a factor of 4.21 billion. However, the storage capability has received insufficient attention and has become a new bottleneck. Therefore, HPC should focus on storage capability. Big data, AI, and other technologies are gradually merging with HPC, allowing valuable data to be exploited. The data being analyzed is gradually changing from cold data, to a large amount of dynamic and constantly changing hot data. IDC predicts that by 2021, the global market for HPC storage will reach \$14.8 billion. The emerging HPDA and HPC-based AI scenarios are predicted to grow rapidly at an annualized rate of 17% and 29.5%.

As data volumes continue to explode, the industry is evolving from compute-intensive HPC, to data-intensive HPDA. HPC faces changes in three areas. First, the expansion of massive data volumes requires storage to evolve to the EB level as soon as possible, and application scenarios are becoming more complex. To meet these challenges, data-intensive supercomputers should adopt more efficient high-density technologies, and provide total cost of ownership (TCO) optimization. The main challenge for data storage is maintaining more valuable data per byte; and high-density

technology will be the best solution. In addition, data-intensive supercomputers will have to focus on optimizing software algorithms—such as compression algorithms—to improve the efficiency of storage. Workloads will become more diverse, and move towards mixed loads. For example, in oil and gas exploration, seismic data requires a higher bandwidth for processing and a higher IOPS for interpretation, each with a different load, which requires a storage design that can support multiple loads simultaneously. Data-intensive supercomputers should also support more data types to accelerate the transfer of data.

In order to cope with the increasing demand for storage capacity, data-intensive supercomputer architects need to design efficient storage management methods—in particular, when the number of objects to be managed is as high as ten billion or even a trillion. When the data size of the mass storage reaches the EB level, the cluster file system needs to manage more than ten billion objects. In particular, when the data are stored in small files, the metadata to be managed can reach trillions in size.

6.3 Enhanced Bandwidth and I/O Performance

High bandwidth, low latency, and high IOPS are important performance indicators for traditional supercomputers, and are also important for data-intensive supercomputers. Bandwidth, latency, and IOPS are key concerns for supercomputers. A higher bandwidth means that the computer can have more processing power and can perform more jobs. Latency is also an important indicator of computer performance. The lower the latency, the better the overall performance.

For data-intensive computing architectures, meeting the bandwidth requirements between computing and storage systems has become a major challenge in handling large amounts of data. In traditional supercomputers, data must pass through a relatively long I/O path from memory to the computing devices. To complete a portion of the computational process, a series of read and write operations are required. When the volume of data exceeds one petabyte, the overhead spent on the storage system will be very high, and it is difficult to meet the needs of large-scale data processing. Therefore, traditional computers tend to use coupled computing and storage architectures. This architecture can effectively shorten the read/write path between computation and storage, reduce system overhead, provide fast data access, and achieve a storage balance between computation and read/write operations. In a storage-computing-coupled architecture, data are distributed on each node responsible for computation and storage, which significantly reduces the system overhead of data movement. However, while this architecture optimizes the overhead between the computing and storage systems, it reduces the overall flexibility of the system, and makes it difficult to solve the problems

of complex data types. Moreover, this architecture is not suitable for all applications, and many aspects need to be optimized and redesigned.

In the case of data-intensive computing, customers and data centers usually interoperate with each other over the internet. Their communication protocol is usually HTTP based. Owing to network bandwidth and stability limitations, and the fact that the requested service usually requires a large amount of data analysis, data-intensive supercomputing should try to achieve a high-speed response and incremental processing of data to meet application requirements.

Data-intensive supercomputing is still cluster-based, and data analysis is very important, so cluster-based supercomputers need to consider how to solve the transmission bottleneck problem. The transfer of large amounts of data between different processors and between different nodes, as well as the detection and processing of errors in parallel processing by the computing system, are issues that require a more robust I/O performance to address. To support efficient data-intensive computing, supercomputers will continue to require more intensive research in hardware, software systems, and algorithms to improve their IOPS.

How to meet the different demands for I/O performance for various scientific computing applications and big data applications should also be one of the issues to be considered in the development of data-intensive supercomputers. In some scientific research, the data processing and computation modes are very demanding in terms of computational performance—modes such as WRF, ARP, and others. The amount of communication between individual. The CPUs when running these software packages are also very high. In addition, because of the large number of users involved and the large number of file read and write operations, this places high demands on the IOPS of the supercomputer.

