Woong Shin, Ph.D.

Research Scientist | High-Performance Computing | AI/ML/Operational Data Analytics & Energy Efficiency

me@woongshin.com | https://woongshin.com

Research Scientist with over 15+ years of combined industry and research experience, spanning mission-critical enterprise software engineering and 8+ years of HPC operational data analytics, data-driven modeling, and sustainable supercomputing. A core contributor to the Oak Ridge Leadership Computing Facility (OLCF)—the U.S. HPC center behind Frontier, the first exascale supercomputer—and its predecessor, Summit. Recipient of multiple best paper awards in HPC sustainability and advanced storage research, as well as a UT-Battelle research accomplishment, now evolving these methods with machine learning and AI to further optimize supercomputing performance, power, and reliability at scale.

Education

Ph.D. in Electrical Engineering & Computer Science (MS & Ph.D. Integrated Course) -
Seoul National University, South Korea (2017) - Dissertation: "OS I/O Stack Optimizations for Flash SSDs"
B.S. in Computer Science
Korea University, South Korea (2003)

Professional Experience

Oak Ridge National Laboratory (ORNL) – Oak Ridge, TN, USA
Research Scientist, Analytics & AI Methods at Scale (AAIMS) Group (2020~Present)

- Designed LLM-based systems to enable predictive analytics and streamline operational data queries
- Spearheaded energy efficiency initiatives for exascale systems, improving power usage insights and enabling sustainable HPC operations.
- Advanced research in data-driven and simulation based modeling techniques (machine learning, digital twins) for HPC operations to increase operational efficiency of ORNL's supercomputers.

Technology Integration Group (2017–2020)

- Developed scalable realtime data pipelines and analytics frameworks for energy and thermal monitoring on Summit and Frontier supercomputers, significantly enhancing operational decision-making.

Seoul National University – South Korea
Research Assistant, Distributed Computing Laboratory (2010 ~ 2017)

- Pioneered cross-layer SSD optimization techniques for HPC storage, resulting in reduced latency and improved throughput.

TmaxSoft | Samsung Networks – South Korea
Software Engineer (2003–2010)

- Delivered middleware monitoring innovations systems solutions for mission-critical enterprise software.

Select Projects

- **AI for HPC Operations (2024–Present)**: Engineered generative AI models enabling natural language querying and automation of machine learning workflows for HPC operators.
- **HPC Application Energy Efficiency (2022–Present)**: Developed advanced tools, models, and analytics dashboards to enhance energy efficiency and sustainability, in collaboration with key vendors.
- **Frontier Digital Twin (2022–Present)**: Co-designed "ExaDigiT," digital twin framework and integrated real-time data with advanced modeling to drive operational excellence and energy conservation.
- **Power & Energy Monitoring and Analytics for Frontier Supercomputer (2021 ~ 2024):** Architected end-to-end data analytics systems for monitoring and analyzing power, energy, and thermal data across OLCF's HPC infrastructure.

Woong Shin, Ph.D. me@woongshin.com | https://woongshin.com

Honors & Awards

- **Best Paper Award**, SC'21 and IEEE BigComp'15
- **Research Accomplishment Award**, UT-Battelle (2022)
- **Technology Innovation Award**, TmaxSoft (2010)

Publications

Selected works:

- Shin et al., **"Towards Sustainable Post-Exascale Leadership Computing"**, in SC24 Sustainable Supercomputing Workshop, November 2024
- Shin et al., **"Navigating Exascale Operational Data Analytics: From Inundation to Insight"**, in SC24 Sustainable Supercomputing Workshop, November 2024
- Brewer et al., **"A Digital Twin Framework for Liquid-cooled Supercomputers as Demonstrated at Exascale"**, SC24, November 2024
- Shin et al., **"Revealing Power, Energy and Thermal Dynamics of a 200PF Pre-Exascale Supercomputer"**, SC21, November 2021, ***Acceptance rate 23.5%*** **(Best paper award)**
- Ott et al., **"Global Experiences with HPC Operational Data Measurement, Collection and Analysis"**, Energy Efficient HPC State of Practice Workshop at CLUSTER20, September 2020
- Shin et al., **"Providing QoS through Host Controlled Flash SSD Garbage Collection and Multiple SSDs"**, BigComp15, 2015, **(Best paper award)**

Skills

- **Innovation & Cross-functional leadership**: **Strategic planning & execution**, **public speaking** (keynotes, panels), **technical writing** (publications, white papers), and driving innovation through **interdisciplinary and global collaboration**
- **Programming & Development:** C-based **systems programming** (Linux, AIX, Solaris, HP-UX), Python for **data pipelines, APIs, automation**; test-driven development (**Pytest, Robot Framework**).
- **HPC & Distributed Systems:** HPC job scheduling (**Slurm, LSF, Torque Moab**), scalable data engineering (**Kafka, Apache Spark, Dask, Parquet**), and **workflow orchestration (Apache Airflow)**.
- **Cloud & Infrastructure: Kubernetes** application deployment (Kustomize, Helm), **GitOps CI/CD** (GitLab, ArgoCD), **Ansible** for automation, high-availability services (**Redis, PostgreSQL, RabbitMQ, Kafka, MinIO, Apache Druid, Apache Spark Cluster**).
- **Data Center & HPC Monitoring: Prometheus + Grafana**, industrial telemetry (**BACnet/IP, Modbus/TCP**), **LD_PRELOAD-based system instrumentation** (Oracle OCI, Tibco, OS file I/O).
- **AI & Machine Learning:** Applied **ML in HPC** for **automation and operational analytics**. Hands-on with generative AI (**LangChain, LiteLLM / Agents, RAG**); collaborated with data scientists to integrate AI into HPC workflows
- **Storage & Embedded Systems:** Linux **kernel & device drivers**, NVMe & SSD optimization, **Xilinx FPGA (Zynq-7000, Virtex-7)**, high-performance storage tuning.

Professional & Public Activities

- **Vice Chair**, Technical Program Proceedings, SC25
- **Member**, Energy-Efficient HPC Working Group (2019 ~ Present)
- **Speaker,** SC24, SC22, and numerous HPC venues

Interests

- Keyboard, Audio engineering, Mountain biking, Ski, Inline skating, Camping, Productivity, Deep Life