

# The Impact of Aurora Supercomputer in Scientific Area and the Development Analysis

<sup>1)</sup> **Ngozi Lilian Okafor**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia

E-Mail: [nlii0001@student.uniji.ac.id](mailto:nlii0001@student.uniji.ac.id)

<sup>2)</sup> **Stella Putri Gunawan**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[sgun0002@student.uniji.ac.id](mailto:sgun0002@student.uniji.ac.id)

<sup>3)</sup> **Rafdah Zhafirah**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[rzha0001@student.uniji.ac.id](mailto:rzha0001@student.uniji.ac.id)

<sup>4)</sup> **Devlin Wen Sujatmiko**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[dwen0002@student.uniji.ac.id](mailto:dwen0002@student.uniji.ac.id)

<sup>5)</sup> **Amalia Shifa Aldila**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[amalia.shifa@uniji.ac.id](mailto:amalia.shifa@uniji.ac.id)

<sup>6)</sup> **Miranti Andhita Scantya**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[miranti.scantya@uniji.ac.id](mailto:miranti.scantya@uniji.ac.id)

<sup>7)</sup> **Lawrence Adi Supriyono**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[lawrence.supriyono@uniji.ac.id](mailto:lawrence.supriyono@uniji.ac.id)

<sup>8)</sup> **Kartiko Eko Putranto**

Universitas Jakarta Internasional, Jl. Letjen S.Parman 1AA Jakarta Barat, DKI Jakarta, Indonesia E-Mail:

[kartiko.putranto@jic.ac.id](mailto:kartiko.putranto@jic.ac.id)

## ABSTRACT

Supercomputers are crucial in solving complex scientific and industrial computing due to its tremendous computational power in enabling large-scale simulations, scientific research, and advancement in various scientific fields. This study is conducted on Aurora supercomputer, a powerful exa-scale supercomputer designed for intricate computing tasks, such as climate modeling, intense simulations, and AI and machine learning. By making use of the literature review approach, we analyze the capabilities and impact of Aurora on the scientific environment. Our research suggests that Aurora is capable of significantly enhancing performance on processing data, surpassing supercomputers such as Frontier and Fugaku. Furthermore, we discuss Aurora's impact on driving groundbreaking research across multiple scientific domains and its real-world applications such as drug discovery driven by AI and machine learning. The result highlights that Aurora marked a remarkable milestone in revolutionizing computational research and further research can show the true power of the Aurora supercomputer.

**Keyword:** Aurora, supercomputers, exascale, machine learning, artificial intelligence

## INTRODUCTION

Aurora is one of the world's first exa-scale supercomputers that can achieve at least a quintillion calculations per seconds performing tasks such as large-scale simulations, machine learning/artificial intelligence applications, establishing design models and techniques, all within a single, integrated system [1]. With around 60,000 GPUs (Graphics Processing Units) and around 85,000 network endpoints, making this

design the largest supercomputer installations to this date. This exa-scale system built by The U.S Department of Energy's (DOE) Argonne National Laboratory and developed by Intel and Hewlett Packard Enterprise (HPE), was designed to provide a flexible platform for groundbreaking innovations. Sweeping advancements across scientific fields that include global climate modelling, materials science, energy storage, and fusion energy research [1].

This system enables researchers to address intricate problems quicker with greater precision and accuracy. With Aurora, years of joint effort between science and engineering professionals, and industry leaders, demonstrated an upsurge in U.S. 's computing and scientific research infrastructure [1].

Aurora is an exascale supercomputer developed at the Argonne Leadership Computing Facility (ALCF), a part of the US Department of Energy (DOE). Aurora is designed to push the boundaries of computational science, helping researchers to solve complex problems in areas such as climate modeling, drug discovery, astrophysics, and artificial intelligence. This proves there is a significant leap in high-performance computing (HPC), providing scientists with powerful computational power [4].

Aurora's architecture is built using technologies from Intel and Hewlett Packard Enterprise (HPE). It is powered by Intel Xeon Max Series CPUs (Sapphire Rapids with High Bandwidth Memory) and Intel Data Center GPU Max Series (Ponte Vecchio architecture). It also utilizes the HPE Cray EX Supercomputer System, which allows for extreme parallelism and computational efficiency. This combination makes Aurora one of the first exascale systems in the world, capable of handling massive AI workloads and large-scale simulations [25], [26].