6.4 Efficient Data Analysis and Management Capabilities

Collecting and maintaining data should be the capability provided by the platform rather than the user. Data-intensive supercomputing should provide efficient data analysis and management capabilities for data-intensive tasks with large data volumes, and high demands on data processing.

Moreover, data-intensive supercomputers must provide efficient data processing and analysis capabilities. The data processing methods for big data are complex, with the most frequently used data processing methods being graph and batch processing—therefore, good cooperation between different files and systems must be obtained. This shows that for data-intensive supercomputers, a single-file system does not meet the requirements for supercomputer operation. Therefore, multiple file systems should be combined for supercomputer design.

On the other hand, a data-intensive supercomputer should provide as rich, comprehensive, and adaptable data management capabilities as possible—such as supporting users to query or update the programs and data they need to use according to their own authority and content; providing functions such as adding, deleting, updating, and checking based on database software; organizing, visualizing, and analyzing the data processed by users in a simple way. Further, the data management functions should be optimized using indexing, striping, and intelligent aggregation technologies. This enables the transfer of complex and additional data processing management and analysis operations to be performed outside the platform by users of traditional supercomputing platforms. And this is to be completed within the platform.

6.5 More Convenient Programming Models and Heterogeneous Programming

Traditional supercomputers tend to use low-level programming primitives to maximize the usage of system resources—often requiring the use of assembly or C languages to manipulate the underlying build the components of supercomputers such as processor cores, memory, and registers. Achieving high performance requires time-consuming and tedious optimization, and requires the user to have very high programming skills and a good understanding of the underlying supercomputing platform.

In the current era of big data, the programming model should meet the basic principles of simplicity and ease of use. In addition, it should improve performance as much as possible to ensure that the programmer's work is faster and smoother. In this process, the programmer does not need to be concerned with the state or condition of the program, but only solve problems that arise.

Big data applications also pose new challenges for programming models. These challenges arise mainly from conflicts between programming simplicity and performance. For the sake of programming simplicity, the programming model should provide developers with a clean programming method that allows them to easily write programs that can execute in parallel on a cluster, where they only need to focus on the problem to be solved and not on the specifics of how the program will run on a large cluster. This can include how the data is distributed, and how it is scheduled. The programming model should minimize the programming burden on developers by hiding specific details from them, making the code easier to understand, reuse, and maintain. As for the performance optimization of the programming model, the programming model usually manages the resources of the entire cluster to ensure that the application can run efficiently on a large-scale cluster, but also enables multiple jobs to run concurrently and share resources to improve the

resource utilization and job throughput of the entire cluster. Therefore, designing a simple, efficient, and reliable parallel programming model for big data applications is a major challenge.

With the development of data-intensive applications represented by intelligent applications, existing intelligent algorithm implementations often use the latest high-level languages such as Python3, Java, Julia, R, and Go. They often hide the underlying device interfaces from users and provide rich APIs and simple programming environments to meet the diverse algorithmic needs of users. Therefore, data-intensive supercomputers need to support these latest high-level programming languages to efficiently map the functionality implemented in their programs to the machine. At the same time, high-level programming languages can support more complex data structures such as 'dictionaries', 'tensors', etc., which are useful in AI tasks, and support for complex data structures is a must for data-intensive supercomputing.

In addition, we must ensure that different types of jobs and tasks can run concurrently on the cluster, which can greatly increase the effective utilization of resources within the cluster. However, with this model, we also face many new problems, such as how to deal effectively with heterogeneity. To improve efficiency, supercomputers generally choose to apply a parallel programming model, which not only improves the efficiency of the workforce, but also contributes to the effective implementation of big data applications to a large extent.

The programming language widely used today for supercomputing is still the C/C++ language invented more than 40 years ago, and still mainly uses the CPU for computational tasks. In recent years, however, supercomputers have been using GPU-assisted computing, which can handle several similar computational tasks simultaneously. Data-intensive supercomputers should therefore also consider how heterogeneous programming can be implemented better.

Heterogeneous computing refers to the collaboration between CPUs and GPUs or other devices (e.g., FPGAs). For example, the computing power of CPUs, GPUs, and even APUs (fusion of CPUs and GPUs) can be combined to increase the computational speed of a system. The use of heterogeneous systems is becoming increasingly common.

