Building a COTS Benchmark Baseline for Graph Analytics

PageRank acceleration for large graphs with scalable hardware and two-step SpMV

Research Problem

PageRank suffers from poor performance and efficiency due to notorious memory access behavior. More importantly, when graphs become bigger and sparser, PageRank applications are inhibited as most solutions strongly rely on large random access fast memory, which is not scalable.

PageRank vector:

$$x_i = \underbrace{\alpha\, x_i^{\mathrm{T}} \mathrm{A}}_{\text{SpMV}} + (1-\alpha)\, x_i^{\mathrm{T}} \frac{ee^{\mathrm{T}}}{\mathrm{N}}$$

Target Graphs

- Very large (~billion nodes)
- Highly sparse (average degree<10)
- No exploitable non-zero pattern

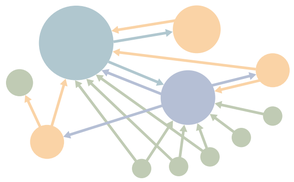

Two-Step SpMV Algorithm

Baseline SpMV: 80M nodes, Avg. degree 3

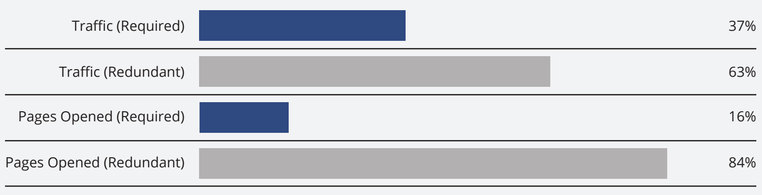

Two-step algorithm conducts SpMV in two separate steps. It requires blocking of the matrix and the source vector as shown below.

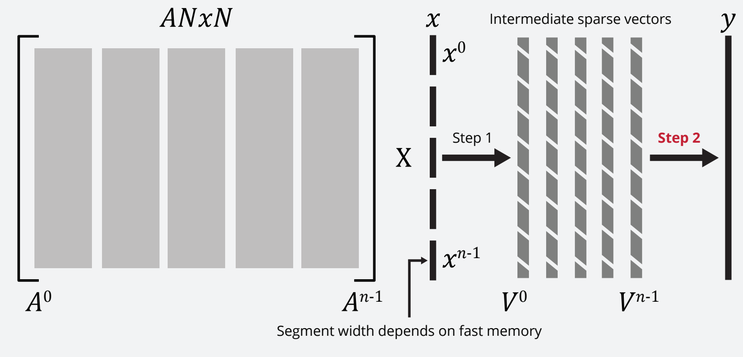

- Guarantees full DRAM streaming access
- Reduces off-chip traffic and enables high-bandwidth utilization
- Requires custom hardware for efficient multi-way merge

Proposed Solution

Custom Hardware with 3D HBM

- Efficient SpMV implementation
- Scalable—less fast memory required

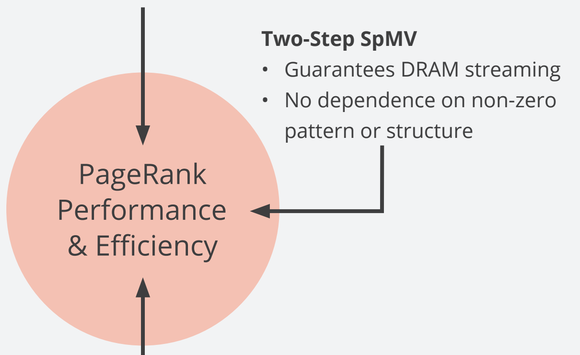

Two-Step SpMV

- Guarantees DRAM streaming
- No dependence on non-zero pattern or structure

Iteration Overlapped PageRank

- Saturates HBM bandwidth
- Reduces off-chip DRAM traffic

16nm FinFET ASIC for PageRank

(can also be realized in COTS FPGA)

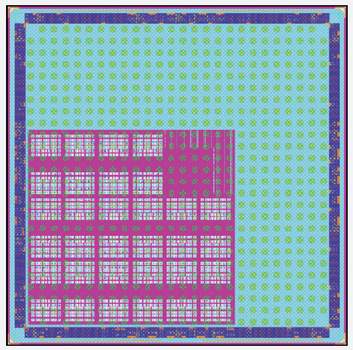

Freq.: 1.4 GHz
Area: 7.5 mm2
Power: 3.11 W

Optimized PageRank by Iteration Overlap (PR_TS_Opt)

Two source vector segment storages in fast memory are required:

1) for computation of Step1 in iteration i+1 and 2) for storing output of Step 2 in iteration i.