With a performance target of over 2 exaFLOPS (2 quintillion calculations per second), Aurora is among the most powerful supercomputers ever built. It is optimized for artificial intelligence, machine learning, and scientific modeling, providing researchers with a tool to accelerate discoveries. This level of computational power enables simulations and analyses that were impossible due to hardware limitations.

Aurora's architecture is built using technologies from Intel and Hewlett Packard Enterprise (HPE). It is powered by Intel Xeon Max Series CPUs (Sapphire Rapids with High Bandwidth Memory) and Intel Data Center GPU Max Series (Ponte Vecchio architecture). It also utilizes the HPE Cray EX Supercomputer System, which allows for extreme parallelism and computational efficiency. This combination makes Aurora one of the first exascale systems in the world, capable of handling massive AI workloads and large-scale simulations [25], [26].

Aurora's architecture is built using technologies from Intel and Hewlett Packard Enterprise (HPE). It is powered by Intel Xeon Max Series CPUs (Sapphire Rapids with High Bandwidth Memory) and Intel Data Center GPU Max Series (Ponte Vecchio architecture). It also utilizes the HPE Cray EX Supercomputer System, which allows for extreme parallelism and

computational efficiency. This combination makes Aurora one of the first exascale systems in the world, capable of handling massive AI workloads and large-scale simulations [25], [26].

With a performance target of over 2 exaFLOPS (2 quintillion calculations per second), Aurora is among the most powerful supercomputers ever built. It is optimized for artificial intelligence, machine learning, and scientific modeling, providing researchers a tool to accelerate discoveries. This level of computational power enables simulations and analyses that were impossible due to hardware limitations.

However, Aurora had multiple delays and challenges. Originally planned for 2018 as a pre-exascale system, Aurora went through many redesigns due to hardware development and supply chain issues. Eventually, it was reconfigured as a full exa-scale system and officially launched in 2023. Despite these setbacks, the upgraded system now meets the latest demands in computational science, making it a valuable resource for researchers worldwide.

Aurora supports groundbreaking research across a wide range of fields. For example similarly to climate science, where it helps model extreme weather patterns and predict global climate changes. In biomedicine, it simulates protein structures to help in drug discovery and medical research. It also plays a key role in particle physics, enhancing data analysis for institutions like CERN. Moreover, Aurora contributes to fusion energy research by simulating plasma physics, and in cosmology, it helps model the evolution of the universe [7].



**Figure 1.** Aurora Supercomputers

Aurora is more than just a powerful machine, it is also a representative to a major step to the next generation in scientific computing. Aurora's capabilities will allow scientist or researchers to solve some of the most challenging problem in science or engineering, giving out innovation and expanding our understanding of the world.

## METHODS

This research was primarily done through a combination of literature review and secondary data approach. By gathering data through trusted and accurate websites, diving deep down on the intricate details of Aurora, would lead to extensive

knowledge and understanding of how powerful an exa-scale system can be. As technology will keep advancing, the development of supercomputers such as Aurora would have a profound effect in the scientific field and crucial in aiding researchers in numerous fields of study. This research can be used as a basis to understand Aurora and there is still room for improvements to fully utilize a machine this complex.

### Aurora Supercomputer Specifications

**Table 1.** Aurora Specifications [8], [10]

Part	Specifications
Compute Node (Processor)	2 Intel Xeon CPU Max Series processors: 64GB HBM on each, 512GB DDR5 each; 6 Intel Data Center GPU Max Series, 128GB HBM on each, RAMBO cache on each; Unified Memory Architecture; 8 SlingShot 11 fabric endpoints [8].
kenapaCores	9,264,128 [8]
Software Stack	HPE Cray EX supercomputer software stack, Intel enhancements, data and learning [8].
GPU Architecture	63,774 Intel Data Center Max 1550 Series; Tile-based chiplets, HBM stack, Foveros 3D integration, 7nm
CPU - GPU Interconnect	CPU-GPU: PCIe; GPU-GPU: Xe Link
Processors	21,248 Intel Xeon CPU Max Series
System Interconnect	Slingshot 11; Dragonfly topology with adaptive routing; Peak Injection bandwidth 2.12 PB/s; Peak Bisection bandwidth 0.69 PB/s
Network Switch	25.6 Tb/s per switch, from 64–200 Gbs ports (25 GB/s per direction)
System Performance	Exascale
High-Performance Storage	230 PB, 31 TB/s, 1024 nodes (DAOS)
Programming Models	Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++
System Size	10,624 nodes
Platform	HPE Cray EX supercomputer

