

# GPU Computing

Weijie Zhao

02/20/2024

# HW2: ANN on DAG

For each test case, at least 50% queries should be correctly answered

Input:

V D E L K A B C M Q

K lines:  $X[i]$  ( $K \leq V * D$ ,  $0 \leq X[i] < M$ )

E lines:  $u[j]$   $v[j]$  ( $u[j] < v[j]$ )

Q lines: start\_point max\_hop  $q_0$   $q_1$   $q_2 \dots q_{D-1}$

For  $i = K$  to  $V * D - 1$ :  $X[i] = (A * X[i - 1] + B * X[i - 2] + C) \% M$  (potential overflow)

Output:

Q lines: vertex id with the smallest L2 distance to q within max\_hop from start\_point. Ties are broken by ids.

V #vertices

$V \leq 10^5$

$Q \leq 10^4$

max\_hop  $\leq 10$

D #dimensions

$D \leq 10^3$

$M \leq 10^2$

E #edges

$E \leq 10^6$

A, B, C non-negative 32-bit int

L maximum out-degree

$L \leq 63$

Due: 2/26/2024 11:59pm EST

Correct GPU Solution will get 5 pts bonus

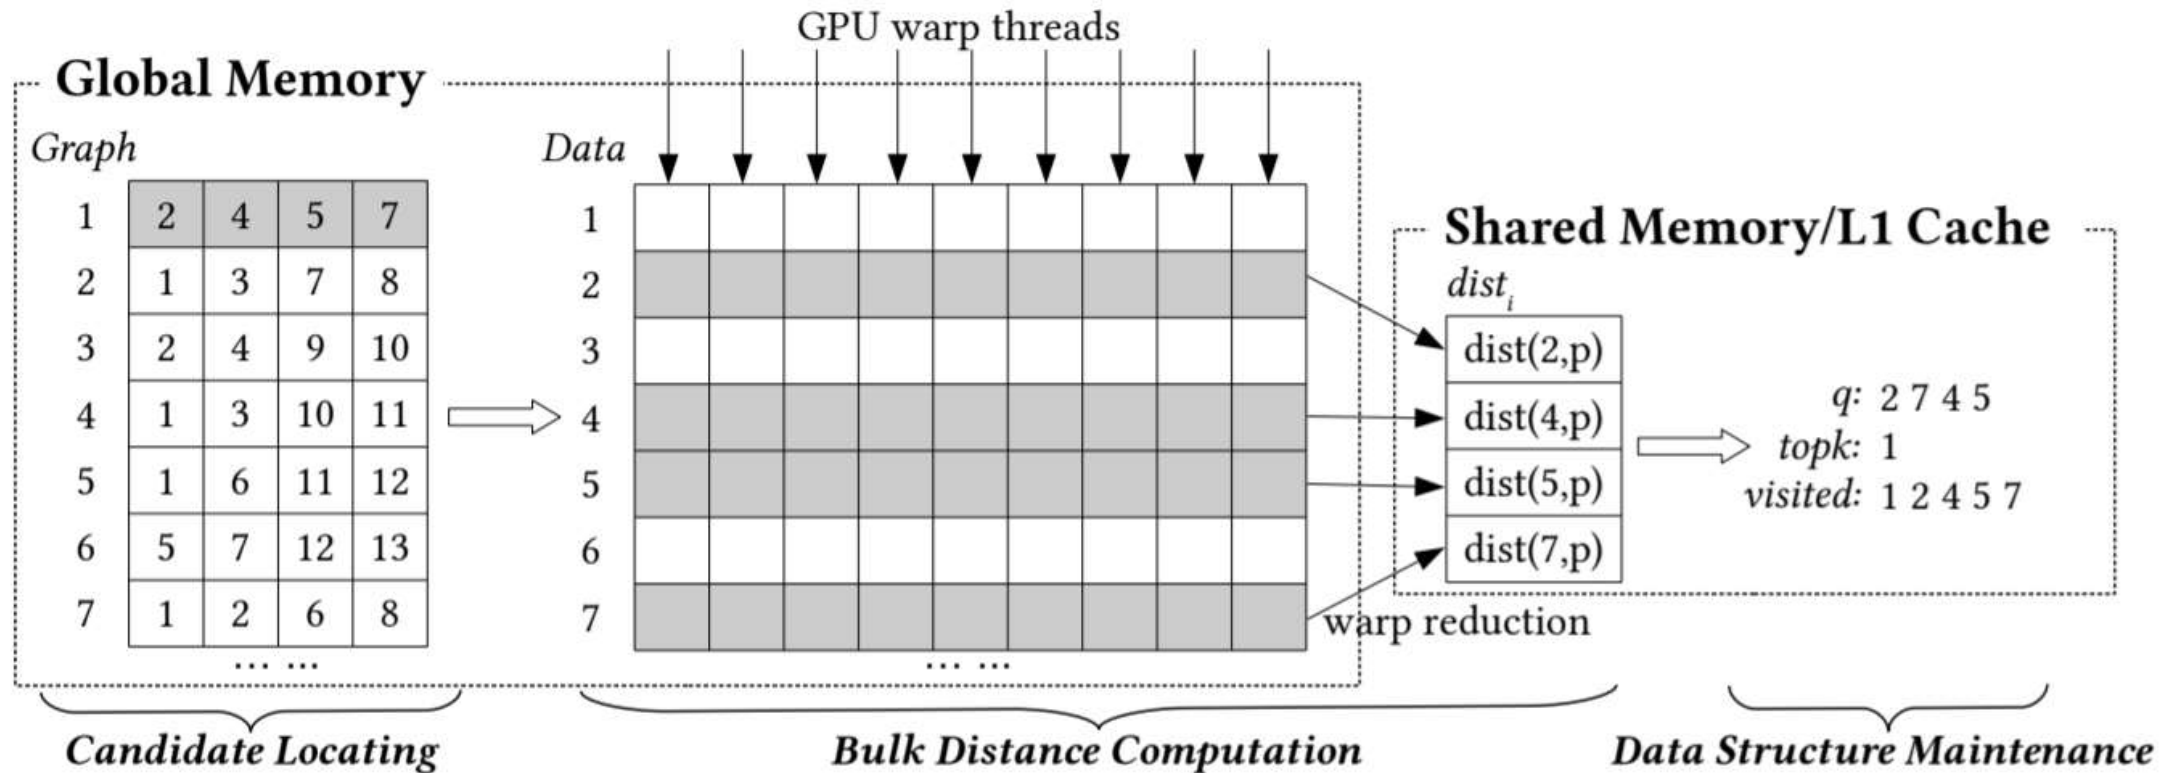
# Test Environment

- [granger.cs.rit.edu](http://granger.cs.rit.edu)
- [weasley.cs.rit.edu](http://weasley.cs.rit.edu)
- [lovegood.cs.rit.edu](http://lovegood.cs.rit.edu)
  
- 8 CPU threads and 1 GPU
- Time limit:
  - 120 seconds compilation time
  - **60 seconds** for each test case

```
52     edges = new int[V * (L + 1)];
53     for(int i = 0; i < V; ++i){
54         edges[i * (L + 1)] = 0;
55     }
56     for(int i = 0; i < E; ++i){
57         int u, v;
58         fscanf(fin, "%d%d", &u, &v);
59         int degree = edges[u * (L + 1)];
60         edges[u * (L + 1) + degree + 1] = v;
61         ++edges[u * (L + 1)];
62     }
```

```
17 int nearest_id(int start_point, int max_hop, int* query_data){
18     std::queue<std::pair<int, int>> q;
19     q.push(std::make_pair(start_point, 0));
20     int min_d = std::numeric_limits<int>::max();
21     int min_id = -1;
22     while(!q.empty()){
23         auto now = q.front();
24         q.pop();
25         int id = now.first;
