

# GPU Accelerated Vessel Segmentation Using Laplacian Eigenmaps


Lin Cheng, Hyunsu Cho and Peter A. Yoon

Trinity College

# Problem

## **Image segmentation**

Partition pictures of vessels into **segments**




# Laplacian eigenmap [1]

Local info embedded in **high dimensional space**

**Project** local info onto low-dimensional plane

**Optimize** the projection to preserve essential characteristics

**Cluster** the projected data points into segments



[1] Tziakos, Laskaris, and Fotopoulos

# Segmentation process





# Build graph of local info

Store the resulting graph in a **weight matrix**

**Edges** reflect variations among different regions (global variation)

*Patch grid*

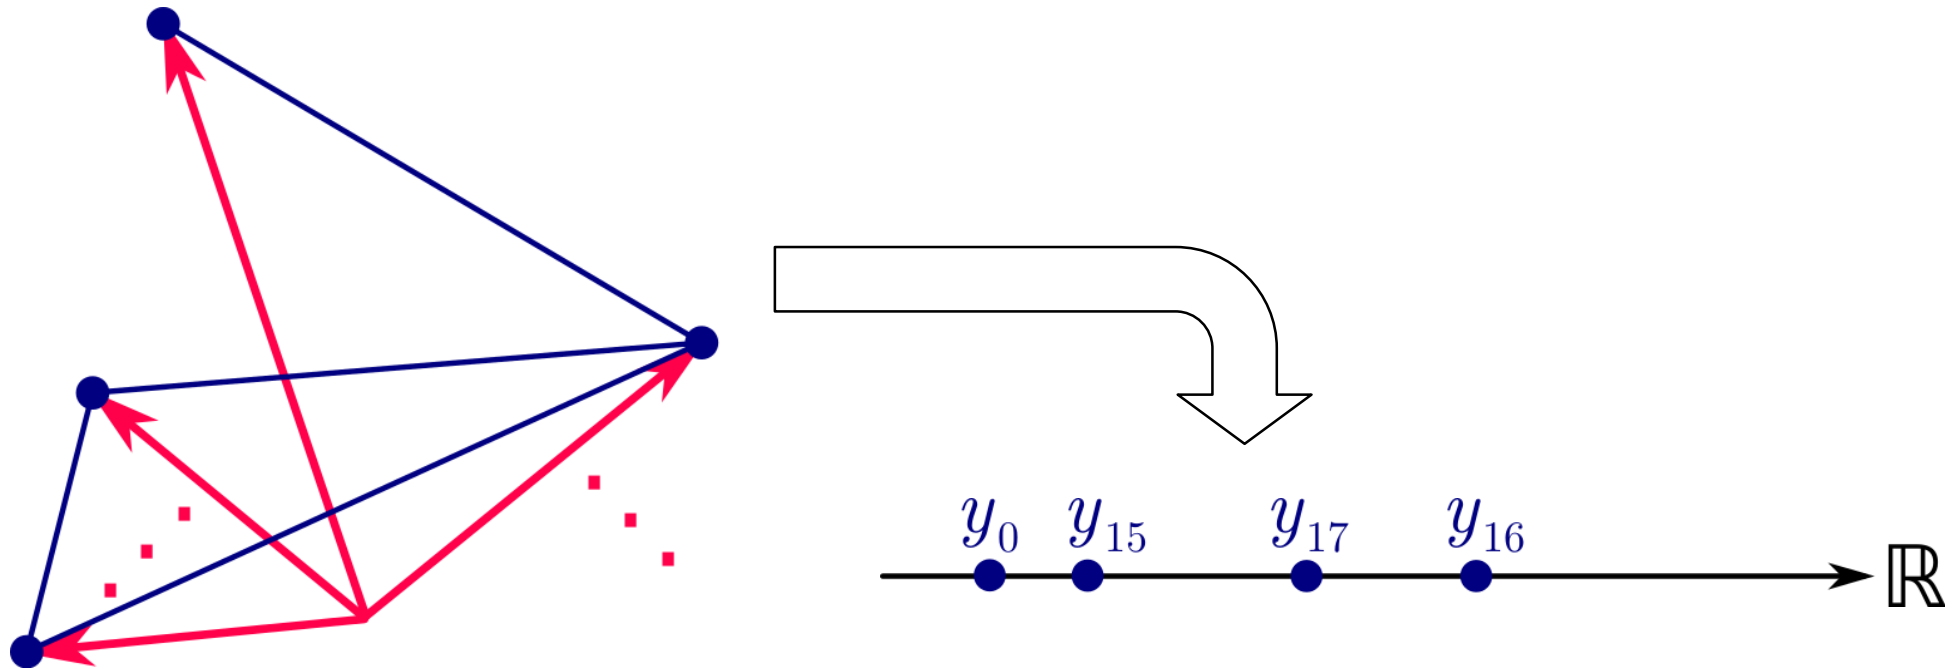
0	1	2	3	4	5
6	7	8	9	10	11
12	13	14	15	16	17
18	19	20	21	22	23



