


## CASE STUDY

High Performance Computing (HPC)  
Intel® Omni-Path Architecture  
Intel® Xeon® Scalable Processor



# KISTI Pushes the Boundaries of Science and Technology with NURION

## Nurion Supercomputer Highlights

- The 13<sup>th</sup> fastest supercomputer in the world as of the November 2018 TOP500 list<sup>2</sup>
- Equipped with both Intel® Xeon® Scalable processors and Intel® Xeon Phi™ processors and utilizing Intel® Omni-Path Architecture, it is the largest supercomputer in South Korea
- Designed to provide the resources to achieve scientific breakthroughs for a wide array of increasingly-complex, data-intensive challenges across modeling, simulation, analytics, and AI



## Executive Summary

No longer strictly focused on computationally intensive workloads, modern HPC centers need performant yet general-purpose systems that can address the many challenging and conflicting resource demands required to achieve scientific breakthroughs across a wide array of increasingly complex memory and data-intensive research projects. Further, world-class supercomputers such as the [Korea Institute of Science and Technology Information \(KISTI\) NURION system](#) are also flagship technology tools procured by an organization to provide for the future—be it in science or to meet the economic needs of a region.

According to Dr. Hee-yoon Choi (KISTI President), “KISTI will grow with the industry, academy and institute community as a central organization to support the dynamic science and technology data ecosystem which shares data and creates value, laying a foundation for Korea’s innovation growth”.<sup>1</sup> Equipped with Intel® Xeon® Scalable and Intel® Xeon Phi™ processors linked via an Intel® Omni-Path Architecture (Intel® OPA) communications fabric, the NURION 146-rack Cray\* CS500 cluster was procured to expand and increase the pace of innovative R&D. It is the largest supercomputer in South Korea and currently the 13<sup>th</sup> fastest supercomputer in the world.<sup>2</sup>

## Challenge

Scalability and the need to solve large-scale PDE problems which involve sparse matrix operations were key technology motivators in the KISTI procurement of a powerful new leadership class supercomputer. Very simply, researchers had outgrown and needed to move beyond the existing decade old TACHYON-II cluster.

Materials research is one of the application areas that KISTI has focused on as a leading HPC R&D institute, since it has the strong potential to lead advanced semiconductor device design that is important for national competitiveness of South Korea. In particular, KISTI has pursued the ability to simulate large-scale solid atomic structures with HPCs.

Dr. Soonwook Hwang (General Director and Principal Researcher, Division of National Supercomputing at KISTI) explains, “Electronic structure simulation of realistically sized solid structures is quite critical to help experimentalists who work on designs of new materials or advanced electronic devices. With large-scale simulations, we expect to cover design factors for nanoscale devices with large-scale simulations that can predict physical behaviors of solid structures having up to several million atoms.”

## Approach

Efficiently utilizing the large amount of many- and multi-core processors at scale as well as chip-level vector parallelism requires both detailed scientific and engineering knowledge. While KISTI has been firmly keeping the leadership of HPC R&D in South Korea during last decade with Tachyon-II cluster, the new NURION

introduced new levels of technology. Dr. Hwang explains, “Our Intel® Parallel Computing Center (Intel® PCC) project has served as a great opportunity for us to better understand and utilize the many- and multi-core Intel® processors. With the NURION system, now we are ready to broaden the leadership of HPC R&D in the Republic of Korea.”

## Results

The Intel PCC collaborative effort has paid off with quick returns as KISTI researchers have already achieved significant success even though NURION was just recently installed and is just starting to be made available to public users.


The Intel PCC project has focused on developing a software package for tight-binding simulations of large-scale electronic structures. Dr. Hoon Ryu (Intel PCC Lead and Principal Researcher, Center for Applied Scientific Computing at KISTI) notes, “The code is useful for advanced semiconducting devices, which is a key national business of South Korea”. KISTI was the first Intel PCC in the Asia-Pacific area starting in 2013.

Dr. Ryu continues, “This work basically needs to solve a Schrödinger equation that normally involves nanostructures consisting of tens of millions of atoms, which are numerically described with system matrices of a billion degrees of freedom. As a result, scalable processors are definitely needed with parallelization of core numerical operations including eigenvalue problems involving large-scale system matrices. With Intel Xeon Phi processors, we are able to drive a huge reduction of end-to-end simulation times for millions of atomic systems.”