Part	Specifications
Aggregate System Memory	20.4 PB
Power Consumption	38,698.36 kW (level 2)
Operating System	SUSE Linux Enterprise Server 15 SP4

### Setup Overview

Aurora GPU Architecture is made up of 2 Stacks also referred to as Tiles, where each stack features Xe Vector Engine or EU which is an execution unit that holds up to 512 units per stack, making 448 units active and supporting up to 8 threads and 512b SIMD. Also a Xe Matrix Engine which is the systolic part of the execution unit has the same number of units as one stack where each thread equips up to 128 registers of 512 bits each. It also includes Xe core which is also referred to as sub-slices where they are units of 8 execution units(Eu). The Xe cores hold up to 64 units per stack and 128 GPU. Its memory resources are provided by an L1 cache of up to 218 KiB with a last level cache which is referred to as the Rambo cache with up to 384 MiB per GPU [14].

**Table 2.** Aurora Computer Node [14]

Node Component	Description	Per Node	Aggregate
Processor	2000 MHz	2	21.248
Core/Threads	Intel Xeon CPU Max 9470C Series	104/208	1.104.896 / 2.209.729
CPU HBM	HBM2e	64x2 GiB	1.328 PiB
CPU DRAM	DDR5	512x2 GiB	10.375 PiB
GPUs	Intel Data Center Max 1550 Series	6	63.744
GPU HBM	HBM2e	768 GiB	7.968 PiB

### Node Performance Overview

A single aurora node is a computing unit that is high in performance that consists of six GPU - Intel Data Center GPU max series with each holding a memory of 768 GB GPU HBM Memory resulting in a peak bandwidth of 19.66 GPU HBM BW (TB/s). It also has two CPUs - Intel Xeon CPU Max Series with each having a memory of 128 GB CPU HBM and a 1025 GB DDR5 memory with peak bandwidth of 2.87 CPU HBM BW (TB/s) and 0.56 DDR5 BW (TB/s). The power of the node does not exceed at least 130 teraflops and it

supports eight NICs [14].

### Early User Notes and Issues

As Aurora is still going through frequent testing and development due its system not being completely developed. Early users could experience issues like; unexpected outage, scheduled and unscheduled downtimes and troubleshooting issues. Early users are expected to encounter unplanned outages and anytime. Aurora's system is not fully stable but has been undergoing significant improvement over the months. Early users need to be cautious when observing Aurora's performance if it encounters any hangs or actions of instability. During this early phase Aurora makes use of the Flare-Lustre File System which may trigger lagging processes and demeaning its performance when it's in heavy use. This can result in possible crashing of the compute node. DAOS which is a scratch file system, it was initially configured so it has a small number of nodes which results in the allocation of smaller projects. Aurora's network and compute node instability may cause compute nodes to be inaccessible causing MPI ranks to be unreachable. Users are advised to stop and rerun if they encounter these error messages.

```
ping failed on x4707c6s4b0n0: Couldn't
forward RPC ping(24c93b8c-3434-4fb5-a8f0-
53cff4cbbe42) to child
x4707c7s6b0n0.hostmgmt2707.cm.aurora.alcf.
anl.gov: Resource temporarily unavailable.
```

**Figure 2.** Showing ping message errors

Hanging are due to Hardware and software failure. For known hardware or low-level software issues such as ping failures as discussed above, just restart the job. GPU Segfault which is also referred to as page faults. Occurring errors on memory may cause page faults but it can be debugged by gdb-oneapi and DDT, both of which allow debugging into GPU kernel threads and looking at GPU data structures [14].

Other known issues Aurora may encounter several runtime errors when in use due to network interface communication, free memory querying and system Limitations. One issue is the Cassini Event Queue overflow detected errors where it occurs for certain MPI communication and different reasons could be job routing, software and hardware communications and so on. This error indicates that the flow of messages is too fast for one of the network interfaces where it's not able to keep up with processing them while it's receiving it. To solve this error, it is advisable for users to alter these environment variables:

```
FI_CXI_DEFAULT_CQ_SIZE,
FI_CXI_OFLOW_BUF_SIZE,
```

And

```
FI_CXI_CQ_FILL_PERCENT,
```

but if the error still persists then users should increase the value of

```
FI_CXI_DEFAULT_CQ_SIZE [14].
```

Another known issue is the GOTPCREL Relocation Error which occurs when complied with relocation failures. This error can be resolved by linking with `-flink-huge-device-code`. Furthermore, users querying for free memory on a device with the Intel SYCL extension

```
get_info<sycl::ext::intel::info::dev
ice::free_memory>(); need to set
export ZES_ENABLE_SYSMAN=1.
```

Also, advisable to check the

```
aspect::ext_intel_free_memory
```

SYCL device property before making a memory query [14].

