OnionChopper

A Modular Arithmetic Hardware Accelerator for Private Information Retrieval

Georgia Shay, Massachusetts Institute of Technology

Goal

Setup: A client wishes to retrieve a record from a server's database without revealing to the server which record was accessed.

Applications: Private voice calling, media streaming, advertising impressions, online presence indicators, etc

Problem: Slow runtime due to privacy preservation and cryptographic operations

Solution: OnionChopper: a small, fast, and energy efficient hardware accelerator

Private Information Retrieval

A typical Private Information Retrieval (PIR) scheme has the client send a one-hot vector encryption of the location of the desired record in the database.

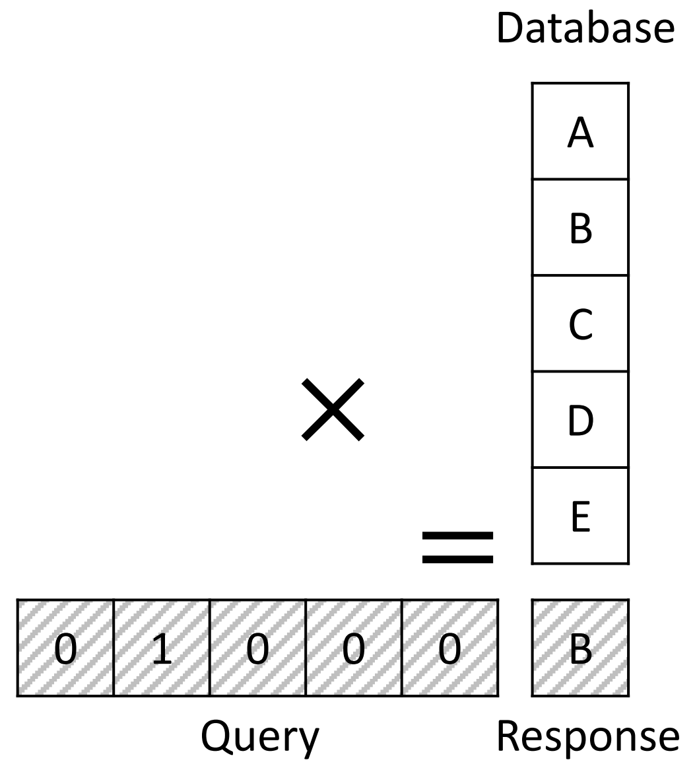

Basic PIR Scheme

After multiplication with the database, the response is an encryption of the desired record.

The database can be reconceptualized as multi-dimensional, and multiple query vectors sent to indicate the row and column.

