Review Paper
J. Astron. Space Sci. 43(3), 117-128 (2026)
https://doi.org/10.5140/JASS.2026.43.3.117

JASS Journal of Astronomy and Space Sciences

Astroparticle Physics and Astronuclear Physics at Exascale: Scientific Opportunities Enabled by the KISTI-6 Supercomputer

Kihyeon Cho†
Korea Institute of Science and Technology Information, Daejeon 34141, Korea

The convergence of theoretical modeling, observational astronomy, and high-performance computing has ushered in a new era in astroparticle physics and astronuclear physics. In this review, we summarize major simulation-based research activities enabled by Korea's national supercomputing infrastructure, with particular emphasis on results obtained using the KISTI-5 supercomputer (Nurion). These include nuclear lattice effective field theory calculations, neutron drip line predictions, nuclear equation-of-state studies, and large-scale simulations supporting the RAON experimental program. The forthcoming KISTI-6 supercomputer represents a transformative transition toward heterogeneous, GPU-accelerated exascale computing, delivering more than 600 peta floating-point operations per second (PFLOPS) of peak performance. We discuss the scientific opportunities unlocked by this architecture, including time-dependent supernova neutrino transport, r-process nucleosynthesis, neutron star crust dynamics, and multi-messenger signal modeling. We also address the associated computational challenges, such as GPU porting of legacy codes and the integration of AI-assisted surrogate models. This review highlights simulations as a central pillar of modern astrophysical research, complementing theory and observation, and outlines how exascale computing will enable a unified, multi-physics description of extreme astrophysical environments in the coming decade.

Keywords: astroparticle physics, astronuclear physics, supercomputer, high-performance computing, exascale computing

1. INTRODUCTION

Astroparticle physics focuses on the origin, propagation, and detection of elementary particles from astrophysical sources, whereas astronuclear physics aims to understand nuclear reactions and nuclear matter under extreme astrophysical conditions. Table 1 shows the comparison between astroparticle physics and astronuclear physics. Astroparticle physics primarily exploits massive event-level parallelism and data-driven inference, whereas astronuclear physics requires high-precision calculations of strongly coupled many-body systems. In this context, supercomputing plays a central role in both domains. In astroparticle physics, large-scale detector and particle transport simulations are typically performed using Geant4 (Shin et al. 2015; Choi et al. 2018; Yeo & Cho 2019, 2020; Sytov et al. 2023; Kim & Cho 2025a). In contrast, astronuclear physics simulations are based on nuclear many-body approaches, such as nuclear lattice effective field theory (NLEFT) and lattice quantum chromodynamics (lattice QCD), to describe nuclear structure and reactions (Kim & Cho 2015; Kim & Cho 2025b).

As supercomputing capabilities advance from the petascale (10^{15} FLOPS) to the exascale (10^{18} FLOPS), these two domains are increasingly converging, enabling unified, time-dependent, multi-physics descriptions of compact-object phenomena. This trend further reinforces the role of simulation as a key component complementing theory and experiment. Together, these complementary approaches provide a unified description of high-energy phenomena and element formation in the universe. Modern astroparticle physics and astronuclear physics aim to understand the

Received 07 APR 2026 **Revised** 08 JUN 2026 **Accepted** 22 AUG 2026
†**Corresponding Author**
Tel: +82-42-869-0722, E-mail: cho@kisti.re.kr
ORCID: https://orcid.org/0000-0003-1705-7399

Table 1. Comparison between astroparticle physics and astronuclear physics

Category	Astroparticle physics	Astronuclear physics
Experiments / Observations	Underground detectors, space-based observatories	Accelerator facilities (e.g., RAON), nuclear experiments
Simulations	Geant4-based detector simulations, cosmic-ray (CR) propagation models	Nuclear many-body methods (NCSM, coupled cluster, IM-SRG, HFB)
Computational Focus	Particle transport, interaction with detectors, signal reconstruction	Nuclear structure, reaction rates, equation of state (EOS)

universe's most extreme environments such as the interiors of neutron stars, the moments surrounding core-collapse supernovae, and the violent mergers of compact objects. These phenomena involve physical conditions that are far beyond what can be reproduced in terrestrial laboratories, making them accessible only through a combination of theoretical models, observational data, and increasingly, high-performance simulations (Cho 2008; Cho & Kim 2009; Cho et al. 2010; Cho et al. 2011; Cho et al. 2012; Cho et al. 2015).

Over the past two decades, simulation-based science has emerged as the "third pillar" of scientific inquiry, complementing traditional theory and observation (Cho 2016a, b; Cho 2017). Fig. 1 shows the convergence of theory, simulation, and observation/experiment for astroparticle physics and astronuclear physics. Numerical simulations serve as a virtual laboratory, enabling the dynamic modeling of multi-scale, multi-physics processes under conditions of extreme density, temperature, and gravitational curvature. They are especially vital in astrophysics where direct experimentation is often impossible (Cho 2017).

Internationally, high-performance computing is undergoing a similar transition toward heterogeneous, accelerator-based exascale architectures. The U.S. Frontier system demonstrated the first sustained exascale performance, followed by El Capitan and Aurora, while Europe has entered the exascale era with JUPITER. These

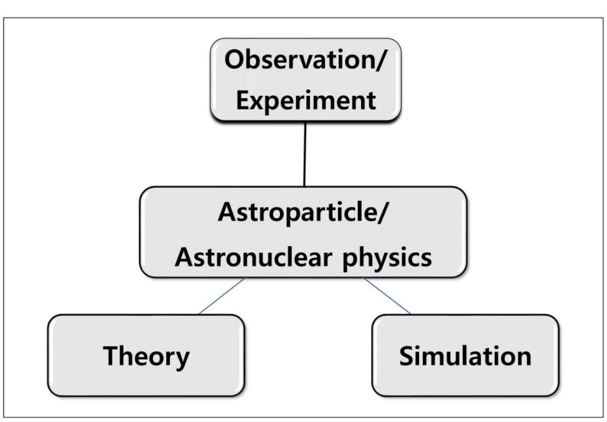

Fig. 1. Convergence of theory, simulation, and observation/experiment in astroparticle physics and astronuclear physics.

developments indicate a global shift from CPU-centric computing toward heterogeneous architectures combining large-scale simulation, accelerator technologies, and AI-assisted scientific workflows (TOP500 2025).

In Korea, the advancement of simulation-based astrophysics has been closely tied to the evolution of its national supercomputing infrastructure. The KISTI-5 supercomputer (Nurion), with 25.7 peta floating-point operations per second (PFLOPS) of theoretical performance and a CPU-centric architecture, has enabled a wide range of research activities (Kim & Cho 2025b). These include the following research areas:

- Lattice simulations of neutron-rich isotopes relevant to r-process nucleosynthesis;
- Theoretical predictions of the neutron drip line using wavefunction-based nuclear models;
- EOS studies for neutron star matter;
- Modeling of cosmic-ray–induced backgrounds for dark matter detection;
- RAON experiment planning and rare isotope reaction network simulations (Kim & Cho 2025a).

These studies have not only deepened our understanding of fundamental nuclear and astrophysical processes but have also informed experimental strategies and data interpretation for large-scale collaborations such as RAON (Park et al. 2026).