Therefore, designing a programming framework for heterogeneous computing should become an important issue in the design of data-intensive supercomputers. Some current heterogeneous programming frameworks, such as CUDA, OpenCL, and OpenACC, are based on the C/C++ language and take more time to program; and it is also difficult for beginners to become proficient in using these programming frameworks in a short time. In addition, existing parallel programming methods make it more difficult to analyze and verify parallel programs, and the uncertainty caused by

the cross-execution of parallel tasks in the program greatly reduces the reliability and predictability of the system.

6.6 Intelligent Scheduling

Existing data-intensive applications use data in increasingly complex ways, including data queries, data updates, simple numerical computations (data pre-processing, normalization, etc.), and complex numerical computations (calculation of neural network weights and reinforcement learning exploration processes). However, data-intensive supercomputers should also provide interactive access. A large number of existing supercomputers still mainly use batch mode to run user-submitted tasks. Although this approach can ensure that the task can be completed even after the terminal is closed or the logging node is down, the user can only obtain the task processing result after the batch task is completed and cannot make corresponding adjustments according to the intermediate outputs of the computation task. Data-intensive supercomputers should also provide an interactive task submission mode, ensuring that users can interact with the program at any point in time by writing code, and should also provide inherent interactive task support—including simple data processing and data visualization.

6.7 High Reliability and High Availability Mechanisms

The availability of supercomputers is an instantaneous concept, defined as a property of a system that indicates that it is ready to be used immediately. In other words, a highly available system works in a timely manner at any given moment. On the other hand, the reliability of a platform is a continuous concept, meaning that the system can continue to operate without failure. For data-intensive supercomputers, both reliability and availability must be guaranteed. Current supercomputers provide reliability primarily by periodically checking the status of programs, and rolling back to the nearest checkpoint when a supercomputer encounters an error. This is an inefficient, non-scalable mechanism that is not suitable for interactive use. Instead, this paper argues that data-intensive supercomputers should employ an uninterrupted reliability mechanism in which all raw and intermediate data are stored in redundant form, and selective recomputation can be performed in the event of component or data failure. Furthermore, the machine should automatically diagnose and disable faulty components, and these should only be replaced if they accumulate enough to impair system performance. Supercomputers should be available all year round, and hardware and software replacements and upgrades should be performed on a running system. This

more distributed and less monolithic architecture is more suitable for use in data-intensive applications than traditional supercomputers.

6.8 Data Privacy and Security Mechanisms

The security and privacy of data are involved in a wide range of applications, such as education, healthcare, and state security. To a certain extent, these data involve private information about individuals and society. If leaked, it could cause great damage to individuals and society. Traditional supercomputers provide technical means to address privacy protection in a closed application environment. With the development of internet technology, the disclosure of information on sensitive data sets poses a significant challenge to privacy security, and will also have an impact on data security in databases and outsourced servers. For this reason, data-intensive supercomputers should provide better data security mechanisms to ensure data privacy, data security, data completeness, and data availability. Most traditional data security technologies use passive defense methods, such as firewalls and encryption. However, with the emergence of new open environments, more active and passive techniques need to be designed to defend against data attacks by unscrupulous people. Data-intensive supercomputers should analyze huge amounts of data, use artificial intelligence to detect data attacks, screen out spam, and provide a more secure and reliable network and data storage base.

6.9 Multi-Protocol Access Interface

Data-intensive supercomputers should provide richer business capabilities, this however also results in a greater variety of data interfaces. Too many data interfaces result in significant conversion overhead and wasted resources. For example, Spectrum Scale provides a large number of interfaces, such as file interfaces (POSIX, NFS, CIFS), object interfaces (S3, SWIFT), big data interfaces, Openstack interfaces, storage service interfaces, etc., to provide rich business access. Traditional supercomputers have complex structures in terms of file management and network structure, and the interfaces are not uniform. For example, the front-end may adopt an FC interface, the storage may adopt an FC interface, and the computing network adopts an Infiniband interface. The complex communication network and various interfaces significantly reduce the conversion efficiency of data, resulting in a large protocol conversion overhead, a serious rate mismatch, a significant reduction in flexibility, an inability to meet the diverse changes in business, and a bottleneck in data transmission and conversion.

6.10 Better Support for Artificial Intelligence Applications

At present, in the fields of life sciences, smart city management, oil surveying, and disaster warning, the combination of artificial intelligence techniques, machine learning, and other technologies with supercomputers, will make traditional industries take on a new luster. Therefore, data-intensive supercomputers need to be considered to provide better performance services for AI-related models and achieve a better combination of AI models and supercomputing.