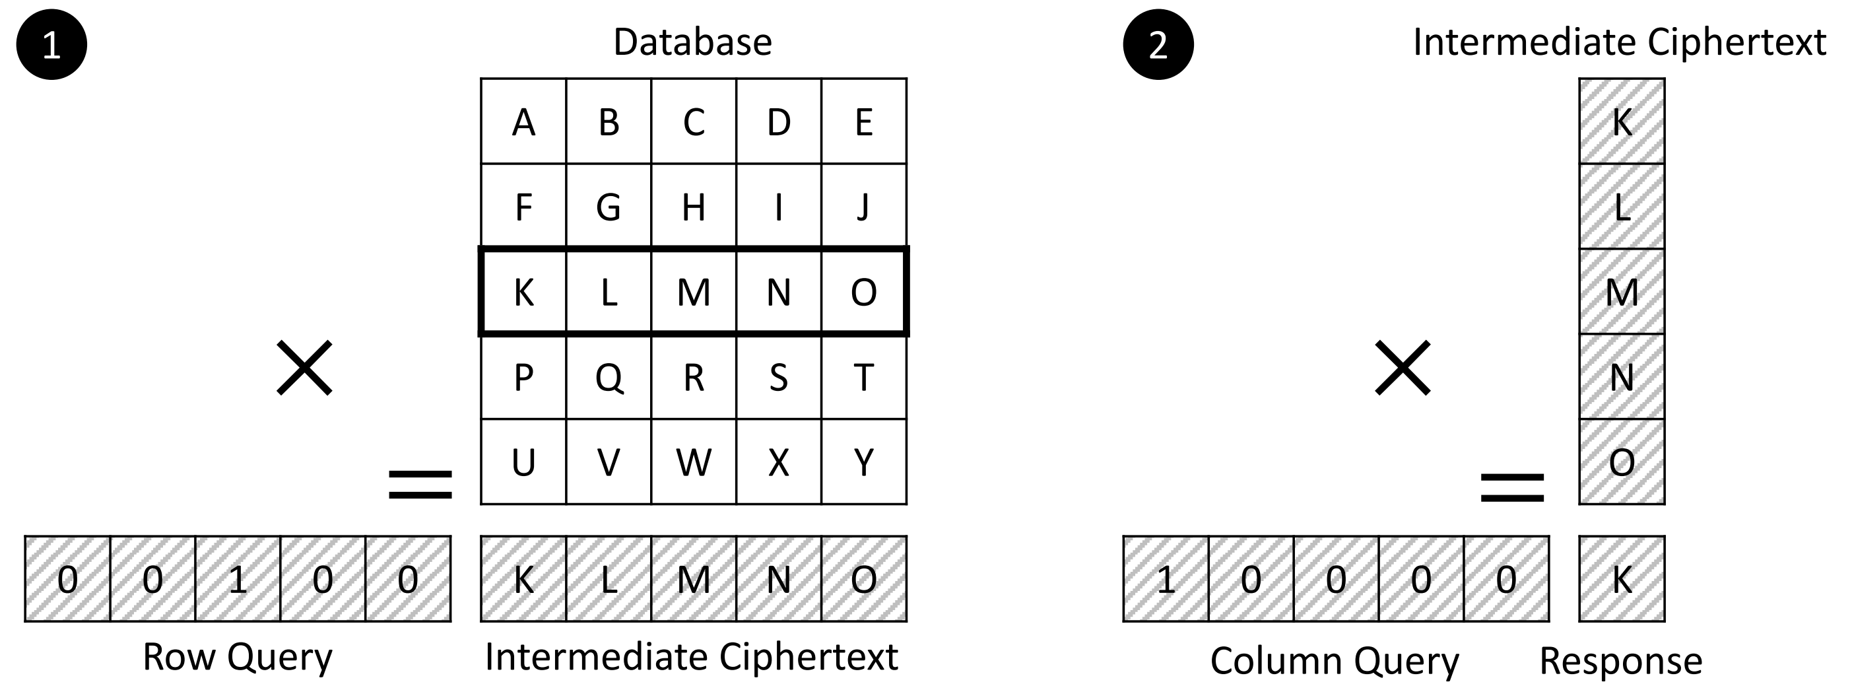

2D PIR Scheme

OnionPIR

OnionPIR uses two different types of ciphertexts in its multiplications to keep response size small. OnionPIR's ciphertexts can be represented as matrices of collections of polynomials.

