ORIE 633 Network Flows	October 25, 2007
Lecture 15	
Lecturer: David P. Williamson	*Scribe: Yu Tung Cheng*

1 Polynomial time algorithms for minimum-cost circulations

1.1 Cost scaling (cont.)

Recall the following algorithm and theorems from the last lecture:

Cost Scaling (Goldberg, Tarjan '90)

Let f be any feasible arc
  $\epsilon \leftarrow C$
  $p_i \leftarrow 0, \forall i \in V$
  While $\epsilon \geq 1/n$
    $\epsilon \leftarrow \epsilon/2$
    $(f, p) \leftarrow$ find-ϵ-opt-circ(f, ϵ, p)

Push/relabel find-ϵ-opt-circ(f, ϵ, p)

$\forall (i,j) \in A_f$ if $c^p_{ij} < 0, f_{ij} \leftarrow u_{ij}$
While $\exists$ active $i \in V$ ($e^f_i > 0$)
  If $\exists j$ s.t. $u^f_{ij} > 0$ and $c^p_{ij} < 0$
    Push $\delta = \min(e^f_i, u^f_{ij})$ flow on (i,j)
  Else
    Relabel $p_i \leftarrow \max_{(i,j)\in A_f} (p_j - c_{ij} - \epsilon)$
Return (f, p)

Last time, we showed the following.

Theorem 1 *Cost scaling takes $O(\min(\log(nC), m \log n))$ iterations.*

The find-ϵ-opt-circ(f, ϵ, p) subroutine does the following:

find-ϵ-opt-circ

- **Input:** 2ϵ-opt circulation f, potentials p s.t. $c^p_{ij} \geq -2\epsilon, \forall (i,j) \in A_f$
- **Goal:** ϵ-opt circulation f', potentials p' s.t. $c^p_{ij} \geq -\epsilon, \forall (i,j) \in A_{f'}$

Last time we described the push/relabel algorithm given above to implement this subroutine. It works by first creating an ϵ-optimal pseudoflow, then gradually changes this to a circulation while maintaining ϵ-optimality. Recall that a pseudoflow is obeys capacity constraints and antisymmetry, but not necessarily flow conservation.

We left the following lemma unproved last time; this time we will finish its proof.

Lemma 2 *For any i, p_i decreases by at most $3n\epsilon$ during the algorithm.*

From this lemma, we derived the following statements about the running time of the algorithm.

Lemma 3 *The total number of relabels is $\leq 3n^2$.*

Lemma 4 *The number of saturating pushes in the above algorithm is at most $O(nm)$.*

Lemma 5 *The number of non-saturating pushes in the algorithm is $O(n^2 m)$.*

Theorem 6 *The* Push/Relabel find-ε-opt-circ *subroutine takes $O(n^2 m)$ time. Furthermore, with a FIFO implementation of Push/Relabel, the subroutine runs in $O(n^3)$ time. With the use of fancy data structures, it runs in $O(mn \log(n^2/m))$ time.*

Corollary 7 *The cost scaling algorithm runs in $O(mn \log(n^2/m) \min(\log(nC), m \log n))$ time.*

Now we proceed to the proof of Lemma 2. We first need the following.

Lemma 8 *Let f be a pseudoflow, f' a circulation. For any i such that $e_i^f > 0$, there exists j such that $e_j^f < 0$ and there exists a path P from i to j with $(k,l) \in A_f, (l,k) \in A_{f'}$ for all $(k,l) \in P$.*

Proof: Claim: It is possible to find P in set of arcs

$$A_< = \{(i,j) : f_{ij} < f'_{ij}\}$$

Note $A_< \subseteq A_f$ since $f_{ij} < f'_{ij}$ implies $f_{ij} < u_{ij}$. Further note that if $(i,j) \in A_<$, then $(j,i) \in A_{f'}$ since then $f'_{ji} < f_{ji} \leq u_{ji}$. Thus given a vertex i such that $e_i^f > 0$, it will be sufficient to find a path in $A_<$ to some j such that $e_j^f < 0$.

To do this, let S be all vertices reachable from i using arcs in $A_<$. Then,

$$\begin{aligned} -\sum_{k \in S} e_k^f &= \sum_{k \in S} \sum_{j:(k,j) \in A} f_{kj} \\ &= \sum_{k \in S, j \notin S, (k,j) \in A} f_{kj} \\ &\geq \sum_{k \in S, j \notin S, (k,j) \in A} f'_{kj} = 0. \end{aligned}$$

The inequality holds because each (k,j) in the sum is not in $A_<$. The last equality holds because f' is a circulation.

Since $e_i^f > 0$, then there must be $j \in S$ such that $e_j^f < 0$. Furthermore, j is reachable from i using arcs of $A_<$. □

Now we can prove Lemma 2.