### Use case: Scaling to 1000k+ atoms

Dr. Min Sun Yeom (Director and Principal Researcher, Center for Applied Scientific Computing at KISTI) says, “With tight-binding simulations of nanostructures having > 1,000,000 atoms on NURION system, we were able to explore the effect of size and structural engineering on band gap energies of physically realizable lead halide perovskite nanostructures within quite reasonable times. We also obtained the preliminary ideas for how to reduce the light-induced phase separation in halide mixtures, which would not be possible with DFT simulations that can normally handle solids consisting of hundreds of atoms.”

Metal halide perovskite is a promising material candidate for optoelectronic devices, and thus provides the motivation for



Connection of experiments and large-scale simulations (a) Experimental image of perovskite (CsPbBr<sub>3</sub>) quantum dots (Nano Letters 15, 3692-3696) (b) Dependency of band gap energies on quantum dot sizes. The KISTI numerical results connect nicely to experiment.

system empirical modelling of large-scale atomic structures. In short, it can provide nice guidelines for device designs such as how to map optical gaps and how to alleviate light-induced phase separation (a bottleneck in LED designs). The best part of empirical modelling is that it can provide direct connections to experiments.

Dr. Ryu points out that the use of Intel® Math Kernel Library (Intel® MKL) helped scale their calculations, “Intel MKL (scalapack packages such as lib\_mkl\_scalapack\_lp64 and libmkl\_blacs\_intelmpi\_lp64) helped a lot to improve the scalability of our Schrödinger solver. We used the LANCZOS algorithm, a well-known iterative method to tackle large-scale eigenvalue problem which has a numerical part that is hard to be MPI-parallelized by users and becomes a performance bottleneck as iterative processes continue. With the Intel MKL subroutines, we were able to reduce the corresponding computing load with improved scalability.”

### Use case: Many-core performance on sparse matrix operations


Leveraging previous work on the first generation Intel Xeon Phi coprocessors, Mr. Kyu Nam Cho (Former Research Associate, Korea University, now Principal Engineer in Samsung Research, Samsung Electronics) says, “The performance of sparse matrix-vector multiplication, which is the core numerical operation needed to solve large-scale electronic structures, was not bad even when we worked with Intel first generation many-core processors (Intel Xeon Phi Co-processors) compared to Intel® Xeon processors V3. The performance on the NURION Intel Xeon Phi nodes is much better, particularly when combined with MCDRAM.” Cho notes that, “Another critical strength of Intel Xeon Phi processor-based systems is their ease of use, particularly if we consider the amount of work that must be performed to port the existing code to run on PCI-E add-in devices.”

The KISTI Intel PCC found that the speedup due to the performance of the Intel Xeon Phi processor’s high bandwidth memory (HBM) meant that a single node could take a larger workload. Dr. Ryu points out that “inter-node scalability is quite nice”. Scalability tests demonstrate a speedup when increasing the number of computing nodes. The KISTI Intel PCC observed a 1.5-3x speedup<sup>3</sup> when they made use of the high bandwidth memory (HBM) packaged with the many-core Intel Xeon Phi processor 7250 nodes. More recently, they successfully ran a 0.4 billion atomic structure in NURION system and checked the strong scalability up to 2,500 computing nodes (170,000 computing cores).

Dr. Ryu points out that “Intel technology matches with the purpose of KISTI HPC”. According to a statistical workload analysis performed at KISTI, approximately 50% of their workloads involve sparse matrix operations. This means the NURION supercomputer should perform well in meeting the needs of KISTI researchers across a wide range of research areas.

## Performance realized

The importance of large-scale simulations for advanced material research to South Korea cannot be underestimated as evidenced by the money spent to procure a world class supercomputer.<sup>5</sup> For this reason, the KISTI Intel PCC critically evaluated the various hardware solutions upon which the NURION procurement could be based—including GPU



Strong scalability of end-to-end simulations (a) Small-scale BMT target was to calculate 5 lowest conduction band states in 27x33x33 nm<sup>3</sup> (~1.5million atoms) Si:P quantum dot.<sup>4</sup> The scalability is tested up to 3 computing nodes (204 cores). (b) Extremely large-scale BMT target was to calculate 3 lowest conduction subbands in 2715x54x54 nm<sup>3</sup> Si:P nanowires (0.4billion atoms). The scalability here is tested up to 2,560 computing nodes (170,000 cores) in NURION system.