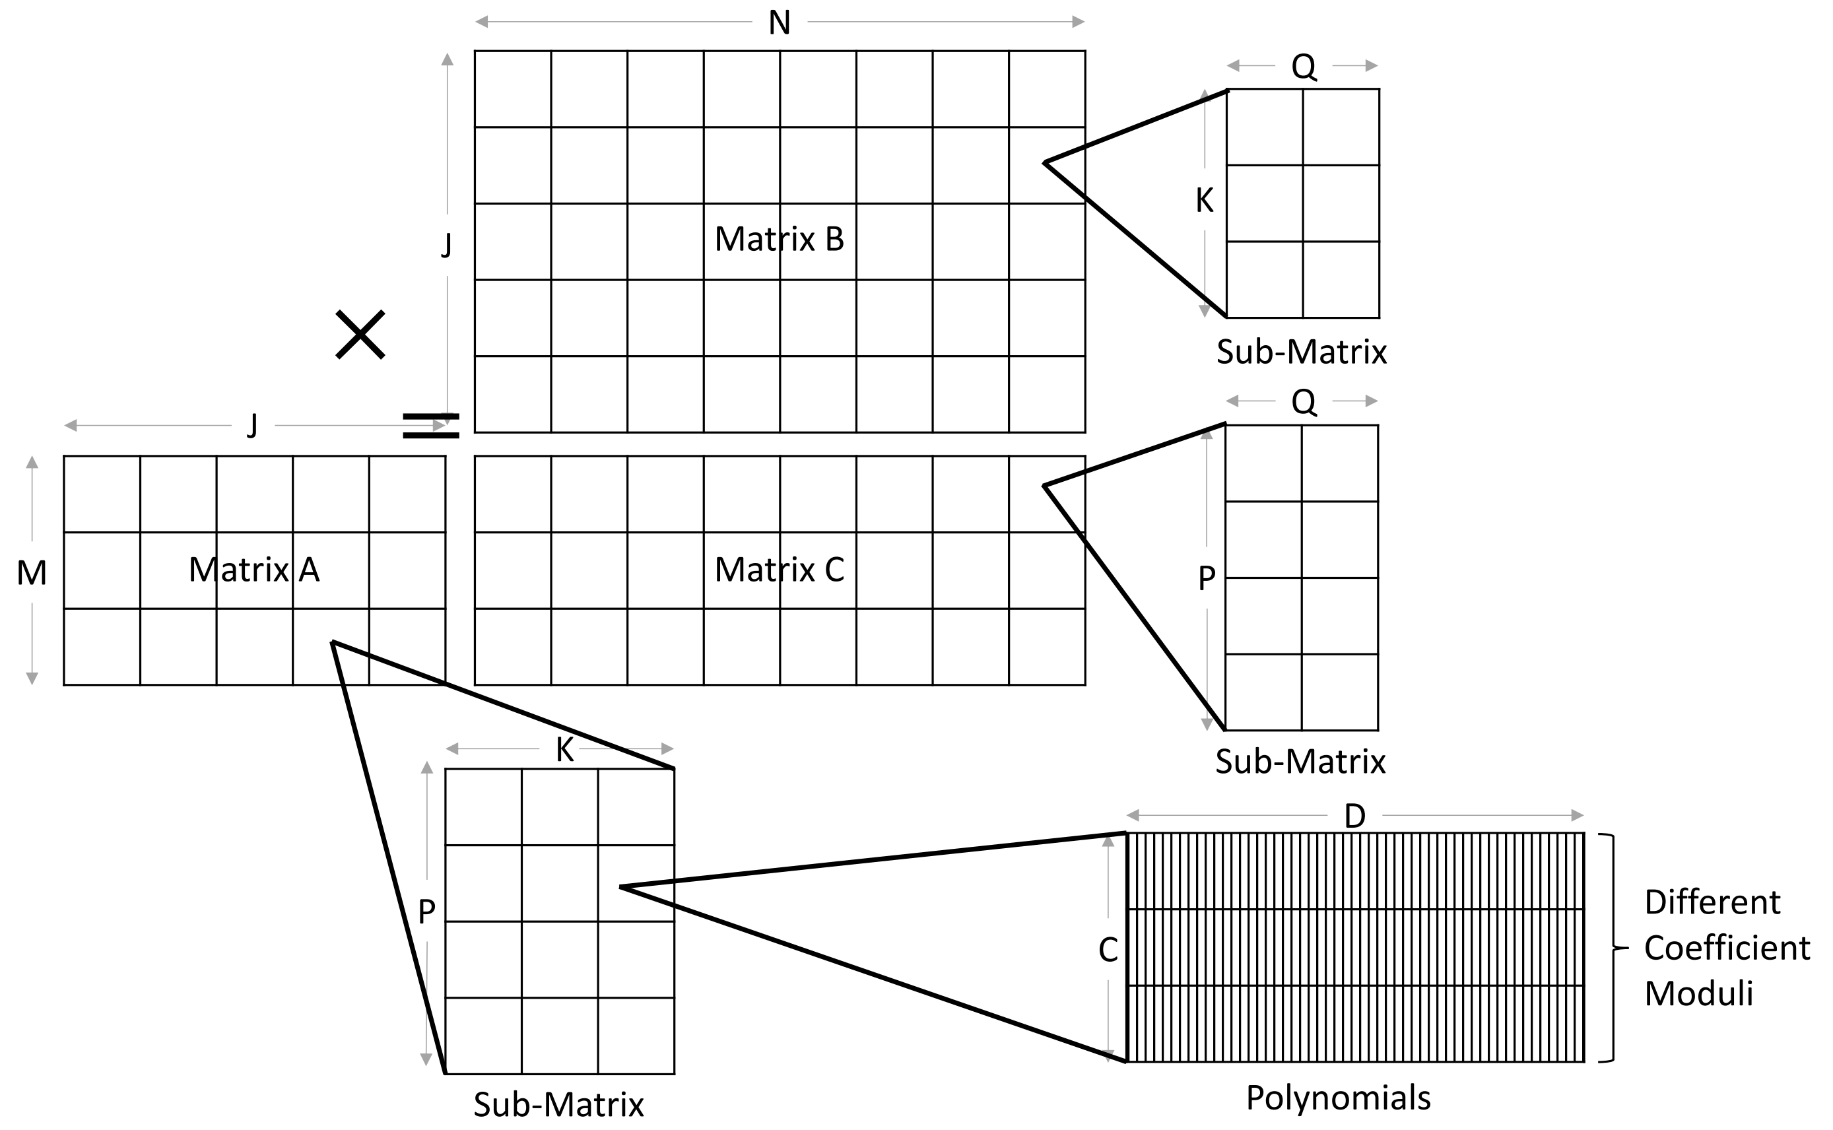

PIR Bottleneck Operation

The same bottleneck operation – a matrix multiplication of submatrices of polynomial collections – is applicable to other PIR schemes.

Hardware Design

Our proposed hardware design is a near-storage accelerator operating at 2GHz located in an SSD unit that contains a DRAM.

Beyond the DRAM will be a storage hierarchy of SRAMs and register files that aims to take advantage of data reuse.

