Pacific Northwest
NATIONAL LABORATORY

FPGA-Accelerated Ripples

October 5, 2021

Reece Neff

Marco Minutoli

Antonino Tumeo

Michela Becchi

U.S. DEPARTMENT OF ENERGY BATTELLE

PNNL is operated by Battelle for the U.S. Department of Energy

Pacific Northwest
NATIONAL LABORATORY

Influence Maximization Problem

- Find most influential nodes in a graph
- Used for network graphs
  - Social Networks
  - Infectious Disease Networks
  - Etc.

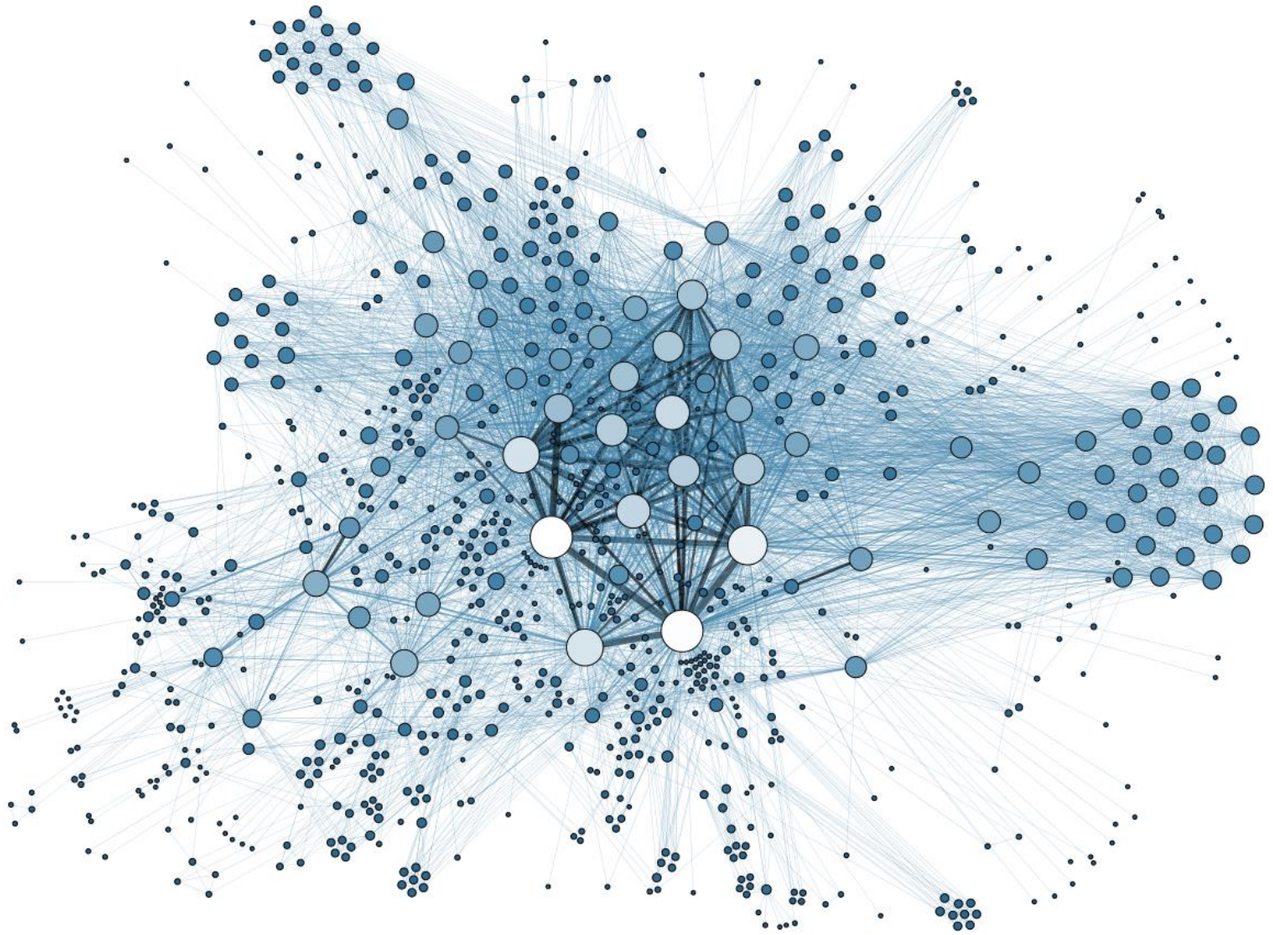

Grandjean, Martin (2014). "La connaissance est un réseau". *Les Cahiers du Numérique* **10** (3): 37-54. DOI:10.3166/LCN.10.3.37-54.