AI applications will become increasingly integrated with supercomputers. Because the demands of AI far exceed the computing power that a single chip can provide, it must be supported using a distributed architecture with many accelerator chips working in tandem. The actual performance of distributed training is highly dependent on the efficiency of the underlying hardware used.

While traditional supercomputers have the computing power and environment to support AI training, the energy consumption and cost effectiveness are not satisfactory. Applying supercomputers to AI requires innovations in architecture and software optimization.

7 Conclusions and Prospects

As supercomputing technology continues to grow in popularity and become more commonplace, it is no longer limited to scientific computing applications only, but is also being widely used in more civil fields. The scale of data that needs to be processed by supercomputers has exploded in both traditional scientific computing and civil fields, such as smart cities. HPC applications are gradually changing from compute-intensive to data-intensive supercomputing. The traditional supercomputing architecture is mainly designed for compute-intensive applications and is inadequate in terms of I/O bandwidth, IOPS, and reliability—and cannot support data-intensive applications well. Therefore, the development of a new generation of secure and reliable data-intensive supercomputers with decoupled storage and computation, and the design of application-optimized unified storage systems for massive data, are key to the sustainable development of data-intensive supercomputers in the future.

References

- Abramson, D., Jin, C., Luong, J., Carroll, J.: A beegfs-based caching file system for data-intensive parallel computing. In: Asian Conference on Supercomputing Frontiers, Springer, Cham, pp. 3–22 (2020)