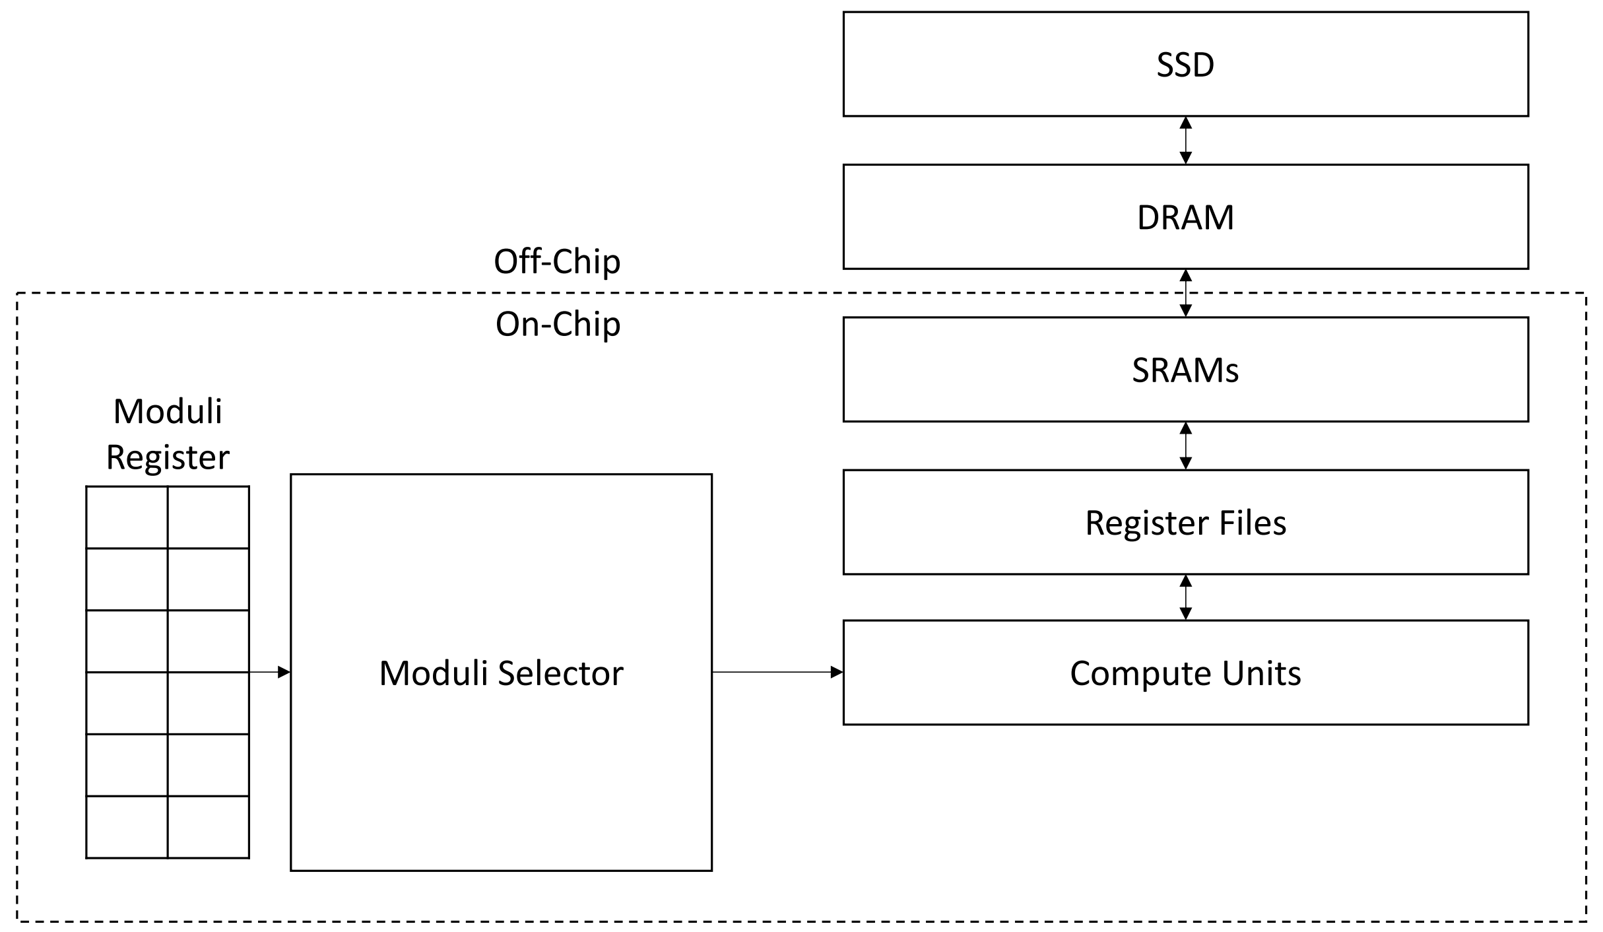

OnionChopper Hardware Design

Optimization

The Timeloop tool can evaluate the best loop ordering and tiling for a tensor calculation on a given architecture, and determine how much area, energy, and time it takes. We defined a range of reasonable values for the storage level sizes and fanouts, and ran the tool on these possible architectures.