Pedro M. Domingos and Matthew Richardson. “Mining the network value of customers”. In: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, San Francisco, CA, USA, August 26-29, 2001. ACM, 2001, pp. 57–66

Pacific Northwest NATIONAL LABORATORY

Influence Maximization Algorithm

```
Input: G, k, ε
Output: S

begin:
  {R,θ}←EstimateTheta(G,k,ε)
  R←Sample(G,θ-|R|,R)
  S←SelectSeeds(G,k,R)
end
```

- **Estimate Theta**
  - Estimate the number of graphs to be generated by the Sampling stage to achieve an approximation factor determined by ε and the number of seeds k
- **Sampling** *(Target for Acceleration)*
  - Generate several graphs determined by Estimate Theta using a given diffusion model.
  - Linear threshold model used for this work
- **Seed Selection**
  - Find the node with the most occurrences in all graphs from the Sampling stage.
  - Remove all graphs with this node
  - Repeat until k seeds reached

Tang, Youze, Yanchen Shi, and Xiaokui Xiao. "Influence maximization in near-linear time: A martingale approach." *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. 2015.

Pacific Northwest
NATIONAL LABORATORY

Linear Threshold Model

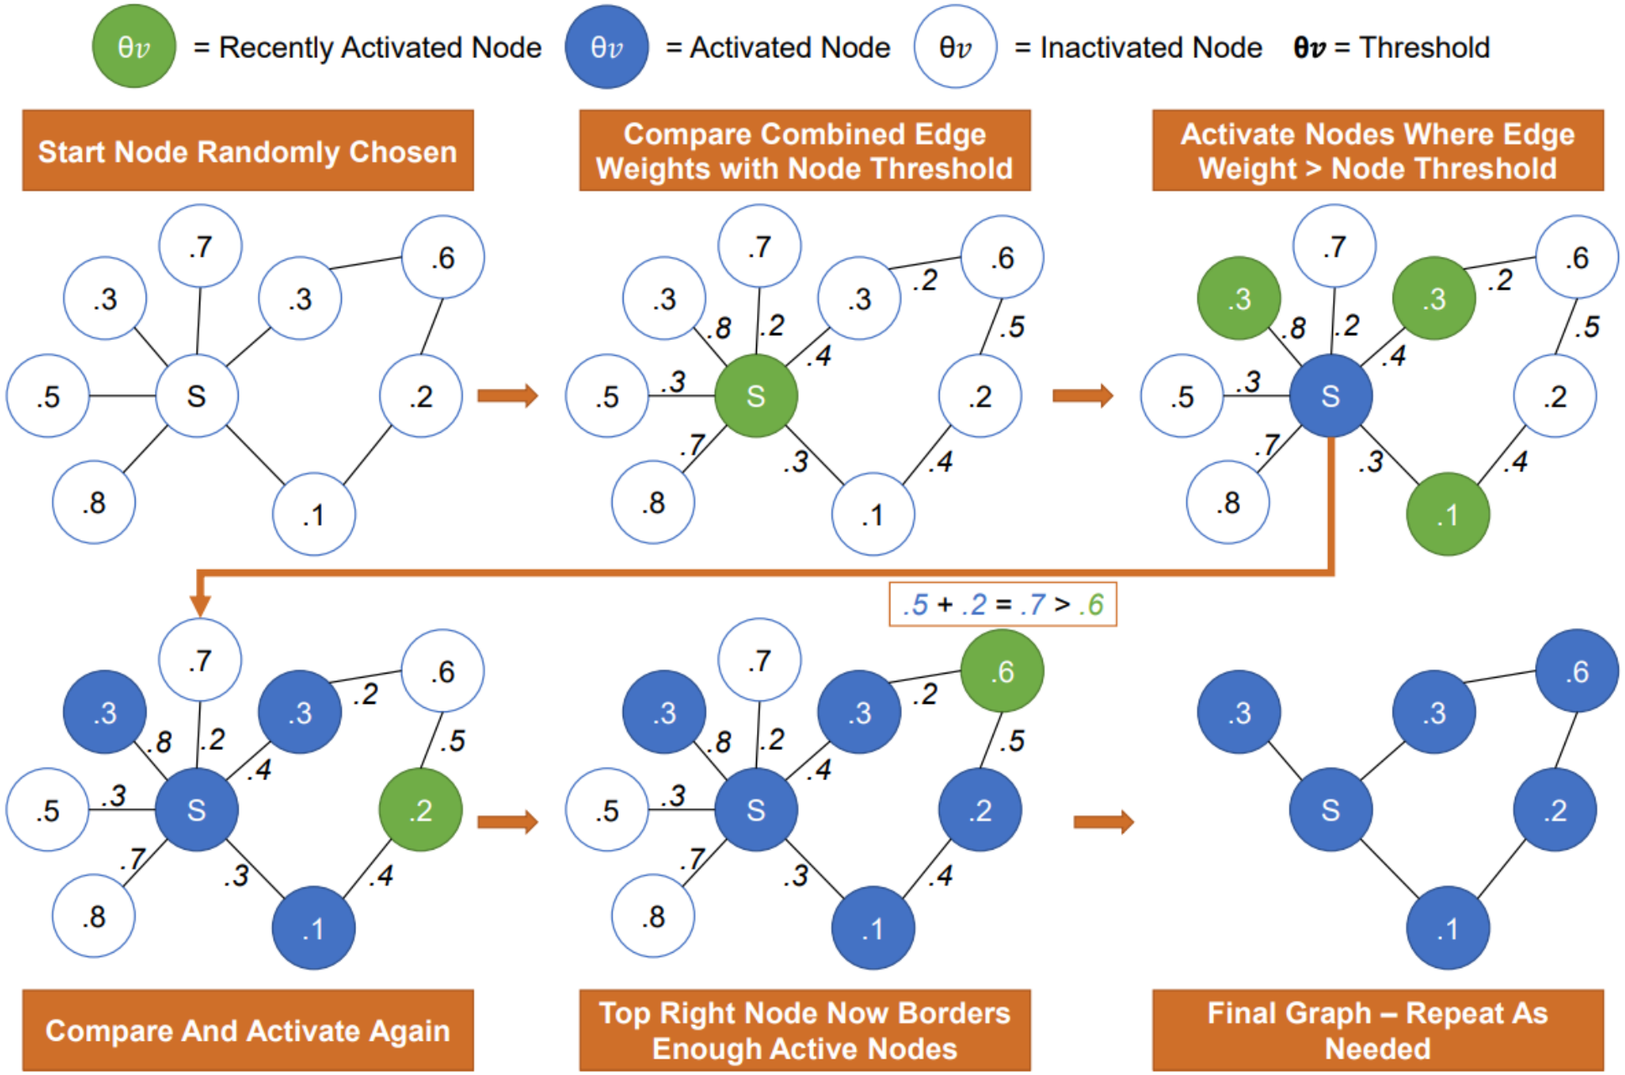

Pacific Northwest
NATIONAL LABORATORY

Acceleration?

- Parallelized in CPU
  - M. Minutoli, M. Halappanavar, A. Kalyanaraman, A. Sathanur, R. Mcclure and J. McDermott, "Fast and Scalable Implementations of Influence Maximization Algorithms," *2019 IEEE International Conference on Cluster Computing (CLUSTER),* 2019, pp. 1-12, doi: 10.1109/CLUSTER.2019.8890991.
- Accelerated with GPU
  - Marco Minutoli, Maurizio Drocco, Mahantesh Halappanavar, Antonino Tumeo, and Ananth Kalyanaraman. 2020. *CuRipples: influence maximization on multi-GPU systems. In Proceedings of the 34th ACM International Conference on Supercomputing (ICS '20).* Association for Computing Machinery, New York, NY, USA, Article 12, 1–11. DOI:https://doi.org/10.1145/3392717.3392750
- FPGA Acceleration
  - Hardware level parallelism and pipelining
  - Lower power consumption