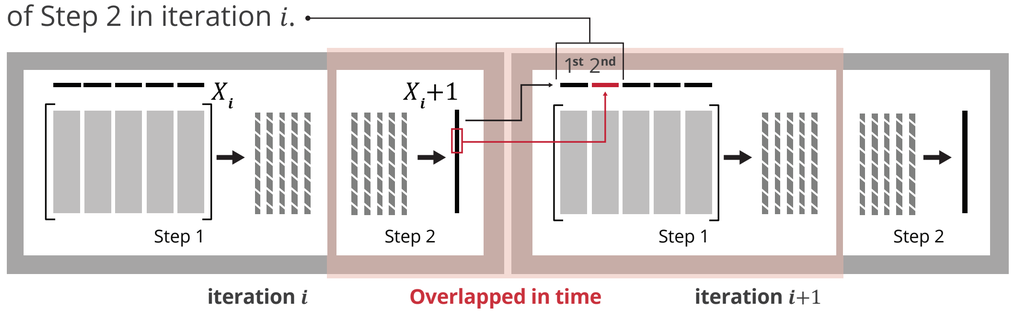

- Step 2 of an iteration runs simultaneously with Step 1 of the next iteration
- Reduces off-chip traffic by eliminating DRAM round trip of both vectors
- Simultaneous Step 1 & 2 doubles the throughput and saturates HBM

Streaming speed PR_TS

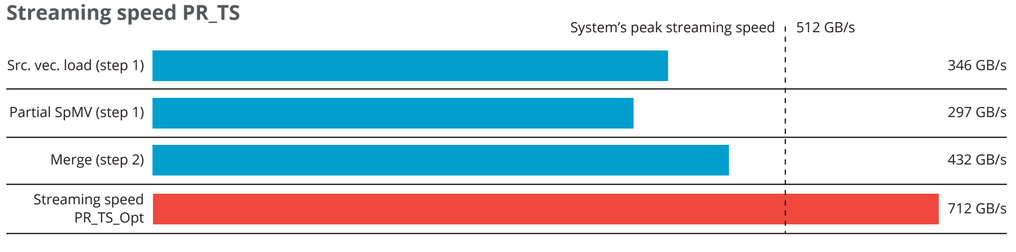

Experimental Results

PageRank Off-chip Traffic Comparison: PR_TS vs Baseline

Matrix 1Bx1B
Avg degree: 3
PageRank: 20 iterations

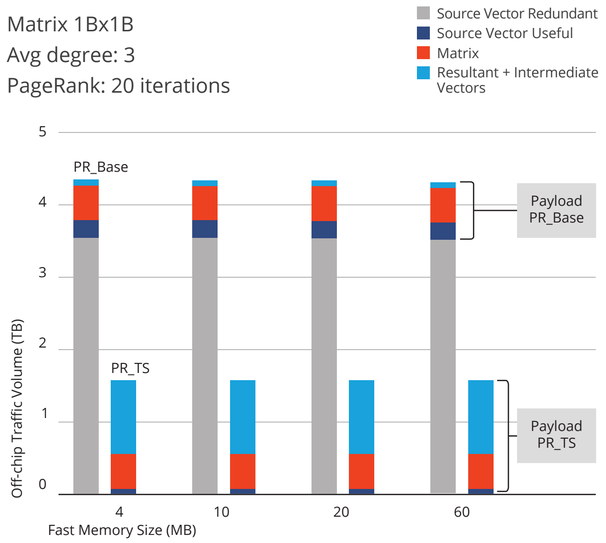

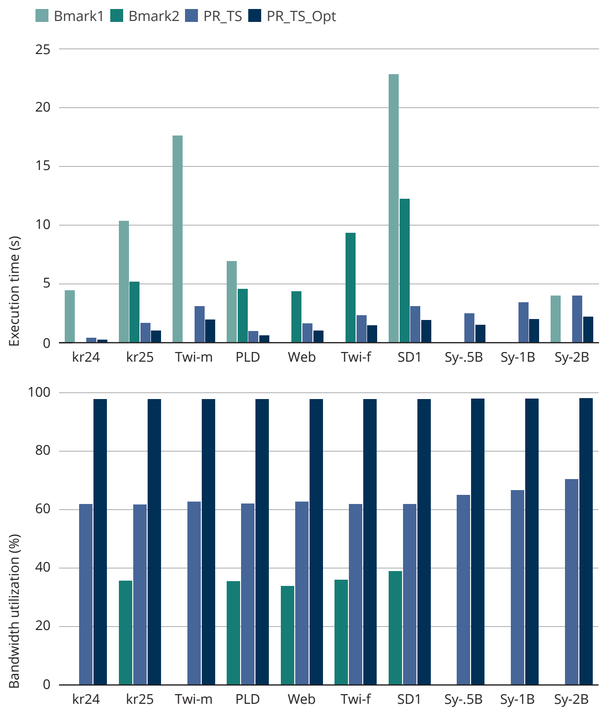

Reference: F. Sadi, J. Sweeney, S. McMillan, T. Z. Low, J. C. Hoe, L. Pileggi, and F. Franchetti, "PageRank Acceleration for Large Graphs with Scalable Hardware and Two-Step SpMV," at *IEEE High Performance Extreme Computing Conference*, Waltham, MA, September 2018.

Dr. Scott McMillan | smcmillan@sei.cmu.edu
with Fazle Sadi, Joe Sweeney, Dr. Tze Meng Low, Dr. James Hoe, Dr. Larry Pileggi, and Dr. Franz Franchetti