Jobs may fail to successfully start at times specifically at higher node counts. If there are no error messages then it is advisable to check the comment field in the full job environment for the command

```
qstat -xfw<JOBID> [14]
```

Some cases that could occur are; user placed the job on hold, user submitted to a queue that is currently not running and user's job tried but failed to start.

### Aurora programming environment

The Aurora Programming Environment (Aurora PE) is the default software stack on the Aurora supercomputer, built around Intel's oneAPI SDK, MPICH, and the Spack PE. It utilizes a hierarchical module system (Lmod) for easy version management and compatibility between its components, allowing users to switch between different versions of compilers and libraries with simple module load commands. The core PE resides in `/opt/aurora` as a read-only squashfs, while `/soft` provides a space for additional, less frequently updated software. However, using software from `/soft` can introduce performance bottlenecks at scale and requires manual conflict resolution due to potential module clashes [14].

The Aurora PE has a planned upgrade cadence on the order of months, with ad-hoc software requests being fulfilled through installations in `/soft` for later consideration in core



PE updates. Users must explicitly add /soft/modulefiles to their module path to access these additional packages. While the module system aims to maintain consistency, users are advised to be cautious of module conflicts, especially when adding modules from non-standard locations, and to prioritize using software from the core PE to maximize performance [14].

### Compilation and Linking

Compilation on Aurora primarily utilizes the oneAPI environment, offering Intel compilers for C, C++, and Fortran, and MPICH for MPI. Compilation can occur on login nodes for CPU-only builds, or compute nodes for GPU-accelerated code. The system supports diverse programming models, including OpenMP, SYCL, and OpenCL for both CPU and GPU parallelism, with specific compiler flags and linking requirements. Notably, GPU compilation allows for Just-in-Time or Ahead-of-Time approaches, with the latter requiring backend specification during linking. The documentation emphasizes the importance of using the same Fortran compiler that MPI was compiled with to prevent incompatibility issues [14].

Practical examples are provided for various programming models, showcasing compilation and execution procedures, including job script usage. A critical aspect highlighted is CPU-GPU affinity management, demonstrated through the ZE\_AFFINITY\_MASK environment variable and a helper script. These examples illustrate how to bind MPI ranks to specific GPU tiles, ensuring efficient resource utilization. The guide also details how to compile and execute example programs, and how to utilize the job scheduler with example scripts [14].

### Build Tools

CMake is a vital component of the Aurora supercomputer's software ecosystem, serving as a versatile, cross-platform build system generator. It operates by interpreting CMakeLists.txt files, which define build instructions in a platform-agnostic way, and then generating native build files tailored to the target system. This abstraction is critical for managing the complexity of Aurora's heterogeneous architecture, which includes both CPUs and GPUs, and for ensuring consistent builds across diverse software dependencies. CMake's flexibility simplifies the handling of complex build processes, dependency management, and custom build rules, making it an indispensable tool for both system software deployment and user application development [14].

On Aurora, CMake facilitates the efficient compilation and linking of software, particularly in scenarios involving complex dependencies and

cross-architecture targeting. Its availability as a module ensures users have access to consistent versions, streamlining the development workflow. Moreover, CMake's ability to handle intricate linker flags, especially those required for Ahead-of-Time (AoT) GPU compilation, is crucial for optimizing performance on Aurora's specialized hardware. By providing a robust and adaptable build system, CMake plays a pivotal role in enabling researchers and developers to fully leverage the computational power of the Aurora supercomputer [14].

### Running Jobs

The Aurora supercomputer at the Argonne Leadership Computing Facility (ALCF) is designed for large-scale computational tasks, utilizing the Portable Batch System (PBS) for job scheduling and management. Users submit jobs via PBS scripts, specifying resource requirements such as node count, walltime, and appropriate queues for optimized scheduling. Given its pre-production status, Aurora may experience node failures. To minimize disruptions, users are encouraged to run MPI jobs interactively, allowing dynamic rerouting around failed nodes. This enhances computational efficiency by ensuring that jobs remain operational despite hardware issues [14].

