


30 Aug 2021

Max bipartite matching

Min-cost bipartite perfect matching

RECAP. For G bipartite and $M \subseteq E(G)$ matching

- a vertex is free if it doesn't belong to an edge of M , else it is matched.
- an M -alternating path (or cycle) is one whose edges alternate between belonging, not belonging to M .
- an M -augmenting path is an M -alternating path from a free vertex to another free vertex.



Algorithm (partly specified) for Max Matching

$M \leftarrow \emptyset$

while \exists M -augmenting path P

$M \leftarrow M \oplus P$

endwhile

output M

how to search efficiently?

(§2) For general graphs (incl. non bipartite) can find P in $O(mn)$ time.

Last time we saw this takes $\leq \frac{n}{2}$ iterations
 ($n = \#$ vertices, $m = \#$ edges)
 and we saw it correctly outputs a
 max matching.



For bipartite G , we can find M -augmenting P
 if it exists, in linear time.

Def. If $G = (L \cup R, E)$ is bipartite
 and M is a matching, the
residual graph G_M has vertices $L \cup R$
 and directed edges

$$\left\{ (u, v) \mid u \in L, v \in R, (u, v) \notin M \right\}$$

$$\cup \left\{ (v, u) \mid u \in L, v \in R, (u, v) \in M \right\}$$

E.g.



{M-augmenting paths in G } $F \triangleq \{\text{free vertices}\}$

|||

{directed paths in G_M from $L \cap F$ to $R \cap F$ }

BFS finds a dir. path in G_M from $L \cap F$ to $R \cap F$ in time $O(m+n)$ if one exists.

$O(n)$ loop iterations, $O(m+n)$ time per iteration
 $\Rightarrow O(mn + n^2)$ time altogether.

Hopcroft & Karp improved this to $O(m\sqrt{n})$.
(§1.4 of notes, to be presented in ≈ 2 weeks)

Minimum Cost Bipartite Perf. Matching.

Input: $G = (L \cup R, E)$ bipartite

$c: E \rightarrow \mathbb{R}$

Convention: $c(u,v) = \infty$ if $(u,v) \notin E$.

If G^+ denotes the complete bipartite graph on vertex sets L, R and M is a matching in G^+

$$c(M) = \sum_{(u,v) \in M} c(u,v).$$

matching with
no free vertices
//

Problem. Find a perfect matching in G^+
with minimum combined edge cost.

Running time will be expressed in terms of

$n = \#$ vertices of G

$m = \#$ edges of G .

Idea: Start with empty matching M_0 .


Build up a sequence of matchings

$M_1, M_2, \dots, M_{n/2}$ using one
augmenting path at a time.

$$M_{k+1} \text{ will be equal to } M_k \oplus P_k$$

where P_k is some (carefully chosen)
 M_k -augmenting path.

How does the cost of a matching change
when we replace M with $M \oplus P$?



$$\text{cost}(M) = 1+1 \quad \text{cost}(M \oplus P) = 3+4+5$$

$$\begin{aligned} \Delta c(P; M) &= \text{cost}(M \oplus P) - \text{cost}(M) \\ &= 3 - 1 + 4 - 1 + 5 \\ &\quad \text{---} \text{---} \text{---} \text{---} \text{---} \end{aligned}$$

If we assign edge costs to directed edges in G_M as follows.

$$\text{cost}(u, v) = \begin{cases} c(u, v) & \text{if } (u, v) \notin M \\ -c(u, v) & \text{if } (v, u) \in M \end{cases}$$

then the bijection

$$\{\text{M-augmenting paths}\} \longleftrightarrow \{\text{directed paths in } G_M \text{ from } L \cap F \text{ to } R \cap F\}$$

is cost-preserving,

GREEDY ALGORITHM

$$M_0 \leftarrow \emptyset$$

for $k = 0, \dots, \frac{n}{2} - 1$

compute G_M with its edge costs defined as above.

let $P_k = \text{min cost directed path from Lof to Rof}$

$$\text{let } M_{k+1} = M_k \oplus P_k$$

end for

output $M_{n/2}$.

↑
efficient search procedure?

Correct?

Yes, but surprisingly subtle to prove.

Induction hypothesis: M_k has the least cost of all matchings with k edges.

(See lecture notes for a self-contained proof of induction step.)

Efficient?

Use Bellman-Ford. Finding P_k takes time $O(mn)$... if G_M has no negative cost cycles!

If $M = M_k$ had a negative cost cycle C in G_M , then consider $M' = M \oplus C$. Note C is an even length M -alternating cycle. So M' is a matching which also has k edges.

$$\text{cost}(C) = \text{cost}(M') - \text{cost}(M).$$

By induction hypothesis above, $M = M_k$ has the least cost among all k -edge matchings.

still waiting to prove this.

$$\therefore \text{cost}(M') \geq \text{cost}(M)$$

$$\therefore \text{cost}(C) \geq 0$$

\therefore we're allowed to use Bellman-Ford to find P_k .

Running time: $O(mn^2)$.

$\frac{n}{2}$ \log iterations, $O(mn)$ per iteration.

Next lecture:

1. Make this faster using Dijkstra in place of Bellman-Ford.
2. Prove the algorithm is correct.

Proving correctness holds the key to speeding up the algorithm.

