

Computing the Pseudoinverse of a Graph's Laplacian using GPUs

Nishant Saurabh

Vrije Universiteit, Amsterdam.
nishants.prmitr7@gmail.com



Dr. Ana Lucia Varbanescu

University of Amsterdam, Amsterdam.
a.l.varbanescu@uva.nl



Dr. Gyan Ranjan

Symantec, CA, USA.
gyan_ranjan@symantec.com




**Workshop on Large-Scale Parallel Processing
IEEE International Parallel and Distributed
Processing Symposium**

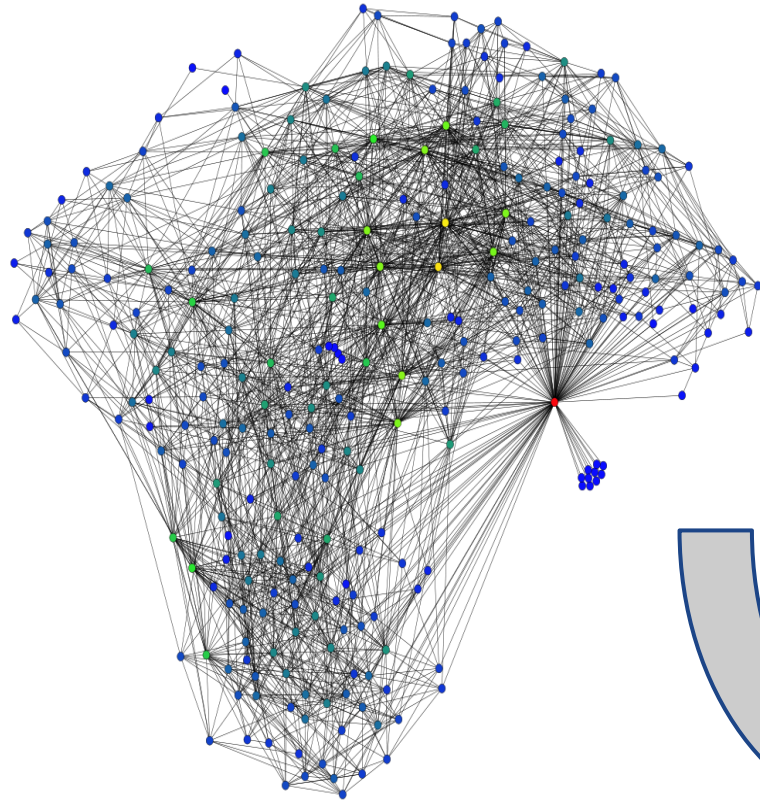


May 25th - 29th, 2015

Graphs are Everywhere !




Graphs are Everywhere !




Small World Characteristics

- Large Scale
- Local World
- Degree Distribution
- Sparse



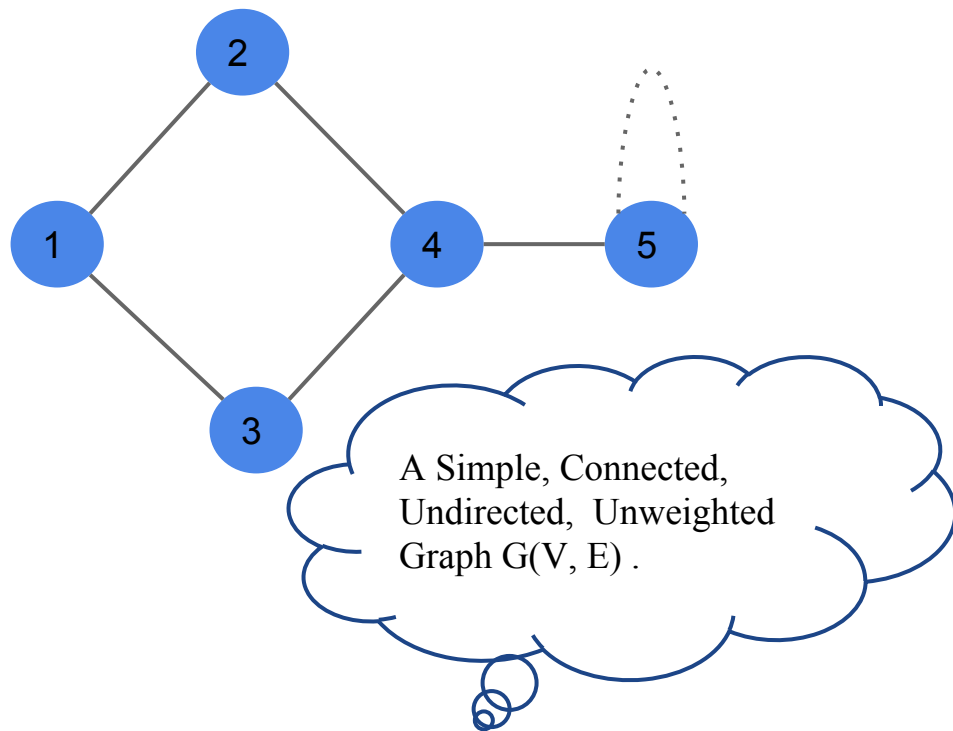
Complex Networks : Define Real World Graphs



Why ?

- Topological Characteristics
- Behavioral Predictability


Graph Formalization & Background



$n = |V|$ is the order of the graph, i.e.,
the number of vertices.

$m = |E|$ is the size of the graph, i.e.,
the number of edges.

Graph Formalization & Background




A is Adjacency Matrix , D is Degree Matrix.

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

$$D = \begin{pmatrix} 2 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Graph Formalization & Background



Degree matrix (D) - Adjacency Matrix (A) = Laplacian Matrix (L)

$$\begin{pmatrix} 2 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 2 & 0 & -1 & 0 \\ -1 & 0 & 3 & -1 & -1 \\ 0 & -1 & -1 & 2 & 0 \\ 0 & 0 & -1 & 0 & 1 \end{pmatrix}$$

Inverse of Laplacian matrix and Eigenspace

*The inverse of Laplacian matrix L is L^{-1}
such that :*

$$\mathbf{L} \mathbf{L}^{-1} = \mathbf{I}$$


where \mathbf{I} is Identity matrix and \mathbf{L} is a square matrix.

Inverse of Laplacian matrix and Eigenspace

*The inverse of Laplacian matrix L is L^{-1}
such that :*

$$L L^{-1} = I$$

where I is Identity matrix and L is a square matrix.




Inverse of Laplacian matrix and Eigenspace

The inverse of Laplacian matrix L is L^{-1} such that :

$$L L^{-1} = I$$

where I is Identity matrix and L is a square matrix.