Optimizing performance on Aurora also involves effective MPI rank and thread binding. Each node contains multiple CPUs and GPUs, making it crucial to map computational tasks efficiently to available resources. Using affinity masks and understanding Aurora's hardware topology helps reduce latency and maximize resource utilization. Proper job placement ensures that tasks are distributed efficiently across nodes, improving execution speed and stability. For detailed guidelines on running jobs effectively, users can refer to the ALCF documentation on job submission, queue selection, and best practices [14].

### Data Management

#### Copper

Copper is a read-only cooperative caching layer developed by the Argonne Leadership Computing Facility (ALCF) to facilitate scalable data loading across a vast number of compute nodes in exa-scale supercomputers. By implementing Copper, ALCF aims to optimize data management strategies within supercomputing environments. The cooperative caching mechanism allows multiple compute nodes to share cached data, reducing redundant data transfers and minimizing latency. Overall, Copper enhances the efficiency and scalability of data-intensive applications running on exa-scale systems [14].

### DAOS

The Distributed Asynchronous Object Storage (DAOS) is a high-performance object store designed to meet the demanding I/O requirements of exascale computing environments, such as those at the Argonne Leadership Computing Facility (ALCF). This versatility allows users to efficiently perform I/O operations and checkpointing on systems like Aurora, seamlessly integrating with the broader compute fabric [14].

### Moving Data to Aurora

Efficient data transfer to the Aurora supercomputer at the Argonne Leadership Computing Facility (ALCF) is facilitated through two primary methods: Globus and the DAOS DataMover. Globus is a widely-used, high-performance data transfer service that enables users to move large datasets securely and reliably between systems. To utilize Globus for transferring data to Aurora, users can follow the detailed instructions provided in the ALCF user guides [14].

### Applications and Libraries

Libraries are used in a way to store and take a collection of pre-written code, in this case involving Cabana, oneDAL, and Spack PE. Aurora provides users with libraries in a way to allow them to stack multiple software and various tools that would affect codes through a large data analysis that supports both CPUs and GPUs [14].

### Cabana

Cabana is a well performance portable library that helps to easily create a particle-based simulation. This application includes molecular dynamics (MD) working with short or long-range atomic interactions following with various other particles like particle-in-cell (PIC) methods including fluid or solid mechanics and plasma physics where cosmology simulations including peridynamics for fracture mechanics.

### oneDAL

oneDAL has also provided many support through batches, online, and distribution yet processing modes that are capable of sending out these datas within its frameworks of other libraries. Hence, during documentation oneDAL showed Algorithms such as Mathematical equations, CPU and GPU supporting, rebuilding and machine learning, Clustering, Covariance, Decomposition, Ensembles, Kernel, Logistic and Regression, Newton-CG Optimizer, Statistics, and etc. Most of these Algorithms go through data management and are then clustered and made into a correlation either through Matrices, Rules, Functions, Quantile, Engines, Metrics, Normalization, and Cholesky Decomposition[21].

### SpackPE

Mentioned in Aurora's Programming Environment, Spack PE is a software that stacks and provides various tools, utilities, and libraries. It contains 2 parts: spack-pe-gcc and spack-pe-oneapi. Spack-pe-gcc is based on a project named E4S Project which increases performance of HPC libraries built in the OneAPI SDK(Software Development Kit) [29]. We can inspect packages as there are several environment variables that are integrated into a package for the user's needs. Like PACKAGE\_ROOT variable, this set contains the path installation prefix to the package similarly to loading cmake. The code to load with cmake written as [14]:

```
$ echo $CMAKE_ROOT
/opt/aurora/23.275.2/spack/gcc/0
.6.1/install/linu x-sles15-
x86_64/gcc-12.2.0/cmake-3.27.7-
mb17dvgbiblpavhu53h5cheyrmpaikdz
$ ls -a $CMAKE_ROOT
. .. bin doc share .spack
```

### Python

Aurora supports Intel's Data Parallel Extensions to allow users to access Intel's Python stack that consist of compilers and libraries for programming varied devices. This DPEP is constructed with three main packages for programming on CPUs and GPUs: dpnp, dpctl, numba-dpex [14].

```
python3 -m venv /path/to/new/venv --system-site-packages
source /path/to/new/venv/bin/activate
```

Figure 3 virtual environment to use Python [14]

### Jupyter

JupyterLab refers to the latest web-based interactive development environment made for notebooks, code, and data. It is the next-generation user interface for Project Jupyter and offers all the familiar building blocks of the classic Jupyter Notebook (such as notebooks, terminals, text editors, and file browsers), but in a more flexible and powerful interface. Its modular and extensible design allows users to arrange multiple documents and activities side by side in tabs or split screens, enabling a highly customizable workflow tailored to individual preferences [31]. JupyterLab supports a wide variety of programming languages (via kernels), with Python being the most commonly used. It is particularly useful in data science, machine learning, and scientific computing environments where interactive code execution and visualization are crucial.

```
ssh
<compute_node_hostname>
source
```

```
<path_to_your_virtual_environment>/
bin/act ivate
jupyter lab --no-browser --
port=9999
```

## Profiling

By understanding and analyzing performance code, users can identify bottlenecks and where to make further optimizations.