26         int hop = now.second;
27         int d = squared_l2_dist(X + id * D, query_data, D);
28         if((d < min_d) || (d == min_d && id < min_id)){
29             min_d = d;
30             min_id = id;
31         }
32         if(hop + 1 <= max_hop){
33             int degree = edges[id * (L + 1)];
34             for(int i = 1; i <= degree; ++i){
35                 int v = edges[id * (L + 1) + i];
36                 q.push(std::make_pair(v, hop + 1));
37             }
38         }
39     }
40     return min_id;
}
```

# GPU Graph Searching Example



# Tensor Computing in Deep Learning

Weijie Zhao

02/20/2024

- Scalar
- Vector
- Matrix
- Tensor
  - Rank
  - Dimension

# Tensor Computing in Deep Learning

Weijie Zhao

02/20/2024



# Matrix

# ~~Tensor~~ Computing in Deep Learning

- Scalar
- Vector
- Matrix
- Tensor
  - Rank
  - Dimension

Weijie Zhao

02/20/2024

# Matrix ~~Tensor~~ Computing in Deep Learning

- Scalar
- Vector
- Matrix
- Tensor
  - Rank
  - Dimension

Weijie Zhao

02/20/2024

- Matrix multiplication
- Non-linear activation
- Gradient descent

# Matrix ~~Tensor~~ Computing in Deep Learning

- Scalar
- Vector
- Matrix
- Tensor
  - Rank
  - Dimension

Weijie Zhao

02/20/2024

•Matrix multiplication

•Non-linear activation

~~•Gradient descent~~

Graduate student descent

# Tensor Operations

- Element-wise add
- Element-wise plus
- Element-wise division
- Hadamard product
- Matrix multiplication
- Batched matrix multiplication
- More linear algebra operations...
- Collect, Scatter, Reduce...

# Libraries

- Numpy
- Blas
- cuBlas
- cuSparse
- MKL
- TensorFlow
- PyTorch
- MXNet
- ...

# Lazy Evaluation and Code Generation

$c = a + b$

$d = c * 2$

for  $i = 1$  to  $n$  do

$c[i] = a[i] + b[i]$

for  $i = 1$  to  $n$  do

$d[i] = c[i] * 2$

for  $i = 1$  to  $n$  do

$d[i] = (a[i] + b[i]) * 2$