The Eigenspace can be formulated as :

$$L v = \lambda v$$


where L is Laplacian Matrix, v is Eigenvector and λ is Eigenvalue.

Inverse of Laplacian matrix and Eigenspace

The inverse of Laplacian matrix L is L^{-1} such that :

$$L L^{-1} = I$$

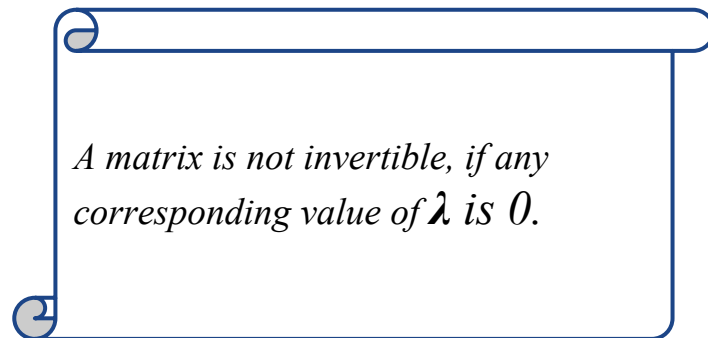
where I is Identity matrix and L is a square matrix.




The Eigenspace can be formulated as :

$$L v = \lambda v$$

where L is Laplacian Matrix, v is Eigenvector and λ is Eigenvalue.




Eigenvalues & Eigenvector



$$L = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 2 & 0 & -1 & 0 \\ -1 & 0 & 3 & -1 & -1 \\ 0 & -1 & -1 & 2 & 0 \\ 0 & 0 & -1 & 0 & 1 \end{pmatrix}$$

Eigenvalues & Eigenvector



$$L = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 2 & 0 & -1 & 0 \\ -1 & 0 & 3 & -1 & -1 \\ 0 & -1 & -1 & 2 & 0 \\ 0 & 0 & -1 & 0 & 1 \end{pmatrix}$$

Non - Invertible !

$$\lambda =$$

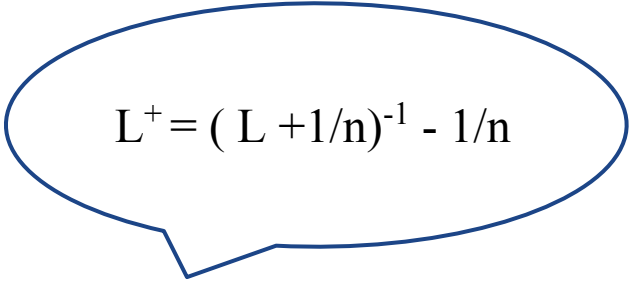
$$\left\{ \boxed{0}, 0.82, 2, 2.68, 4.4812 \right\}$$

Moore Penrose pseudo-inverse

To calculate the inverse for a rank deficient matrix ($L =$ laplacian matrix, of order n):

Moore Penrose pseudo-inverse

To calculate the inverse for a rank deficient matrix ($L = \text{laplacian matrix, of order } n$):


$$L^+ = (L + 1/n)^{-1} - 1/n$$

Moore Penrose pseudo-inverse

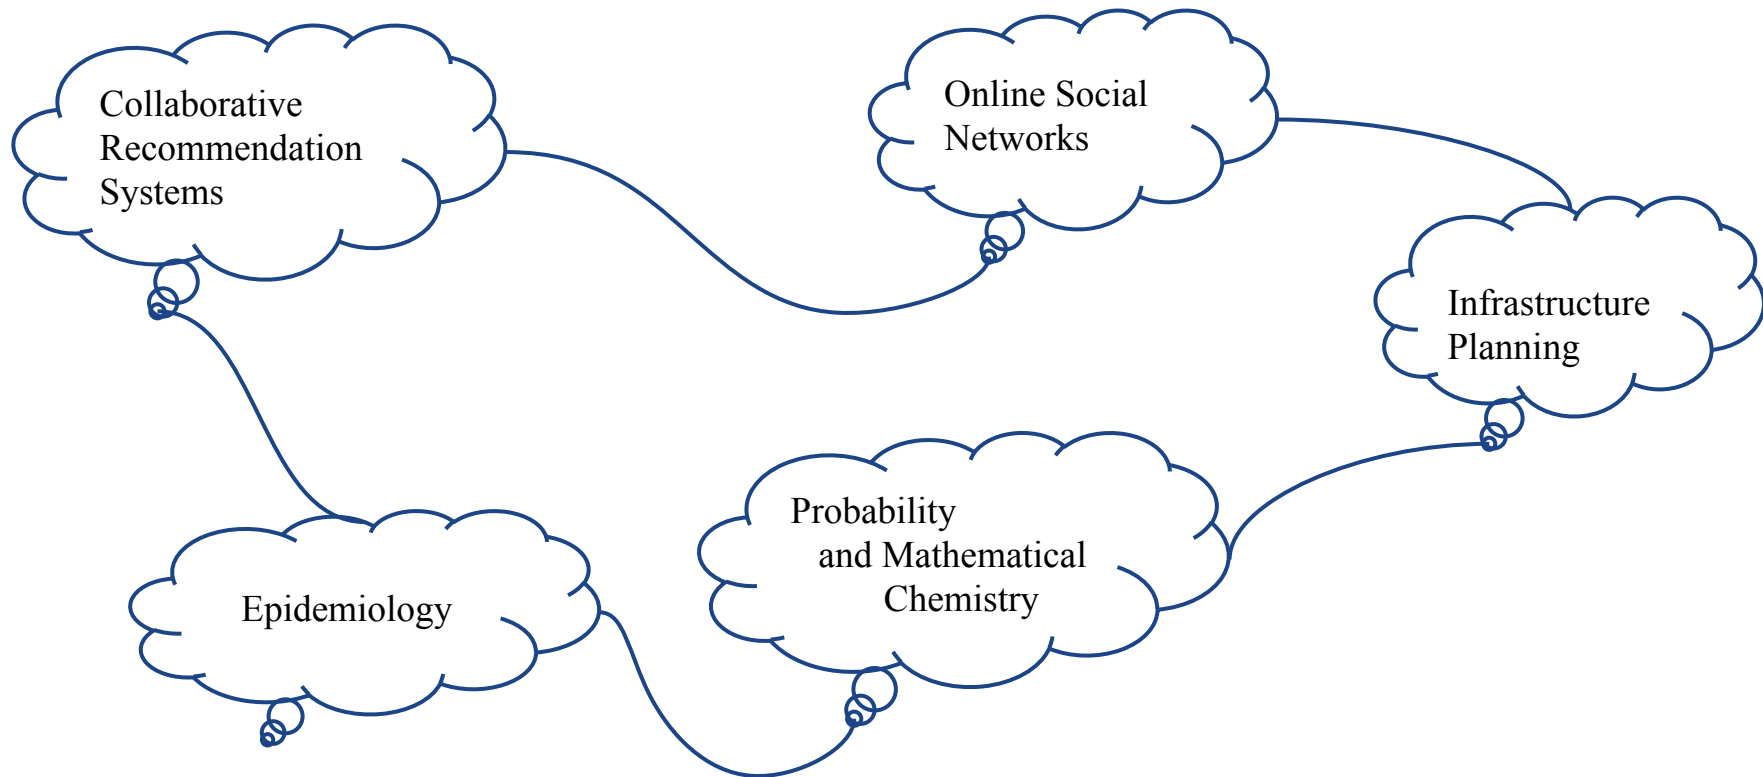
To calculate the inverse for a rank deficient matrix ($L = \text{laplacian matrix, of order } n$):

