Minimum Cost Flows

Definitions

- G=(V,E) is a directed graph
- capacity u(v,w) for every v,w $\in$V: If (v,w) $\notin$ E then u(v,w) = 0
- Cost c(v,w) for every (v,w) $\in$ E, c(w,v) = -c(v,w)
- Two distinguished vertices s and t.

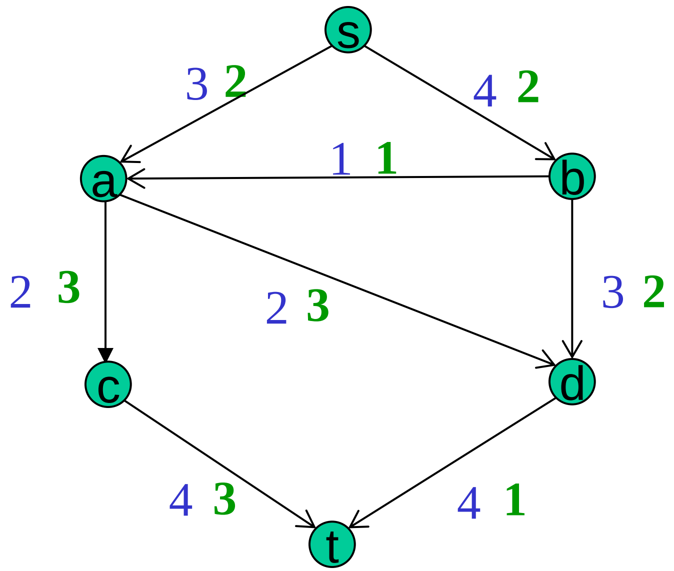

A Minimum Cost Flow

A flow is a function on the edges which satisfies the following requirements

- $f(v,w) = -f(w,v)$ skew symmetry
- $f(v,w) \leq u(v,w)$
- For every v except s and t $\Sigma_w f(v,w) = 0$

The value of the flow $|f| = \Sigma_w f(s,w)$

The cost of f: $\text{cost}(f) = \frac{1}{2} \sum_{(v,w)\in E} c(v,w) \cdot f(v,w)$

The minimum cost flow problem is to find a maximum flow of minimum cost

A Minimum Cost Circulation

A circulation is a function on the edges which satisfies the following requirements

- $f(v,w) = -f(w,v)$ skew symmetry
- $f(v,w) \leq u(v,w)$
- For every v $\Sigma_w f(v,w) = 0$

The cost of f: $\text{cost}(f) = \frac{1}{2} \sum_{(v,w)\in E} c(v,w) \cdot f(v,w)$

