

# OPENCL

## Episode 5 - Questions and Answers

David W. Gohara, Ph.D.

Center for Computational Biology

Washington University School of Medicine, St. Louis

email: [sdg0919@gmail.com](mailto:sdg0919@gmail.com)

# THANK YOU




# Q&A

- GPU Layout
- Bank Conflicts



# Q&A

- 10 TPCs
- 30 SMs
- 8 SPs (cores or FPU)/SM
- 240 cores/FPU/SPs
- 30 DPU
- 60 SFUs




# Q&A

- Bank Conflicts - Local Memory
  - 16 KB
  - 16 banks
  - Each entry is 32 bits wide
  - 4096 entries total
- Successive 32-bit words assigned to successive banks
- Whether reading or writing two or more simultaneous access to same bank results in serialization



# MATRIX TRANSPOSE





# MATRIX TRANSPOSE

Read from global mem

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1,1 | 1,2 | 1,3 | 1,4 | 2,1 | 2,2 | 2,3 | 2,4 | 3,1 | 3,2 | 3,3 | 3,4 | 4,1 | 4,2 | 4,3 | 4,4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

Write to shared mem

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1,1 | 1,2 | 1,3 | 1,4 | 2,1 | 2,2 | 2,3 | 2,4 | 3,1 | 3,2 | 3,3 | 3,4 | 4,1 | 4,2 | 4,3 | 4,4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

Once in shared memory any thread (work item) in a thread block  
(work group) can read the data fast