- Allen, R.M.: Transforming earthquake detection? *Science* **335**(6066), 297–298 (2012)
- Amin, M.S., Ahn, H.: Earthquake disaster avoidance learning system using deep learning. *Cognit. Syst. Res.* **66**, 221–235 (2021)
- Anbuvizhi, R., Balakumar, V.: Credit/debit card transaction survey using map reduce in hdfs and implementing syferlock to prevent fraudulent. *Int. J. Comput. Sci. Netw. Security (IJCSNS)* **16**(11), 106 (2016)
- Anh Khoa, T., Quang Minh, N., Hai Son, H., Nguyen Dang Khoa, C., Ngoc Tan, D., VanDung, N., Hoang Nam, N., Ngoc Minh Duc, D., Trung Tin, N.: Wireless sensor networks and machine learning meet climate change prediction. *Int. J. Commun. Syst.* **34**(3), e4687 (2021)
- de Assuncao, M.D., da Silva, Veith A., Buyya, R.: Distributed data stream processing and edge computing: a survey on resource elasticity and future directions. *J. Netwk. Comput. Appl.* **103**, 1–17 (2018)
- Belair, S., Carrera, M.L., Abrahamowicz, M., Alavi, N., Badawy, B., Shahabadi, M.B., Bilodeau, B., Charpentier, D., Deacu, D., Durnford, D., et al.: Spaceborne l-band radiometry in environment and climate change canada (eccc)'s numerical analysis and prediction systems. In: *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, pp. 7526–7528 (2019)
- Braam, P.J., Zahir, R.: Lustre technical project summary. Attachment A to RFP B514193 Response (2001)
- Chang, C.C., Nicholson, A.N., Rinaldi, E., Berkowitz, E., Garron, N., Brantley, D.A., Monge-Camacho, H., Monahan, C.J., Bouchard, C., Clark, M.A., et al.: A per-cent-level determination of the nucleon axial coupling from quantum chromodynamics. *Nature* **558**(7708), 91–94 (2018)
- Chen, Q., Chen, K., Chen, Z.N., Xue, W., Ji, X., Yang, B.: Lessons learned from optimizing the sunway storage system for higher application i/o performance. *J. Comput. Sci. Technol.* **35**(1), 47–60 (2020)
- Chien, S., Bashir, R., Nerem, R.M., Pettigrew, R.: Engineering as a new frontier for translational medicine. *Sci. Trans. Med.* **7**(281), 281fs13 (2015)
- Dai, Y., Yan, J., Tang, X., Zhao, H., Guo, M. Online credit card fraud detection: A hybrid framework with big data technologies. In: *2016 IEEE Trustcom/BigDataSE/ISPA*, IEEE, pp. 1644–1651 (2016)
- Dong, W., Li, K., Kang, L., Quan, Z., Li, K.: Implementing molecular dynamics simulation on the sunway taihulight system with heterogeneous many-core processors. *Concurrency Comput.* **30**(16), e4468 (2018)
- Duan, X., Gao, P., Zhang, T., Zhang, M., Liu, W., Zhang, W., Xue, W., Fu, H., Gan, L., Chen, D., et al.: Redesigning lammgs for petascale and hundred-billion-atom simulation on sunway taihulight. In: *SC18: International conference for high performance computing, networking, storage and analysis*, IEEE, pp. 148–159 (2018)
- Forough, J., Momtazi, S.: Ensemble of deep sequential models for credit card fraud detection. *Appl. Soft Comput.* **99**, 106883 (2021)
- Fu, H., Liao, J., Yang, J., Wang, L., Song, Z., Huang, X., Yang, C., Xue, W., Liu, F., Qiao, F., et al.: The sunway taihulight supercomputer: system and applications. *Sci. China Inform. Sci.* **59**(7), 1–16 (2016)
- Gao Z.-Y., Zhang L.-M., Duan: A quantum machine learning algorithm based on generative models. *Sci. Adv.* (2018)
- Gao, J., Zheng, F., Qi, F., Ding, Y., Li, H., Lu, H., He, W., Wei, H., Jin, L., Liu, X., et al.: Sunway supercomputer architecture towards exascale computing: analysis and practice. *Sci. China Inform. Sci.* **64**(4), 1–21 (2021)
- Gao, P., Duan, X., Zhang, T., Zhang, M., Yang, G.: Millimeter-scale and billion-atom reactive force field simulation on sunway taihulight. *IEEE Transactions on Parallel and Distributed Systems* **PP**(99), 1–1 (2020)
- Gianinetto, M., Frassy, F., Marchesi, A., Maianti, P., De Paulis, R., Biffi, P.G., Nodari, F.R.: Mapping large-scale microseepage signals for supporting oil and gas exploration in new ventures. In: *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, IEEE, pp. 5430–5433 (2016)
- Guo, S., Qiao, W., Chen, B., Wang, B.: Prediction and abnormality analysis of climate change based on pca-arma and pcc. In: *2020 IEEE International Conference on Networking, Sensing and Control (ICNSC)*, IEEE, pp. 1–6 (2020)
- Hager, G., Wellein, G.: *Introduction to High Performance Computing for Scientists and Engineers*. CRC Press (2010)
- Harchol-Balter, M.: *Performance Modeling and Design of Computer Systems: Queuing Theory in Action*. Cambridge University Press (2013)
- Henz, B.J., Elliot, L., Barton, M., Shires, D.: High-performance computing for the next generation combat vehicle. Tech. rep., US Army Research Laboratory Aberdeen Proving Ground, United States (2018)
- Hernández, B., Somnath, S., Yin, J., Lu, H., Eaton, J., Entschew, P., Kirkham, J., Ronaghi, Z.: Performance evaluation of python based data analytics frameworks in summit: Early experiences. In: *Smoky Mountains Computational Sciences and Engineering Conference*, Springer, pp. 366–380 (2020)
- Hohman, F., Park, H., Robinson, C., Chau, D.H.P.: Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Trans. Vis. Comput. Graphics* **26**(1), 1096–1106 (2019)
- Hong, H.J., Chuang, J.C., Hsu, C.H.: Animation rendering on multimedia fog computing platforms. In: *2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, IEEE, pp. 336–343 (2016)
- Hu, Y., Yang, H., Luan, Z., Gan, L., Yang, G., Qian, D.: Massively scaling seismic processing on sunway taihulight supercomputer. *IEEE Trans. Parallel Distributed Syst.* **31**(5), 1194–1208 (2019)
- Hush, M.R.: Machine learning for quantum physics. *Science* **355**(6325), 580–580 (2017)
- Iannone, F., Ambrosino, F., Bracco, G., De Rosa, M., Funel, A., Guarnieri, G., Migliori, S., Palombi, F., Ponti, G., Santomauro, G., et al.: Cresco enea hpc clusters: a working example of a multifabric gpus spectrum scale layout. In: *2019 International Conference on High Performance Computing & Simulation (HPCS)*, IEEE, pp. 1051–1052 (2019)
- Ichimura, T., Fujita, K., Yamaguchi, T., Naruse, A., Wells, J.C., Schulthess, T.C., Straatsma, T.P., Zimmer, C.J., Martinasso, M., Nakajima, K., et al.: A fast scalable implicit solver for nonlinear time-evolution earthquake city problem on low-ordered unstructured finite elements with artificial intelligence and transprecision computing. In: *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, pp. 627–637 (2018)
- Joubert, W., Weighill, D., Kainer, D., Climer, S., Justice, A., Fagnan, K., Jacobson, D.: Attacking the opioid epidemic: Determining the epistatic and pleiotropic genetic architectures for chronic pain and opioid addiction. In: *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, pp. 717–730 (2018)
- Jun, S.P., Yoo, H.S., Choi, S.: Ten years of research change using google trends: from the perspective of big data utilizations and applications. *Technol. Forecasting Soc. Change* **130**, 69–87 (2018)
- Kahle, J.A., Moreno, J., Dreps, D.: 2.1 summit and sierra: designing ai/hpc supercomputers. In: *2019 IEEE International Solid-State Circuits Conference-(ISSCC)*, IEEE, pp. 42–43 (2019)
- Kappe, C.P., Böttinger, M., Leitte, H.: Analysis of decadal climate predictions with user-guided hierarchical ensemble clustering. In: *Computer Graphics Forum*, Wiley Online Library, vol 38, pp. 505–515 (2019)

