第 28 回先進スーパーコンピューティング環境（ASE）研究会実施報告

近 藤 正 章

東京大学情報基盤センター

　東京大学情報基盤センターASE 研究会（Advanced Supercomputing Environment）は内外からの講演者をお招きして不定期に開催している。2017 年 5 月 22 日（月）に実施された第 28 回 ASE 研究会[1]では，Global Extensible Open Power Manager（GEOPM）の開発者である Jonathan Eastep 博士（Intel）をお招きして、スーパーコンピュータシステムの電力管理を行うランタイムフレームワークについてご講演頂いた。また、本センター教員が実施しているプロジェクトで開発中の電力管理手法や本部門の紹介を実施した。プログラムは表 1 の通りである。学内外から 9 名の出席があり、活発な議論が行われた。

写真: 講演する Jonathan Eastep 博士

表 1　プログラム

時間帯	講演者	題　目
17:00 - 17:05	Masaaki Kondo (The University of Tokyo)	Welcome & Opening
17:05 - 18:00	Jonathan Eastep (Intel, USA)	Global Extensible Open Power Manager: A Vehicle for HPC Community Collaboration on Co-Designed Energy Management Solutions
The power scaling challenge associated with Exascale systems is a well-known issue. In this talk, we provide an overview of the Global Extensible Open Power Manager (GEOPM). GEOPM is an open source power management runtime framework which is being contributed to the HPC community to foster collaboration on new power management runtime techniques to address Exascale power challenges or enhance performance and power efficiency on today's systems as well. Through GEOPM's plug-in extensible architecture, it enables rapid prototyping of new runtime algorithms. This talk will cover GEOPM's architecture and usages then discuss opportunities for collaboration. For additional information, please visit: https://geopm.github.io/geopm/		
18:00 - 18:15	Thang Cao (The University of Tokyo)	Cooling-Aware Job Scheduling and Node Allocation for Overprovisioned HPC Systems
Limited power budget is becoming one of the most crucial challenges in developing supercomputer systems. Hardware overprovisioning which installs a larger number of nodes beyond the limitations of the power constraint is an attractive way to design future supercomputers. In air-cooled HPC centers, it is said that about half of the total power is consumed by cooling facilities. Reducing cooling power and effectively utilizing power resource for computing nodes are important challenges. In this talk, we present a cooling and node location-aware job scheduling strategy which tries to optimize job-to-node mapping while improving the total system throughput under the constraint of total system (compute nodes and cooling facilities) power consumption.		
18:15 - 18:25	Kengo Nakajima (The University of Tokyo)	Overview of Supercomputing Research Division, Information Technology Center, the University of Tokyo
18:25 - 18:30	Masaaki Kondo (The University of Tokyo)	Closing

---

[1] http://www.cc.u-tokyo.ac.jp/event/ase/28.html