Pacific Northwest
NATIONAL LABORATORY

Acceleration Challenges

- Irregular memory accesses
- Large, complex loops inhibit FPGA pipelining

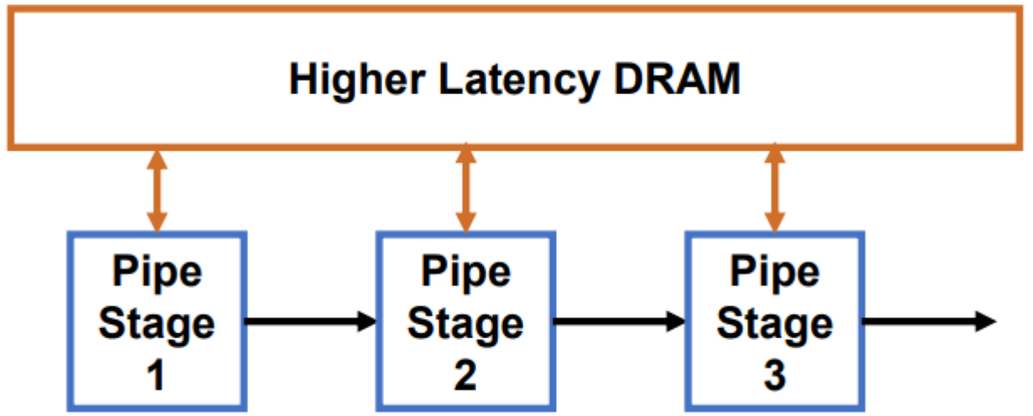

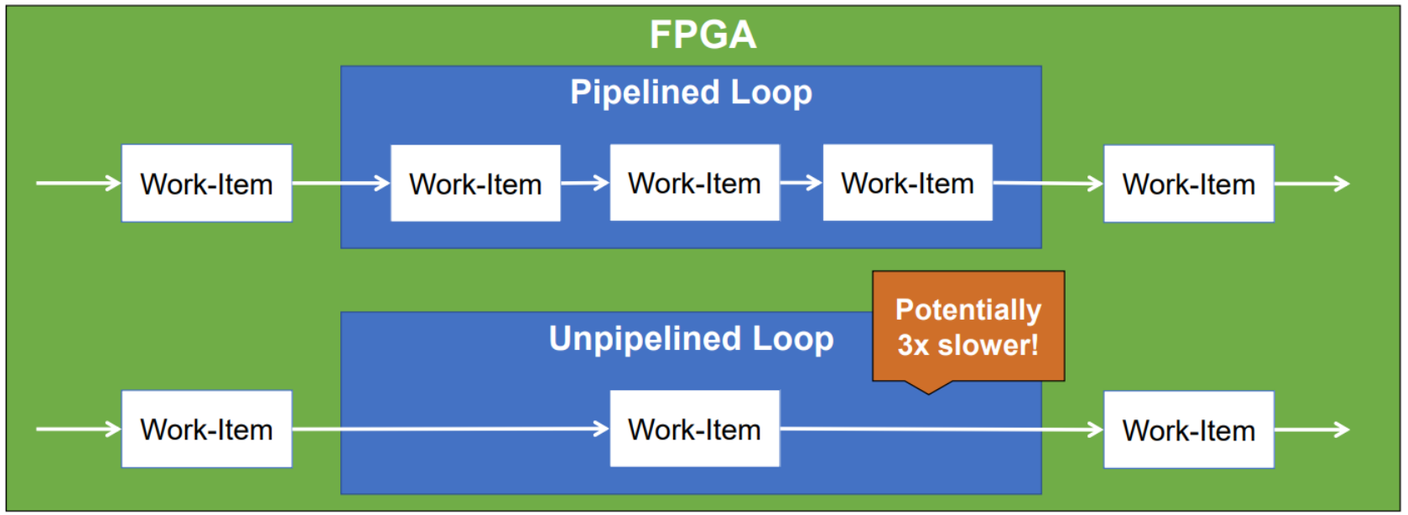

Pacific Northwest
NATIONAL LABORATORY

FPGA Optimizations

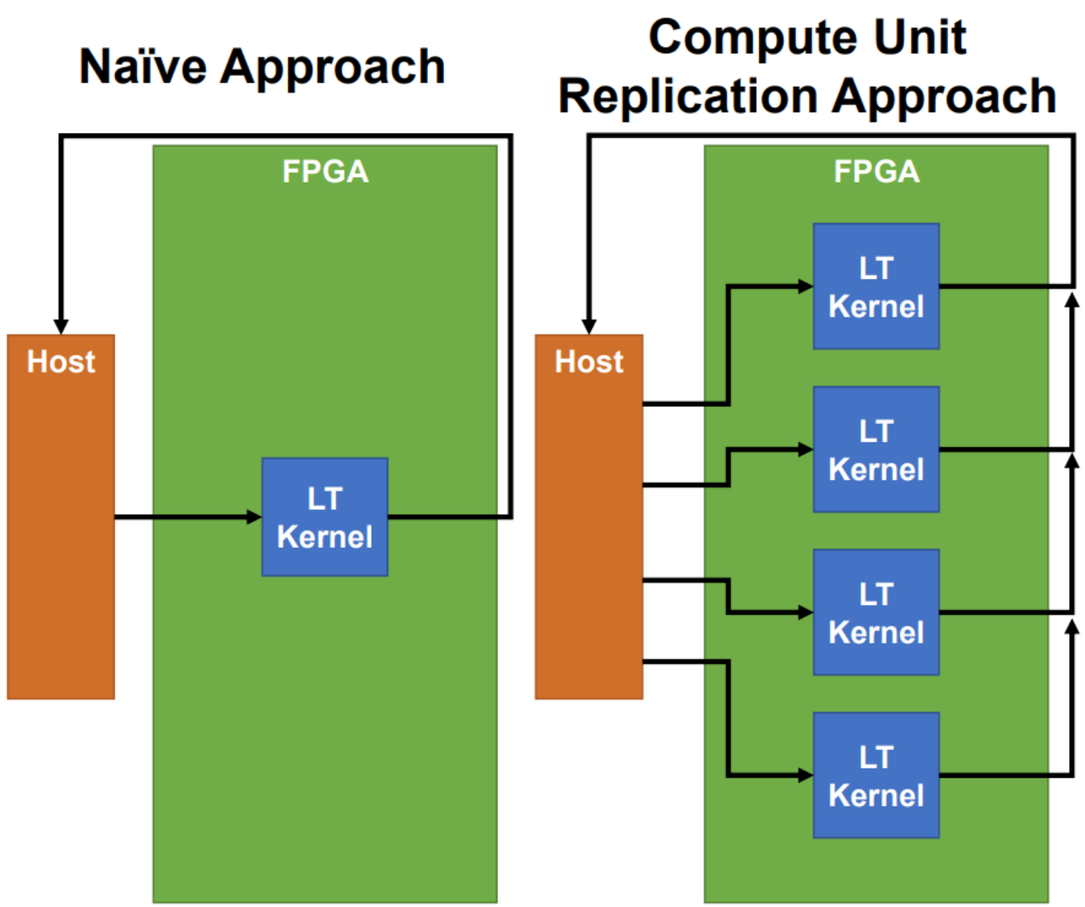

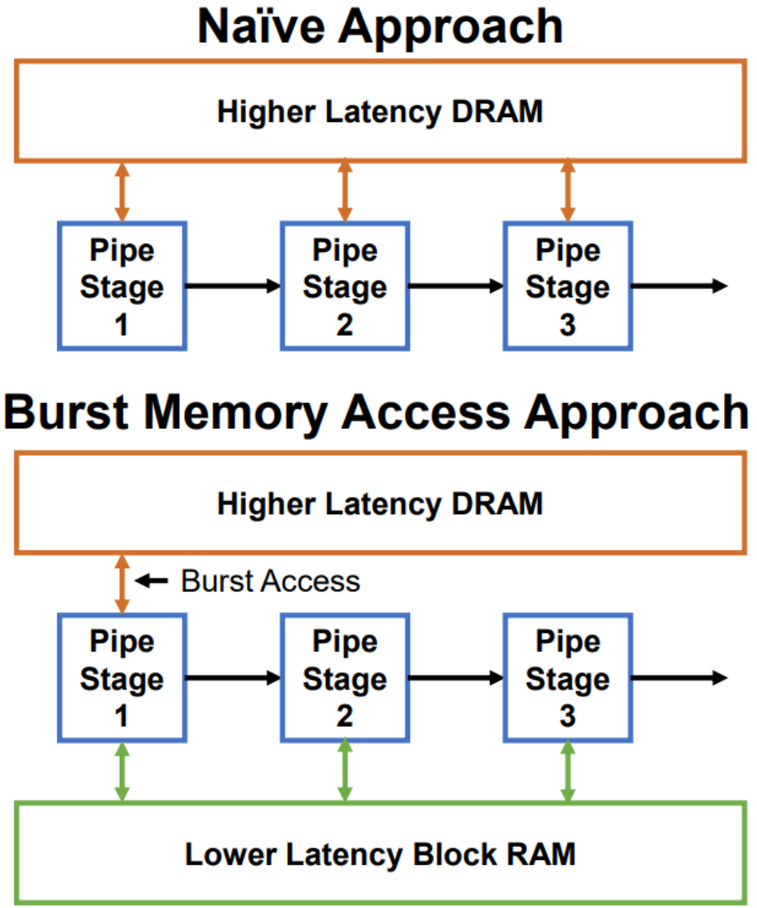

Pacific Northwest NATIONAL LABORATORY

Experimental Setup

Component Type	Component
CPU	Intel Xeon E5-2367 w/ 8x32GB 2400MHz DDR4 RAM
FPGA	Xilinx Alveo U250 w/ 4x16GB DDR4 RAM (1x16GB used)
Input Graph	Selections from SNAP[2]
HLS Tool	Vitis Unified Software Platform 2020.2 (OpenCL to FPGA)

Graph	Nodes	Edges	Avg. Degree	Max Degree
cit-HepTh	27,770	352,807	12.70	2,468
web-BerkStan	685,230	7,600,595	22.18	84,290
web-Google	875,713	5,105,039	11.66	6,353
soc-pokec-relationships	1,632,803	30,622,564	37.51	20,518
wiki-topcats	1,791,489	28,511,807	31.83	3,907
com-Orkut	3,072,441	117,185,083	76.28	33,313
soc-LiveJournal1	4,847,571	68,993,773	28.47	22,889

