



Distributed Bipartite Matching

Brian A. Page
bpage1nd.edu
December 10, 2018

The College of Engineering
at the University of Notre Dame




Bipartite Matching




A Maximum Matching

- Matching (M) is set of edges such that $E(u,v)$
- Vertices incident to only one edge in M


Alternating/Augmenting Paths




Current Matching



Find Alternate Paths



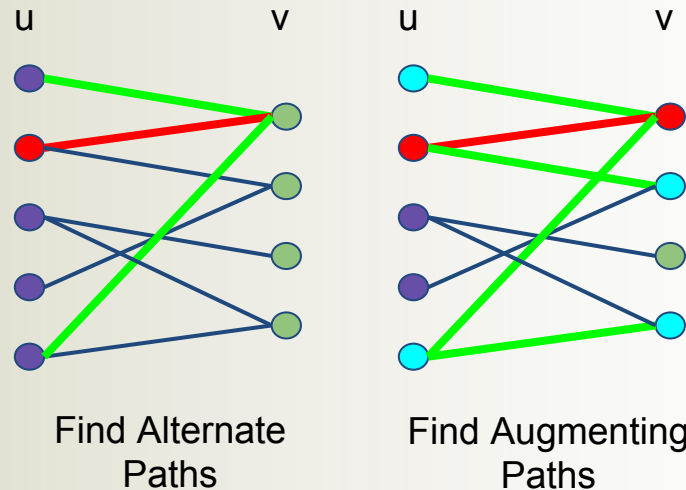
Find Augmenting Paths



Flip Matching Selection

- Many flow networks rely on augmenting path algorithms
- Can be used to find a more optimal matching
- Runtime influenced by number of potential augmenting paths needing verification

MBM: Multithreaded



- Utilizes Hopcroft-Karp algorithm
 - OpenMP for alternating/augmenting path checks
 - Runs in $O(|E|\sqrt{|V|})$
- Requires atomic operations for matching reversal
- **Slower than sequential**

Hopcroft-Karp

```
204 // Returns size of maximum matching
205 ▽ int Bipartite::hopcroftKarp() {
206     pairU = new int[uVertices.size() + 1];
207     pairV = new int[vVertices.size() + 1];
208     dist = new int[uVertices.size() + 1];
209
210     // Initialize NIL as pair of all vertices
211 ▽ for (int u = 0; u < uVertices.size(); u++) {
212         pairU[u] = NIL;
213     }
214 ▽ for (int v = 0; v < vVertices.size(); v++) {
215         pairV[v] = NIL;
216     }
217
218     // Initialize result
219     int result = 0;
220
221     // Keep updating the result while there is an augmenting path.
222 ▽ while (bfs()) {
223         // Find a free vertex
224 ▽ for (int u=1; u<=uVertices.size(); u++) {
225             // If current vertex is free and there is
226             // an augmenting path from current vertex
227 ▽ if (pairU[u] == NIL && dfs(u)) {
228                 result++;
229             }
230         }
231     }
232     return result;
233 }
```

1. BFS for alternate path frontier
2. DFS for augmenting verification
3. Check if matching increases if augmenting path chosen
4. stop when there are no more augmenting paths possible.



BFS

- Breadth First Search checks for vertices adjacent to a matched vertex
 - Adjacent vertices are candidates for alternating paths
- Parts of BFS can be parallelized using OpenMP
 - matched vertices can be checked simultaneously for adjacencies
- Initial experiments have used:
#pragma omp for

```
128 // Returns true if there is an augmenting path, else returns false
129 bool Bipartite::bfs(){
130     queue<int> Q; //an integer queue
131
132     // First layer of vertices (set distance as 0)
133     for (int u = 1; u <= uVertices.size(); u++) {
134         // If this is a free vertex, add it to queue
135         if (pairU[u] == NIL){
136             // u is not matched
137             dist[u] = 0;
138             Q.push(u);
139         }
140
141         // Else set distance as infinite so that this vertex
142         // is considered next time
143         else dist[u] = INF;
144     }
145
146     // Initialize distance to NIL as infinite
147     dist[NIL] = INF;
148
149     // Q is going to contain vertices of left side only.
150     while (!Q.empty()) {
151         // Dequeue a vertex
152         int u = Q.front();
153         Q.pop();
154
155         // If this node is not NIL and can provide a shorter path to NIL
156         if (dist[u] < dist[NIL]){
157             // Get all adjacent vertices of the dequeued vertex u
158             //list<int>::iterator i;
159             for (int i = 0; i < uVertices[u].edgelist.size(); ++i){
160                 int v = uVertices[u].edgelist[i].v;
161
162                 // If pair of v is not considered so far
163                 // (v, pairV[V]) is not yet explored edge.
164                 if (dist[pairV[v]] == INF) {
165                     // Consider the pair and add it to queue
166                     dist[pairV[v]] = dist[u] + 1;
167                     Q.push(pairV[v]);
168                 }
169             }
170         }
171     }
172
173     // If we could come back to NIL using alternating path of distinct
174     // vertices then there is an augmenting path
175     return (dist[NIL] != INF);
176 }
```