$$L^+ = (L + 1/n)^{-1} - 1/n$$

$L^+ = \text{pinv}(L)$ in Matlab

$L^+ = \text{numpy.linalg.pinv}(L)$ in Python

Applications to Computation of L^+



The Goal




Algorithm

The Goal



Algorithm

+




GPU
THOUSANDS OF CORES

The Goal

Algorithm

+



GPU
THOUSANDS OF CORES

=

Speedup ?


Divide and Conquer Approach to compute L^+

Divide and Conquer Approach to compute L^+

Three Steps



➤ Partition

Divide : Partition



A simple, connected, unweighted,
undirected Graph G .

Divide : Partition



A simple, connected, unweighted, undirected Graph G .


Connected Subgraph G_1 and G_2
Compute $L^+_{G_1}$ and $L^+_{G_2}$

Divide and Conquer Approach to compute L^+

Three Steps

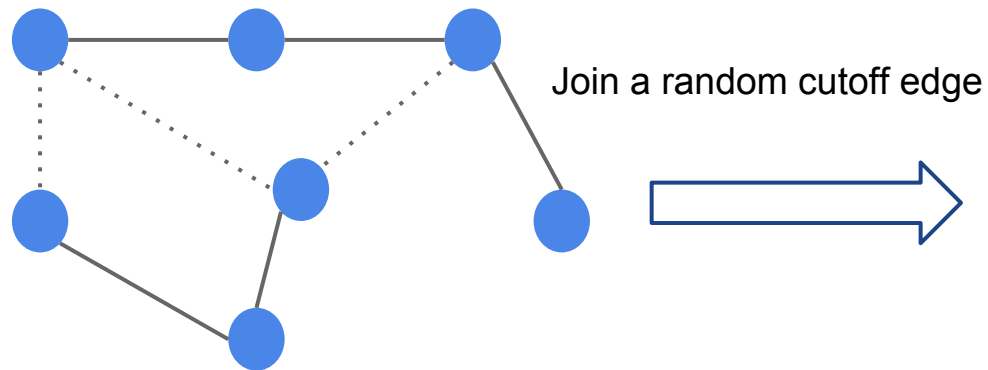
- Partition
- First Join

Conquer : First Join




Dotted lines represent minimized cutoff edges during Partition.

Conquer : First Join



Dotted lines represent minimized cutoff edges during Partition.




We get a new connected Graph G_3
Using $L^+_{G_1}$ and $L^+_{G_2}$, compute $L^+_{G_3}$