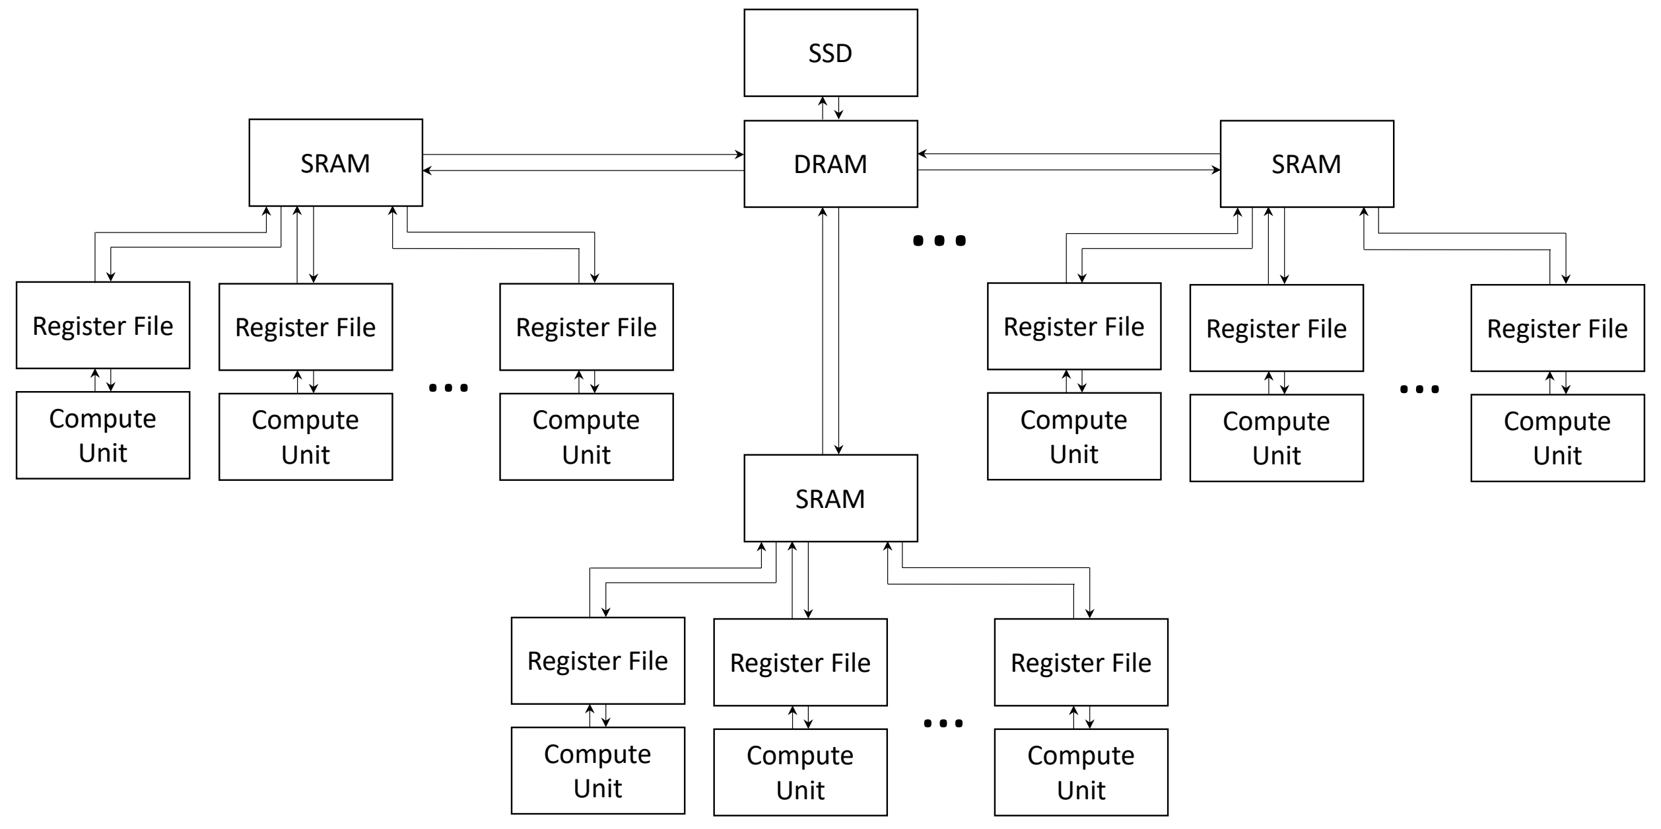

Storage Hierarchy

The fastest architecture with the best area-energy product had two 4KB SRAMs and two 32-register register files. It takes up 6.87mm^2 in area, largely from the compute units, which are heavily pipelined.

Results

On a 64GB database, OnionPIR's main operation takes 6.4 seconds on OnionChopper, compared to 29 minutes on a general-purpose processor.

