

Minimum spanning tree (MST)

Graph $G = (V, E, w)$

MST: T subgraph of G (connected) - it spans G i.e. all vertices are covered $V(T) = V(G)$


It has minimum weight.

$$w(T) = \sum_{e \in E(T)} w(e)$$

exactly $(n-1)$ edges if there are n vertices.

Prim's algorithm

root some vertex with which algorithm starts.



$T = \{b\};$

$d[] = [1 \ 0 \ 5 \ 1 \ \infty]$ // distance of each vertex from T

pick the edge (b, d)

$T = \{b, d\};$

$d[] = [1 \ 0 \ 2 \ 1 \ \infty]$

pick the edge (b, a)

$T = \{b, d, a\};$

$d[] = [1 \ 0 \ 2 \ 1 \ 3]$

pick the edge (d, c)

~~pick the edge~~ $T = \{b, d, a, c\};$

$d[] = [1 \ 0 \ 2 \ 1 \ 3]$

pick the edge (a, e)

$T = \{b, d, a, c, e\};$

keep it at the old value
change for the new
vertices which
are not in T .

PRIM_MST_Seq (V, E, W, r)

{

T = {r};

d[r] = 0;

for all v ∈ V \ T:

if (r, v) ∈ E : d[v] = w(r, v)

else : d[v] = ∞ // something appropriate

for (n-1) iterations // in each iteration 1 new edge and 1 new vertex are added

n/p writes to each processor

combine across processors in log p steps.

⇒ find u s.t. d[u] = min{d[v] | v ∈ V - T}

T = T ∪ {u};

for all v ∈ V - T

d[v] = min(d[v], w(u, v));

no edge w(u, v) = ∞

}

}

OpenMP parallelization

for (n-1) iterations # pragma omp for (u: reduction(min))

find u s.t. d[u] = min{d[v] | v ∈ V - T}

any reasonable openMP implementation will create local copy of u - reduce across processors (log p) - useful if many many processors, but in shared memory have moderate number of threads, so openMP might just use critical section.

how to get u for which d[u] holds.

find the vertex having smallest index for which d[u] = min

reduction (u: min)

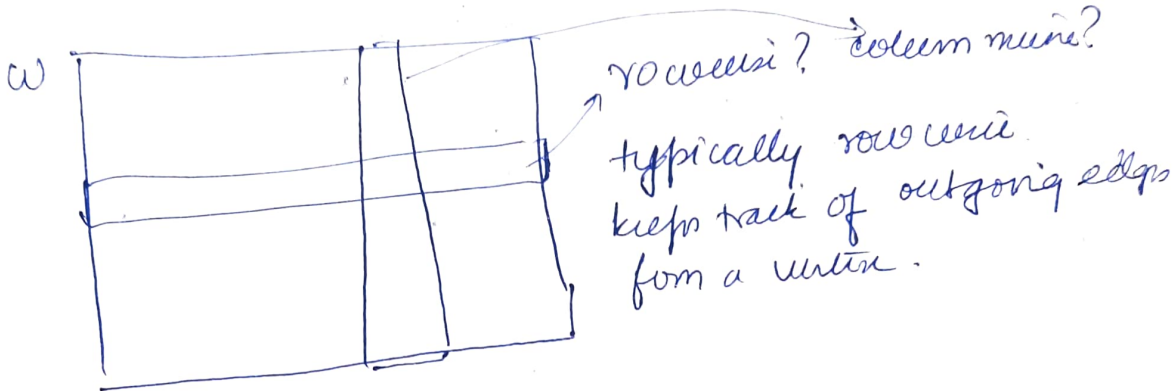
for all v ∈ V - T # pragma omp for

MPI implementation

$\text{reduce} \rightarrow \text{MPI-Reduce}$
 MPI-allreduce } similar to open MP reduction

careful about which vertex is assigned to which processor?