Divide and Conquer Approach to compute L^+


Three Steps

- Partition
- First Join
- Edge Firing


Conquer : Edge Firing



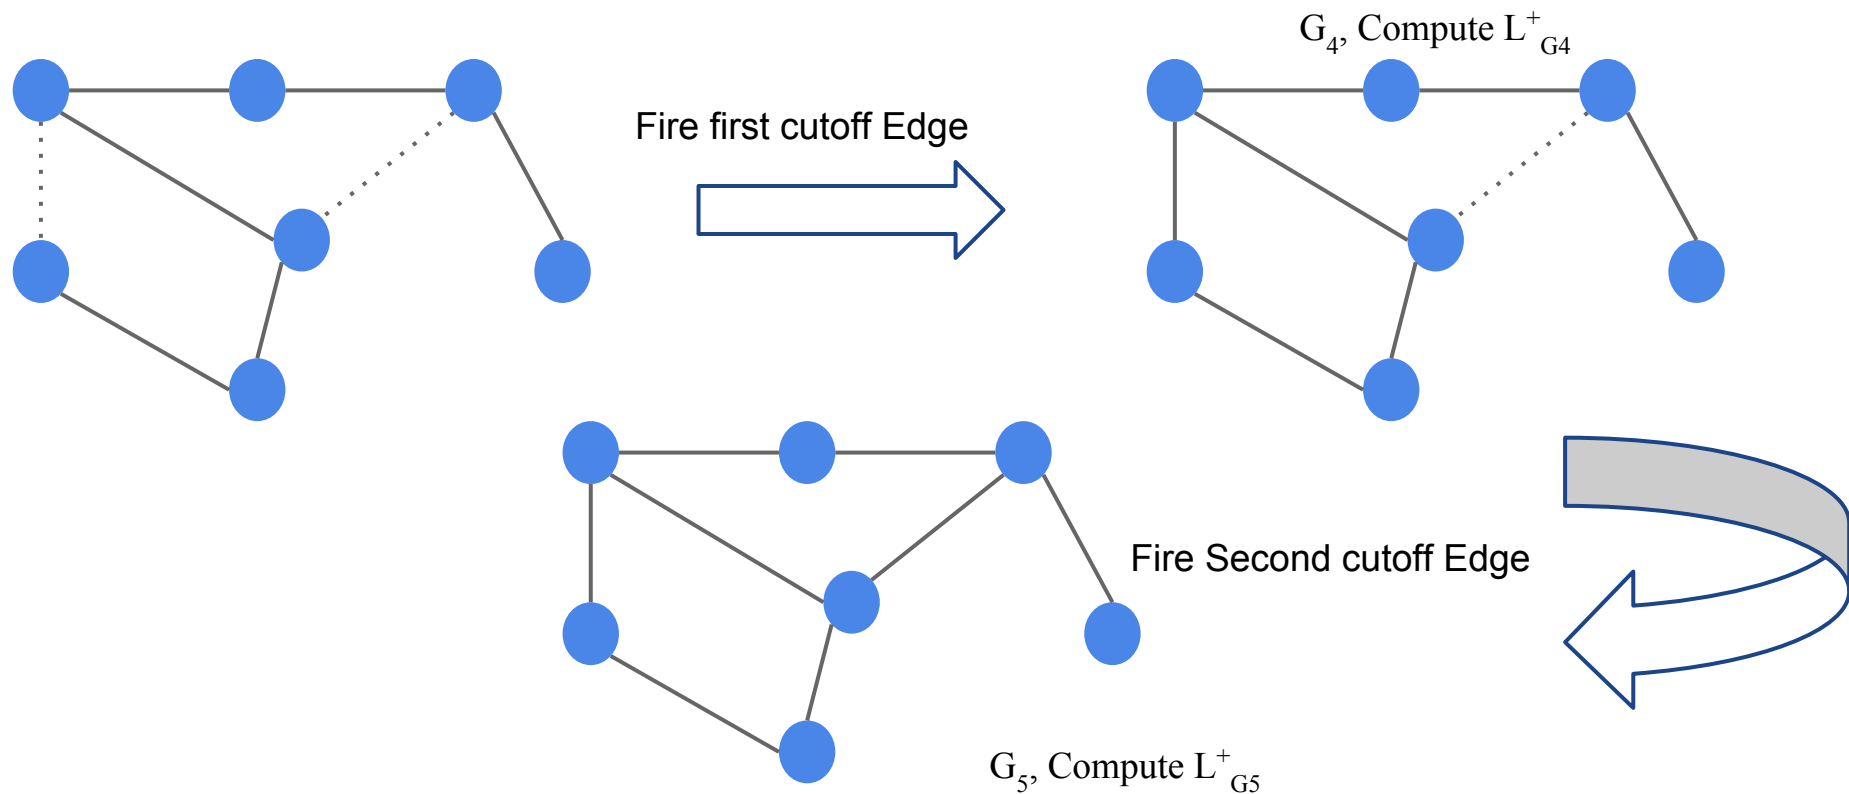
Conquer : Edge Firing



Fire first cutoff Edge




Conquer : Edge Firing



Methodology

- Abstract a simple Partition Approach.
- Use GPU to compute L^+ of the subgraphs.
- Apply element-wise computation on First Join and Edge Firing using GPU.
- Minimize Data Transfer.

Implementation Approach - Representation




A simple, connected, unweighted, undirected Graph G.

$$\begin{pmatrix} 1, & 2, & 1 \\ 2, & 1, & 1 \\ 1, & 5, & 1 \\ 5, & 1, & 1 \\ 1, & 7, & 1 \\ 7, & 1, & 1 \\ 2, & 3, & 1 \\ 3, & 2, & 1 \\ 3, & 6, & 1 \\ 6, & 3, & 1 \\ 3, & 7, & 1 \\ 7, & 3, & 1 \\ 4, & 7, & 1 \\ 7, & 4, & 1 \\ 4, & 5, & 1 \\ 5, & 4, & 1 \end{pmatrix}$$

Sparse Representation

Implementation Approach - Partition




Node	Degree
1	3
3	3
7	3
2	2
4	2
5	2
6	1

A simple, connected, unweighted, undirected Graph G.


Implementation Approach - Partition

Remove edges of node 1




Largest Component has order 6

Remove edges of node 3



Largest Component has order 3 , which is $\leq n / 2$ ($n = 7$). But we wanted two connected components, and there are some isolated ones.

Implementation Approach - Recombine



We get a bipartition, with two simple, connected components

Parallel Implementation - First Join & Edge Firing

1. *CPU Based Parallel Implementation :*

- ★ Parallelisation using pthreads
- ★ *inverse* $(L + 1/n) - 1/n$
- ★ *dgetri.f* Blas routine
- ★ First Join and Edge Firing
- ★ 4 threads generated

Parallel Implementation - First Join & Edge Firing

1. *CPU Based Parallel Implementation :*

- ★ Parallelisation using pthreads
- ★ $inverse(L + 1/n) - 1/n$
- ★ *dgetri.f* Blas routine
- ★ First Join and Edge Firing
- ★ 4 threads generated

2. *Matlab based GPU Implementation :*

- ★ Parallel Computing Toolbox
- ★ *inv()* , GPU enabled function
- ★ First Join and Edge Firing
- ★ *bsxfun()*

Parallel Implementation - First Join & Edge Firing

1. *CPU Based Parallel Implementation :*

- ★ Parallelisation using pthreads
- ★ *inverse* ($L + 1/n$) - $1/n$
- ★ *dgetri.f* Blas routine
- ★ First Join and Edge Firing
- ★ 4 threads generated

2. *Matlab based GPU Implementation :*

- ★ Parallel Computing Toolbox
- ★ *inv()* , GPU enabled function
- ★ First Join and Edge Firing
- ★ *bsxfun()*

3. *CUDA based GPU Implementation :*

- ★ Thrust Library
- ★ *inverse* cuBlas library routine
- ★ First Join - Three Device Kernels
- ★ Edge Firing - Single Kernel
- ★ 256 threads generated per block

Parallel Implementation - First Join & Edge Firing

1. *CPU Based Parallel Implementation :*

- ★ Parallelisation using pthreads
- ★ $inverse(L + 1/n) - 1/n$
- ★ *dgetri.f* Blas routine
- ★ First Join and Edge Firing
- ★ 4 threads generated

2. *Matlab based GPU Implementation :*

- ★ Parallel Computing Toolbox
- ★ *inv()* , GPU enabled function
- ★ First Join and Edge Firing
- ★ *bsxfun()*

3. *CUDA based GPU Implementation :*


- ★ Thrust Library
- ★ $inverse$ cuBlas library routine
- ★ First Join - Three Device Kernels
- ★ Edge Firing - Single Kernel
- ★ 256 threads generated per block

4. *cuBlas Implementation -Baselining the performance :*

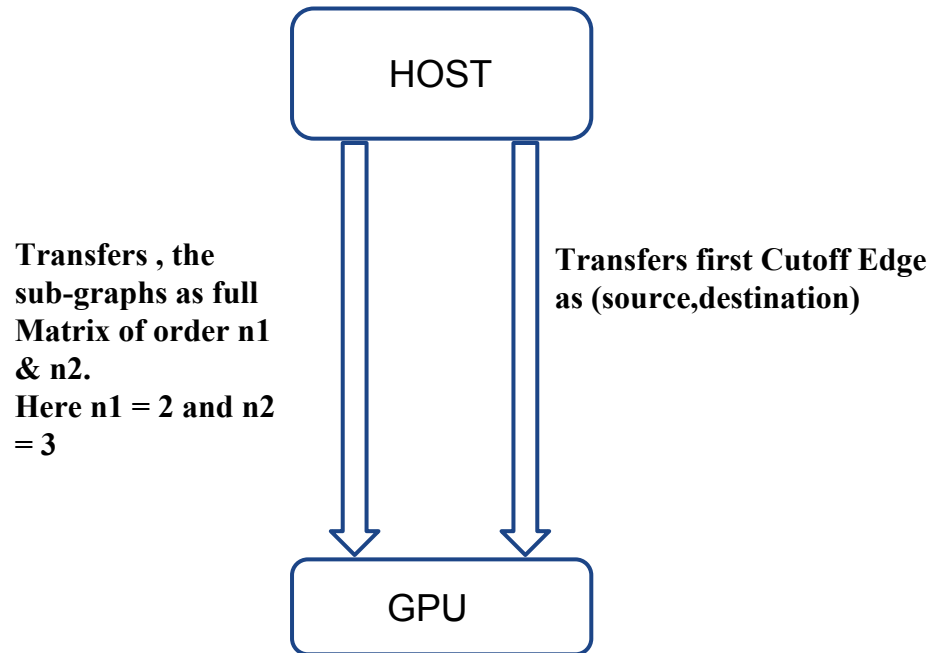
- ★ $inverse(L + 1/n) - 1/n$
- ★ *cublas<t>getriBatched* routine
- ★ *cublas<t>getrfBatched* routine

