Graph analytics on RISC-V GPU
Where are the bottlenecks?

Nimish Shah and Marian Verhelst
nimish.shah@esat.kuleuven.be

Aim

Benchmarking and analyzing Graph Analytics throughput on a RISC-V GPU

Social networks

Road networks

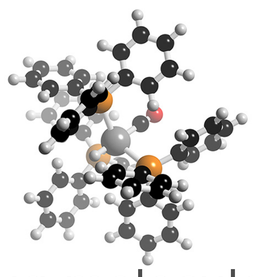

Drug molecules

Biological networks

Real-world graphs

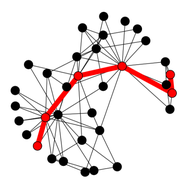

Traversal

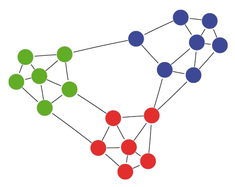

Clustering

Ranking

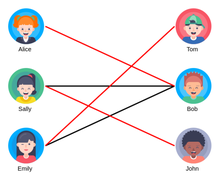

Matching

Analytics

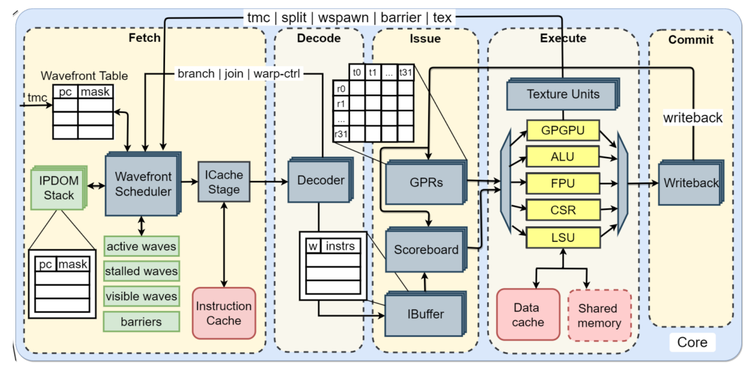

Vortex- a RISC-V GPU [1]

Benchmarking flow

Analytics kernel used: Breadth-first search (BFS)

Target graphs

Name	Domain	Node – Edges
Cora	Citation network of scientific papers	2.7k – 5.4k
Pubmed	Citation network of Pubmed diabetes papers	19k – 44k
CA-HEPh	Collaboration network in high-energy physics	12k – 237k
Cit-HEPh	Citation network in high-energy physics	34k – 421k
Gnutella	Gnutella peer-to-peer network	36k – 88k
Slashdot	Social network	77k – 905k