Looking forward, the upcoming KISTI-6 supercomputer (Hangang) represents a transformative leap in capability. With over 600 PFLOPS of projected peak performance and a heterogeneous architecture integrating thousands of GPUs, the KISTI-6 supercomputer will allow simulations to reach unprecedented resolution, scale, and complexity. It will enable researchers to explore time-sensitive and resource-intensive problems such as 3D neutrino-radiation transport in supernovae, r-process nucleosynthesis during neutron star mergers, and gravitational wave signal modeling with full general relativistic treatment.

However, harnessing this power demands a paradigm shift in computational strategy. Legacy CPU-based simulation codes such as NLEFT, lattice QCD solvers, and Geant4-based detector simulations must be adapted for GPU execution through CUDA. Moreover, artificial intelligence (AI)-assisted surrogate models are increasingly being integrated to accelerate simulations and optimize exploration of high-dimensional parameter spaces (Kim & Cho 2022).

Finally, simulation-based research will not only bridge theory and observation but also serve as a central pillar in Korea's contribution to global astrophysical discovery in the decade ahead.

2. ASTROPHYSICAL AND ASTRONUCLEAR PHENOMENA REQUIRING LARGE-SCALE SIMULATIONS

Fig. 2 shows a simplified and closed-loop view of the stellar life cycle and the associated nucleosynthesis processes. The diagram summarizes how baryonic matter is continuously recycled through successive generations of stars, linking stellar evolution to the chemical enrichment of the universe. Massive stars evolve from molecular clouds to main-sequence objects where hydrogen fusion dominates the energy production. As nuclear fuel is exhausted, advanced burning stages lead to the formation of heavier elements up to the iron group through successive fusion processes. The final fate of many massive stars is core-collapse supernovae which play a central role in dispersing heavy elements into the interstellar medium (Li 2008; Abbott et al. 2017).

These explosive environments provide the necessary conditions for r-process nucleosynthesis, producing neutron-rich nuclei beyond iron. The diagram also highlights the formation of compact remnants, such as neutron stars, whose internal structure and stability are governed by the properties of dense neutron-rich matter. In this context, the EOS is a key ingredient in determining neutron star masses, radii, and tidal deformabilities.

Accurate modeling of these phenomena requires the simultaneous treatment of neutrino transport, nuclear reaction networks, and hydrodynamic evolution. Neutrinos regulate the dynamics of core-collapse supernovae and influence nucleosynthesis pathways, while large-scale reaction networks determine the final isotopic abundances ejected into space.

Fig. 2 further emphasizes the strong connection between theoretical modeling and modern observations. Multi-messenger data from gravitational-wave detectors, neutrino observatories, and gamma-ray telescopes provide direct constraints on stellar explosions, compact objects, and element formation. As such, stellar life cycle simulations serve as a critical bridge between nuclear microphysics and astrophysical observations. Gravitational wave observatories such as LIGO, Virgo, and KAGRA have enabled the detection of neutron star mergers, providing constraints on the equation of state and nucleosynthesis processes (Aso et al. 2013; Abbott et al. 2017; Akutsu et al. 2021). Neutrino observatories including Super-Kamiokande

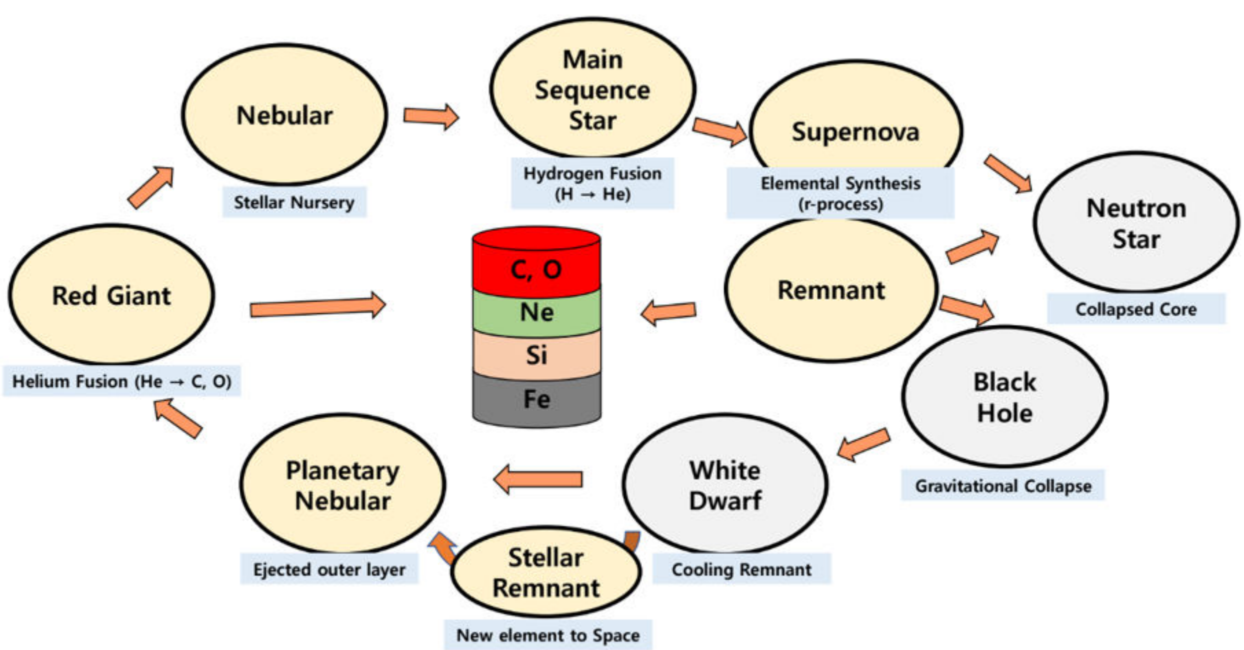

Fig. 2. Stellar life cycle and associated nucleosynthesis processes. Data from Li (2008) and Abbott et al. (2017).

and IceCube probe core-collapse supernovae and high-energy astrophysical neutrinos (Fukuda et al. 1998; IceCube Collaboration et al. 2013; Abe et al. 2014). In addition, large-scale surveys such as the Square Kilometre Array (SKA) and the Vera C. Rubin Observatory (LSST) trace electromagnetic counterparts of transient astrophysical events (Dewdney et al. 2013; Ivezić et al. 2019; Braun et al. 2020).