Implementation - First Join Using GPU

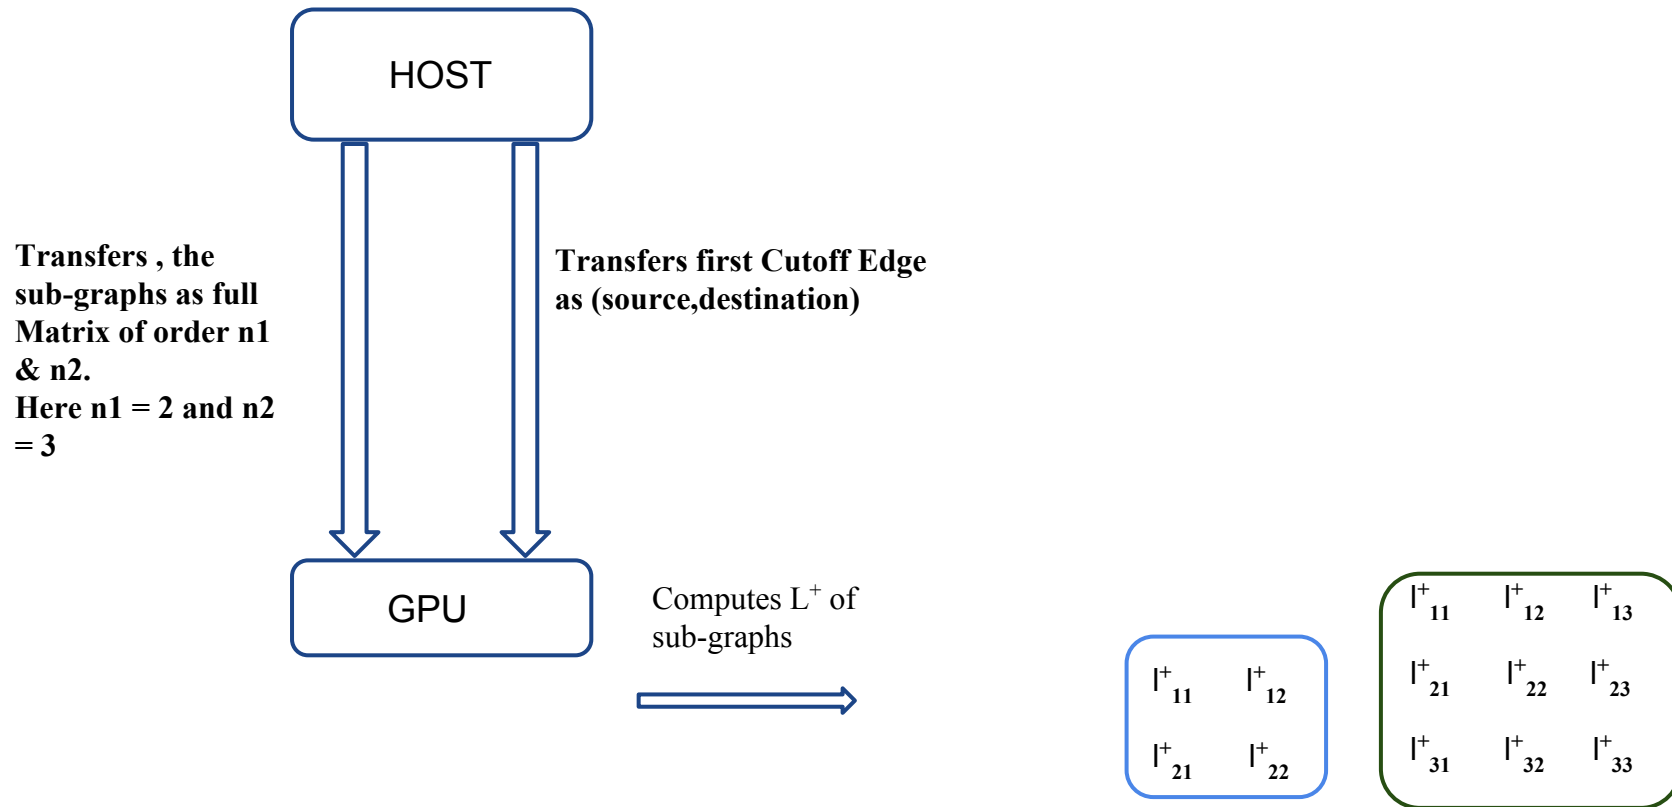
**Transfers , the
sub-graphs as full
Matrix of order n_1
& n_2 .
Here $n_1 = 2$ and n_2
 $= 3$**