- Kleppmann, M.: Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems. "O'Reilly Media, Inc." (2017)
- Kodama, Y., Odajima, T., Arima, E., Sato, M.: Evaluation of power management control on the supercomputer fugaku. In: 2020 IEEE International Conference on Cluster Computing (CLUSTER), IEEE, pp. 484–493 (2020)
- Koppers, L., Wormer, H., Ickstadt, K.: Towards a systematic screening tool for quality assurance and semiautomatic fraud detection for images in the life sciences. *Sci. Eng. Ethics* **23**(4), 1113–1128 (2017)
- Kudo, S., Nitadori, K., Ina, T., Imamura, T.: Implementation and numerical techniques for one eflop/s hpl-ai benchmark on fugaku. In: 2020 IEEE/ACM 11th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA), IEEE, pp. 69–76 (2020)
- Kumar, S., Huang, C., Zheng, G., Bohm, E., Bhatele, A., Phillips, J.C., Yu, H., Kalé, L.V.: Scalable molecular dynamics with namd on the ibm blue gene/l system. *J. Res. Dev.* **52**(1.2), 177–188 (2008)
- Kurth, T., Treichler, S., Romero, J., Mudigonda, M., Luehr, N., Phillips, E., Mahesh, A., Matheson, M., Deslippe, J., Fatica, M., et al.: Exascale deep learning for climate analytics. In: SC18: International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE, pp. 649–660 (2018)
- Li, K., Shang, H., Zhang, Y., Li, S., Wu, B., Wang, D., Zhang, L., Li, F., Chen, D., Wei, Z.: Openkmc: a kmc design for hundred-billion-atom simulation using millions of cores on sunway taihulight. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 1–16 (2019)
- Li, Z., Liu, G., Jiang, C.: Deep representation learning with full center loss for credit card fraud detection. *IEEE Trans. Comput. Soc. Syst.* **7**(2), 569–579 (2020)
- Lin, H., Tang, X., Yu, B., Zhuo, Y., Chen, W., Zhai, J., Yin, W., Zheng, W.: Scalable graph traversal on sunway taihulight with ten million cores. In: 2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS), IEEE, pp. 635–645 (2017)
- Liu, Z., Chu, X., Lv, X., Meng, H., Shi, S., Han, W., Xu, J., Fu, H., Yang, G.: Sunwaylb: Enabling extreme-scale lattice boltzmann method based computing fluid dynamics simulations on sunway taihulight. In: 2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS), IEEE, pp. 557–566 (2019)
- Lu, Y., Qian, D., Fu, H., Chen, W.: Will supercomputers be super-data and super-ai machines? *Commun ACM* **61**(11), 82–87 (2018). <https://doi.org/10.1145/3239556>
- Lu, Y.T., Cheng, P., Chen, Z.G.: Design and implementation of the tianhe-2 data storage and management system. *J. Comput. Sci. Technol.* **35**(1), 27–46 (2020)
- Luo, H., Paal, S.G.: Advancing post-earthquake structural evaluations via sequential regression-based predictive mean matching for enhanced forecasting in the context of missing data. *Adv. Eng. Inform.* **47**, 101202 (2021)
- Luo, L., Straatsma, T.P., Suarez, L.A., Broer, R., Bykov, D., D'Azevedo, E.F., Faraji, S.S., Gottiparthi, K.C., De Graaf, C., Harris, J.A., et al.: Pre-exascale accelerated application development: The orn1 summit experience. *IBM J. Res. Dev.* **64**(3/4), 11–1 (2020)
- Lv, G., Li, M., An, H., Lin, H., Chen, J., Han, W., Xiao, Q., Wang, F., Lin, R.: Distributed deep learning system for cancerous region detection on sunway taihulight. *CCF Trans. High Performance Comput.* **2**(4), 348–361 (2020)
- Mapar, J., Holtermann, K., Legary, J., Mahrous, K., Guzman, K., Heath, Z., John, C.J., Mier, S.A., Mueller, S., Pancerella, C.M., et al.: The role of integrated modeling and simulation in disaster preparedness and emergency preparedness and response: the summit platform. In: 2012 IEEE Conference on Technologies for Homeland Security (HST), IEEE, pp. 117–122 (2012)
- Massonnet, F., Bellprat, O., Guemas, V., Doblus-Reyes, F.J.: Using climate models to estimate the quality of global observational data sets. *Science* **354**(6311), 452–455 (2016)