Database Size vs Speed

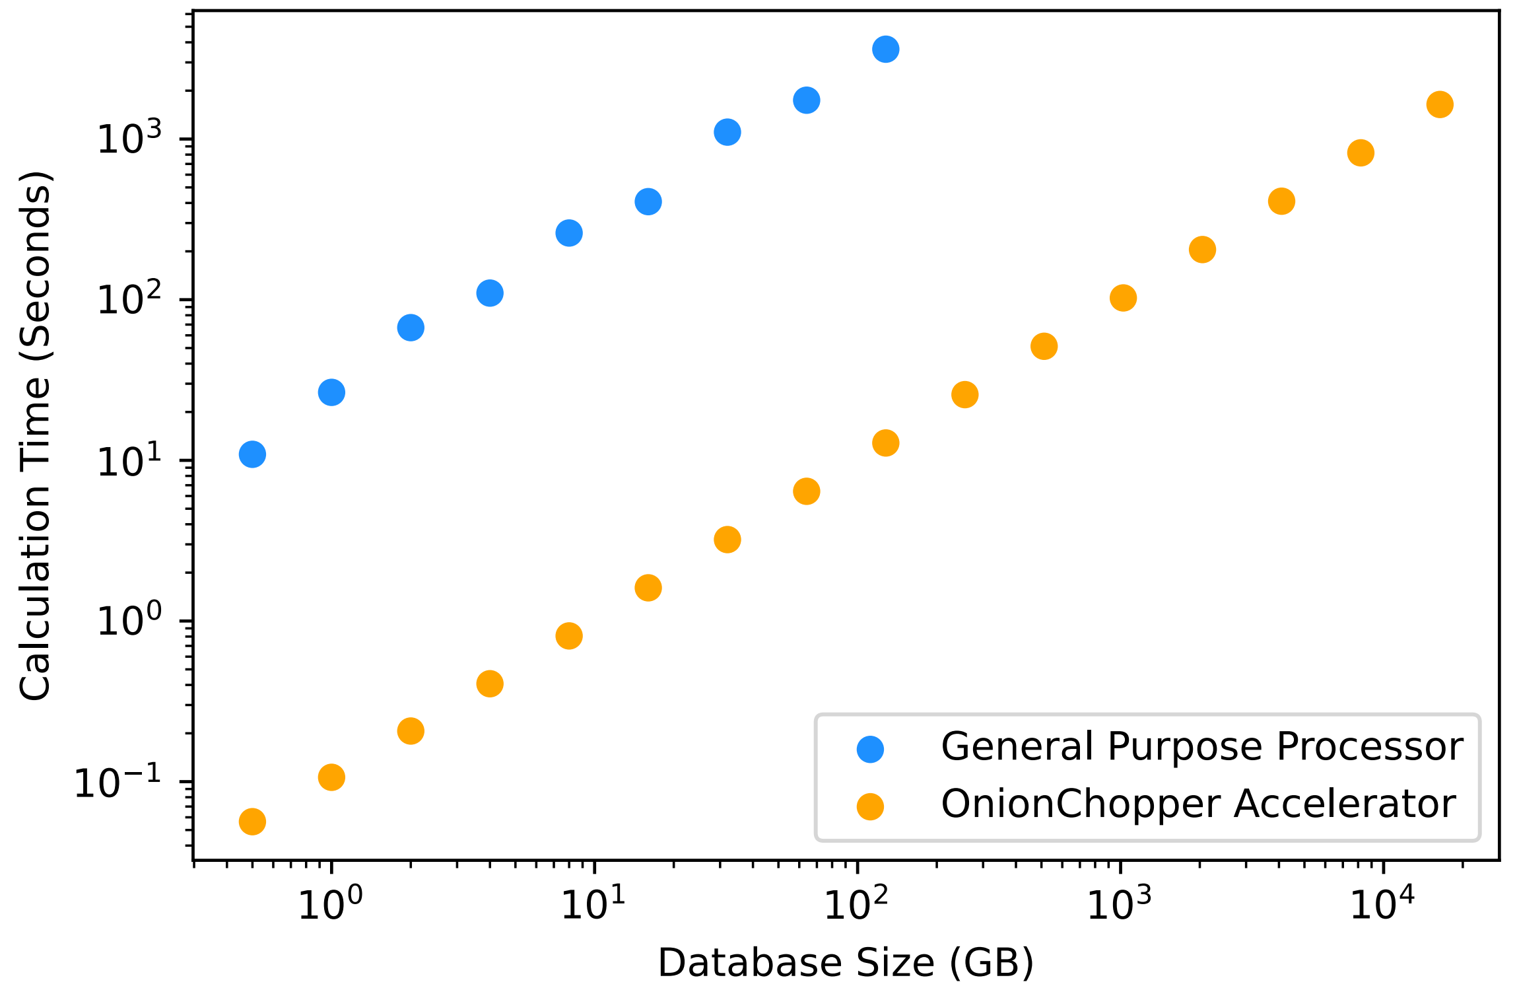

OnionChopper is very extensible to different database sizes, approaching the SSD bandwidth of 10GB/s at large database sizes on OnionPIR and other algorithms.