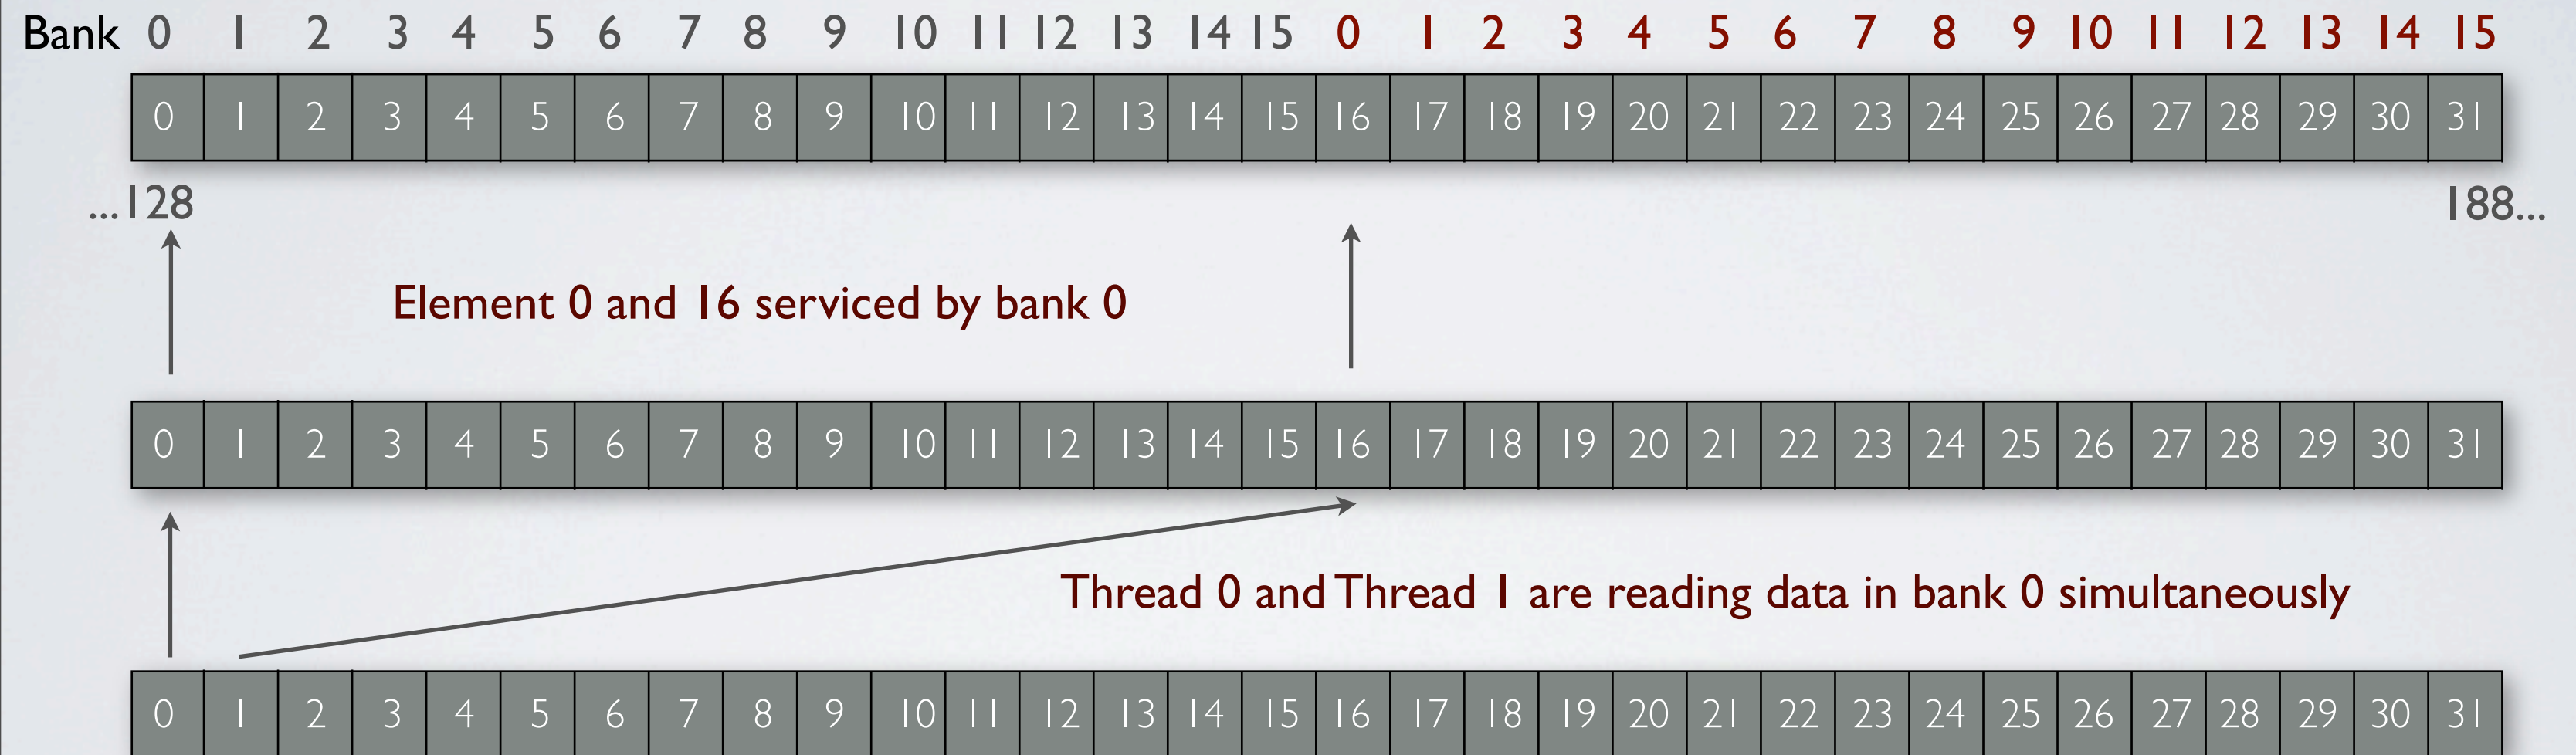
Read data transposed

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1,1 | 2,1 | 3,1 | 4,1 | 1,2 | 2,2 | 3,2 | 4,2 | 1,3 | 2,3 | 3,3 | 4,3 | 1,4 | 2,4 | 3,4 | 4,4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