[illegible]

**Figure 4.** Showing Example Usage of Profiling Deep Learning Application [14]

## RESULT AND DISCUSSION

## Aurora in Drug Discovery

Aurora has had a great impact on the discovery of drugs, by sorting through a large database of chemical compounds in search for medicines that could serve as a treatment for some of the world's most deadly diseases. The ALCF team has enabled Aurora to be able to screen about 11 billion molecules of drug per hour making use of 128 nodes of Aurora then, with double the number of nodes, Aurora is able to screen about 22 billion molecules per hour. However, the ALCF team still has plans in hopes that Aurora will be able to screen about 1 trillion candidates per hour once it's fully complete[11]. Aurora's performance in drug discovery highlights its capabilities in analyzing and handling large datasets in order to develop suitable models that will accelerate development in scientific fields.

## Aurora in Cosmology

The US Department of Energy (DOE) researchers at Argonne National Laboratory are implementing their own computational mining tool that will explain the nature of dark matter. Although there are evolving theories behind dark matter, Argonne's scientists are incorporating supercomputers like Aurora in a project called Dark Sky Mining. This project is made up of creating highly detailed sky maps that include the actual cosmic observation and computer intensive simulation of the entire universe. With contributions from 15 research teams that will assemble different sets of codes and software to match the architecture and scale of Aurora,

making use of its capability to execute over a quintillion calculations per second enabling a rapid simulation of large quantities of physics intense cosmic scenarios. However they will also be integrating the use of AI (artificial Intelligence), and statistics to provide a more reliable theory on the nature of dark matter and its connections to the universe [12].

## Future Potentials and Applications of Aurora Towards Efficiency

Argonne's Aurora supercomputer is an exascale system which impacts supercomputer studies by enabling researchers to do much more actions consisting of; detailed simulations, accelerate AI and machine learning advancements, and tackle previously impossible and non-human calculations across various scientific fields. [6] With its efficiency and to build a larger margin of the system, Aurora's exascale supercomputer attribute towards researching and science ground-breaking makes it more robust in proving many probabilities and theories such as medicinal, dark matter, and AI integrated systems capable of reaching to a fast research of levels that it could contribute many to human life changes or environmental impacts.

## Towards Environmental Footprint

Aurora is designed with energy efficiency in mind, utilising advanced technologies to reduce its environmental impact compared to other and previous year-built supercomputers [6]. Additionally, Aurora is housed in a new state-of-the-art data centre specifically designed to support efficient energy use. Hence, Aurora is capable of achieving an eco-friendly way of energy consumption. [5]

## Towards Long-term Projects

Aurora evolves into an ecosystem-like state in a exa-scale supercomputer in which it aims to unlock much more possibilities for scientific research, its goals include to develop AI-enabled workflows and models consistently producing revolutionised fields for clean energy, universe knowledge, and drug discovery. [4] Not only has it been a work-in-progress, the supercomputer itself has worked to design a next-generation system, Helios. In computer learning and human understanding [4].

## Future of Dark Matter Learning

Aurora is capable of detecting its presence, using Lattice QCD simulations, understanding how this subatomic reaction worked and could determine the potential interactions of dark matter with computational touches [13]. However, it proves that Aurora's detection was capable of sensing these interactions and come into play with its



nature. Though, its future is uncertain as dark matter wasn't much to be studied than that of learning new ways of medicine and detection of earthquakes [13].

### Future of AI and Exa-scale

With the increasing speed of knowledge regarding science, Aurora was capable enough to be one of the three fastest systems below El Capitan at Lawrence Livermore National Laboratory. Even at its development, Aurora had established itself as one of the world's leading systems in AI performance, on the HPL-MxP benchmark in November 2024 [6]. With its world performance, Aurora is also trained to learn large language models for science [11]. Though, its massive datasets are also critical as this requires researchers to perform many data streams at a large-scale for this supercomputer in facilities such as Argonne's Advanced Photon Source and CERN's Large Hadron Collider. Hence, this renewed project was capable of supporting other engineering and science projects of over 70 and increasing [11].