# Apply Laplace operator

Form **Laplacian matrix**  $L = I - D^{1/2}WD^{1/2}$  encoding the Laplace operator.


The operator formulates an **optimization problem**:  
Projections of well-connected nodes should also be tightly clustered.



# Solve optimization problem

Solutions to eigenvalue problem  $L\mathbf{y} = \lambda\mathbf{y}$  are optimal solutions

If solutions are good, we can detect clusters




# Characteristics of GPUs

**Massively parallel** – lots of small cores (workers)

Good for high-throughput, compute-bound tasks

**Separate memory** space from main memory





# Strategy: Reduce memory footprint

On-GPU memory is limited

Reduce memory usage and we can pack in more work into GPU


# Strategy: Reduce memory footprint

Weight matrix generation:

Do not store intermediate results


More entries can be calculated in parallel; **10x faster**

0	1	2	3	0	1	2	3



(0,0)	(0,1)	(0,2)	(0,3)
(1,0)	(1,1)	(1,2)	(1,3)
(2,0)	(2,1)	(2,2)	(2,3)
(3,0)	(3,1)	(3,2)	(3,3)

# Worker allocation




# Strategy: use Lanczos method

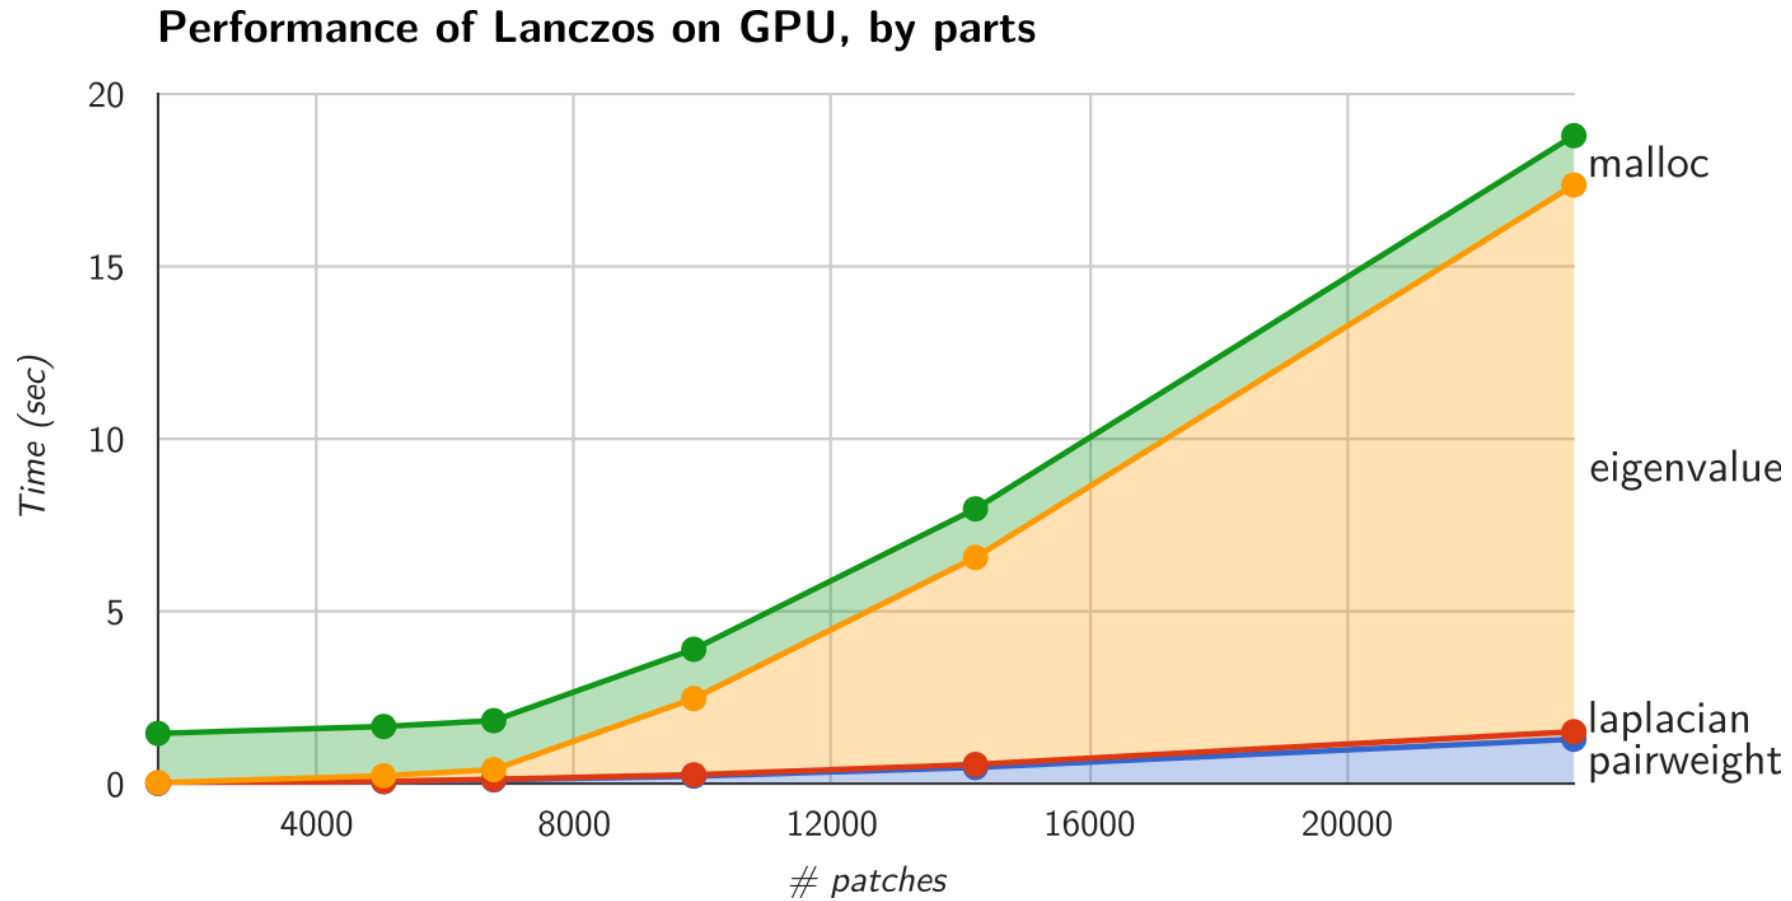
We need only a few smallest eigenvalues of  $L$