Write to global mem

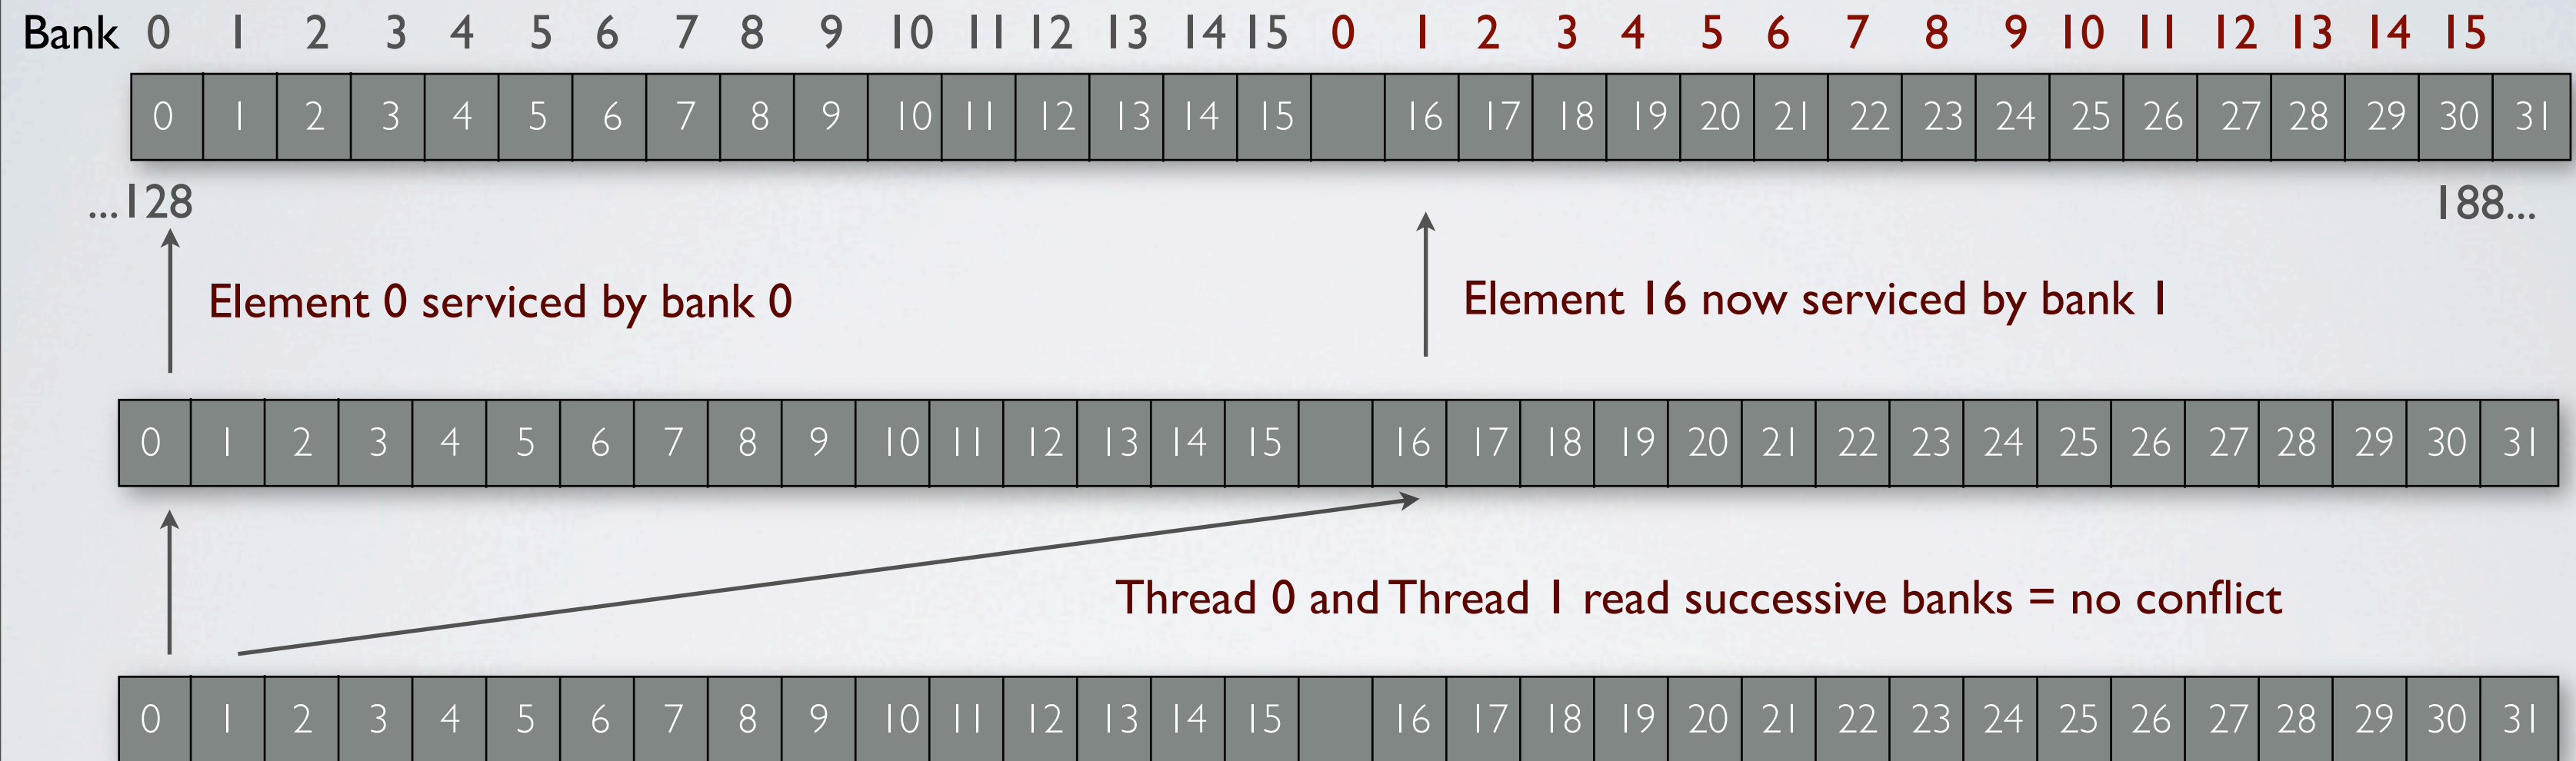
|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1,1 | 2,1 | 3,1 | 4,1 | 1,2 | 2,2 | 3,2 | 4,2 | 1,3 | 2,3 | 3,3 | 4,3 | 1,4 | 2,4 | 3,4 | 4,4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

# MATRIX TRANSPOSE





# MATRIX TRANSPOSE



# MORE INFORMATION

- MacResearch.org
  - OpenCL - <http://www.macresearch.org/openc1>
  - Amazon Store - <http://astore.amazon.com/macreseorg-20>
- NVIDIA Online Seminars
  - [http://developer.nvidia.com/object/gpu\\_computing\\_online.html](http://developer.nvidia.com/object/gpu_computing_online.html)