### CONCLUSION

Aurora is an extremely powerful exa-scale supercomputer developed at the ALCF, designed to push the boundaries of computational science and aid researchers in various scientific fields. With the ability of achieving at least a quintillion calculations per second, complex computational tasks can be executed, Aurora has accomplish remarkable achievements such as advancement in AI and machine learning, impact on the discovery of drugs, uncover mysteries such as the dark matter, and large simulations to predict the climate using climate modelling.

With its immense power, further research can still be attained to understand the true potential of Aurora. This exa-scale supercomputer can further contribute to future research and science ground-breaking discoveries, proving many probabilities involving medicinal, dark matter, and AI integrated systems that can lead to better life expectancy. This makes Aurora a major step in the next generation of scientific computing. It has the capabilities to give major challenges and innovations that can aid researchers in expanding their understanding of the world.

### REFERENCES

- [1] UChicago News, "Argonne National Laboratory opens 'Aurora' supercomputer to researchers," UChicago News, Feb. 11, 2025. [Online]. Available: <https://news.uchicago.edu/story/argonne-national-laboratory-opens-aurora-supercomputer-researchers>
- [2] J. Thomas, "The Aurora supercomputer achieves exascale," Innovation News Network, Sept. 13, 2024. [Online]. Available: <https://www.innovationnewsnetwork.com/the-aurora-supercomputer-achieves-exascale/50825/>
- [3] InsideHPC, "Argonne Releases Aurora Exascale Supercomputer to Researchers," InsideHPC, Jan. 28, 2025. [Online]. Available: <https://insidehpc.com/2025/01/argonne-releases-aurora-exascale-supercomputer-to-researchers/>
- [4] T. Rosenberg, "Aurora supercomputer heralds a new era of scientific innovation," Argonne National Laboratory, May 30, 2024. [Online]. Available: <https://www.anl.gov/article/aurora-supercomputer-heralds-a-new-era-of-scientific-innovation>
- [5] InsideHPC, "HPE delivers 2nd Exascale Supercomputer, Aurora, to Argonne National Laboratory," InsideHPC, May 15, 2024. [Online]. Available: <https://insidehpc.com/2024/05/hpe-delivers-2nd-exascale-supercomputer-aurora-to-argonne-national-laboratory/>
- [6] J. Collins, "Argonne's Aurora supercomputer breaks exascale barrier," Argonne National Laboratory, May 13, 2024. [Online]. Available: <https://www.anl.gov/article/argonnes-aurora-supercomputer-breaks-exascale-barrier>
- [7] N. Heinonen, "From cells to the cosmos: the research campaigns already underway on Aurora," Argonne Leadership Computing Facility, May 8, 2024. [Online]. Available: <https://www.alcf.anl.gov/news/cells-cosmos-research-campaigns-already-underway-aurora>
- [8] Argonne Leadership Computing Facility, "Aurora," Argonne National Laboratory. [Online]. Available: <https://www.alcf.anl.gov/aurora>
- [9] NASA Scientific Visualization Studio, "Supercomputing the Climate," Jun. 2, 2010. [Online]. Available: <https://svs.gsfc.nasa.gov/10563/>
- [10] TOP500.org, "TOP500 List - November 2024," TOP500, Nov. 2024. [Online]. Available: <https://top500.org/lists/top500/2024/11/>
- [11] K. Jackson, "Aurora proves AI supremacy, snags top AI spot from Frontier," HPCwire, May 13, 2024. [Online]. Available: <https://www.hpcwire.com/2024/05/13/aurora-proves-ai-supremacy-snags-top-ai-spot-from-frontier/>
- [12] J. Spizzirri, "Mining the dark sky with the



