


SIMD-X: Programming and Processing of Graph Algorithms on GPUs



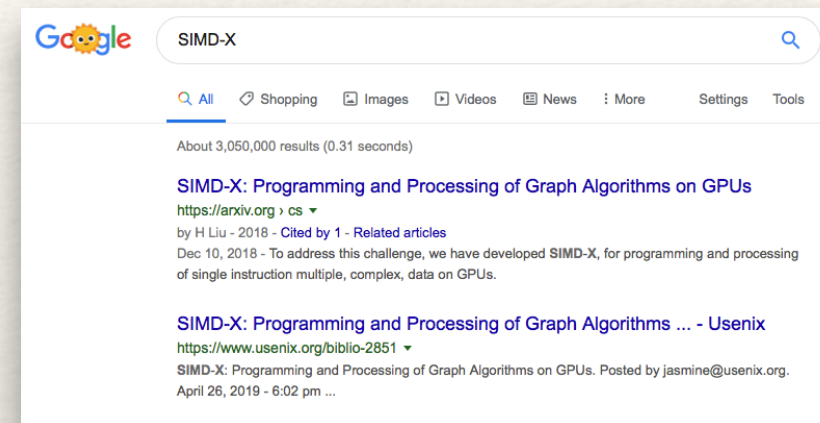
Hang Liu

H. Howie Huang

University of Massachusetts Lowell


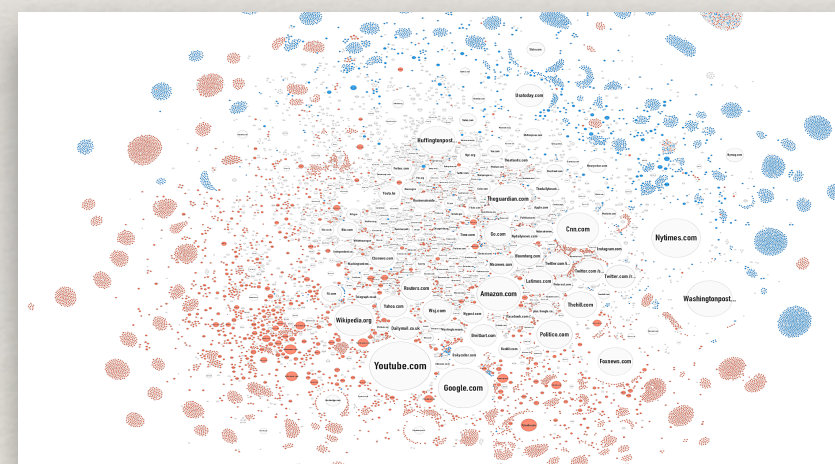
George Washington University

Graphs are Everywhere ...





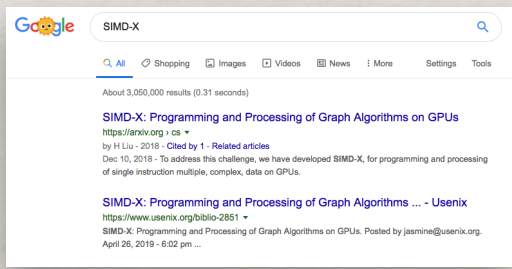

...

Our Daily Life



...

Graph Algorithms are Insightful

Algorithm	Functionality	Graph
Triangle completing	Friend recommendation	
Shortest path	Navigation	
PageRank	Webpage searching	
Matrix factorization	Product recommendation	