Astrophysical environments are governed by extreme conditions such as high densities, intense gravitational fields, rapid rotations, and strong magnetic forces that are difficult or impossible to reproduce in terrestrial laboratories. Consequently, numerical simulations have become indispensable for exploring the physical processes that shape the cosmos (Cho 2016b). Among the most critical targets for computational astrophysics are:

- r-process nucleosynthesis: The rapid neutron capture process is responsible for producing approximately half of the elements heavier than iron in the universe. It is believed to occur during explosive astrophysical events such as core-collapse supernovae and neutron star mergers. However, the detailed nuclear reaction pathways involve hundreds of unstable, neutron-rich isotopes that are inaccessible to current experimental techniques. Accurate simulations that couple hydrodynamics, neutrino transport, and large nuclear reaction networks are essential to understand where and how the r-process operates (Li 2008).
- Neutrino-driven winds: In the aftermath of a core-collapse supernova, neutrinos escaping the proto-neutron star deposit energy into the surrounding matter, driving outflows known as neutrino-driven winds. These winds are promising sites for nucleosynthesis, particularly of lighter heavy elements like Sr, Y, and Zr. Modeling these winds requires solving the coupled neutrino-radiation-hydrodynamics equations under relativistic conditions (IceCube Collaboration et al. 2013; Abe et al. 2014). Moreover, the interplay of flavor oscillations and weak interactions in such environments remains an open challenge that can only be addressed by high-resolution simulations.
- Neutron-rich Matter, Nuclear EOS, and Neutron Drip Line Simulations: The properties of neutron-rich matter play a central role in determining the nuclear EOS, which governs the structure and dynamics of neutron stars, including their radii, maximum masses and tidal deformabilities. A key link between finite nuclei and infinite nuclear matter is provided by the neutron drip line, which marks the boundary beyond which additional neutrons can no longer be bound to a nucleus. The location of the drip line is strongly sensitive to the density dependence of the nuclear symmetry energy and many-body correlations (Li 2008; Qu 2013), both of which are fundamental ingredients of the EOS at sub- and near-saturation densities. Consequently, neutron drip line simulations serve as a critical benchmark for EOS models applied to neutron star crusts and neutron-rich astrophysical environments.

Fig. 3 shows the nuclear chart, illustrating the connection between nuclear structure and cosmic evolution (Kortelainen et al. 2010; Erler et al. 2012). The proton-rich region (upper-left) is associated with proton-rich nucleosynthesis pathways relevant to the early evolution of the universe. The valley of stability represents nuclei described within the framework of the Standard Model and experimentally accessible regions. The neutron-rich region (right) corresponds to the evolution of the universe, including r-process nucleosynthesis in extreme astrophysical environments such as neutron star mergers and core-collapse supernovae. The progression from ab initio approaches to density functional theory indicates the theoretical frameworks required to describe nuclei across the nuclear landscape (Navrátil et al. 2000).

From a theoretical perspective, different nuclear many-body approaches are required to describe various regions of the chart. Ab initio methods are applicable to light nuclei, while shell-model and density functional theory approaches extend the description toward medium and heavy nuclei (Navrátil et al. 2000; Saxena et al. 2018). The neutron-rich frontier thus represents a critical boundary where nuclear theory, astrophysical observations, and large-

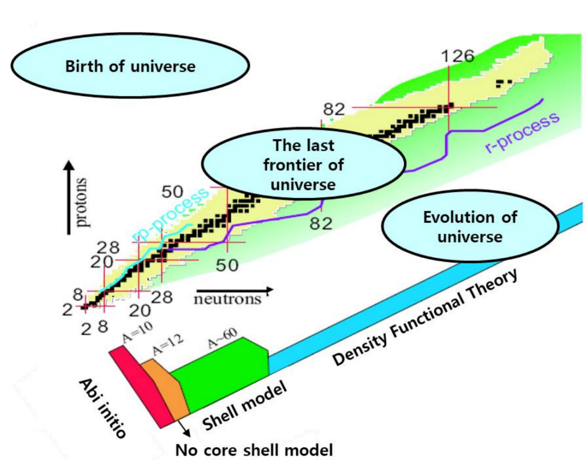

Fig. 3. Nuclear chart. Data from Kortelainen et al. (2010) and Erler et al. (2012).

scale simulations converge, highlighting the importance of advanced computational resources for exploring the limits of nuclear existence and their implications for the evolution of the universe.

The neutron-rich frontier of the nuclear chart represents the last unexplored region that connects nuclear structure with the astrophysical evolution of the universe.

3. ASTROPARTICLE AND ASTRONUCLEAR SIMULATIONS ON THE KISTI-5 SUPERCOMPUTER

The KISTI-5 supercomputer has played a pivotal role in advancing astronuclear simulations. With a theoretical peak performance of 25.7 PFLOPS and over 560,000 cores, Nurion enabled several grand challenge projects that pushed the frontier of computational astrophysics.

3.1 Nuclear Lattice Effective Field Theory and Rare Isotope Structure

A central thrust of these efforts was the NLEFT simulations for the structure of exotic, neutron-rich nuclei. These ab initio simulations employed Monte Carlo techniques and auxiliary field methods to describe the clustering and excitation properties of isotopes such as $^{18–26}$O and $^{12–22}$C, which are relevant to both r-process nucleosynthesis and observational studies of astronuclear physics. By calibrating low-energy constants from chiral effective field theory and performing large-scale ensemble averaging, these studies yielded charge radii and binding energies in agreement with experimental data, including upcoming rare isotope beam experiments at RAON (Kim & Cho 2025b).

3.2 Neutron Drip Line and Equation of State Constraints

Large-scale simulations performed on the KISTI-5 supercomputer have enabled systematic investigations of neutron-rich isotopes near the drip line using advanced many-body approaches. These studies combined ab initio-inspired interactions with wave-function fitting and large sparse-matrix techniques, requiring massive parallel resources to handle the computational complexity. The resulting predictions of binding energies, separation energies, and charge radii for neutron-rich nuclei provide quantitative constraints on the symmetry energy and its slope, thereby anchoring EOS models used in astrophysical simulations. In particular, these results establish a consistent connection between finite nuclear systems and neutron-rich matter relevant to neutron star crusts, laying a robust foundation for next-generation EOS studies and neutron star merger simulations on forthcoming exascale platforms such as the KISTI-6 supercomputer (Elhatisari et al. 2024).

Neutron drip line predictions based on ab initio calculations of neutron-rich nuclei provide direct constraints on the density dependence of the symmetry energy, thereby anchoring the nuclear equation of state relevant to neutron-star crust and neutron-rich matter (Hebeler et al. 2010; Hagen et al. 2016; Drischler et al. 2019).

3.3 RAON Experiment Simulations and Nuclear Reaction Networks

Furthermore, large-scale computational modeling of RAON-based experiments was conducted. Given the limited beam time at rare isotope facilities, numerical simulations were used to optimize target-projectile combinations, predict cross sections, and guide the design of detectors. These simulations also helped identify promising isotopes for measuring key astrophysical reaction rates, such as (n,γ), (p,γ), and (α,γ) channels critical to the rp- and r-processes as shown in Fig. 4. The nuclear reaction pathways start from hydrogen and progress through helium, carbon, oxygen, silicon, and beyond, including upper and lower branch reactions (e.g., C-12 → O-16, O-14 → F-15). This figure visually represents the flow of rare isotope reactions that can be observed or experimentally studied at RAON (Jeon et al. 2014; Chung et al. 2023; Choi et al. 2025).

3.4 Dark Matter and Background Simulations

Dark matter has been extensively studied in high-energy collider experiments, where missing energy signatures are used to search for physics beyond the Standard Model (Yeo & Cho 2018). Complementary to these efforts, underground direct detection experiments aim to observe rare interactions between dark matter particles and detector materials. In this context, accurate modeling of

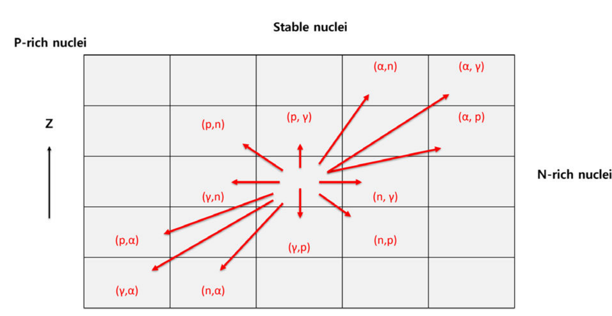

Fig. 4. Nuclear reaction network relevant to astronuclear physics.

environmental and cosmogenic backgrounds becomes essential.