- Aurora exascale supercomputer," Argonne National Laboratory, Jan. 23, 2025. [Online]. Available: <https://www.anl.gov/article/mining-the-dark-sky-with-the-aurora-exascale-supercomputer>
- [13] J. Sykes, "Argonne's Aurora supercomputer to help scientists advance dark matter research," Argonne National Laboratory, Dec. 11, 2023. [Online]. Available: <https://www.anl.gov/article/argonnes-aurora-supercomputer-to-help-scientists-advance-dark-matter-research>
- [14] Argonne Leadership Computing Facility, "ALCF User Guides." [Online]. Available: <https://docs.alcf.anl.gov>
- [15] Microsoft Research, "DeepSpeed," Microsoft. [Online]. Available: <https://www.microsoft.com/en-us/research/project/deepspeed/>
- [16] NVIDIA, "PyTorch," NVIDIA. [Online]. Available: <https://www.nvidia.com/en-au/glossary/pytorch/>
- [17] N. Shiledarbaxi, "Guide to GPyTorch: A Python Library for Gaussian Process Models," Analytics India Magazine, Mar. 10, 2021. [Online]. Available: <https://analyticsindiamag.com/ai-trends/guide-to-gpytorch-a-python-library-for-gaussian-process-models/>
- [18] F. J. Hwang, "A visual exploration of Gaussian processes," Distill, 2019. [Online]. Available: <https://distill.pub/2019/visual-exploration-gaussian-processes/>
- [19] TensorFlow, "TensorFlow," TensorFlow.org. [Online]. Available: <https://www.tensorflow.org/>
- [20] ECP-CoPA, "Cabana Wiki," GitHub. [Online]. Available: <https://github.com/ECP-copa/Cabana/wiki>
- [21] UXL Foundation, "oneDAL," UXL Foundation GitHub. [Online]. Available: <https://uxlfoundation.github.io/oneDAL/index.html>
- [22] Argonne Leadership Computing Facility, "Preparing Applications for Aurora: OpenACC and OpenMP Migration Tool." [Online]. Available: <https://www.alcf.anl.gov/support-center/training/preparing-applications-aurora-openacc-openmp-migration-tool>
- [23] OpenMP Architecture Review Board, OpenMP Application Programming Interface Version 6.0, OpenMP, 2024. [Online]. Available: <https://www.openmp.org/wp-content/uploads/OpenMP-RefGuide-6.0-OMP60SC24-web.pdf>
- [24] OpenMP Architecture Review Board, OpenMP Seminar Material, OpenMP, 2024. [Online]. Available: <https://www.openmp.org/wp-content/uploads/OMP-seminar-material.pdf>
- [25] Intel Corporation, "Argonne's Aurora: An Exascale Supercomputer," 2024. [Online]. Available: <https://www.intel.com/content/www/us/en/customer-spotlight/stories/argonne-aurora-customer-story.html>
- [26] Hewlett Packard Enterprise, "Aurora at Argonne National Laboratory," HPE, 2024. [Online]. Available: <https://www.hpe.com/us/en/compute/hpc/cray/argonne-national-laboratory.html>
- [27] Innovation News Network, "Aurora Supercomputer Breaks Exascale Computing Threshold," 2024. [Online]. Available: <https://www.innovationnewsnetwork.com/aurora-supercomputer-breaks-exascale-computing-threshold/47488/>
- [28] Spack, "Spack Package Manager Documentation," 2024. [Online]. Available: <https://spack.readthedocs.io/en/latest/index.html>
- [29] Argonne Leadership Computing Facility, "Spack PE for Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/applications-and-libraries/libraries/spack-pe/>
- [30] Argonne Leadership Computing Facility, "Aurora Containers and Modules," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/containers/containers/#load-modules-on-compute-node>
- [31] Jupyter, "Project Jupyter," 2024. [Online]. Available: <https://jupyter.org/>
- [32] Argonne Leadership Computing Facility, "ADIOS Workflows on Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/workflows/adios/>
- [33] Argonne Leadership Computing Facility, "Balsam Workflows on Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/workflows/balsam/>
- [34] Argonne Leadership Computing Facility, "libEnsemble Workflows on Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/workflows/libensemble/>
- [35] Argonne Leadership Computing Facility, "Parsl Workflows on Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/workflows/parsl/>
- [36] Argonne Leadership Computing Facility, "SmartSim Workflows on Aurora," 2024. [Online]. Available: <https://docs.alcf.anl.gov/aurora/workflows/smart-sim/>

- [Online]. Available:  
<https://docs.alcf.anl.gov/aurora/workflows/smartsim/>
- [37] Intel, "Bringing HIP to oneAPI on Aurora," 2024. [Online]. Available:  
<https://simplecore.intel.com/oneapi-io/wp-content/uploads/sites/98/Brice-V-HIP-on-Aurora-Bringing-HIP-to-oneAPI.pptx.pdf>
- [38] CrayLabs, "SmartSim-Zoo," GitHub Repository, 2024. [Online]. Available:  
<https://github.com/CrayLabs/SmartSim-Zoo>
- [39] Kitware, "ParaView - Open Source Scientific Visualization," 2024. [Online]. Available: <https://www.paraview.org/>

