第 27 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

伊田　明弘

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招いて不定期に開催している。2017 年 4 月 7 日（金）に実施された第 27 回 ASE 研究会[1]では，Faisal Shahzad 博士（Erlangen-Nuremberg 大学）に，SPPEXA-ESSEX プロジェクトで開発が進められている CRAFT(自動フォルトトレランスライブラリ)ついて，ご講演いただくとともに，関連分野についてセンター教員から 2 件の発表を実施した。学内外から合計 12 名の出席者があり，活発な議論が行われた。

写真：講演する Shahzad 博士

表 1　プログラム

時間帯	講演者	題　目
16:00-16:05	Akihiro Ida (ITC, The University of Tokyo)	Welcome & Opening
16:05-16:55	Edmond Chow (University of Erlangen-Nuremberg, Germany)	CRAFT: A library for application-level Checkpoint/Restart & Automatic Fault Tolerance
In order to efficiently use future generations of supercomputers, fault tolerance and power usage are two of the prime challenges anticipated by the High Performance Computing (HPC) community. A significant share of faults in HPC systems constitute of hard failures, which in many cases lead to process(es) and eventually job failure. In this talk, we will present our fault tolerance approach developed in the scope of SPPEXA-ESSEX project. We have developed a Checkpoint/Restart and Automatic Fault Tolerance (CRAFT) library that serves two purposes. First, it provides a framework that significantly reduces the effort needed for the implementation of application-level checkpoint/restart methods in a program. The user can extend the library to add more user-specific data-types, making them ‘checkpointable’ for future use. Secondly, it provides an easier interface for dynamic process recovery, thus enabling applications to recover automatically after process failures. For this purpose, we have used User-Level Failure Mitigation (ULFM), which is a prototype implementation of fault tolerant MPI. We have significantly reduced the complexity of failure detection and application recovery mechanism. Both of these functionalities of CRAFT can either be used separate as well as combined. CRAFT-library features, optimizations and limitations will be discussed in detail.		
16:55-17:25	Toshihiro Hanawa (ITC, The University of Tokyo)	Tightly Coupled Accelerators Architecture: A Very Low Latency　Communication System on GPU Cluster
17:25-17:55	Eishi Arima (ITC, The University of Tokyo)	Efficient Big Data Processing through Page Table Walk Aware Cache Management
17:55-18:00	Akihiro Ida	Closing

[1] http://www.cc.u-tokyo.ac.jp/event/ase/27.html