Another line of research supported on Nurion focused on dark matter background modeling (Park & Cho 2021a, b; Park et al. 2022; Park et al. 2024). Detailed Monte Carlo simulations were employed to estimate cosmic-ray–induced neutron and muon backgrounds in underground laboratories. Such background estimates are crucial for interpreting signals in next-generation direct detection experiments, where sub-keV precision is required.

Finally, Nurion was also utilized for lattice QCD calculations and nuclear EOS construction at finite density and temperature, contributing to the modeling of neutron star matter and supernova core dynamics. These inputs serve as essential components in multi-messenger models of compact object mergers observed by gravitational wave detectors. Taken together, these achievements illustrate the power of HPC in tackling the most pressing questions in astronuclear physics. Nurion enabled a synergistic integration of nuclear theory, astrophysical modeling, and experimental planning that is indispensable in the era of precision astrophysics.

4. TRANSITION TO EXASCALE: THE KISTI-6 SUPERCOMPUTER

Today, modern science is increasingly driven by big data and deep learning, enabled by the integration of experiments, theory, and computational simulations (Cho 2008; Cho & Kim 2009; Cho et al. 2010; Cho et al. 2011; Cho et al. 2012; Cho et al. 2015). Fig. 5 illustrates the convergence of computational paradigms in astroparticle physics and astronuclear physics, facilitated by the KISTI-6 GPU-based supercomputer through AI, big data, and advanced computational techniques.

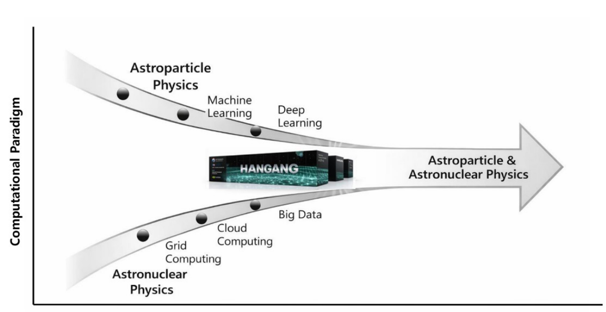

Fig. 5. Computational paradigms in astroparticle physics and astronuclear physics. The KISTI-6 GPU-based supercomputer (Hangang) enables advances through artificial intelligence, big data, and advanced computational techniques.

The evolution of supercomputing architectures has progressed from traditional CPU-based systems to more advanced frameworks incorporating Grid and cloud computing (Jeung et al. 2009), and is now transitioning toward heterogeneous computing. While the KISTI-5 supercomputer is a CPU-only system, the KISTI-6 supercomputer adopts a GPU-based heterogeneous architecture. In this paper, we present recent developments in astroparticle physics and astronuclear physics and discuss the scientific opportunities enabled by the KISTI-6 supercomputer in the following sections.

The transition from traditional CPU-based architectures to heterogeneous HPC platforms represents a major evolution in computational science for astroparticle and astronuclear physics. The forthcoming KISTI-6 supercomputer is designed to deliver an aggregate peak performance exceeding 600 PFLOPS by integrating thousands of next-generation GPUs with high-bandwidth interconnects. This leap represents more than a 20-fold increase in computational capability compared to its predecessor, Nurion. Table 2 compares the architectures and specifications of KISTI-5 (Nurion) and KISTI-6 (Hangang) (Kim & Cho 2025b). Table 3 summarizes representative computational limitations encountered on Nurion together with the expected scientific and computational capabilities enabled by the KISTI-6 supercomputer.

Recent studies on GPU-enabled platforms have demonstrated the feasibility of GPU acceleration for nuclear many-body simulations. In particular, the NLEFT code has been ported to the KISTI Neuron system equipped with NVIDIA V100 GPUs using CUDA and OpenMPI, while the reference CPU platform is the KISTI-5 Nurion system based on Intel Xeon Phi 7250 processors. Fig. 6 compares the relative computational performance of a single NVIDIA V100 GPU with that of a single Intel Xeon Phi 7250 processor for representative NLEFT calculations as a function of nuclear mass number A. The single-GPU/single-processor performance ratio increases from approximately 55% at A = 6 to 60% at A = 12 and 70% at A = 24. Thus, the figure is not intended to represent a conventional speedup factor greater than unity, but rather to demonstrate that the relative computational efficiency of the GPU improves as the nuclear mass number increases. These benchmark results also demonstrate the successful porting of the CPU-based NLEFT code to the GPU environment and suggest that GPU acceleration becomes increasingly advantageous for computationally more demanding nuclear systems.

Although these benchmarks were performed using a single NVIDIA V100 GPU and a single Intel Xeon Phi 7250 processor,

Table 2. Comparison of the KISTI-5 (Nurion) and KISTI-6 (Hangang) supercomputers

Feature	KISTI-5 (Nurion)	KISTI-6 (Hangang)
Peak performance (theoretical)	25.7 PFLOPS	> 600 PFLOPS
Architecture	CPU-only (KNL/Skylake)	GPU (NVIDIA GH200) + CPU
Memory	797 TB (KNL)	1.5 PB (GPU) + 0.75 PB (CPU)
Storage	21 PB disk, 10 PB tape	20 PB flash + 180 PB HDD + 40 PB tape
Interconnect	100 Gbps	400 Gbps
Service period	2018–present	2026–future

PFLOPS, peta floating-point operations per second.
Data from Kim & Cho (2025b).

Table 3. Computational limitations of KISTI-5 and expected improvements with KISTI-6

Category	KISTI-5 (Nurion) limitation	KISTI-6 expected capability
NLEFT simulations	Practical limit around A ≈ 50	Extended toward A ≈ 100
NCSM basis size	Restricted Nmax due to memory	Larger Nmax and model spaces
Sparse matrix calculation	Communication bottleneck at large node count	GPU-accelerated sparse operations
Monte Carlo statistics	Limited ensemble size	Larger ensemble and uncertainty quantification
Runtime	Long wall-clock time (~3000 s cases)	Reduced execution time via GPU acceleration
Parallel scalability	Early saturation for communication-intensive jobs	Improved heterogeneous scaling
Parameter-space exploration	Computationally expensive	AI-assisted surrogate exploration
Multi-messenger analysis	Offline post-processing	Near-real-time inference workflows
Detector simulation	CPU-dominated Geant4 workflows	GPU-accelerated event processing
Space weather / Magnetohydrodynamic (MHD)	Limited resolution and coupling	Multi-scale real-time simulation potential

the computational advantage of KISTI-6 is expected to emerge at the system level through its large number of next-generation NVIDIA GH200 GPUs and high-bandwidth interconnect. The substantially enhanced computational and memory capabilities of the GH200 architecture, combined with large-scale GPU parallelism, are expected to provide a more favorable environment for GPU-optimized NLEFT calculations than the CPU-centric KISTI-5 system. The higher interconnect bandwidth of KISTI-6 is also expected to mitigate communication bottlenecks in large multi-GPU calculations. The projected extension toward A ≈ 100 shown in Table 3 should therefore be interpreted as a system-level expectation for the KISTI-6 supercomputer rather than as a direct extrapolation from the benchmark points in Fig. 6. However, the achievable performance will ultimately depend on application-specific GPU optimization, memory requirements, and multi-node scaling efficiency.