Implementation - First Join Using GPU




Implementation - First Join Using GPU



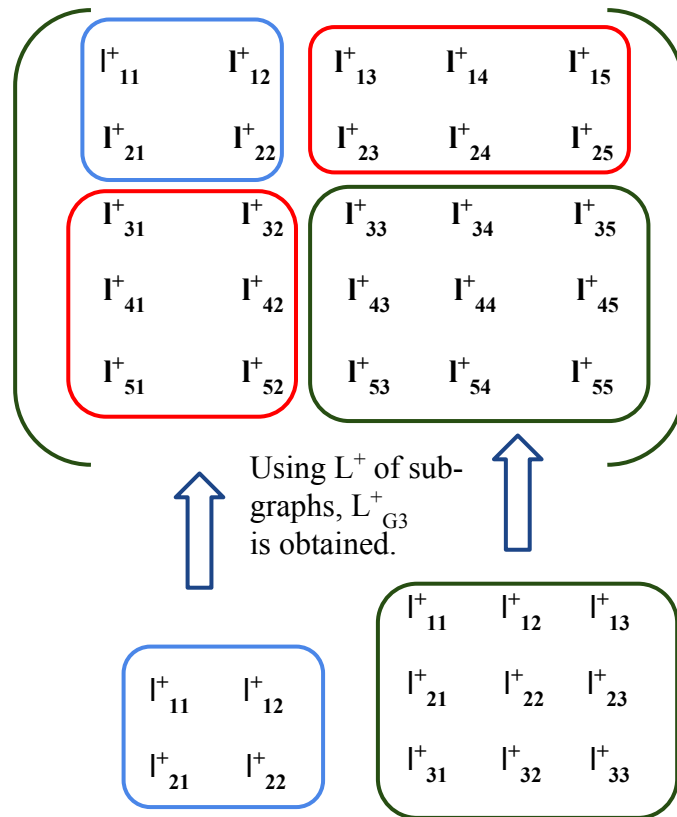
Implementation - First Join Using GPU

Transfers, the sub-graphs as full Matrix of order n_1 & n_2 .
Here $n_1 = 2$ and $n_2 = 3$




Transfers first Cutoff Edge as (source,destination)

Computes L^+ of sub-graphs




Implementation - Edge Firing Using GPU

Remaining cutoff edges transferred as (source, destination), N number of times. N is the number of edges fired after first Join.



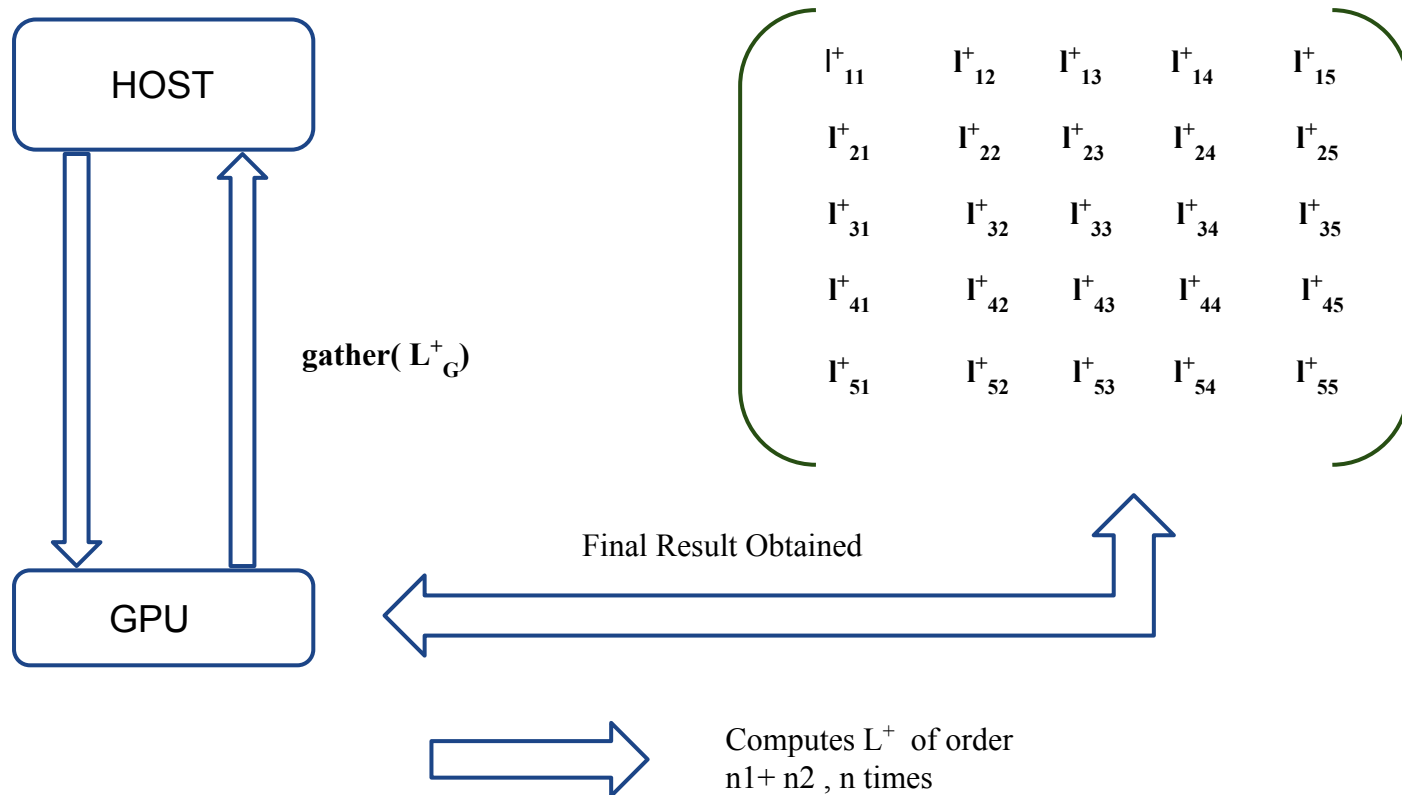
I_{11}^+	I_{12}^+	I_{13}^+	I_{14}^+	I_{15}^+
I_{21}^+	I_{22}^+	I_{23}^+	I_{24}^+	I_{25}^+
I_{31}^+	I_{32}^+	I_{33}^+	I_{34}^+	I_{35}^+
I_{41}^+	I_{42}^+	I_{43}^+	I_{44}^+	I_{45}^+
I_{51}^+	I_{52}^+	I_{53}^+	I_{54}^+	I_{55}^+



Computes L^+ of order $n_1 + n_2$, n times

Implementation - Edge Firing Using GPU

Remaining cutoff edges transferred as (source, destination), N number of times. N is the number of edges fired after first Join.




Hardware Platform :

- ★ NVIDIA Tesla K20m GPU from DAS4 , a six-cluster wide-area distributed system designed and used by 5 research institutions in The Netherlands.
- ★ 5GB of GPU global memory, Memory bandwidth of 208 GB/sec, and a peak performance of 3520 GFlops (single precision).
- ★ CUDA 5.5 , Matlab R2014a
- ★ CPU experiments (sequential and parallel) - performed on DAS4 computing node, using dual-quad-core 2.4 GHz CPU configuration and 24GB memory.