1-D distribution — each processor n/p vertices.
 w matrix distributed among processors



A processor already knows $d[v]$ if we do row-wise distribution, then $w[u, v]$ will be with owner of u .

* So instead of row-wise, we distribute the w matrix column-wise \rightarrow incoming edges to v .

Time

sequential time $\Theta(n^2)$

$$T_p = \Theta\left(\frac{n^2}{p}\right) \text{ [computation]}$$

$+ \Theta(\log p)$ [communication for reduction each iteration] $\log p$, so for n iterations $n \log p$

$$\text{Speedup} = \frac{\text{sequential time}}{\text{parallel time}} = \frac{\Theta(n^2)}{\Theta\left(\frac{n^2}{p}\right) + \Theta(n \log p)}$$

$$\text{efficiency} = \frac{\text{Speedup}}{p} = \frac{1}{\frac{\Theta\left(\frac{1}{p}\right) + \Theta\left(\frac{\log p}{n}\right)}{p \log p}} \rightarrow \text{Theoretical}$$

How many processors to have ideal efficiency of 1? $p = \frac{n}{\log n}$

Single source shortest Path

Dijkstra


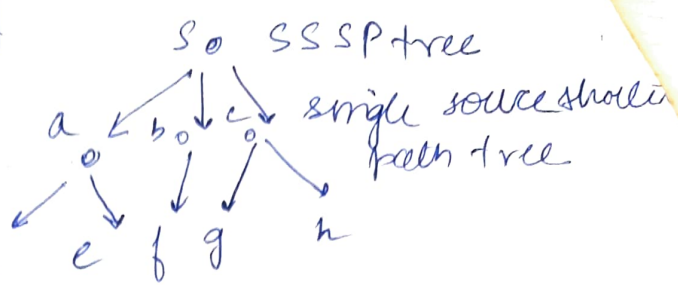
$$V = \{v\};$$

Shortest distance from every vertex d

$$d = \begin{bmatrix} s & a & b & c & d \\ 0 & \infty & \infty & \infty & \infty \end{bmatrix}$$

tentative shortest path length discovered so far

when terminal, has the shortest path distance for each vertex. can find the actual path with backtracking.



Relaxation on edge (u, v)

$$d[v] \leq d[u] + w(u, v)$$

$$\text{then } d[v] = d[u] + w(u, v)$$

Every edge from the vertices in V

$$d = \begin{bmatrix} s & a & b & c & d \\ 0 & 2 & 4 & \infty & \infty \end{bmatrix}$$

Similar to Prim's, I pick the minimum and include in V .

$$d = \begin{bmatrix} s & a & b & c & d \\ 0 & 2 & 3 & \infty & \infty \end{bmatrix}$$

SSSP - Seq (V, E, W, π)

$$\{ T = \{v\};$$

$$d[v] = 0;$$

for all $v \in V/T$ if $(r, v) \in E$; $d[v] = w(r, v)$
else ; $d[v] = \infty$

for $(n-1)$ iterations:

$$\text{find } u \text{ s.t. } d[u] = \min \{d[v] \mid v \in V - T\}$$

<p>In Prim</p> <p>$d[u]$: tentative shortest edge from T to V/T</p> <p>$d[u]$: tentative shortest path from $s(r)$ to u</p>

$$T = T \cup \{u\};$$

for all $v \in V - T$

$$d[v] = \min \{d[v], d[u] + w[u, v]\}$$

Parallelization Technique is same as Prim's.

Speedup

$$\frac{p}{\Theta(1) + \Theta\left(\frac{p \log p}{n}\right)}$$

$$p = \frac{n}{\log n}$$

Sparse graph: each vertex has constant number of neighbors

$\Theta(n) \rightarrow$ for $(n-1)$ iteration

$\Theta(n) \rightarrow$ find u s.t. $d[u] = \min \{d[v] \mid v \in V - T\}$

Linear search for min

$$T = T \cup \{u\};$$

for all $v \in V - T$

$$d[v] = \min \{d[v], d[u] + w[u, v]\}$$

$\Theta(n)$

for dense graph

but $\Theta(1)$ for sparse graph (as only neighbors of u can be updated)

①

Use Priority Queue data structure to find minimum

②

$$T_{seq} = \Theta(n \log n) \quad T_{parallel} = n \left(\log \frac{n}{p} + \log p \right) + \frac{c}{p} \log \frac{n}{p}$$

PQ = V

for each vertex u that is settled/frozen i.e. added to SSSP tree add distance along shortest path known \rightarrow removed from PQ.

