Homework 4

Algorithms on Directed Graphs, Winter 2018/9

Due: 23.11.2018 by 16:00

1 Disjoint Paths and k-linkage

In this note we consider two related problems on directed graphs: *disjoint paths*, and *k-linkage*. In the following sections, we describe three reductions with implications for the computational complexity of disjoint paths and k-linkage.

Let $G = (V(G), E(G))$ be a directed graph, and let $A, B \subseteq V(G)$ be disjoint subsets with $|A| = |B| = k$. For concreteness we write

$$A = \{a_1, a_2, \ldots, a_k\} \quad \text{and} \quad B = \{b_1, b_2, \ldots, b_k\}$$

We say that (G, A, B) satisfies the ***disjoint paths property*** if there exist k pair-wise vertex disjoint directed paths $P_1, P_2, \ldots, P_k$ such that the first vertex of each P_i is in A and the last vertex of each P_i is in B. We say that the paths $P_1, P_2, \ldots, P_k$ form a ***k-linkage*** if additionally for each $i = 1, 2, \ldots, k$, P_i is a path from a_i to b_i. If such a family of paths exist, we say that (G, A, B) satisfies the ***k-linkage property***.

2 Reducing Disjoint Paths to Max Flow/Min Cut

We describe a reduction from disjoint paths to the max flow/min cut problem. Let (G, A, B) be an instance of the disjoint paths problem. We construct the graph G' as follows. The vertex set $V(G')$ contains vertices as follows:

- For each $a_i \in A$ and $b_i \in B$, $i \in [k]$, we have $a_i, b_i \in V(G')$.
- For each $v \in V(G) \setminus (A \cup B)$, there are two vertices $v', v'' \in V(G')$.
- There are two additional vertices $s, t \in V(G')$.

The edge set of G' and edge capacities are constructed as follows:

- For each $a_i \in A$, $b_i \in B$, we have $(s, a_i), (b_i, t) \in E(G')$ with $c((s, a_i)) = c((b_i, t)) = 1$.

- For each $v \in V(G) \setminus (A \cup B)$, $(v', v'') \in E(G')$ with $c((v', v'')) = 1$
- For each $(u, v) \in E(G)$ we have
  - $(u'', v') \in E(G')$ if $u, v \notin A \cup B$
  - $(u, v') \in E(G')$ if $u \in A$, $v \notin A \cup B$
  - $(u'', v) \in E(G')$ if $u \notin A \cup B$, $v \in B$
  - $(u, v) \in E(G')$ if $u \in A$, $v \in B$.

  The capacities of all edges above are ∞.

There are no other edges in $E(G')$. In particular, G' does not contain any edges of the form (u, a) for $a \in A$, $u \neq s$, nor does it contain edges of the form (b, w) for $w \neq t$.

Exercise 1. Prove that the graph G' constructed above has max flow/min cut value of k if and only if (G, A, B) satisfies the disjoint paths property.

Exercise 2. Use the construction above and the max flow/min cut theorem to prove Menger's theorem: the size of the minimum vertex cut[1] separating A and B is equal to the number of vertex-disjoint paths between A and B.

3 Reducing 3-SAT to k-linkage

Let $x_1, x_2, \ldots, x_n$ be Boolean variables, and let Φ be a 3-CNF formula over $x_1, x_2, \ldots, x_n$. That is,

$$\Phi = \phi_1 \wedge \phi_2 \wedge \cdots \wedge \phi_m,$$

where each ϕ_j is of the form

$$\phi_j = y_j^1 \vee y_j^2 \vee y_j^3$$

and each y_j^ℓ is a literal (equal to some x_i or its negation $\neg x_i$). The 3-SAT problem is to determine if there exists an assignment of the x_i to true or false such that Φ evaluates to true—in this case we say that Φ is ***satisfiable***. 3-SAT is one of the classical NP-complete problems.

Given a 3-CNF formula Φ as above, we construct a directed acyclic graph $G = G(\Phi)$ such that G satisfies the k-linkage property if and only if Φ is satisfiable. We build the graph G as follows.

- For each literal x_i, $V(G)$ contains two vertices s_i, t_i along with two vertex-disjoint paths from s_i to t_i. We label the two paths by T_i and F_i respectively. Initially T_i and F_i each consists of a single edge from s_i to t_i, but the edges will be sub-divided as the construction proceeds.

[1] A ***vertex cut*** separating A and B is a set $X \subseteq V(G)$ such that removing the vertices in X from G disconnects $A \setminus X$ and $B \setminus X$. In particular, the cuts $X = A$ and $X = B$ separate A and B, so that a min cut has size at most k.

- For each clause ϕ_j in Φ, $V(G)$ contains two vertices u_j and w_j along with three vertex-disjoint paths from u_j to w_j. The paths correspond to the three literals in ϕ_j. For each literal y in ϕ_j:
  - if $y = x_i$ appears as a positive literal in ϕ_j, subdivide the path F_i by adding a new vertex v to this path (between s_i and t_i),
  - if $y = \neg x_i$ appears as a negative literal in ϕ_j, subdivide the path T_i by adding a new vertex v to this path.

  Then add the edges (u_j, v) and (v, w_j) to $E(G)$ so that they form a path from u_j to w_j.

Exercise 3. For the graph G constructed as above, prove that there exists a k-linkage from $A = \{s_1, s_2, \dots, s_n\} \cup \{u_1, u_2, \dots, u_m\}$ to $B = \{t_1, t_2, \dots, t_n\} \cup \{w_1, w_2, \dots, w_m\}$ $(k = n + m)$ if and only if Φ is satisfiable.

