Session on

Data Storage Research

Attended by 22+7+1 participants

Devesh Tiwari, Yong Chen, Carlos Maltzahn, Sudharshan Kannan, Erez Zadok, Anduo Wang, Jason Liu, Geoff Kuenning, Weijun Xiao, Mai Zheng, Satish Puri, Shrin Bidokhti, David Du, Ethan Miller, Shannon Beck, Raju Rangaswami, Krishna Kant, **Ali R. Butt**

Research Thrusts

1. **Tools and techniques to manage variety/heterogeneity of storage**
   a. Design methods to deal with deep hierarchy in storage systems
   b. Employ new technologies, e.g., **NVRAM**, technologies on the horizon, e.g., **DNA storage**, and revived old technology, e.g., **tapes**
   c. Explore hardware-software solutions, e.g., FPGAs for storage stack
   d. Bring data and computation closer, e.g., in-situ processing
2. **Cross-stack optimizations for storage**
   a. Consider all aspects of the application-to-storage device hierarchy
   b. Exploit specialized applications/uses, e.g., Edge computing, ML, DL
3. **Long-term and archival storage**
   a. Design new storage systems for massive archival data
   b. Address issues of provenance, metadata, indexing, and searching

Research Thrusts

4. **Support representative storage testbeds & benchmarks**
   a. Design/support public testbeds for storage systems research
   b. Make traces and benchmarks available
   c. Support reproducibility
5. **Evolve storage APIs and abstractions**
   a. Rethink POSIX
   b. Explore new models such as key-value stores
6. **Suggestions to NSF**
   a. Support specialized programs to support **Storage+X** research a la *FMitF*
   b. Develop transition to practice track
   c. Grow “competitions”/ events to raise awareness of storage
   d. Leverage NSF Workshop on Data Storage Research Vision 2025