Skysong Innovations
Advancing the Arizona State University Knowledge Enterprise

1475 N. Scottsdale Road, Suite 200
Scottsdale, AZ 85287-3538
Phone: 480 884 1996 Fax: 480 884 1984

Case ID:M22-101P^
Published: 12/27/2022

Inventors

Deliang Fan
Fan Zhang
Shaahin Angizi

Contact

Physical Sciences Team

Fast and Efficient Max/Min Searching in DRAM

-In the era of big data, min/max searching from bulk data arrays is one of the most important and widely used fundamental operations in data-intensive applications such as sorting, ranking, bioinformatics, data mining, graph processing, and route planning. Online news and social media require real-time ranking using fast min/max searching from massive data stores to evaluate trending information to display on their sites.

The process of min/max searching is a time-consuming computation in many large-scale graph processing algorithms. However, implementing fast and efficient min/max searching for big data faces significant challenges in conventional computer systems with respect to memory architecture and computing algorithms. Within memory architectures, the well-known 'memory-wall' challenge causes significant issues, like long off-chip memory access latency, data congestion due to limited memory bandwidth and two orders higher energy consumption in data movement than data processing. Within computing algorithms, min/max searching is in general a comparison-based algorithm, where the CPU needs to compare every element serially of colossal amounts of raw data. Such computing properties causes min/max searching to demand ultra-high computing resources and power. There is a need for efficient min/max searching algorithms and supporting hardware for bulk data storage that greatly minimizes the time and cost associated with the computing.

Researchers at Arizona State University (ASU) have developed a min/max-in-memory algorithm to find the minimum/maximum of an array stored in dynamic random-access memory (DRAM). Additionally, the ASU researchers have developed a hardware that hosts and computes the data via in-DRAM computing. The algorithm and hardware support parallel in-memory searching for minimum and maximum values of bulk data stored in DRAM as unsigned and signed integers, fixed-point and floating-point numbers.

Related publication: Max-PIM: Fast and Efficient Max/Min Searching in DRAM

Potential Applications:

- Accelerate min/max searching within big data widely used by:
  - cloud computing
  - social media
  - online news
  - data center providers
  - bioinformatics

Benefits and Advantages:

- Optimized with a one-cycle fast XNOR logic in-DRAM operation and in-

memory data transpose

- Example experiments utilizing the algorithm and hardware in big data sorting and graph processing applications produced speeds up to ~50x and ~1000x faster than GPU and CPU while only consuming 10% and 1% of the energy, respectively
- Compared to in-DRAM computing platforms, i.e., Ambit and DRISA, ASU's algorithm and hardware increases speed of computations by ~3-10x