Fig. 7 illustrates several conceptual scaling scenarios for heterogeneous computing architectures. The effective overhead parameter α is introduced as a qualitative measure of aggregate communication, synchronization, and memory-access costs. The value α = 0.15 was empirically

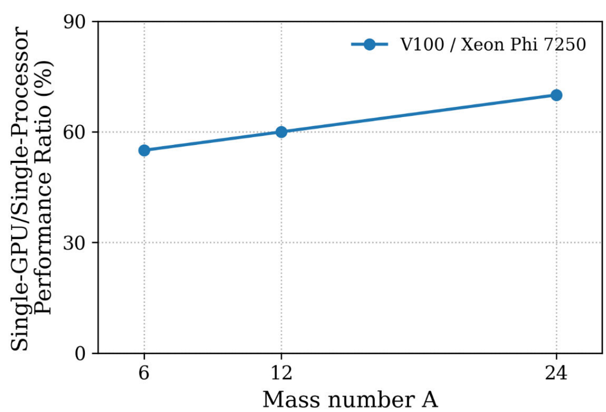

Fig. 6. Single-GPU/single-processor performance ratio for nuclear lattice effective field theory (NLEFT) calculations using NVIDIA V100 and Intel Xeon Phi 7250. The relative GPU performance increases with nuclear mass number A.

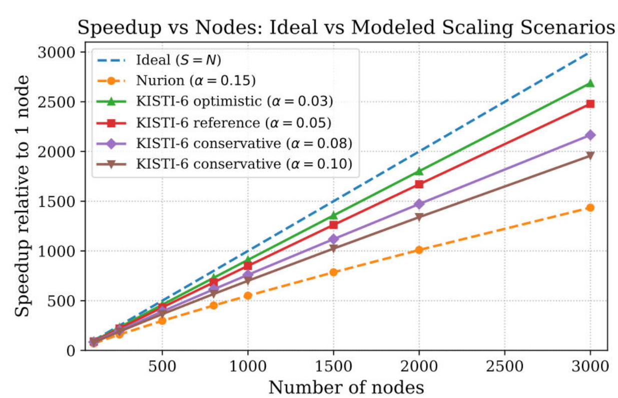

Fig. 7. Conceptual scaling scenarios for the KISTI-6 heterogeneous architecture. Different effective overhead parameters (α = 0.03–0.10) represent optimistic, reference, and conservative GPU-scaling assumptions, compared with ideal linear scaling and the Nurion CPU-based system.

selected to represent the communication-limited behavior typically observed in CPU-centric architectures such as Nurion, while smaller values (α = 0.03–0.10) illustrate possible scaling scenarios for future heterogeneous GPU-accelerated systems. Because production-scale benchmark data for the KISTI-6 supercomputer are not yet available, these curves are intended to illustrate qualitative scaling behavior rather than predictive performance estimates.

5. EXASCALE APPLICATIONS IN ASTROPARTICLE PHYSICS AND ASTRONUCLEAR PHYSICS

5.1 Convergence of Astroparticle Physics and Astronuclear Physics at Exascale

Advances in HPC have played a decisive role in shaping modern research in both astroparticle physics and astronuclear physics. Although these two fields have traditionally been distinguished by their physical targets and computational requirements, the rapid increase in available computing power from the petascale to the exascale regime has progressively narrowed this distinction. As supercomputing capability grows, the scientific scope accessible to each domain not only expands but also increasingly overlaps, ultimately enabling a unified description of compact astrophysical phenomena. The two research domains differ primarily in their physical focus and numerical requirements, as summarized in Table 4. For astroparticle physics, the dominant theories are quantum field theory (QFT), Standard Model and Beyond the Standard Model (Cho & Nam 2013; Cho et al. 2015).

At exascale performance levels, the traditional boundary between astroparticle physics and astronuclear physics becomes increasingly blurred. While the former has historically benefited from massive event-level parallelism and statistical inference, and the latter has demanded extreme numerical precision for strongly correlated many-body systems, GPU-accelerated architectures enable both paradigms to coexist within a unified computational framework. This convergence allows microscopic nuclear physics to be directly connected with macroscopic astrophysical observables, such as neutrino emission and gravitational-wave signals, in a single, time-dependent simulation pipeline.

Furthermore, the KISTI-6 supercomputer supports a sophisticated software stack, including containerization environments (Docker, Singularity), GPU-tuned numerical libraries, and a range of job schedulers for managing resource-intensive workflows. This ecosystem facilitates reproducible and scalable computational experiments, which are increasingly critical in the context of multi-messenger astrophysics and data-driven discovery.

Fig. 8 shows the progress in HPC for astroparticle physics and astronuclear physics. The advent of exascale HPC architectures, such as the KISTI-6 supercomputer, significantly expands the horizon for multi-physics and multi-scale simulations in computational astrophysics. Several challenging applications, previously deemed infeasible due to computational limitations, are now within reach.

The transition from the KISTI-5 supercomputer to the GPU-accelerated KISTI-6 system is expected to produce not only higher raw computational performance but also qualitatively new scientific capabilities. On CPU-centric systems such as Nurion, large-scale ab initio calculations were often restricted by memory capacity, sparse-matrix communication overhead, and practical wall-clock limitations. The shell-model calculations involving O (10^{13}–10^{14}) sparse-matrix elements required up to approximately 1,000 computational nodes, while large NLEFT simulations utilized as many as 1,500 exclusive nodes and more than 100,000 CPU cores simultaneously. The KISTI-6 heterogeneous architecture is expected to substantially extend accessible model spaces, Monte Carlo ensemble sizes, and parameter-space coverage for future nuclear and astrophysical simulations (Kim & Cho 2025b).

5.2 Key Astrophysical Applications Enabled by Exascale HPC

One such application is multi-dimensional supernova neutrino transport. The precise modeling of neutrino-

Table 4. Comparison of astroparticle physics and astronuclear physics from a computational perspective

Category	Astroparticle physics	Astronuclear physics
Primary targets	Neutrinos, dark matter, cosmic rays, gravitational waves	Extreme nuclear matter, nuclear reactions, neutron stars, supernova nuclei
Dominant theory	Quantum field theory, Standard Model and beyond Standard Model	Nuclear forces (EFT), many-body nuclear structure
Computational nature	Large-scale statistical simulations, event-based MC, ML	High-precision many-body calculations, matrix and lattice methods
Parallel characteristics	Event/sample independence → strong embarrassingly parallel scaling	Strong inter-particle coupling → communication-intensive scaling