- Mazzucco, W., Pastorino, R., Lagerberg, T., Colotto, M., d'Andrea, E., Marotta, C., Marzuillo, C., Villari, P., Federici, A., Ricciardi, W., et al.: Current state of genomic policies in healthcare among eu member states: results of a survey of chief medical officers. *Euro. J. Public Health* **27**(5), 931–937 (2017)
- Mehmood, Y., Ahmad, F., Yaqoob, I., Adnane, A., Imran, M., Guizani, S.: Internet-of-things-based smart cities: Recent advances and challenges. *IEEE Commun. Mag.* **55**(9), 16–24 (2017)
- Middleton, A.M.: Data intensive supercomputing solutions. In: Big Data Technologies and Applications, Springer, pp. 257–306 (2016)
- Minson, S.E., Meier, M.A., Baltay, A.S., Hanks, T.C., Cochran, E.S.: The limits of earthquake early warning: timeliness of ground motion estimates. *Sci. Adv.* **4**(3), eaaq0504 (2018)
- Musser, G.: One of quantum physics' greatest paradoxes may have lost its leading explanation. *Science* (2020)
- Nakao, M., Ueno, K., Fujisawa, K., Kodama, Y., Sato, M.: Performance evaluation of supercomputer fugaku using breadth-first search benchmark in graph500. In: 2020 IEEE International Conference on Cluster Computing (CLUSTER), IEEE, pp. 408–409 (2020)
- Odajima, T., Kodama, Y., Tsuji, M., Matsuda, M., Maruyama, Y., Sato, M.: Preliminary performance evaluation of the fujitsu a64fx using hpc applications. In: 2020 IEEE International Conference on Cluster Computing (CLUSTER), IEEE, pp. 523–530 (2020)
- Orhan, A.E.: Robustness properties of facebook's resnext wsl models. arXiv preprint [arXiv:1907.07640](https://arxiv.org/abs/1907.07640) (2019)
- Puertas-Martín, S., Banegas-Luna, A.J., Paredes-Ramos, M., Redondo, J.L., Ortigosa, P.M., Brovarets', O.O., Pérez-Sánchez, H.: Is high performance computing a requirement for novel drug discovery and how will this impact academic efforts? *Expert Opin. Drug Discovery* **15**(9), 981–985 (2020)
- Rajak, R.: A comparative study: Taxonomy of high performance computing (hpc). *Int. J. Electr. Comput. Eng.* **8**(5), 3386 (2018)
- Sapountzi, A., Psannis, K.E.: Social networking data analysis tools & challenges. *Future Generat. Comput. Syst.* **86**, 893–913 (2018)
- Schmidt, B., Hildebrandt, A.: Next-generation sequencing: big data meets high performance computing. *Drug Discovery Today* **22**(4), 712–717 (2017)
- Seal, S.K., Lim, S.H., Wang, D., Hinkle, J., Lunga, D., Tsaris, A.: Toward large-scale image segmentation on summit. In: 49th International Conference on Parallel Processing-ICPP, pp. 1–11 (2020)
- Sejdic, E., Malandraki, G.A., Coyle, J.L.: Computational deglutition: Using signal-and image-processing methods to understand swallowing and associated disorders [life sciences]. *IEEE Signal Process. Mag.* **36**(1), 138–146 (2018)
- Shimizu, T.: Supercomputer fugaku: Co-designed with application developers/researchers. In: 2020 IEEE Asian Solid-State Circuits Conference (A-SSCC), IEEE, pp. 1–4 (2020)
- Shipilova, E., Barret, M., Bloch, M., Boelle, J.L., Collette, J.L.: Simultaneous seismic sources separation based on matrioshka orthogonal matching pursuit, application in oil and gas exploration. *IEEE Trans. Geosci. Remote Sens.* **58**(7), 4529–4546 (2020)
- Shuman, C.A., Steffen, K., Box, J.E., Stearns, C.R.: A dozen years of temperature observations at the summit: central greenland automatic weather stations 1987–99. *J. Appl. Meteorol.* **40**(4), 741–752 (2001)
- Stavriniades, G.L., Karatza, H.D.: The impact of data locality on the performance of a saas cloud with real-time data-intensive applications. In: 2017 IEEE/ACM 21st International Symposium on Distributed Simulation and Real Time Applications (DS-RT), IEEE, pp. 1–8 (2017)