DFS

```
178 // Returns true if there is an augmenting path beginning with free vertex u
179 ▽ bool Bipartite::dfs(int u){
180 ▽   if (u != NIL){
181     //list<int>::iterator i;
182 ▽   for (int i = 0; i < uVertices[u].edgeList.size(); ++i){
183     int v = uVertices[u].edgeList[i].v;
184
185     // Follow the distances set by BFS
186 ▽   if (dist[pairV[v]] == dist[u]+1)           {
187     // If dfs for pair of v also returns
188     // true
189 ▽   if (dfs(pairV[v]) == true){
190     pairV[v] = u;
191     pairU[u] = v;
192     return true;
193   }
194 }
195 }
196
197 // If there is no augmenting path beginning with u.
198 dist[u] = INF;
199 return false;
200 }
201 return true;|
202 }
```

- Depth First Search checks to see if the alternating paths are augmenting paths
- Like BFS parts are easily parallelizable using OpenMP
 - check each path(edge)



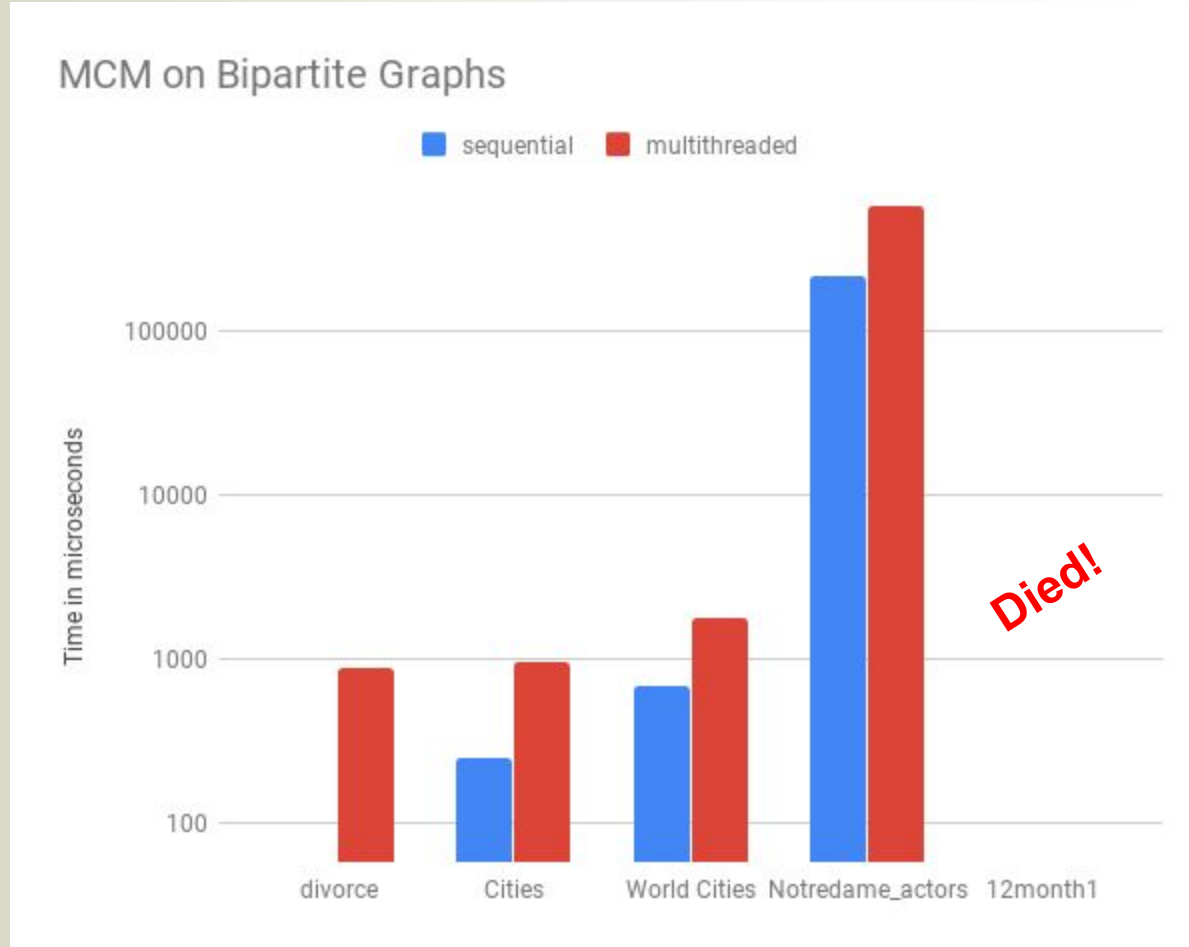
Benchmark Graphs

<u>Graph Name</u>	<u>Rows</u>	<u>Columns</u>	<u>Edges</u>
divorce	50	9	225
Cities	50	46	1342
World Cities	315	100	7518
Notredame_actors	392400	127823	1470404
12month1	12471	872622	22624727

- Suite sparse matrix/graph collection



Multithreaded Results



Distributed MBM: Scaling


- Use MPI for interprocess communication
- Hope to decrease run-time as process count scales
- **Allow for streaming of edge/vertex changes**

Challenges

- Optimal workload distribution necessary for reduced communication
- Dynamic Graph optimization




Partitioning



- Vertices sorted by degree
- Vertices assigned based on optimal packing of edge count
- Vertices have associated edge list
- Outing edges “owned” by source vertex
- Two copies of each edge
 - Aids alternating path search

BFS and DFS



- Process performs BFS with regard to locally owned vertices