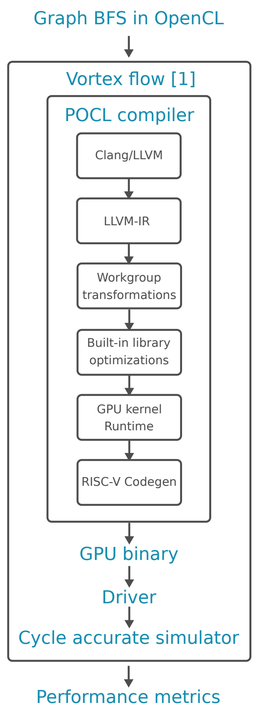

Graph BFS achieves low throughput

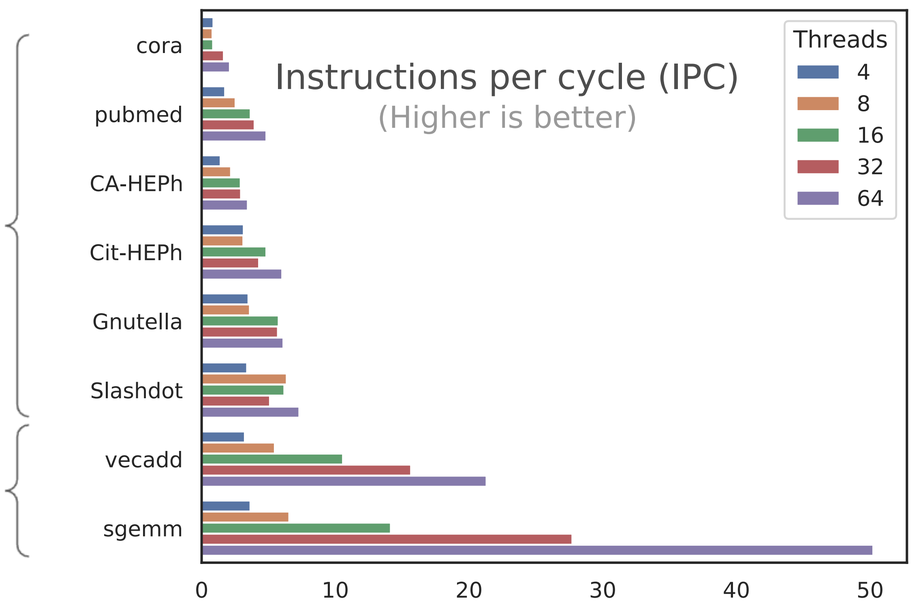

Improper memory coalesing due to irregular memory requests

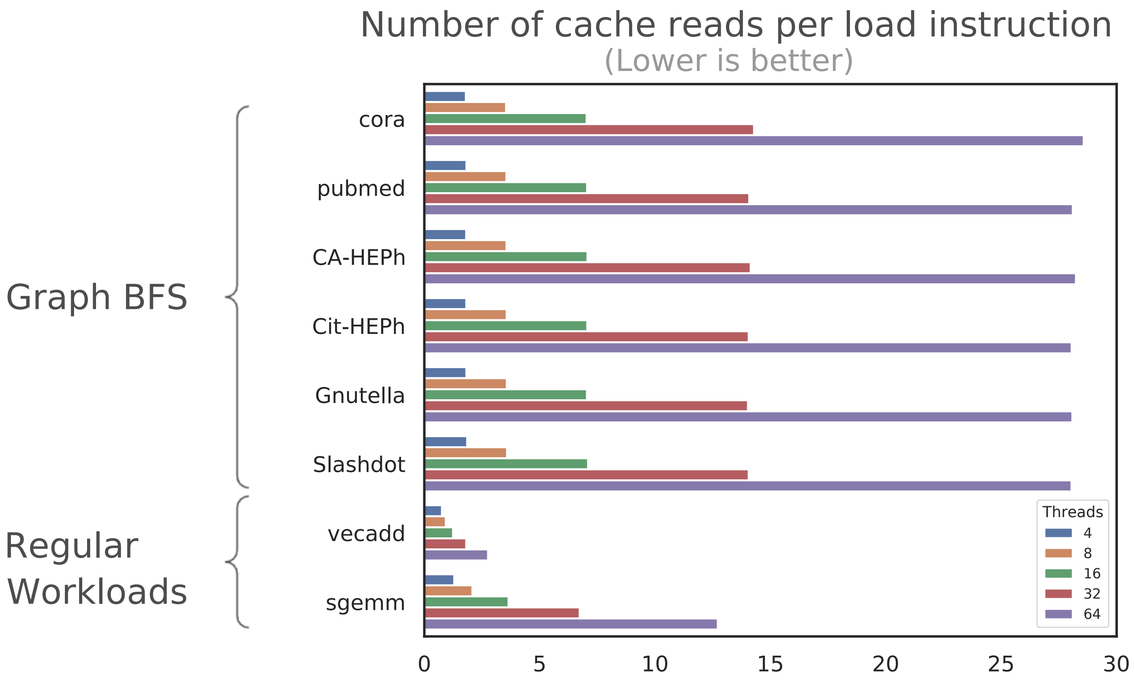

Cache bandwidth underutilized due to irregular memory requests

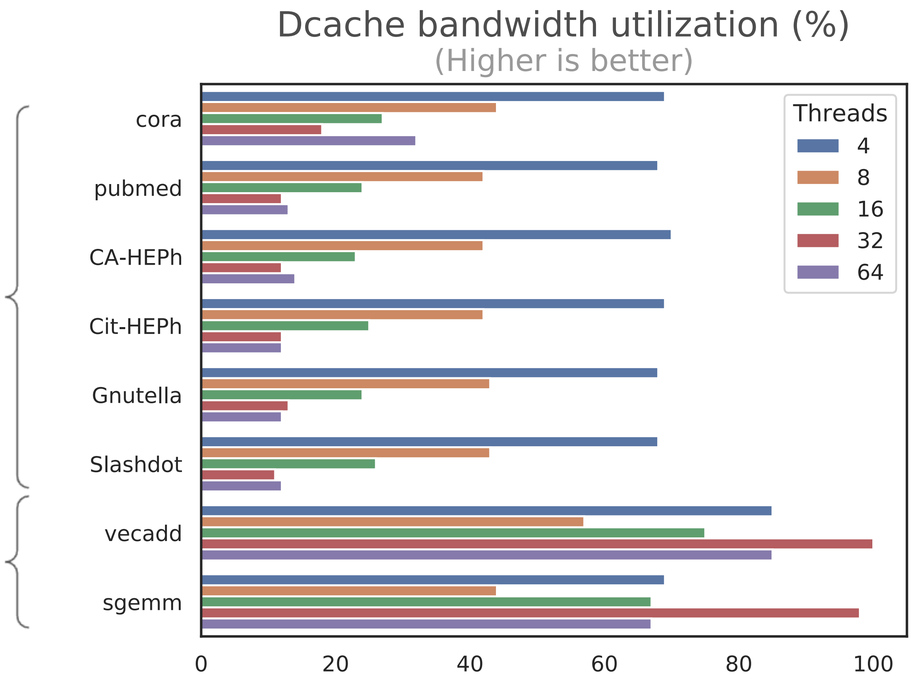

Conclusions

- Graph analytics throughput does not scale with more parallel hardware threads in GPU
- GPU memory hierarchy remains underutilized due to irregular memory requests
- Future work: Upgrade the memory heirarchy to minimize the impact of irregular requests

[1] Tine et al., ”Vortex: Extending the RISC-V ISA for GPGPU and 3D-Graphics”, MICRO, 2021.