accelerated systems. Their results have been published in the literature for Intel processors<sup>6, 7, 8</sup> and GPUs.<sup>9</sup> They present solid technical evidences to show why the choice for NURION was an Intel based system that delivers 25.7 PFlop/s (Rpeak), 13.9 PFlop/s (Rmax),<sup>3</sup> ranking it at #13 on the November 2018 TOP500 list.<sup>2</sup> Dr. Ryu is developing a white paper to tell the full CPU vs. GPU story in an article to be published later this year.<sup>10</sup>

However the story does not stop with the NURION system as the KISTI Intel PCC is evaluating the use of FPGAs for large-scale electronic structure calculations. In particular, the Intel Scalable processor family provides a pathway towards future FPGA acceleration.<sup>11</sup> As with the GPU and Intel processor evaluations, the KISTI Intel PCC has been publishing their work on FPGAs as well.<sup>12</sup>



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of November 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.


Intel, the Intel logo, Xeon and Phi are trademarks of Intel Corporation in the U.S. and/or other countries. \*Other names and brands may be claimed as the property of others.

© Intel Corporation

0119/RJM/J/PDF



338512-001US



KISTI people who enabled scalable simulations of extremely large electronic structures in NURION system: (From left) Dr. Hoon Ryu, Dr. Ji-Hoon Kang (Principal Researcher, Center for Applied Scientific Computing), Mr. Taeyoung Hong (NURION Operation Team Lead and Senior Researcher, Supercomputing Service Center)

#### Endnotes

- <https://kisti.re.kr/about/pageView/245> Configuration: Intel Xeon Phi 7250 nodes; 68 cores/node using 2 MPI processes + 32 threads per node; Quad / Flat memory mode; 100G network connectivity. 2500 Intel Xeon Phi nodes, a total of 68x2500 computing cores were used for the benchmark test of KISTI's in-house code. BIOS: S72C610.86B.01.03.0018.C0001.012420182107; Memory: 96GB DDR4-2400 memory + 16GB 7.2GT/s MCDRAM; Networking and Storage: Intel Omni-Path Architecture, 100Gb network connectivity; OS and Kernel details: CentOS Linux Release 7.3, Linux kernel 3.10.0-514.26.2.el7.x86-64; Application software: Quantum simulation tool for Advanced Nanoscale Devices; Tested by KISTI in November, 2018
- Currently according to the November [2018 TOP500 list](https://www.top500.org/lists/2018/11/)
- Test performed by KISTI in November 2018. Rmax is maximal LINPACK performance achieved; Rpeak is theoretical peak performance per TOP500.org. Configuration: Intel Xeon Phi 7250 nodes; Up to 272 (68x4) cores/node using 4 MPI processes + 68 threads per node; Quad/Flat memory mode; 10 G network connectivity
- Si:P alloy structures have been popularly studied to build Si-based qubit systems. See Nature Nanotechnology 9, 430-435, and Nano Letters 15, 1, 450-456
- <https://www.datacenterdynamics.com/news/cray-to-build-south-koreas-largest-supercomputer-for-kisti/>
- [http://itoc.sjtu.edu.cn/wp-content/uploads/2018/05/201805\\_IPC-AsiaSummit\\_HoonRyu.pdf](http://itoc.sjtu.edu.cn/wp-content/uploads/2018/05/201805_IPC-AsiaSummit_HoonRyu.pdf)
- [https://www.researchgate.net/publication/307613037\\_Time-efficient\\_simulations\\_of\\_tight-binding\\_electronic\\_structures\\_with\\_Intel\\_Xeon\\_PhiTM\\_many-core\\_processors](https://www.researchgate.net/publication/307613037_Time-efficient_simulations_of_tight-binding_electronic_structures_with_Intel_Xeon_PhiTM_many-core_processors)
- Ji-Hoon Kang, Oh-Kyoung Kwon, Jinwoo Jeong, Kyunghun Lim, Hoon Ryu: Performance Evaluation of Scientific Applications on Intel Xeon Phi Knights Landing Clusters. HPCS 2018: 338-341
- GPU results were published in "Fast, energy-efficient electronic structure simulations for multi-million atomic systems with GPU devices" by Hoon Ryu and Oh-Kyoung Kwon in Journal of Computational Electronics (2018) 17:698-706, <https://doi.org/10.1007/s10825-018-1138-4>.
- Please check Dr. Ryu's publications list to see the article when it appears: [https://www.researchgate.net/profile/Hoon\\_Ryu3](https://www.researchgate.net/profile/Hoon_Ryu3)
- <https://www.intel.com/content/www/us/en/programmable/solutions/acceleration-hub/platforms.html>
- <https://ieeexplore.ieee.org/document/8048988>