Proof of Lemma 2: Let f' be the initial 2ϵ-optimal circulation, and p' initial potentials. We consider the last point in the algorithm during which p_i is relabelled. Note that if p_i is relabelled, then $e_i^f > 0$. By the previous lemma, we know there is $j \in V$ such that $e_j^f < 0$ and there is a path P from i to j in A_f, with the reverse of the path in $A_{f'}$.

First, observe that f being ϵ-optimal implies

$$-|P|\epsilon \leq \sum_{(k,l)\in P} c_{kl}^p = \sum_{(k,l)\in P} (c_{kl} + p_k - p_l) = \left(\sum_{(k,l)\in P} c_{kl}\right) + p_i - p_j$$

Next, observe that since f' is 2ϵ-optimal and the reverse of P from j to i is in $A_{f'}$ implies that

$$-2\epsilon|P| \leq \sum_{(k,l)\in P} c_{lk}^{p'} = \sum_{(k,l)\in P} c_{lk} + p'_j - p'_i$$

Finally, observe that by our definition of costs $\sum_{(k,l)\in P} c_{kl} = -\sum_{(k,l)\in P} c_{lk}$. Thus by adding the previous inequalities, we get

$$-3\epsilon|P| \leq (p_i - p'_i) + (p'_j - p_j).$$

Because $e_j^f < 0$, the node j must not have been relabelled to this point in the algorithm, and thus $p_j = p'_j$. Therefore we have that

$$p_i - p'_i \geq -3n\epsilon.$$

Since we assumed that this was the last point in the algorithm during which i was relabelled, the lemma statement follows. □

We close our performance analysis of the cost-scaling algorithm with two open questions. First, is a minimum-cost circulation problem solvable with $O(\min(m \log n, \log(nC)))$ iterations of any maximum flow algorithm? It looked like the push/relabel algorithm could be used as the *find-ε-opt-circ* subroutine with only minor modifications; this is also the case for the blocking flow variant of this subroutine.

More to the point, can the Goldberg-Rao maximum flow algorithm be used for this subroutine? This would then give us a minimum-cost circulation algorithm that runs in $O(\Lambda m \log n(\log(mU))(\log(nC)))$ time, which would be the fastest known algorithm.

1.2 Capacity scaling

We consider one last algorithm for the minimum-cost circulation problem. So far we've considered algorithms that have used ϵ-optimality to measure their progress toward optimality. Today we'll start looking at an algorithm that maintains 0-optimality but has a pseudoflow and works towards primal feasibility.

To make the algorithm work, we need to ensure that we can push any amount of flow from any node to any other node. To do this, we can add arcs to the graph of infinite capacity but very high cost, high enough that any optimal flow would never use them.

We want to maintain a pseudoflow with potentials p, and a parameter Δ ($= U$ initially), where

$$A_f(\Delta) = \{(i,j) \in A_f : u_{ij}^f \geq \Delta\}$$

$$S(\Delta) = \{i \in V : e_i^f \geq \Delta\}$$

$$T(\Delta) = \{i \in V : e_i^f \leq -\Delta\}$$

The idea is to enforce $c_{ij}^p \geq 0$ for all $(i,j) \in A_f(\Delta)$. Then we repeatedly move Δ units of flow from $S(\Delta)$ to $T(\Delta)$ until either $S(\Delta) = \emptyset$ or $T(\Delta) = \emptyset$ (which will imply that there is not too much excess left). Then divide Δ by 2 and repeat.

We claim that the algorithm will do what we want.

Claim 9 *When $\Delta < 1$, f is a feasible circulation and it is optimal.*

Proof: First we show that the circulation is feasible. At the end of iteration, $\Delta < 1$. This means $\{i : e_i \geq 1\} = \emptyset$ or $\{i : e_i \leq -1\} = \emptyset$. Since $\sum e_i = 0$, by the integrality of flow, this implies that $e_i = 0$ for all $i \in V$. Hence f is feasible.

The circulation is optimal since $c_{ij}^p \geq 0$ for all $(i,j) \in A_f(\Delta)$, and when $\Delta < 1$, $A_f(\Delta) = A_f$. By the optimality conditions for circulations we showed some time ago, this implies optimality. □

At the beginning of an iteration, if $c_{ij}^p < 0$ and $(i,j) \in A_f(\Delta)$, we saturate edge (i,j). Then we will have a pseudoflow such that $c_{ij}^p \geq 0$ for all $(i,j) \in A_f(\Delta)$ initially. How do we maintain this as we push flow on paths in $A_f(\Delta)$? It's a bit tricky since if we push flow on any edge (i,j) such that $c_{ij}^p > 0$, then the arc (j,i) will enter the residual graph and have $c_{ji}^p < 0$. To resolve this, we'd like to push flow on arcs (i,j) such that $c_{ij}^p = 0$. We'll see how we can accomplish this in the next lecture.