The minimum cost circulation problem is to find a circulation of minimum cost

Min Cost Flow is eq. to Min Cost Circulation

can solve MCF ➔ Can solve MCC

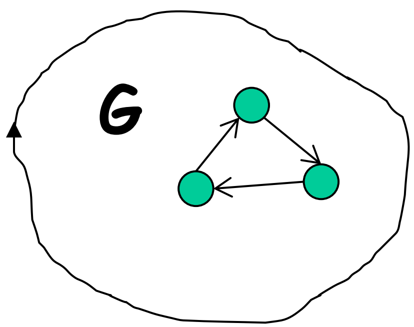

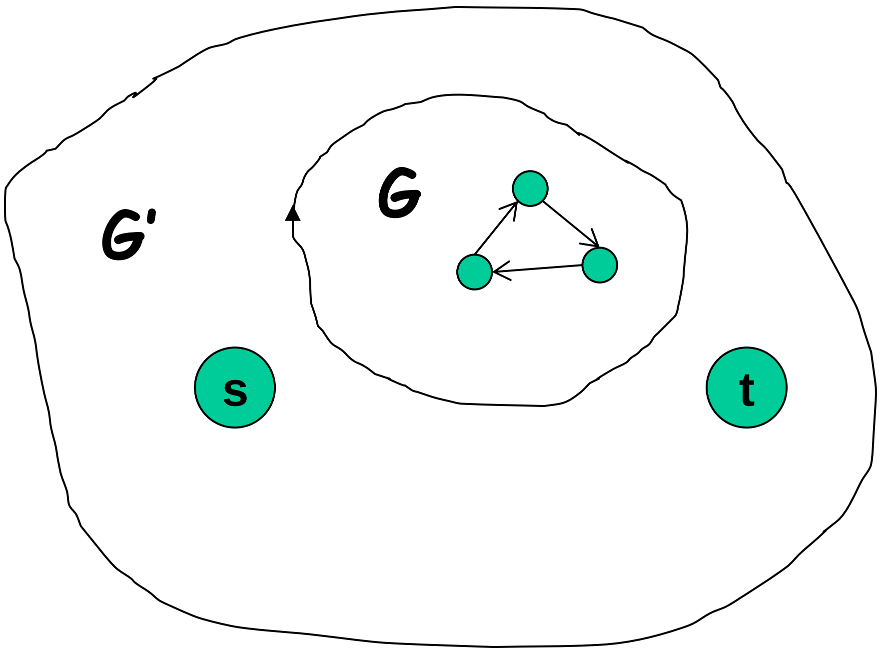

Min Cost Flow is eq. to Min Cost Circulation

can solve MCC ➔ Can solve MCF

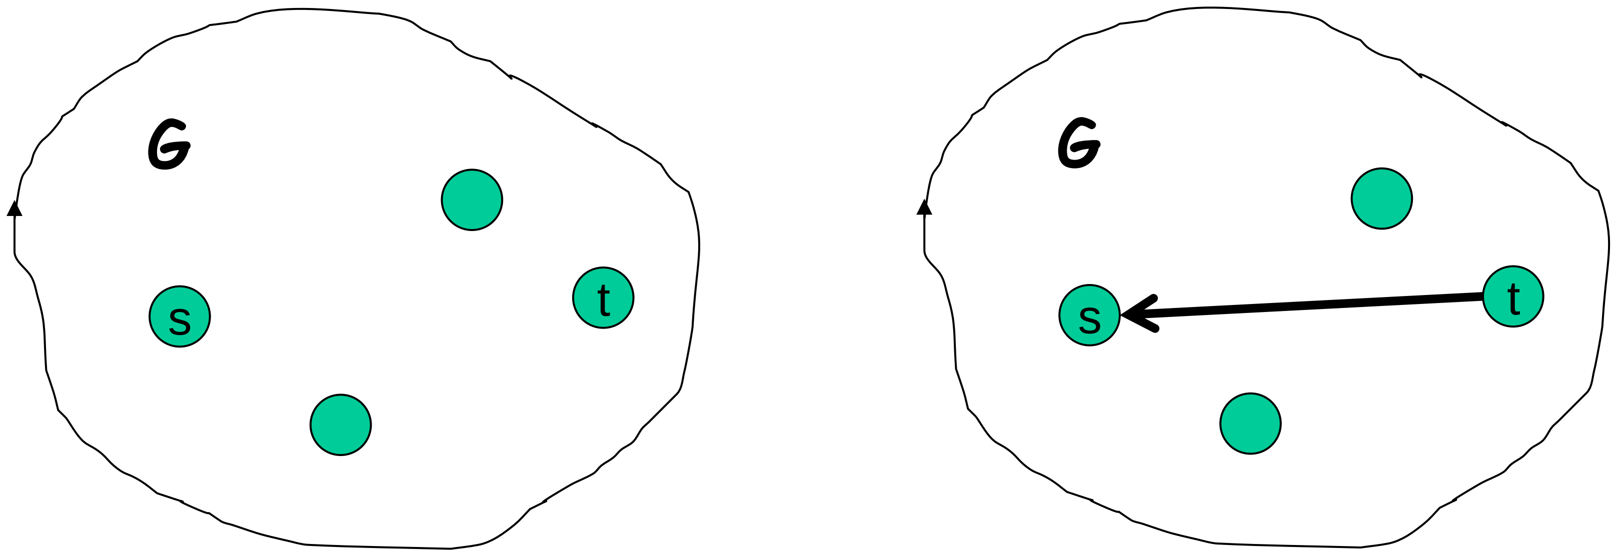

$$u(t,s) = mU$$

$$c(t,s) = -(C+1)n$$

More definitions

The residual capacity of a flow is a function r on the edges such that

$r(v,w) = u(v,w) - f(v,w)$

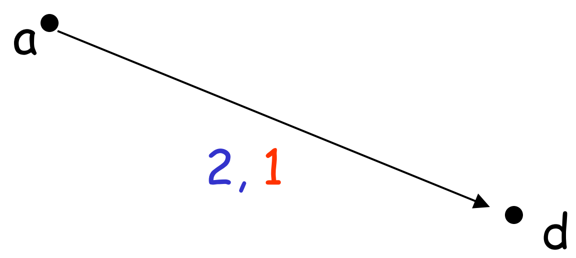