Graph computation is NOT easy

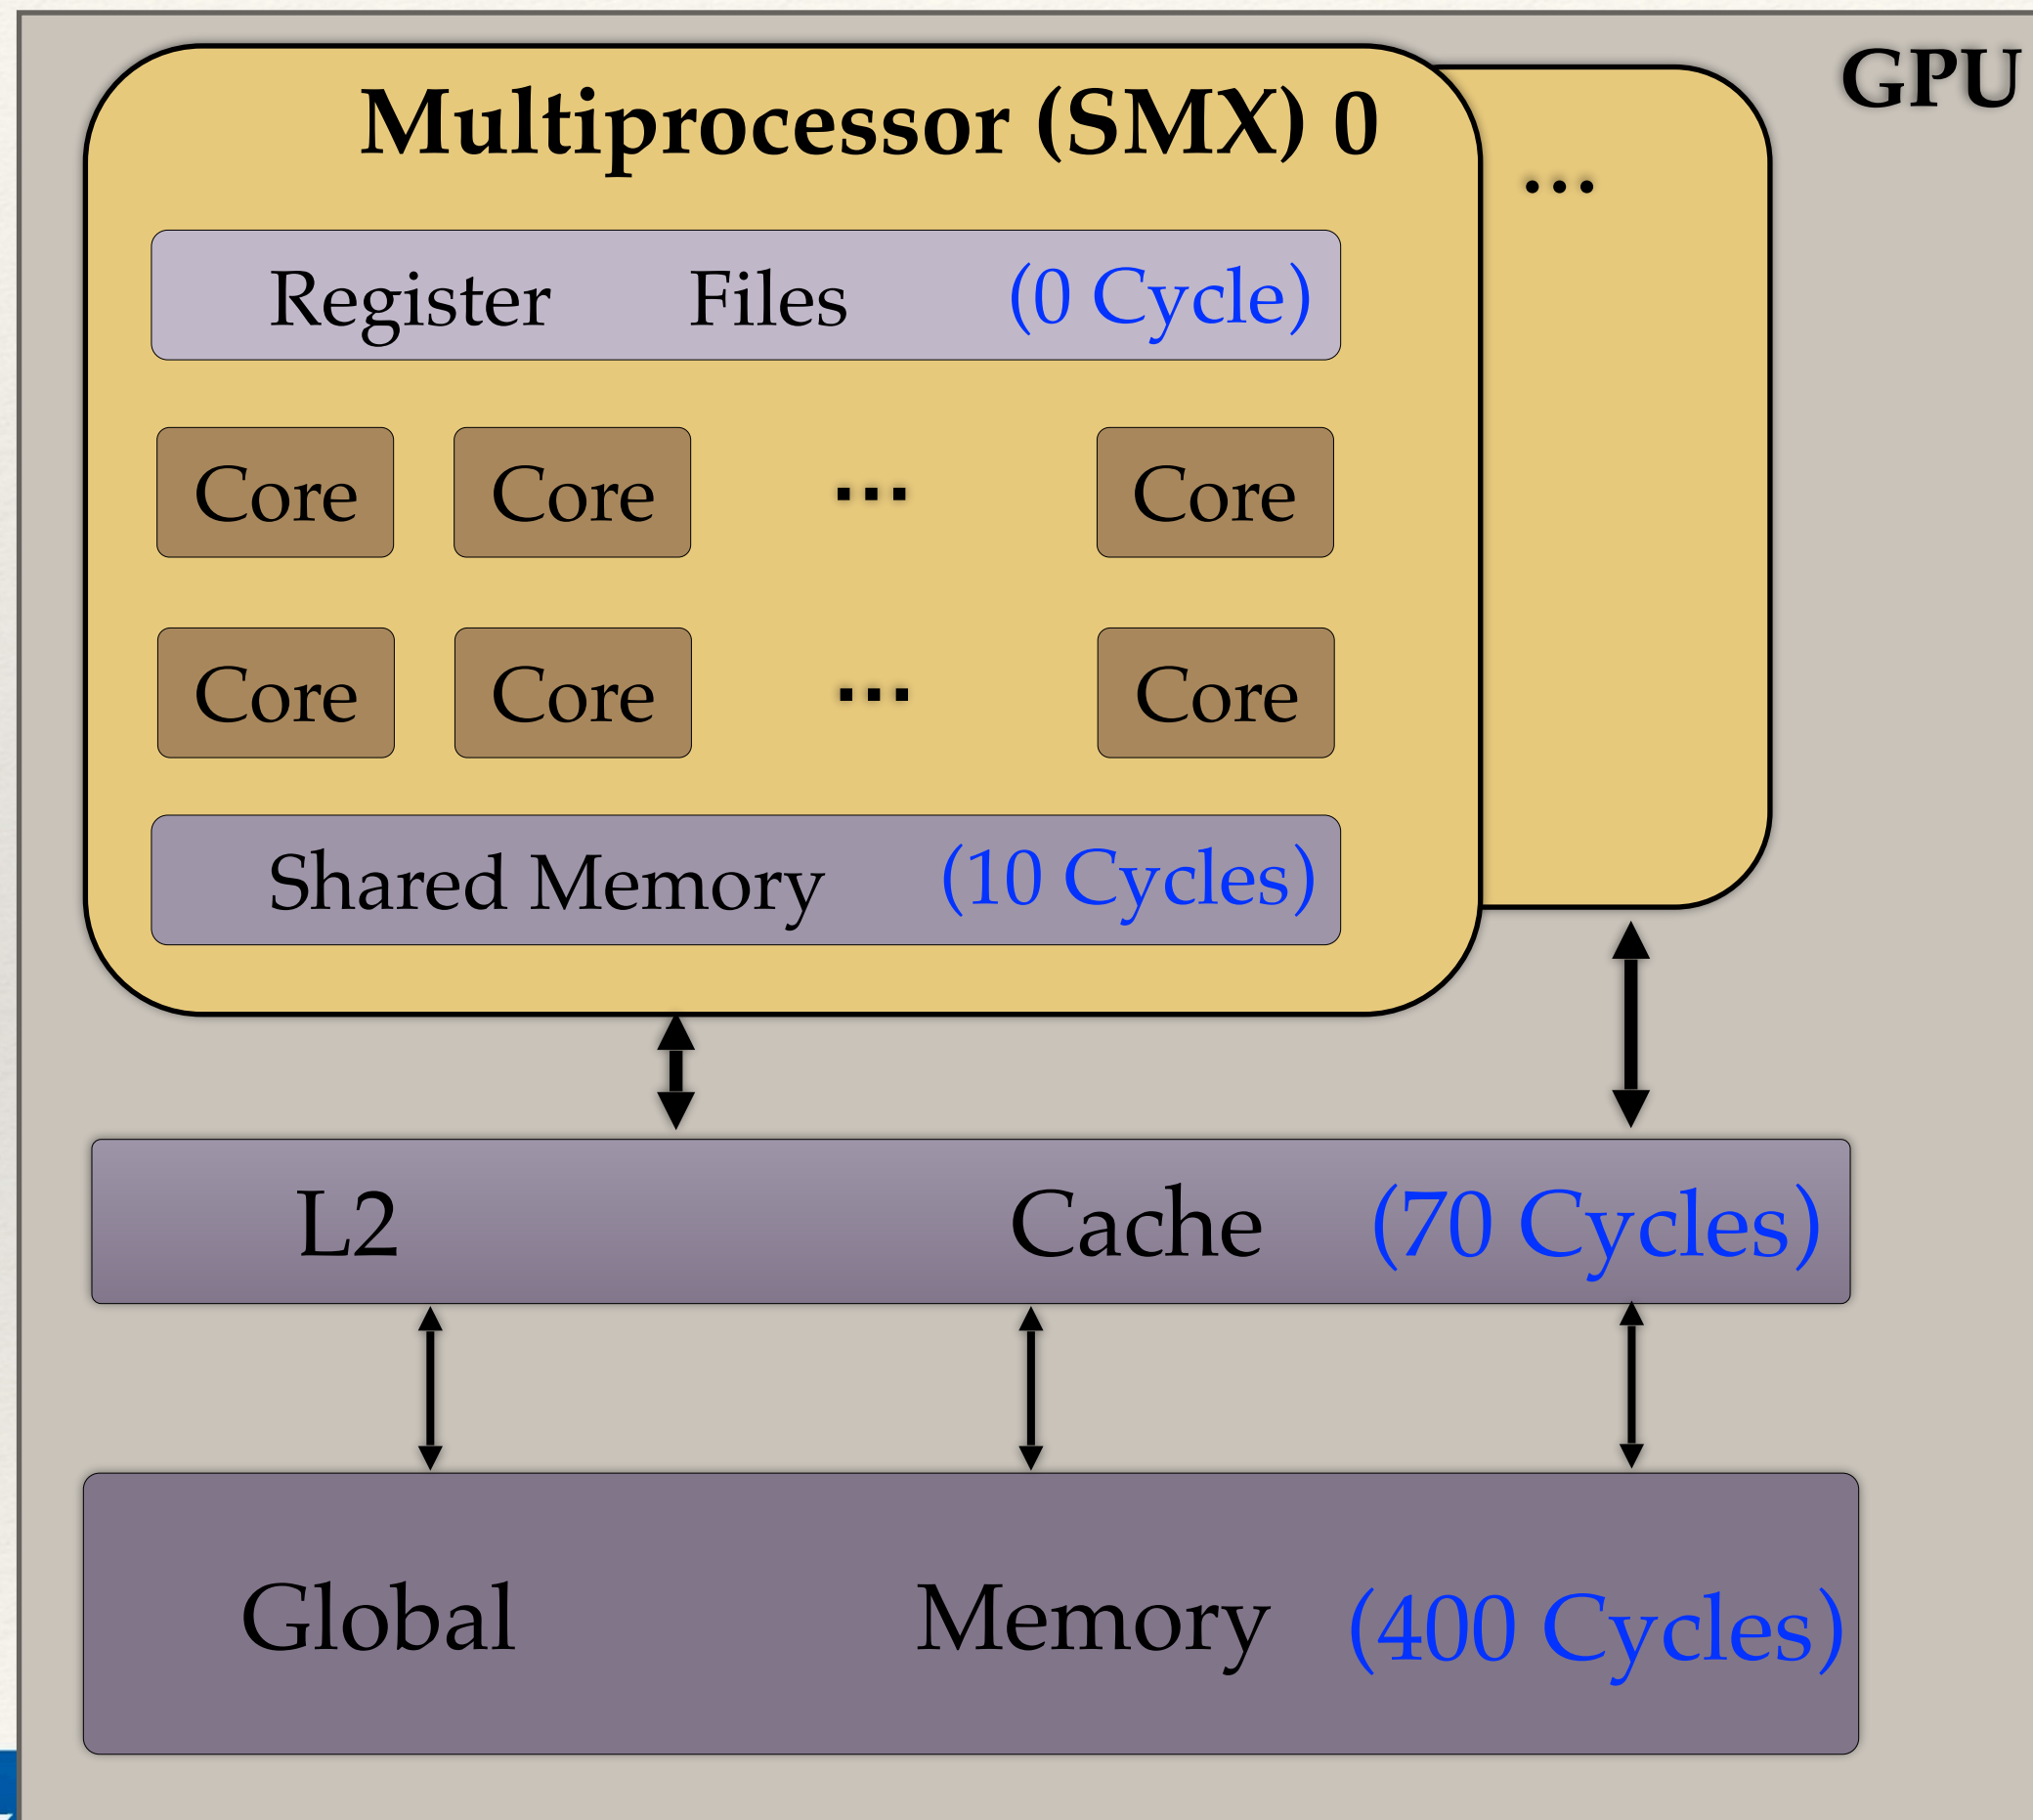
Algorithm	Complexity	Runtime
Triangle completing	$ E * D_{\max}^2$	20 mins
Shortest path	$ E + V \sim V * E $	17 mins
PageRank	$n * E $	41 mins
Matrix factorization	$n * L * E $	2 hours