4 Reducing k-linkage for DAGs to Connectivity

Let G be a directed acyclic graph (DAG), and let $A, B \subseteq V(G)$ with $A \cap B = \varnothing$ and $|A| = |B| = k$. Without loss of generality, we assume that there are no edges of the form (u, a) for $a \in A$, nor any edges of the form (b, w) for $b \in B$.[2] We construct a graph G' from G in the following manner. The vertex set $V(G')$ is the the set of k-tuples of pair-wise distinct vertices in G:

$$V(G') = \{(v_1, v_2, \dots, v_k) \,|\, v_1, v_2, \dots, v_k \in V(G), i \neq j \implies v_i \neq v_j\}.$$

The edge set $E(G')$ is defined as follows. Since G is a DAG, for every $\mathbf{v} = (v_1, v_2, \dots, v_k) \in V(G')$, there is some index r such that v_r is not reachable from any of the other v_i $(i \neq r)$. That is, there is no directed path from any v_i to v_r. For this index r (choosing one arbitrarily if there is more than one), and each edge $(v_r, w) \in E(G)$ with $w \neq v_1, v_2, \dots, v_n$, we add the edge $(\mathbf{v}, \mathbf{v}_r(w)) \in E(G')$ where $\mathbf{v}_r(w)$ is $\mathbf{v}$ with v_r replaced by w:

$$\mathbf{v}_r(w) = (v_1, v_2, \dots, v_{r-1}, w, v_{r+1}, \dots, v_k).$$

We will first show that if G' contains a directed path P' from $\mathbf{a} = (a_1, a_2, \dots, a_k)$ to $\mathbf{b} = (b_1, b_2, \dots, b_k)$, then G contains a k-linkage from A to B. To this end, suppose P' consists of vertices

$$\mathbf{v}_1(= \mathbf{a}), \mathbf{v}_1, \dots, \mathbf{v}_\ell, \mathbf{v}_{\ell+1}(= \mathbf{b}).$$

From the definition of G', for each edge $(\mathbf{v}_j, \mathbf{v}_{j+1})$ along P', there is a unique index r_j such that $\mathbf{v}_j$ and $\mathbf{v}_{j+1}$ differ only in the entry with index r_j. Let u_j and w_j denote the r_j-th entry of $\mathbf{v}_j$ and $\mathbf{v}_{j+1}$, respectively. Further, by the definition of $E(G')$, we must have $(u_j, w_j) \in E(G)$. If $r_j = i$, we call (u_j, w_j) an i-edge.

[2] No edge of the form (u, a) or (b, w) can appear in any k-linkage, so removing such edges will not change whether or not G satisfies the k-linkage property.

Claim. The set of i-edges induced by P' forms a path P_i from a_i to b_i in G.

To see the claim, consider a fixed i, and let consider the i-edges in the order that they appear in P'. Consider the first i-edge (u_j, w_j), and let $(\mathbf{v}_j, \mathbf{v}_{j+1})$ be the corresponding edge in P'. Since this the first i-edge, the i-th entry of $\mathbf{v}_j$ is the same as the i-th entry of $\mathbf{v}_1 = \mathbf{a}$. Therefore, $u_j = a_i$, so the first i-edge is an edge out of a_i: (a_i, w_j). Similarly, the next i-edge is an edge from w_j to some other vertex. Continuing in this way, i-edges form a path from a_i. Since the final vertex in P' is $\mathbf{v}_{\ell+1} = \mathbf{b}$, the final i-edge in P' must change the i-th entry of the corresponding $\mathbf{v}_j$ to b_i, so that P_i does indeed terminate at b_i.

By the claim, the path P' induces k paths $P_1, P_2, \ldots, P_k$ where each P_i is a path from a_i to b_i. We now show that the P_i are pair-wise vertex disjoint, thus forming a k-linkage from A to B. To this end, suppose towards a contradiction that there exist $i \neq j$ and a vertex w appearing in both P_i and P_j. Let u_i denote the vertex before w in P_i and let u_j denote the vertex before w in P_j. Thus (u_i, w) is and i-edge and (u_j, w) is a j-edge. Suppose the edges in P' corresponding to (u_i, w) and (u_j, w) are the t_i-th and t_j-th edges respectively. Assume without loss of generality that $t_i < t_j$. Let u'_j be the j-th entry of $\mathbf{v}_{t_i}$—i.e., u'_j is the j-th entry of $\mathbf{v}$ when the i-edge (u_i, w) appears in P'.

Observe that j-edges induced by the sub-path $\mathbf{v}_{t_i}, \mathbf{v}_{t_i+1}, \ldots, \mathbf{v}_{t_j}$ of P' form a path from u'_j to u_j (and then to w). Thus, w is reachable from u'_j in G. Therefore (by the definition of G') the i-th entry of $\mathbf{v}_{t_j}$ cannot have changed from w in this sub-path. But this implies that in $\mathbf{v}_{t_j}$, both the i-th and j-th entries are equal to w, contradicting the definition of G'! Therefore, the paths P_i and P_j do not intersect, as desired.

Exercise 4. Assume (without loss of generality) that G does not contain edges of the form (u, a) for $a \in A$, nor any edges (b, w) for $b \in B$. Prove that if G contains a k linkage from A to B, then there exists a path P' in G' from $(a_1, a_2, \ldots, a_k)$ to $(b_1, b_2, \ldots, b_k)$.