Threading Opportunities in High-Performance Flash-Memory Storage

Craig Ulmer — Sandia National Laboratories, California
Maya Gokhale — Lawrence Livermore National Laboratory

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000

Sandia National Laboratories

Revolutionary Storage Technologies

- Storage-Intensive Supercomputing (SISC) at LLNL
  - System architectures for applications with ***massive datasets***
  - New technologies: processing elements, networks, and storage

- NAND-Flash storage in high-performance computing
  - Flash chips have great potential: 100x access times, 10x bandwidth
  - However, few commercial products have delivered performance

- Exception: Fusion-io's ioDrive
  - PCIe x4 card with 80-320 GB of flash
  - Theoretical read speed of 700 MB/s
  - Hardware allows many IOPs to be in-flight concurrently

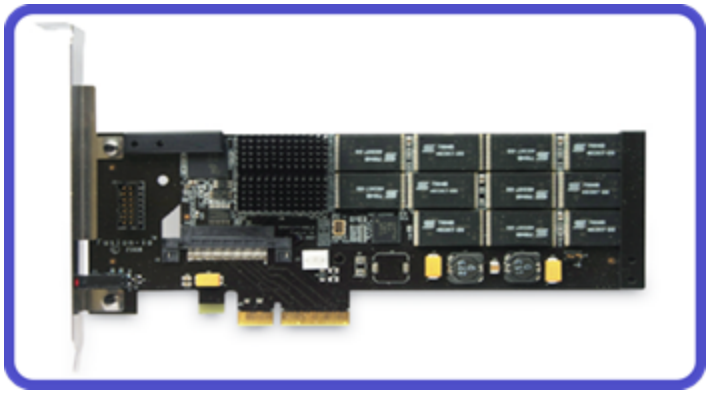

Sandia National Laboratories

Threaded I/O Microbenchmarks

- Observation: Increasing number of IOPs improves performance
  - Opposite of what we expect from hard drives
  - Due to flash memory packaging: chip is actually a die stack

- Implemented a set of I/O microbenchmarks to investigate
  - Threaded with mixed I/O characteristics (mostly read-only)
  - Currently: Block transfer, kNN, external sort, binary search

- Example: k-Nearest Neighbors (kNN)
  - Stream through all training vectors and find k vectors that are most similar to each input vector
  - Each thread works on portion of training data

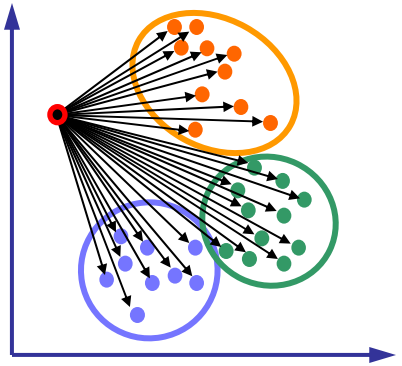

Sandia National Laboratories

kNN Results

- *Single* ioDrive vs. *Three* SATA hard drives in RAID0
  - ioDrive provided a 3x improvement to end application
  - Small number of threads can have large impact with flash memory

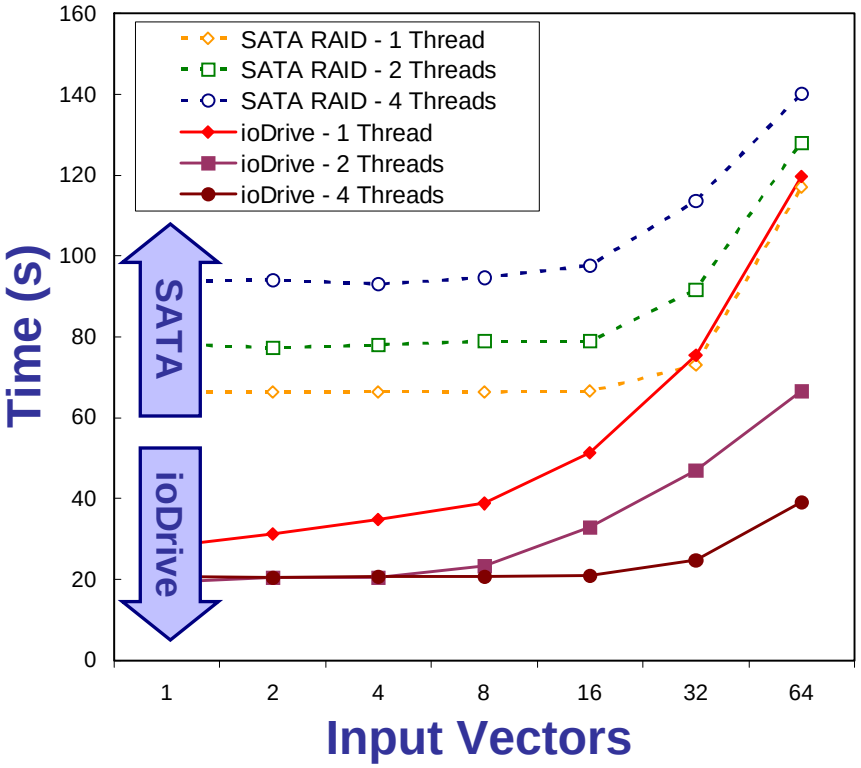

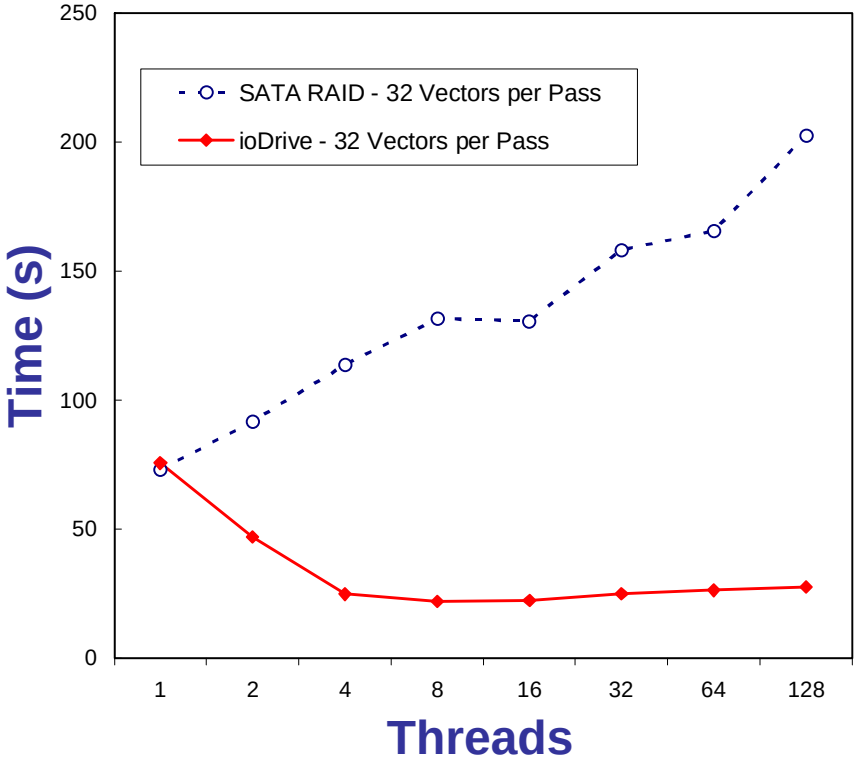

Sandia National Laboratories