* V and E: vertex and edge count
n: #iterations to converge
L: #latent factors

* PowerGraph [OSDI '12]
on Twitter dataset
❖ 53M vertices
❖ 2B edges

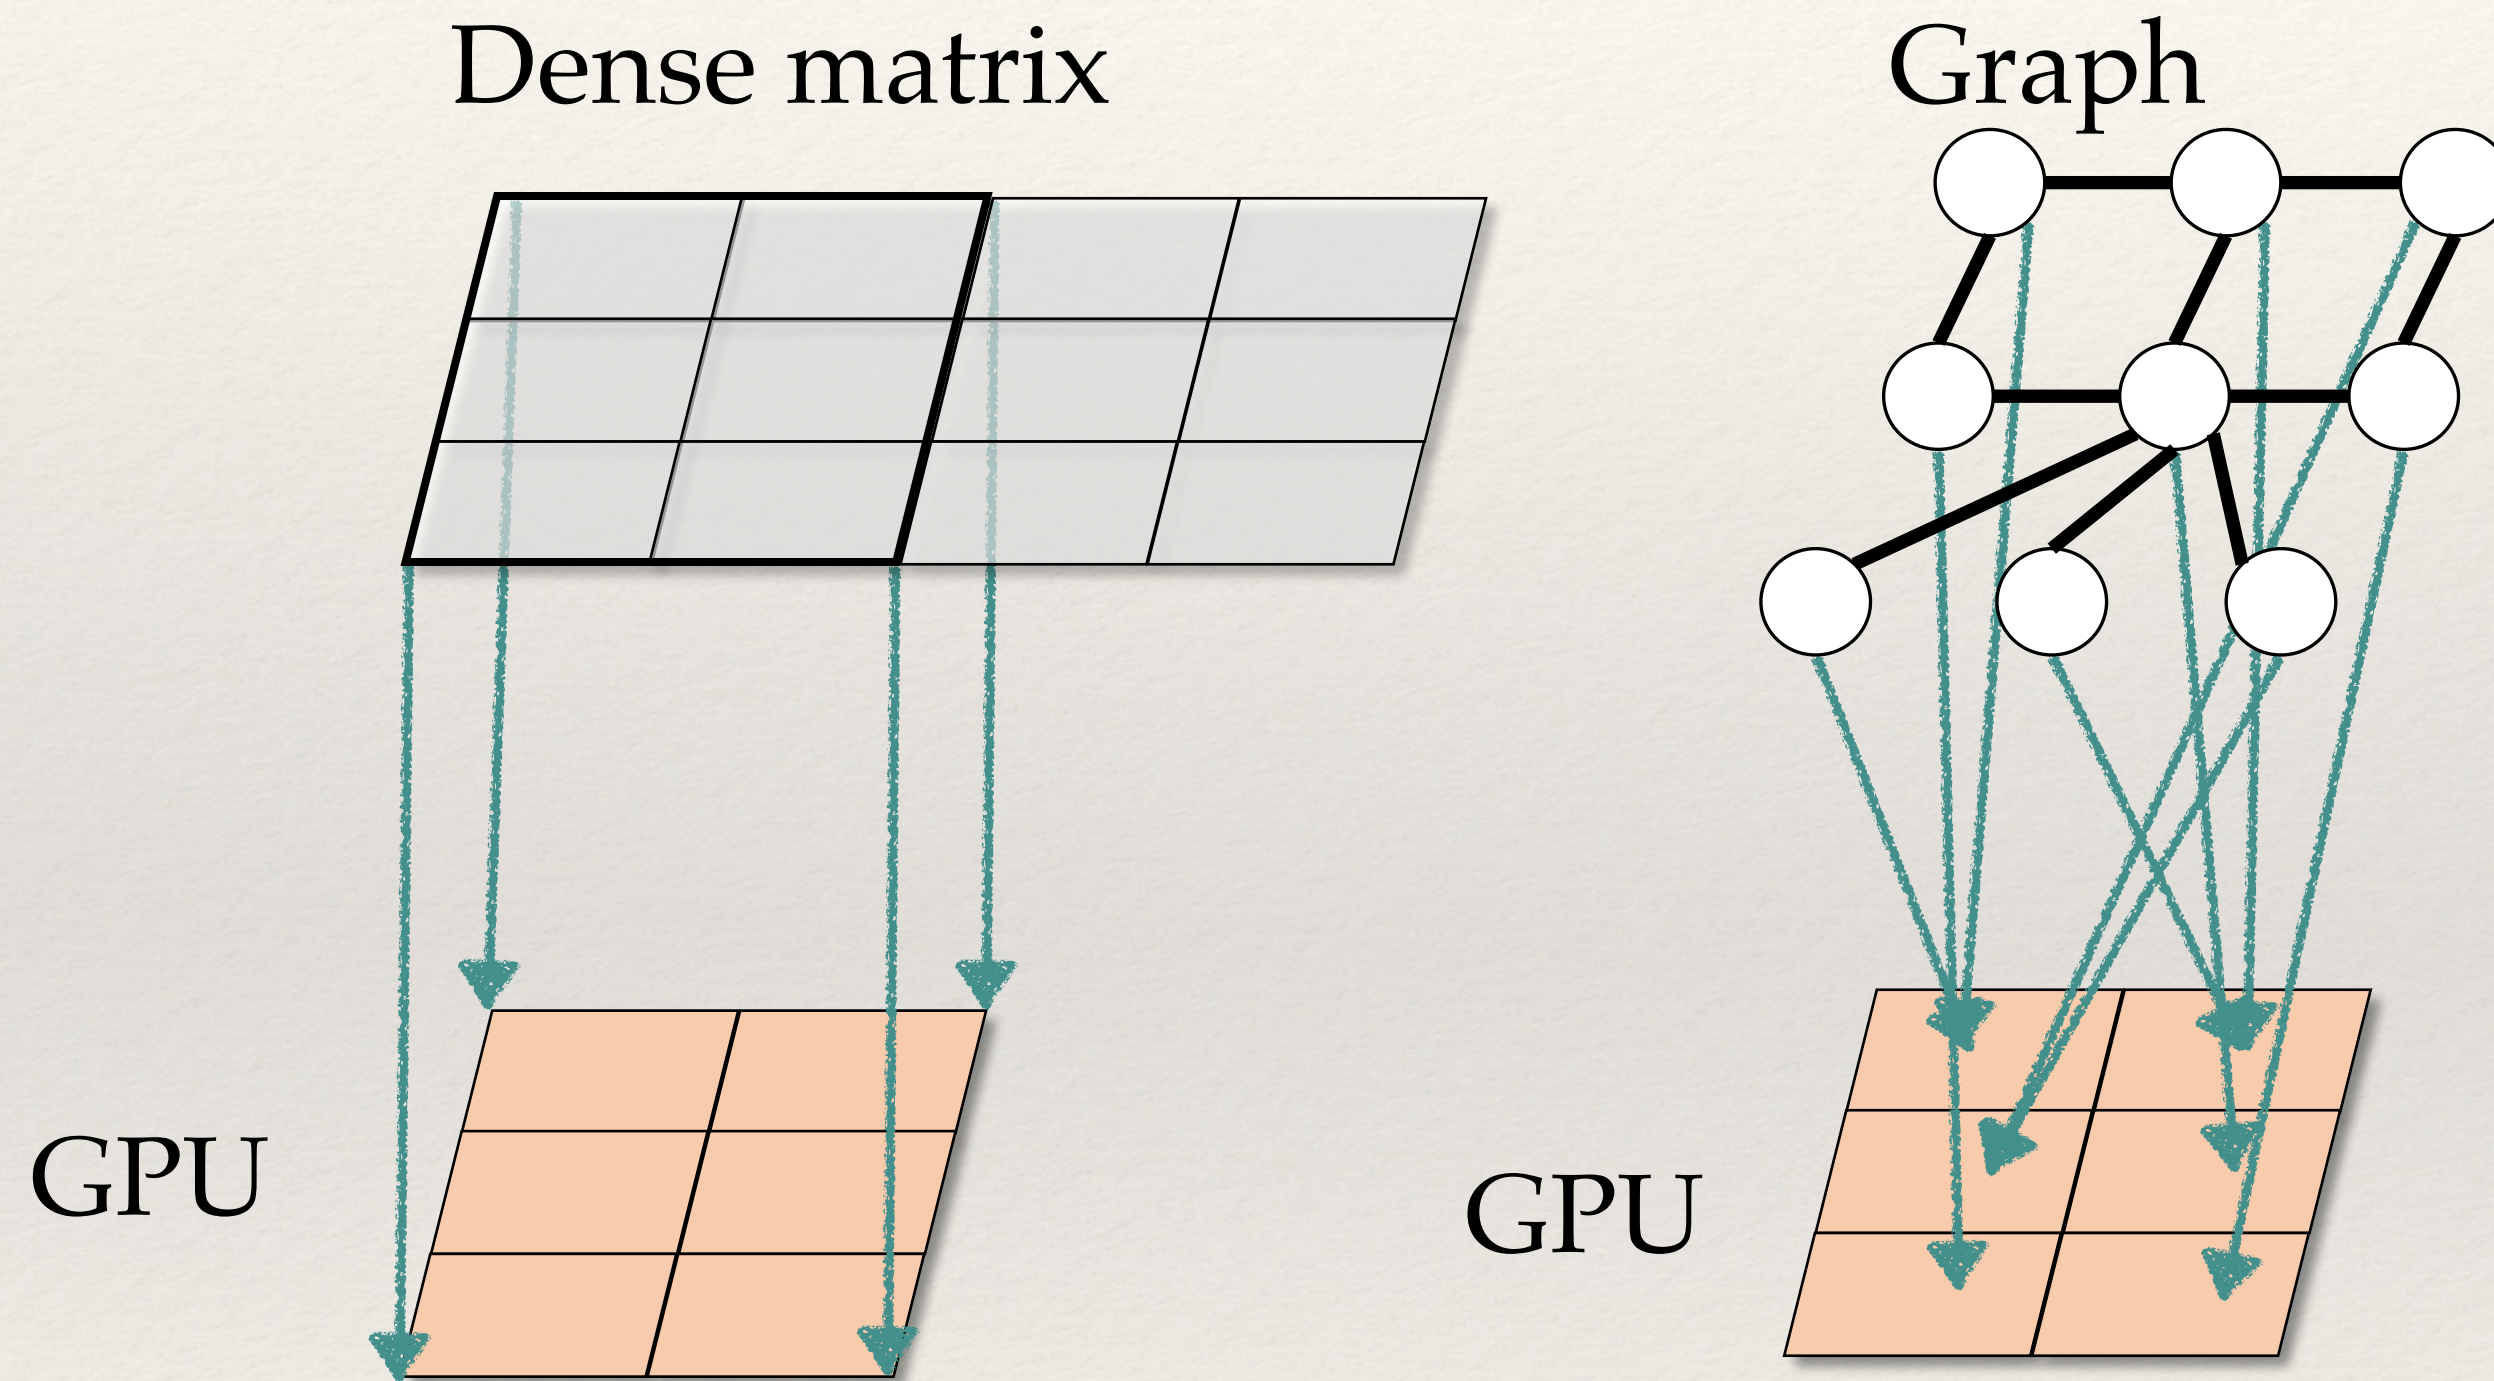


General Purpose GPUs (Tesla V100)



- ❖ **Massive parallelism and high bandwidth**
 - ❖ **22.5/120** TFLOPS and **900** GB/s.
 - ❖ 5,120 cores, supporting ~millions of threads.
- ❖ **Thread granularity**
 - ❖ **Warp**: consecutive 32 threads, executed in **SIMD**.
 - ❖ **Lack of inter-SMX communication support.**
- ❖ **Memory access pattern**
 - ❖ 900 GB/s: **consecutive** threads access **adjacent** data.
 - ❖ **Random / stride access is ~10x slower.**

Mapping Graph Computing on GPUs



❖ GPU: **SIMD**


❖ Graph: Comple**X**

❖ **SIMD-X** bridges the gap!

SIMD-X Approaches

- ❖ SIMD-X — up to 10x faster than state-of-the-art “Gunrock” [PPoPP ’16]
 - ❖ Data-parallel ACC programming abstraction
 - ❖ Just-in-time task management
 - ❖ Push-pull based kernel fusion

SIMD-X: Programming and Processing of Graph Algorithms on GPUs



Hang Liu

H. Howie Huang

University of Massachusetts Lowell

George Washington University