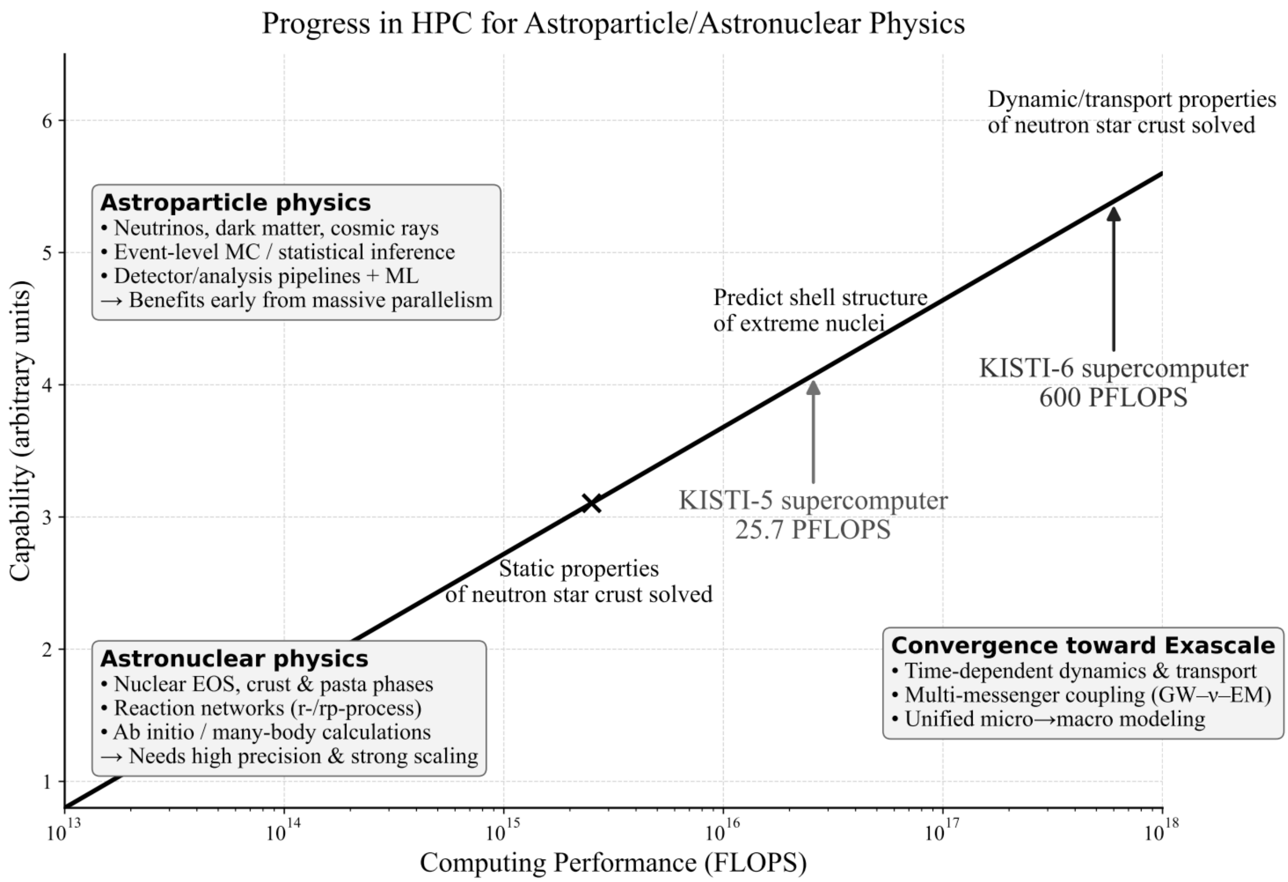

Fig. 8. The progress in HPC for astroparticle physics and astronuclear physics. Data from Kim & Cho (2025b).

matter interactions in core-collapse supernovae is crucial for understanding the explosion mechanism and the subsequent synthesis of heavy elements. Solving the Boltzmann neutrino transport equation in full 3D with energy and angle dependence requires massive parallelization and optimized GPU kernels. Exascale platforms enable these simulations with increased fidelity, resolution, and physical realism.

Another transformative application is the modeling of gravitational wave signals from binary neutron star mergers. These events are not only sources of gravitational waves but also sites of r-process nucleosynthesis and short gamma-ray bursts. By combining general relativistic hydrodynamics with detailed microphysical inputs such as nuclear EOS and neutrino radiation, realistic gravitational waveform predictions can be generated and compared with LIGO/Virgo observations. Machine-learned emulators trained on simulation data sets may also assist in interpreting multi-messenger signals in real time.

Moreover, Geant4-based detector simulations are increasingly moving to large-scale GPU environments to meet the needs of upcoming astrophysical and cosmological experiments. These include simulations of radiation backgrounds for underground neutrino detectors, optimization of shielding and veto systems for dark-matter experiments, and cosmic ray propagation studies for space-based instruments. GPU readiness differs among the major simulation frameworks considered here. For lattice QCD, GPU acceleration is already well established through libraries such as QUDA, including mixed-precision and multi-GPU solvers (Babich et al. 2010). In contrast, standard Geant4 workflows remain primarily CPU-oriented, although GPU-based particle transport is being actively developed through frameworks such as AdePT (Amadio et al. 2023). This distinction highlights the different migration strategies required for future KISTI-6 applications.

Fig. 9 illustrates a hybrid AI–HPC workflow expected to become increasingly important in GPU-accelerated exascale environments. Large-scale first-principles simulations performed on GPU-accelerated supercomputers generate training data for scientific machine learning models, including physics-informed neural networks (PINNs) and neural emulators. These AI-assisted surrogate models enable accelerated parameter exploration, uncertainty

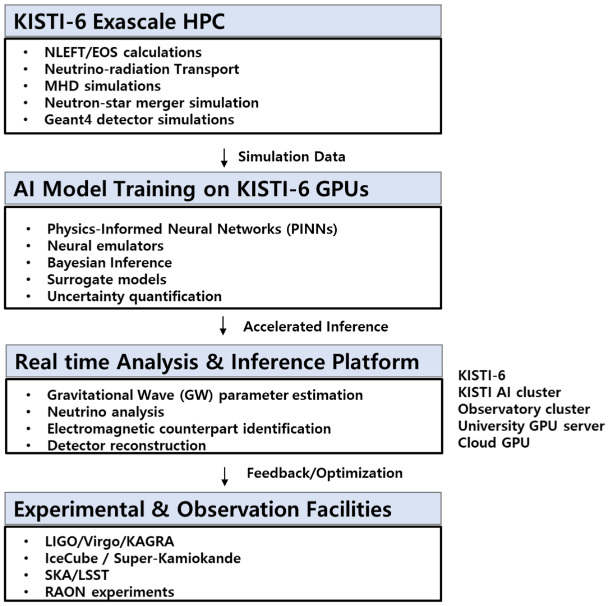

Fig. 9. Hybrid AI–HPC workflow for exascale astrophysical simulations.

quantification, and real-time multi-messenger analysis while preserving physical consistency through embedded conservation laws and physics constraints. In this framework, AI does not replace first-principles simulations but complements HPC workflows by accelerating parameter exploration and uncertainty quantification (Field et al. 2014; Radovic et al. 2018; Raissi et al. 2019; Cuomo et al. 2022).

Beyond astroparticle and astronuclear physics, exascale GPU architectures are also expected to enable large-scale three-dimensional MHD simulations for space-weather forecasting. Such simulations can couple solar-wind dynamics, magnetospheric responses, and kinetic plasma processes across multiple spatial and temporal scales. Future heterogeneous systems such as KISTI-6 may support near-real-time prediction of extreme space-weather events through GPU-accelerated MHD and kinetic simulations. These capabilities would further broaden the scientific impact of exascale computing across the wider space-science community (White et al. 2016).