Database Size vs Speed

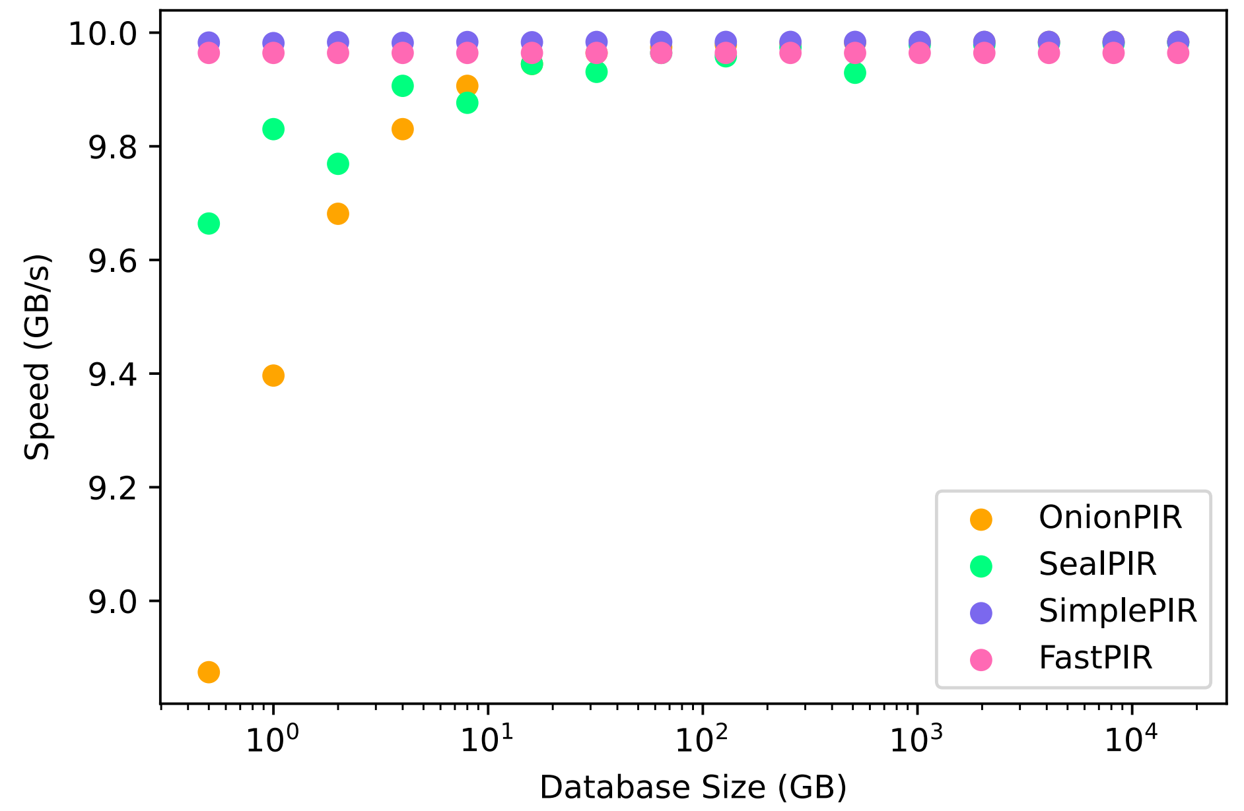

The on-chip accelerator consumes 169J, or 26W during the calculation, roughly one-third of which is from the compute units and two-thirds from the SSD. Similar consumption is found at large database sizes and for other protocols.