**Lanczos method** iteratively solve for the eigenvalues needed

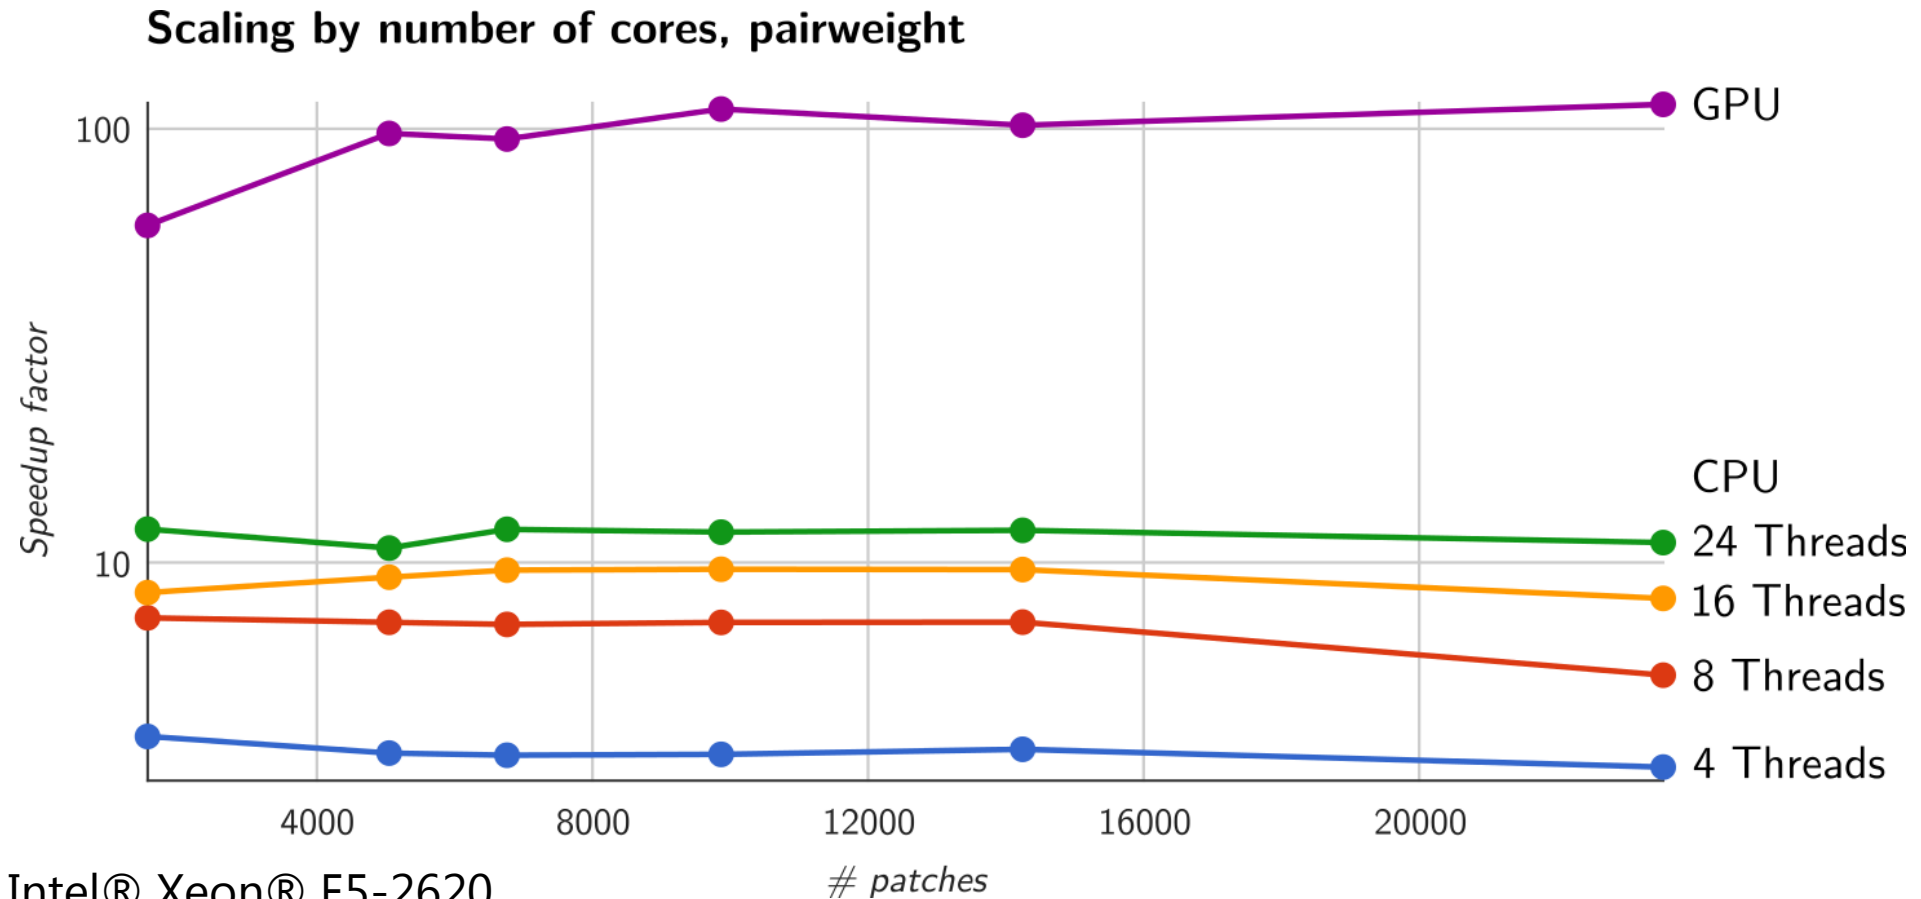
Takes 1/28 time of conventional method



# Performance



# Performance: vs. multicore CPUs



CPU: two Intel® Xeon® E5-2620

GPU: one Nvidia Tesla® K20c

# Acknowledgement

Trinity College, Student Research Program

Nvidia Corporation, CUDA Teaching Center Program