- Tang, B., Chen, Z., Hefferman, G., Pei, S., Wei, T., He, H., Yang, Q.: Incorporating intelligence in fog computing for big data analysis in smart cities. *IEEE Trans. Ind. Inform.* **13**(5), 2140–2150 (2017)
- Torbicki, M.: Longtime prediction of climate-weather change influence on critical infrastructure safety and resilience. In: 2018 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM), IEEE, pp. 996–1000 (2018)
- Usman, S., Mehmood, R., Katib, I.: Big data and hpc convergence for smart infrastructures: a review and proposed architecture. *Smart Infrastructure and Applications*, pp. 561–586 (2020)
- Wang, A., Zhang, A., Chan, E.H., Shi, W., Zhou, X., Liu, Z.: A review of human mobility research based on big data and its implication for smart city development. *ISPRS Int. J. Geo-Inform.* **10**(1), 13 (2021)
- Wang, D., Yuan, F.: High-performance computing for earth system modeling. *High Performance Computing for Geospatial Applications*, pp. 175–184 (2020)
- Wang, R., Tobar, R., Dolensky, M., An, T., Wicenc, A., Wu, C., Dulwich, F., Podhorszki, N., Anantharaj, V., Suchyta, E., et al.: Processing full-scale square kilometre array data on the summit supercomputer. In: SC20: International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE, pp. 1–12 (2020)
- Womble, D.E., Shankar, M., Joubert, W., Johnston, J.T., Wells, J.C., Nichols, J.A.: Early experiences on summit: Data analytics and ai applications. *IBM Journal of Research and Development* **63**(6), 2–1 (2019)
- Yu, Z., Zhu, K., Hattori, K., Chi, C., Fan, M., He, X.: Borehole strain observations based on a state-space model and apne analysis associated with the 2013 lushan earthquake. *IEEE Access* **9**, 12167–12179 (2021)
- Zhang, J., Zhou, C., Wang, Y., Ju, L., Du, Q., Chi, X., Xu, D., Chen, D., Liu, Y., Liu, Z.: Extreme-scale phase field simulations of coarsening dynamics on the sunway taihulight supercomputer. In: SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE, pp. 34–45 (2016)
- Zhang, K., Su, H., Zhang, P., Dou, Y.: Optimization and performance modeling of stencil computations on arm architectures. In: 2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS), IEEE, pp. 113–121 (2020)
- Zhang, P., Yin, D., Atkinson, P.M.: Future extreme climate prediction in western jilin province based on statistical downscaling model. In: IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium, IEEE, pp. 9886–9889 (2019a)
- Zhang, X., Wang, Y., Wang, Q., Zhao, X.: A new approach to double i/o performance for ceph distributed file system in cloud computing. In: 2019 2nd International Conference on Data Intelligence and Security (ICDIS), IEEE, pp. 68–75 (2019b)
- Zhang, Y., Zhu, Z., Cui, H., Dong, X., Chen, H.: Small files storing and computing optimization in hadoop parallel rendering. *Concurrency Comput.* **29**(20), e3847 (2017)
- Zhong, X., Yang, H., Luan, Z., Gan, L., Yang, G., Qian, D.: swtensor: accelerating tensor decomposition on sunway architecture. *CCF Trans. High Performance Comput.* **1**(3), 161–176 (2019). <https://doi.org/10.1007/s42514-019-00017-5>