- If adjacent vertices are not “owned” notify owner of traversal
- Similar behavior for DFS phase

Distributed MBM

Pros

- Partitioning possible and working as designed
 - Near uniform edge distribution

Cons

- Off node/process notifications dont work properly

```
[mbm(G)] = [cswarmfe.crc.nd.edu:mpi_rank_0][error_sighandler] Caught error: Segmentation fault (signal 11)

=====
= BAD TERMINATION OF ONE OF YOUR APPLICATION PROCESSES
= PID 5706 RUNNING AT cswarmfe.crc.nd.edu
= EXIT CODE: 11
= CLEANING UP REMAINING PROCESSES
[= YOU CAN IGNORE THE BELOW CLEANUP MESSAGES
=====
YOUR APPLICATION TERMINATED WITH THE EXIT STRING: Segmentation fault (signal 11)
```



Lessons Learned

- Scaling Bipartite Matching is hard!
 - race condition potential
 - sequential nature in native form
- Multithreading works, but due to some required atomic had worse performance
- I believe the use of futures might alleviate much of the off node communication issues (not necessarily the overhead)



Can This Be Done?!

- Yes. However it is non-trivial!
- Best methods to date still require heavy communication volumes



References

- Jeremy Kepner and John Gilbert. 2011. Graph Algorithms in the Language of Linear Algebra. Soc. for Industrial and Applied Math., Philadelphia, PA, USA.
- https://en.wikipedia.org/wiki/Bipartite_graph
- [https://en.wikipedia.org/wiki/Matching_\(graph_theory\)#Bipartite_matching](https://en.wikipedia.org/wiki/Matching_(graph_theory)#Bipartite_matching)
- <https://www.geeksforgeeks.org/ford-fulkerson-algorithm-for-maximum-flow-problem/>
- https://en.wikipedia.org/wiki/Hungarian_algorithm#The_algorithm_in_terms_of_bipartite_graphs
- Ariful Azad and Aydin Buluc. 2016. Distributed-Memory Algorithms for Maximum Cardinality Matching in Bipartite Graphs. Lawrence Berkeley National Labs.