6. SUMMARY

Over the past decade, HPC infrastructure in Korea has played a crucial role in advancing astroparticle and astronuclear physics. The KISTI-5 supercomputer enabled a broad range of simulation-based studies, from nuclear structure and neutron drip line predictions to dark matter background modeling and experimental support for RAON. These efforts demonstrated the essential role of large-scale computation in bridging theory, experiment, and observation.

The deployment of the KISTI-6 supercomputer marks a decisive step toward the exascale era. Its heterogeneous CPU–GPU architecture and advanced software ecosystem will allow simulations of unprecedented scale, accuracy, and physical realism. In particular, the KISTI-6 supercomputer will enable time-dependent, multi-physics studies of supernovae, neutron star mergers, and dense nuclear matter, while supporting real-time multi-messenger data analysis through AI-assisted workflows.

As simulations increasingly serve as the third pillar of scientific discovery, alongside theory and observation, continued investment in HPC infrastructure, algorithm development, and international collaboration will be essential. With the KISTI-6 supercomputer, we are well positioned to make leading contributions to global efforts aimed at understanding the most extreme environments in the universe.

ACKNOWLEDGMENTS

We thank Kyungho Kim, Youngman Kim, Young-Ho Song, Ik Jae Shin and Myugkuk Kim. This research is supported by the major institutional R&D program, KISTI (K26L1M2C3, N25NM055-26). This study was also supported by National Supercomputing Center with supercomputing resources including technical support (KSC-2022-CHA-0003, KSC-2023-CHA-0005, KSC-2024-CHA-0001 and KSC-2025-CHA-0004).

ORCID

Kihyeon Cho https://orcid.org/0000-0003-1705-7399

REFERENCES

Abbott BP, Abbott R, Abbott TD, Acernese F, Ackley K, et al., GW170817: observation of gravitational waves from a binary neutron star inspiral, Phys. Rev. Lett. 119, 161101 (2017). https://doi.org/10.1103/PhysRevLett.119.161101

Abe K, Hayato Y, Iida T, Iyogi K, Kameda J, et al., Calibration of the Super-Kamiokande detector, Nucl. Instrum. Methods Phys. Res. A. Accel. Spectrom. Detect. Assoc. Equip. 737, 253-272 (2014). https://doi.org/10.1016/j.nima.2013.11.081

Akutsu T, Ando M, Arai K, Araki S, Araya A, et al., Overview of KAGRA: detector design and construction history, Prog. Theor. Exp. Phys. 2021, 05A101 (2021). https://doi.org/10.1093/ptep/ptaa125

Amadio G, Apostolakis J, Buncic P, Cosmo G, Dosaru D, et al., Offloading electromagnetic shower transport to GPUs, J. Phys. Conf. Ser. 2438, 012055 (2023). https://doi.org/10.1088/1742-6596/2438/1/012055

Aso Y, Michimura Y, Somiya K, Ando M, Miyakawa O, et al., Interferometer design of the KAGRA gravitational wave detector, Phys. Rev. D 88, 043007 (2013). https://doi.org/10.1103/PhysRevD.88.043007

Babich R, Clark MA, Joó B, Parallelizing the QUDA library for multi-GPU calculations in lattice quantum chromodynamics, Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC10), New Orleans, LA, 13-19 Nov 2010.

Braun R, Bonaldi A, Bourke T, Keane EF, Wagg J, et al., Anticipated SKA1 science performance, Publ. Astron. Soc. Aust. 37, e026 (2020). https://doi.org/10.1017/pasa.2020.19

Cho K, e-Science for high energy physics in Korea, J. Korean Phys. Soc. 53, 1187-1191 (2008). https://doi.org/10.3938/jkps.53.1187

Cho K, Computational science and the search for dark matter, New Phys. Sae Mulli 66, 950-956 (2016a). https://doi.org/10.3938/NPSM.66.950

Cho K, e-Science paradigm for astroparticle physics at KISTI, J. Astron. Space Sci. 33, 63-67 (2016b). https://doi.org/10.5140/JASS.2016.33.1.63

Cho K, Computational science-based research on dark matter at KISTI, J. Astron. Space Sci. 34, 153-159 (2017). https://doi.org/10.5140/JASS.2017.34.2.153

Cho K, Kim H, Heavy flavor physics through e-science, J. Korean Phys. Soc. 55, 2045-2050 (2009). https://doi.org/10.3938/jkps.55.2045

Cho K, Kim H, Jung M, Cyberinfrastructure for high energy physics in Korea, J. Phys. Conf. Ser. 219, 072032 (2010). https://doi.org/10.1088/1742-6596/219/7/072032

Cho K, Kim J, Nam SH, Collider physics based on e-Science paradigm of experiment-computing-theory, Comput. Phys. Commun. 182, 1756-1759 (2011). https://doi.org/10.1016/j.cpc.2010.12.019

Cho K, Bae T, Kim J, Kim Y, A research on the theory-experiment-simulation paradigm for probing the Standard Model, New Phys. Sae Mulli 62, 1143-1150 (2012). https://doi.org/10.3938/NPSM.62.1143

Cho K, Nam SH, Right-handed current contributions in $B \to K\pi$ decays, Phys. Rev. D. 88, 035012 (2013). https://doi.org/10.1103/PhysRevD.88.035012

Cho K, Kim J, Kim JH, Research and development of the evolving architecture for beyond the Standard Model, J. Phys. Conf. Ser. 664, 072011 (2015). https://doi.org/10.1088/1742-6596/664/7/072011

Choi S, Chung YS, Hong IS, Lee J, Jung Y, et al., The RAON facility: an overview, J. Korean Phys. Soc. 87, 447-454 (2025). https://doi.org/10.1007/s40042-025-01435-1

Choi W, Cho K, Yeo I, Performance profiling for brachytherapy applications, Comput. Phys. Commun. 226, 180-186 (2018). https://doi.org/10.1016/j.cpc.2017.12.022

Chung YS, Hong SW, Jang HM, Ki TK, Kim HJ, et al., Commissioning preparation of RAON rare isotope accelerator facility, J. Korean Phys. Soc. 82, 613-621 (2023). https://doi.org/10.1007/s40042-023-00715-y

Cuomo S, Di Cola VS, Giampaolo F, Rozza G, Raissi M, et al., Scientific machine learning through physics-informed neural networks: where we are and what's next, J. Sci. Comput. 92, 88 (2022). https://doi.org/10.1007/s10915-022-01939-z

Dewdney PE, Lazio J, Diamond P, SKA1 system baseline design, SKA Organization Technical Document, SKA-TEL-SKO-DD-001(2013).