Experiments & Results

Execution Time in milliseconds: Comparison between ParallelCPU, cuBlas, Cuda and Matlab

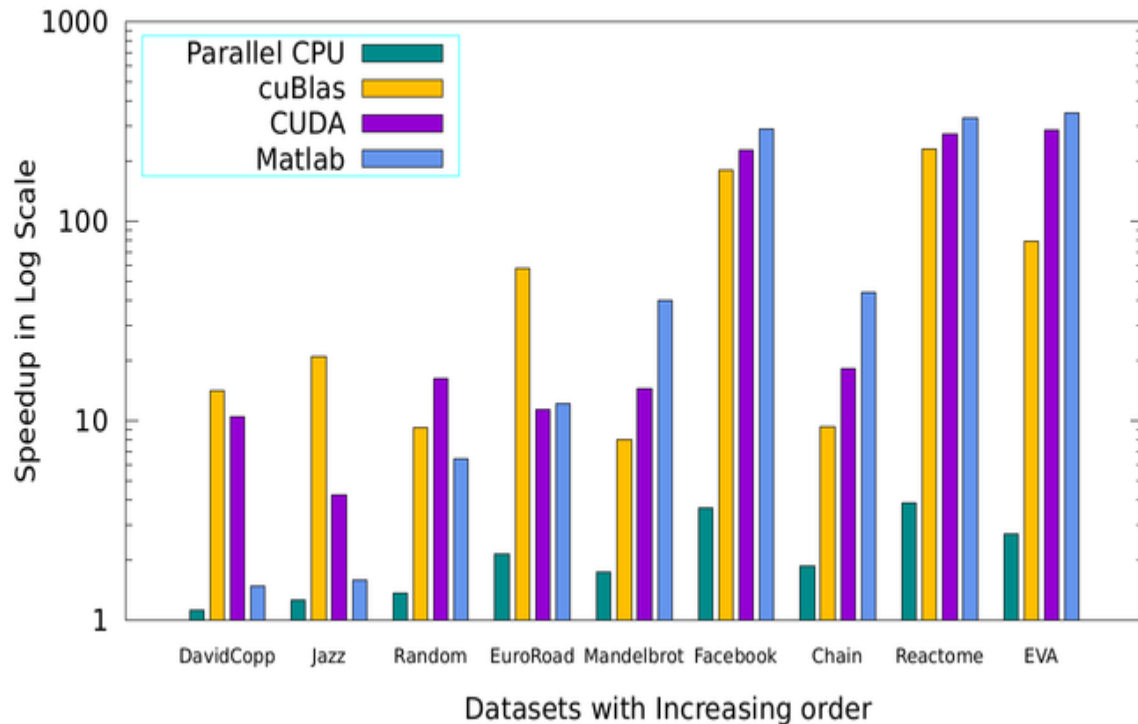


Finding :

- Different behavior for different graphs.
- cuBlas and CUDA better for smaller graphs.
- Matlab suitable for graphs of large order.
- Divide and Conquer approach versus $(L + 1/n)^{-1} - 1/n$

Experiments & Results

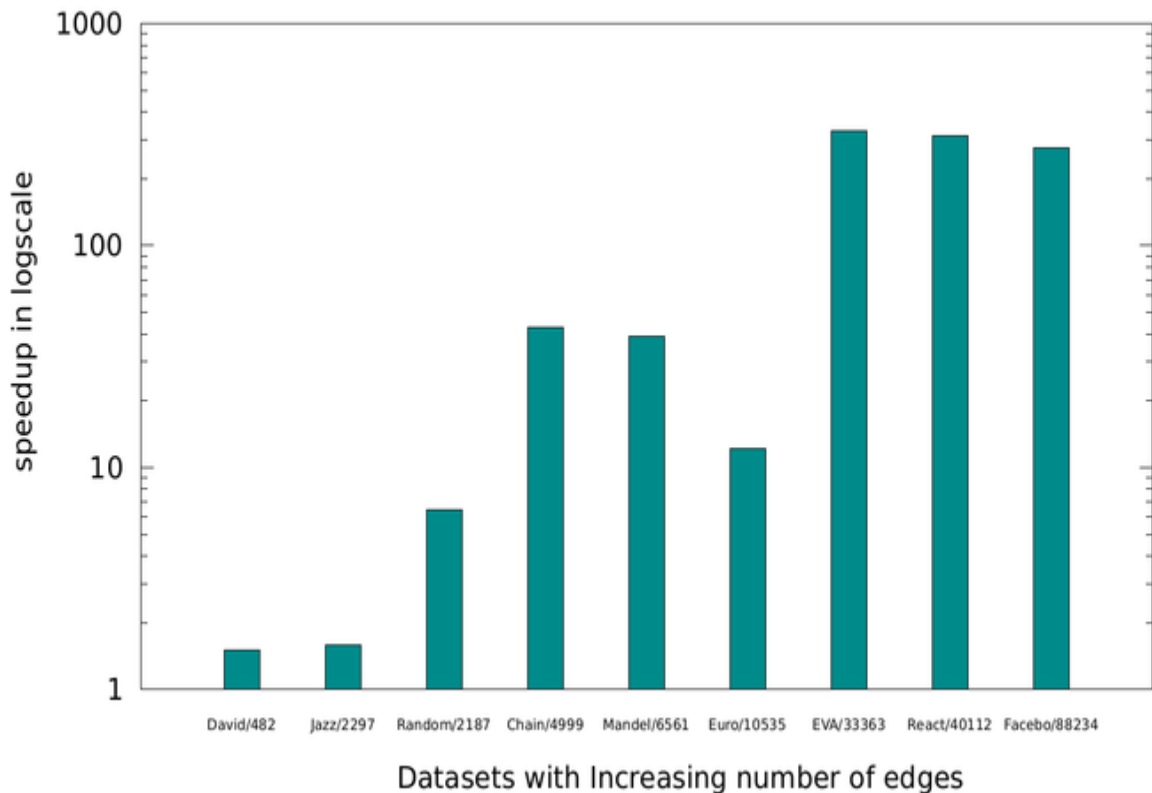
Speedup Comparison between ParallelCPU, cuBlas, Cuda and Matlab



Finding :

- Speedup achieved up-to 300 times.
- Matlab - Speedup - better for large order graphs.