Preliminary Results

CPU Only = 8 CPU Cores
FPGA = 7 CPU Cores + 1 FPGA Device

Performance

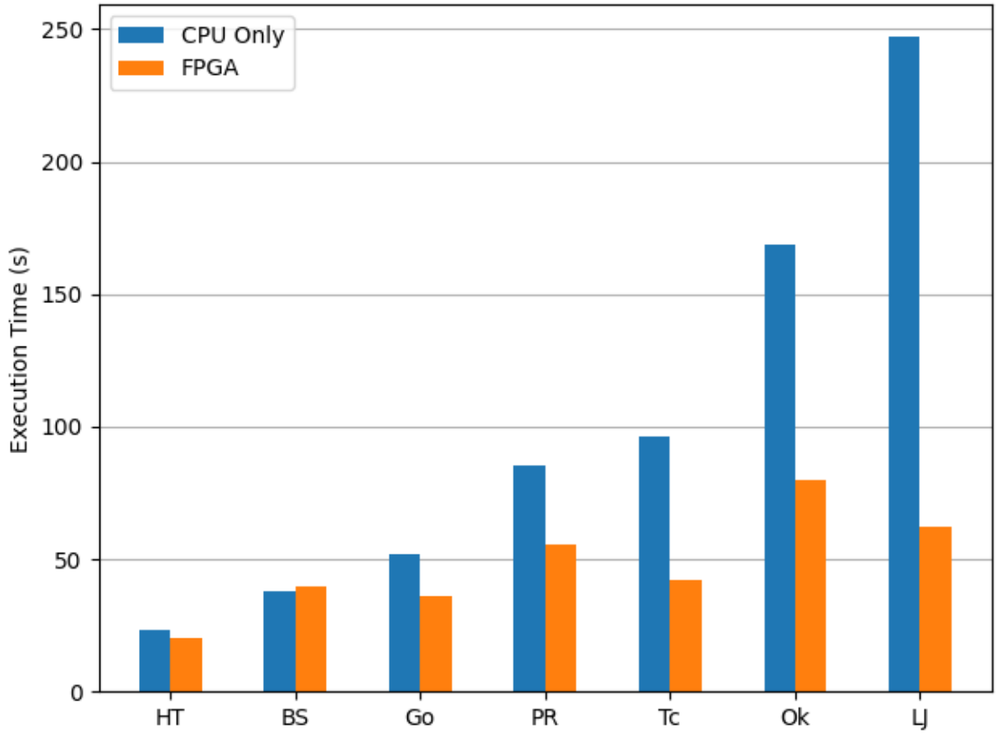

Energy Consumption

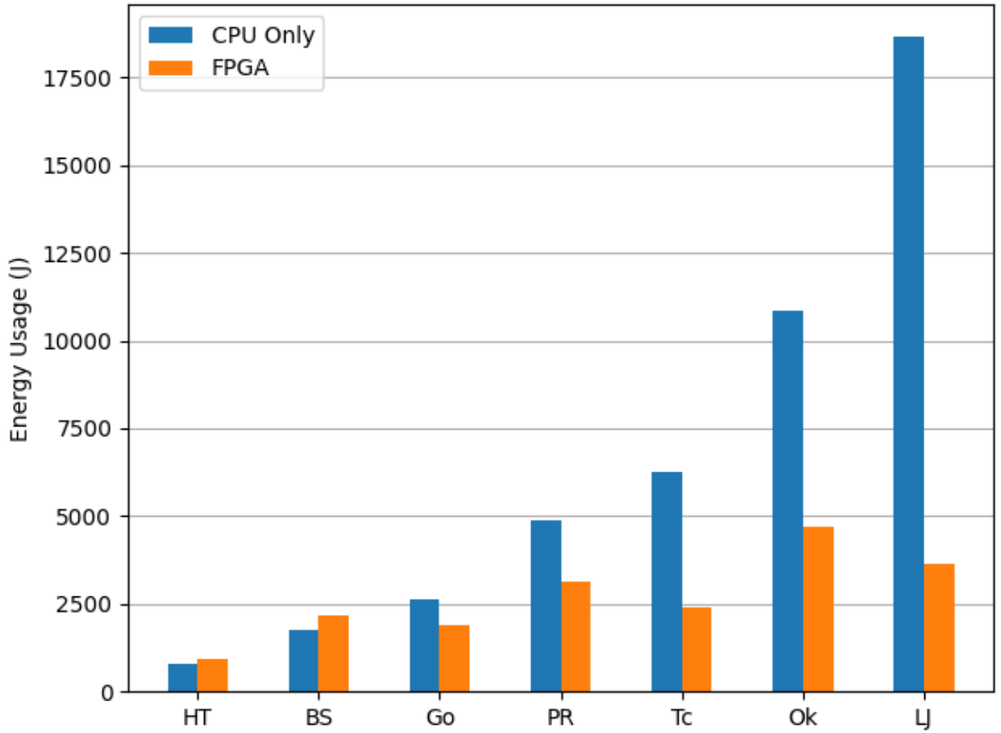

Resource Usage (% of total)

Kernel	FFs	LUTs	BRAMs	DSP
Graph Initialization	5,763 (0.17%)	4,544 (0.28%)	4 (0.16%)	0
Linear Threshold	43,345 (1.32%)	32,323 (1.99%)	24 (0.95%)	23 (1.88%)
Total	179,143 (5.44%)	133,836 (8.25%)	100 (3.96%)	924 (7.52%)

Pacific Northwest NATIONAL LABORATORY

Preliminary Results Tables

Execution Time (s)	CPU Only	FPGA	Speedup
cit-HepTh	23.21	20.01	1.16
web-BerkStan	37.81	39.97	0.95
web-Google	52.15	35.92	1.45
soc-pokec-relationships	85.17	55.37	1.54
wiki-topcats	96.08	42.04	2.29
com-Orkut	168.46	80.02	2.11
soc-LiveJournal1	247.04	62.32	**3.96**

Energy Usage (J)	CPU Only	FPGA	Savings
cit-HepTh	774.08	935.85	0.83
web-BerkStan	1739.35	2190.06	0.79
web-Google	2616.09	1877.08	1.39
soc-pokec-relationships	4880.21	3117.28	1.57
wiki-topcats	6250.00	2416.80	2.59
com-Orkut	10846.58	4702.91	2.31
soc-LiveJournal1	18645.45	3623.19	**5.15**

Pacific Northwest
NATIONAL LABORATORY

Conclusion and Future Work

- Implemented Linear Threshold Sampling on FPGA
  - Added several optimizations to improve the naïve port
- Achieved up to 3.96x Speedup and 5.15x Energy Savings over CPU baseline

- Future Work
  - Multiple Super Logic Regions
  - IC Diffusion Model
  - Seed Selection

Pacific Northwest
NATIONAL LABORATORY

Thank you

Questions?
Email me at reece.neff@pnnl.gov