Database Size vs Power Consumption

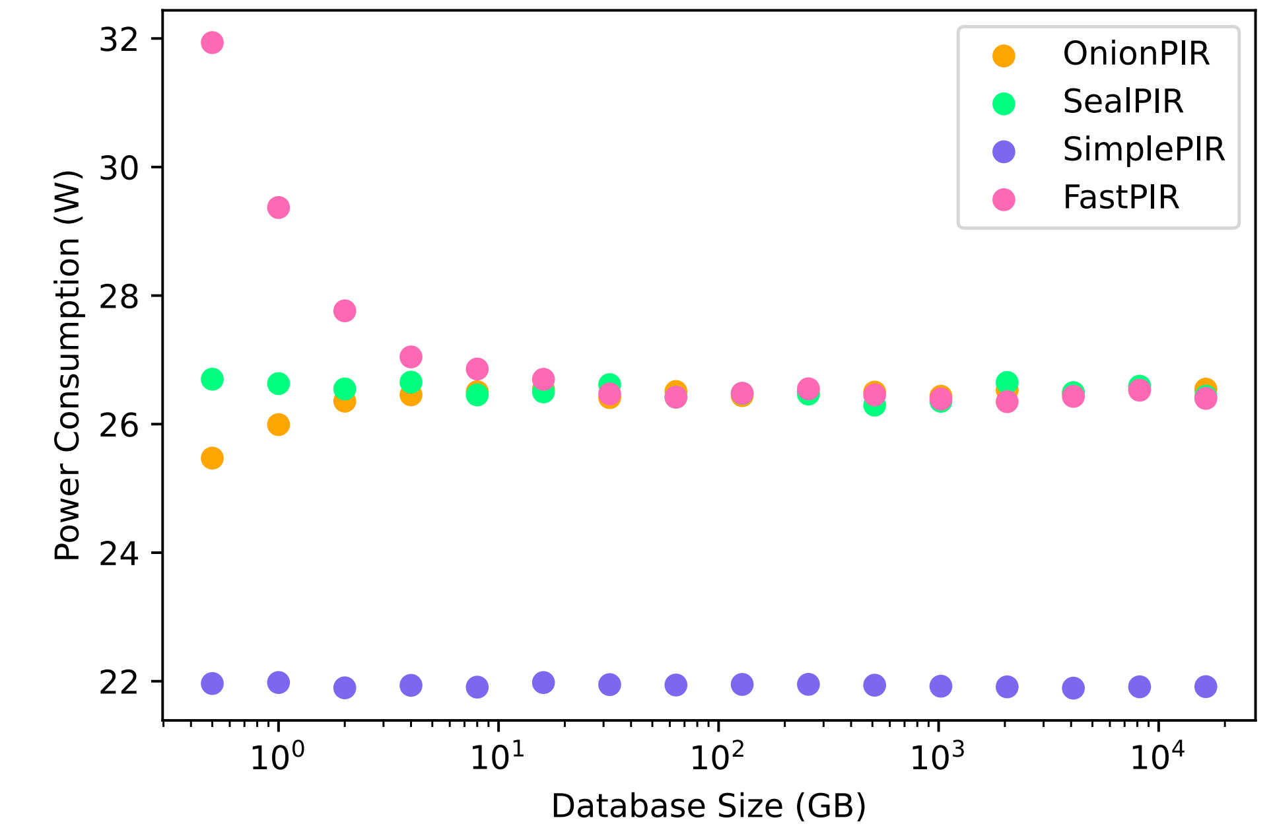

Database Size vs Energy Consumption

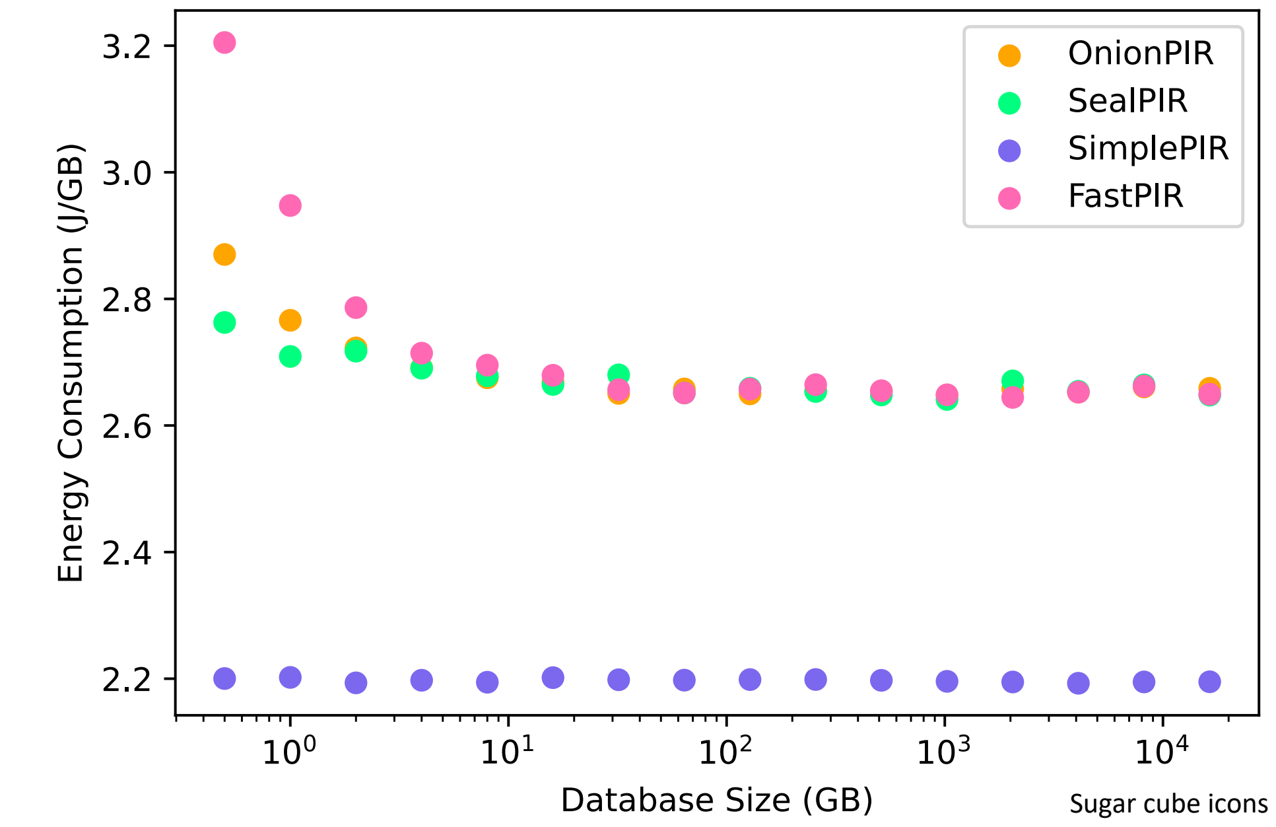

Sugar cube icons created by Freepik – Flaticon
Onion icon by Microsoft