$\Theta(\log n)$ asymptotically similar

Johnson's algorithm

How to parallelize if we need to use Priority Queue?

each processor owns n/p vertices. (Partition V into p disjoint sets V_1, V_2, \dots, V_p)

each processor i maintains priority queue

Priority Queue is own vertices which are not yet finalized w.r.t. shortest path.

Initialize PQ with V_i .

Maintain d as before for its own vertices.

Extracted node is settled locally and may get activated again only global mean it can be settled.

Look at all neighbors \rightarrow send messages to owners of all neighbors of the extracted min vertex

for all neighbors v of u : send (owner(v), (u, v) , $d[u] + w[u, v]$)

Receiver will examine (u, v) , $d[u] + w[u, v]$


owner of v checks if $d[v] > d[u] + w[u, v]$ then update.

* relaxation for p vertices instead of 1

iterations = diameter of the graph.


for $(n-1)$ iterations \leftarrow attack this.

SSSP-Tree



P0	P1	P2
S	a	b, c
{S}	{ }	{ }

$p = n$

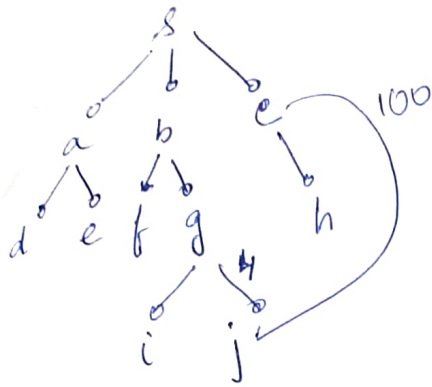


how many iterations?

depth of the tree

P0 \rightarrow sends messages to P1, P2

Priority Queue P0 { }
 P1 {a} \downarrow neighbors of a send msg
 P2 {b, c} \downarrow neighbors of b send msg



$P_8 \{b, fg\}$ $P_{10} \{c, j\}$

itcu1 { }

{ }

itcu2 {b} extract

{c} extract

itcu3 {f, g}

{j} extract

"thinks settled j"

itcu4 extract

Temporary settlement

but when P_{10} gets message from P_8 realize j has a shorter path from s .
activated again

Johnson - SSSP - ID (V, E, w, r)

{ if ($owner(r) == myRank$) {

$\bar{d}(r) = 0;$

insert r into PQ; }

$d \rightarrow$ stable value

$\bar{d} \rightarrow$ tentative value

use in PQ

for all $v \in V$.

$d(v) = \infty$

while (True)

{ $u = \text{extract_min}(PQ)$

$d[u] = \bar{d}[u]$

for all neighbors $v \in \text{Adj}[u];$

send $owner(v), (u, v) (d[u] + w[u, v])$

buffer for each proc send together

Receive messages

for each received msg $(x, v) (d[x] + w[x, v])$

if u in PQ

{ if $d[x] + w[x, v] < \bar{d}(v)$

update $\bar{d}(v) \leftarrow \left(\frac{\lg n}{p} \right)$ adjust priority queue

}

else // v not in PQ

if $(d[x] + w[x, v]) < d[v]$ →

$$d[v] = d[x] + w[x, v]$$

insert v into PQ ↗ readjust

replacemⁿ
of a settled
vertic^e.
(we thought
it is already
settled, but
not).

in priority queue d value
outside priority queue d value

check for termination

everyone checks size of PQ

all reduce

size 0

exit