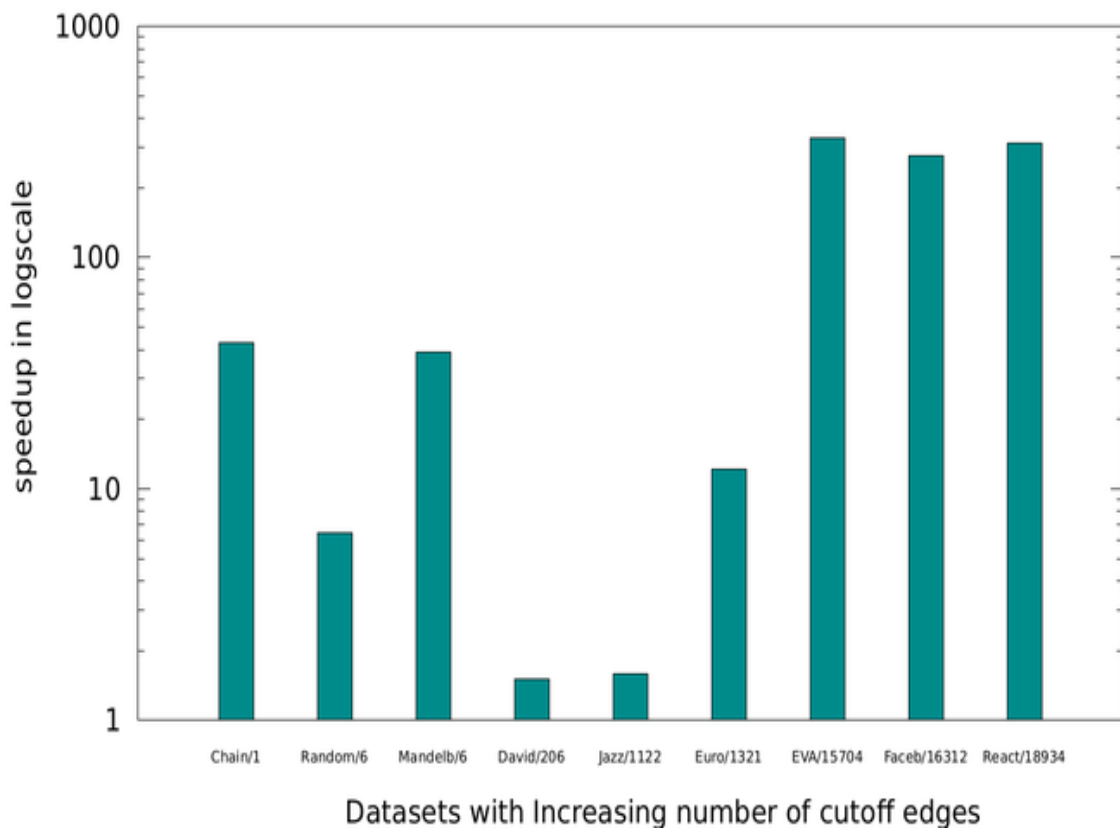
Experiments & Results



No Correlation Found !

Investigate more parameters !

Experiments & Results



No Correlation Found !

Future Work : Investigate more parameters.

Contributions

- ★ Designed a parallel version of the Divide-and-Conquer computation of the Moore-Penrose pseudo-inverse of the Laplacian .
- ★ Designed a GPU-enabled version of this parallel solution, and implemented it in Matlab, with significant speedup.
- ★ Implemented three other parallel versions, one using CUDA, one using cuBLAS, and a pThreads-based version.
- ★ Empirical evidence that the performance of three GPU-enabled versions is heavily dependent on the input graph properties.

Conclusion & Future Work

Conclusion :

- ❑ cuBlas and CUDA - small order graphs.
- ❑ Matlab implementation outperforms cuBlas and CUDA for large order graphs.
- ❑ Divide and Conquer Approach - Large Graphs - Significant Performance.
- ❑ Matlab GPU Computing - Productivity and Performance.
- ❑ Performance Variation - Input Graph

Conclusion & Future Work

Future Work :

- ★ Multiple GPU's , Recursive Partitioning.
- ★ Spanning Trees to Compute L^+ .
- ★ Investigate parameters of the graph affecting performance,

Thank you !

Any Questions ?

Nishant Saurabh
nishants.prmitr7@gmail.com



*Workshop on Large-Scale Parallel
Processing
IEEE International Parallel and Distributed
Processing Symposium*



Back-up Slides

Nishant Saurabh
*Vrije Universiteit
Amsterdam.*



Dr. Ana Lucia Varbanescu
*University of Amsterdam
Amsterdam.*



Dr. Gyan Ranjan
*Symantec
CA USA.*



Workshop on Large-Scale Parallel
Processing
IEEE International Parallel and Distributed
Processing Symposium




Computational Formula

Operations	Ω	L^+
First Join	$x, y \in G_1: \Omega_{xy}^{G3} = \Omega_{xy}^{G1}$	$l_{xy}^{+(1)} - \frac{n_2 n_3 (l_{xi}^{+(1)} + l_{iy}^{+(1)}) - n_2^2 (l_{ii}^{+(1)} + l_{jj}^{+(2)} + \omega_{ij})}{n_3^2}$
	$x, y \in G_2: \Omega_{xy}^{G3} = \Omega_{xy}^{G2}$	$l_{xy}^{+(2)} - \frac{n_1 n_3 (l_{xj}^{+(2)} + l_{jy}^{+(2)}) - n_1^2 (l_{ii}^{+(1)} + l_{jj}^{+(2)} + \omega_{ij})}{n_3^2}$
	$x \in G_1, y \in G_2: \Omega_{xy}^{G3} = \Omega_{xi}^{G1} + \omega_{ij} + \Omega_{jy}^{G1}$	$\frac{n_3 (n_1 l_{xi}^{+(1)} + n_2 l_{jy}^{+(2)}) - n_1 n_2 (l_{ii}^{+(1)} + l_{jj}^{+(2)} + \omega_{ij})}{n_3^2}$

Computational Formula

Operations	Ω	L^+
Edge Firing	$\frac{\Omega_{xy}^{Gl} - [(\Omega_{xj}^{Gl} - \Omega_{xi}^{Gl}) - (\Omega_{jy}^{Gl} - \Omega_{iy}^{Gl})]^2}{4(\omega_{ij} + \Omega_{ij}^{Gl})}$	$l_{xy}^{+(l)} - \frac{(l_{xi}^{+(l)} + l_{xj}^{+(l)})(l_{iy}^{+(l)} + l_{jy}^{+(2)})}{\omega_{ij} + \Omega_{ij}^{Gl}}$

Case Study - Topological Centrality



Γ_{11}^+	Γ_{12}^+	Γ_{13}^+	Γ_{14}^+	Γ_{15}^+
Γ_{21}^+	Γ_{22}^+	Γ_{23}^+	Γ_{24}^+	Γ_{25}^+
Γ_{31}^+	Γ_{32}^+	Γ_{33}^+	Γ_{34}^+	Γ_{35}^+
Γ_{41}^+	Γ_{42}^+	Γ_{43}^+	Γ_{44}^+	Γ_{45}^+
Γ_{51}^+	Γ_{52}^+	Γ_{53}^+	Γ_{54}^+	Γ_{55}^+

Topological Centrality (C_i^*) = $1/\Gamma_{ii}^+$, where Γ_{ii}^+ is general term of L^+ for node i

Experiments & Results