Drischler C, Hebeler K, Schwenk A, Chiral interactions up to next-to-next-to-next-to-leading order and nuclear saturation, Phys. Rev. Lett. 122, 042501 (2019). https://doi.org/10.1103/PhysRevLett.122.042501

Elhatisari S, Bovermann L, Ma YZ, Epelbaum E, Frame D, et al., Wavefunction matching for solving quantum many-body problems, Nature. 630, 59-63 (2024). https://doi.org/10.1038/s41586-024-07422-z

Erler J, Birge N, Kortelainen M, Nazarewicz W, Olsen E, et al., The limits of the nuclear landscape, Nature. 486, 509-512 (2012). https://doi.org/10.1038/nature11188

Field SE, Galley CR, Hesthaven JS, Kaye J, Tiglio M, Fast prediction and evaluation of gravitational waveforms using surrogate models, Phys. Rev. X. 4, 031006 (2014). https://doi.org/10.1103/PhysRevX.4.031006

Fukuda Y, Hayakawa T, Ichihara E, Inoue K, Ishihara K, et al., Evidence for oscillation of atmospheric neutrinos, Phys. Rev. Lett. 81, 1562 (1998). https://doi.org/10.1103/PhysRevLett.81.1562

Hagen G, Ekström A, Forssén C, Jansen GR, Nazarewicz W, et al., Neutron and weak-charge distributions of the 48Ca nucleus, Nat. Phys. 12, 186-190 (2016). https://doi.org/10.1038/nphys3529

Hebeler K, Lattimer JM, Pethick CJ, Schwenk A, Constraints on neutron star radii based on chiral effective field theory interactions, Phys. Rev. Lett. 105, 161102 (2010). https://doi.org/10.1103/PhysRevLett.105.161102

IceCube Collaboration, Evidence for high-energy extraterrestrial neutrinos at the IceCube detector, Science. 342, 1242856 (2013). https://doi.org/10.1126/science.1242856

Ivezić Z, Kahn SM, Tyson JA, Abel B, Acosta E, et al., LSST: from science drivers to reference design and anticipated data products, Astrophys. J. 873, 111 (2019). https://doi.org/10.3847/1538-4357/ab042c

Jeon D, Hong IS, Kim HJ, Kim JW, Bodenstein RM, et al., Design of the RAON accelerator systems, J. Korean Phys. Soc. 65, 1010-1019 (2014). https://doi.org/10.3938/jkps.65.1010

Jeung M, Kim H, Cho K, Byeon O, The data processing of e-Science for high energy physics, J. Korean Phys. Soc. 55, 2067-2071 (2009). https://doi.org/10.3938/jkps.55.2067

Kim J, Cho K, A study on the optimization of finite volume effects of B K in lattice QCD by using the CUDA, J. Korean Phys. Soc. 67, 307-310 (2015). https://doi.org/10.3938/jkps.67.307

Kim K, Cho K, A study on the automatic reconstruction of the particles with machine learning at e+e- collider experiments, New Phys. Sae Mulli. 72, 564-572 (2022). https://doi.org/10.3938/NPSM.72.564

Kim K, Cho K, Numerical analysis of RAON heavy ion beam and target interaction using Geant4, J. Korean Phys. Soc. 87, 843-849 (2025a). https://doi.org/10.1007/s40042-025-01296-8

Kim K, Cho K, Nuclear theory in the KISTI-6 supercomputer era, J. Subatom. Part. Cosmol. 3, 100050 (2025b). https://doi.org/10.1016/j.jspc.2025.100050

Kortelainen M, Lesinski T, Moré J, Nazarewicz W, Sarich J et al., Nuclear energy density optimization, Phys. Rev. C. 82, 024313 (2010). https://doi.org/10.1103/PhysRevC.82.024313

Li BA, Chen LW, Ko CM, Recent progress and new challenges in isospin physics with heavy-ion reactions, Phys. Rep. 464, 113-281 (2008). https://doi.org/10.1016/j.physrep.2008.04.005

Navrátil P, Vary JP, Barrett BR, Properties of ^{12}C in the *Ab Initio* nuclear shell model, Phys. Rev. Lett. 84, 5728 (2000). https://doi.org/10.1103/PhysRevLett.84.5728

Park K, Cho K, A study of dark photon at the electron-positron collider experiments using KISTI-5 supercomputer, J. Astron. Space Sci. 38, 55-63 (2021a). https://doi.org/10.5140/JASS.2021.38.1.55

Park K, Cho K, Study of dark matter at e^+e^- collider using KISTI-5 supercomputer, Int. J. Contents 17, 67-73 (2021b). https://doi.org/10.5392/IJoC.2021.17.3.067

Park K, Kim K, Cho K, A study of double dark photons produced by lepton colliders using high performance computing, J. Astron. Space Sci. 39, 1-10 (2022). https://doi.org/10.5140/JASS.2022.39.1.1

Park K, Kim K, Sytov A, Cho K, Study of dark photons using future electron-positron colliders based on machine learning, J. Korean Phys. Soc. 84, 403-426 (2024). https://doi.org/10.1007/s40042-024-01005-x

Park SH, Suk H, Chung M, Accelerator science in Korea: current challenges and future opportunities, J. Korean Phys. Soc. 88, 617-624 (2026). https://doi.org/10.1007/s40042-025-01546-9

Qu X, Chen Y, Zhang S, Zhao P, Shin IJ, et al., Extending the nuclear chart by continuum: from oxygen to titanium, Sci. China Phys. Mech. Astron. 56, 2031-2036 (2013). https://doi.org/10.1007/s11433-013-5329-5

Radovic A, Williams M, Rousseau D, Kagan M, Bonacorsi D, et al., Machine learning at the energy and intensity frontiers of particle physics, Nature. 560, 41-48 (2018). https://doi.org/10.1038/s41586-018-0361-2

Raissi M, Perdikaris P, Karniadakis GE, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, J. Comput. Phys. 378, 686-707 (2019). https://doi.org/10.1016/j.jcp.2018.10.045

Saxena A, Srivastava PC, Suzuki T, *Ab initio* calculations of Gamow-Teller strengths in the *sd* shell, Phys. Rev. C. 97, 024310 (2018). https://doi.org/10.1103/PhysRevC.97.024310

Shin JW, Min KJ, Ham C, Park TS, Hong SW, Yield estimation of neutron-rich rare isotopes induced by 200 MeV/u 132Sn beams by using GEANT4, Nucl. Instrum. Methods Phys. Res. B Beam Interact. Mater. Atoms. 349, 221-229 (2015). https://doi.org/10.1016/j.nimb.2015.03.005

Sytov A, Bandiera L, Cho K, Cirrone GAP, Guatelli S, et al., Geant4 simulation model of electromagnetic processes in oriented crystals for accelerator physics, J. Korean Phys. Soc. 83, 132-139 (2023). https://doi.org/10.1007/s40042-023-00834-6

TOP500, TOP500 List - November 2025, 66th Edition (2025) [Internet], viewed 2026 Feb 20, available from: https://new.top500.org/lists/top500/2025/11/

White CJ, Stone JM, Gammie CF, An extension of the Athena++ code: framework for GRMHD based on advanced Riemann solvers and staggered-mesh constrained transport, Astrophys. J. Suppl. Ser. 225, 22 (2016). https://doi.org/10.3847/0067-0049/225/2/22

Yeo I, Cho K, Researches on dark matter using e^+e^- collider, J. Astron. Space Sci. 35, 67-74 (2018). https://doi.org/10.5140/JASS.2018.35.2.67

Yeo I, Cho K, Study on Geant4 simulation toolkit using a low-energy physics profiling system, J. Korean Phys. Soc. 74, 923-929 (2019). https://doi.org/10.3938/jkps.74.923

Yeo I, Cho K, Low-energy physics profiling of the Geant4 simulation tool kit on evolving computing architectures, J. Korean Phys. Soc. 76, 1047-1053 (2020). https://doi.org/10.3938/jkps.76.1047