

Liu Liu  
Discussion Leader Cover Letter  
Research Methods Spring 2024

**Title:** [Sieve: Scalable In-situ DRAM-based Accelerator Designs for Massively Parallel k-mer Matching](#)

**General Topic:** Leveraging In-DRAM computing to accelerate genome analysis.

**Specific Behavior:** This paper introduces three configurations of DRAM-based in-situ k-mer matching accelerator designs for different optimization targets. A lightweight matching circuitry for fast pattern matching and an early termination scheme is also proposed to achieve better latency and energy efficiency.

**Research Questions:**

- Can the memory-centric solution be employed to accelerate the k-mer matching process?
- Is it possible to leverage the large memory capacity of DRAM for better scalability?

**Challenges:** The software approaches using conventional hardware (CPU/GPU) are not efficient for k-mer matching tasks due to irregular access patterns and poor cache behavior. Additionally, the novel hardware accelerator needs to be combined with the unique data mapping scheme to fully utilize its advantages.

**Paradigm:** The authors proposed novel hardware designs and conducted extensive simulations on real datasets to validate the functionality and evaluate the performance of their designs.

**Importance:**

- *k-mer* matching is critical for many genome analysis pipelines, which can occupy 65%-80% of the total runtime for various important bioinformatics applications.
- The genome data is generated at an incredible speed, and there are vast needs for precision medicine, disease surveillance, population genetics, etc.

**Claims:** The proposed in-situ DRAM-based accelerator can achieve an average of 326x/32x speedup and 74x/48x energy savings compared to the CPU/GPU baseline on several datasets.

**State of Knowledge:** The authors systematically analyze different bioinformatics applications/workloads, highlight the importance of k-mer matching operations, and summarize the state-of-the-art hardware implementations for k-mer matching.

**Story Structure:** In this paper, the authors first highlight the k-mer matching as a bottleneck in many genomics pipelines. They proposed Sieve, an in-memory accelerator based on DRAM. Furthermore, they proposed different types of implementations to optimize for area, latency, and energy. Finally, they conduct extensive experiments on real-world bioinformatics workloads to evaluate the performance and functionality of their design.