Interpretation: We can push r(v,w) more flow from v to w by increasing f(v,w) and decreasing f(w,v)

Optimality Criteria 1

f is MCC ⇔ The residual network does not have a negative cost cycle

Proof. ➔ Obvious

🡐 Suppose f is not optimal and let f^* be a circ. such that $cost(f^*) < cost(f)$

- f^*-f is a circulation in the residual of f
- $Cost(f^*-f) < 0$

So f^*-f contains a negative cost cycle, a contradiction ☐

Optimality Criteria 1

f is MCC ⇔ The residual network does not have a negative cost cycle

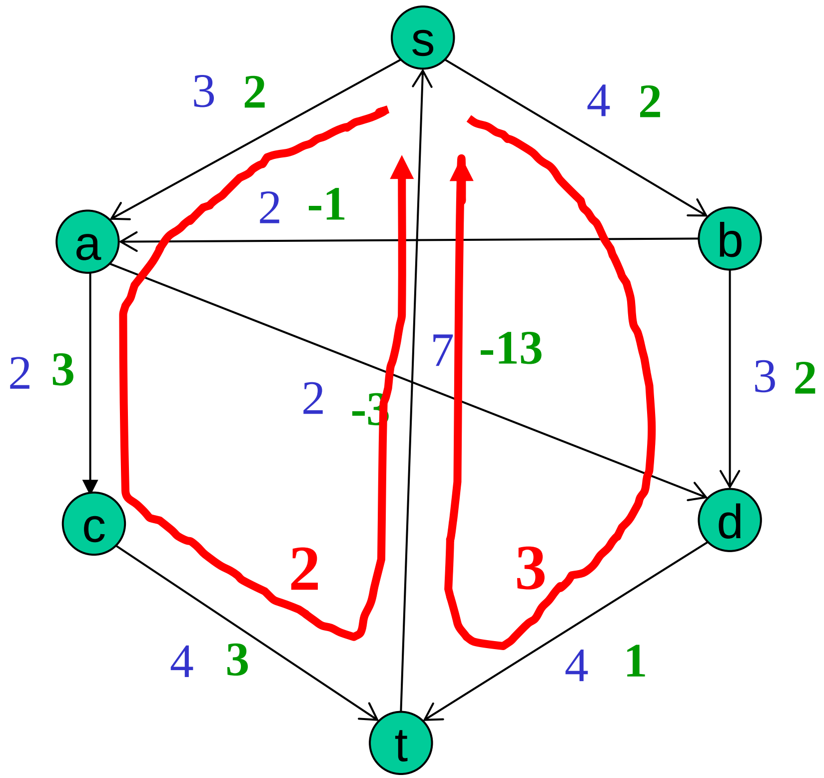

Optimality Criteria 1

f is MCC ⇔ The residual network does not have a negative cost cycle

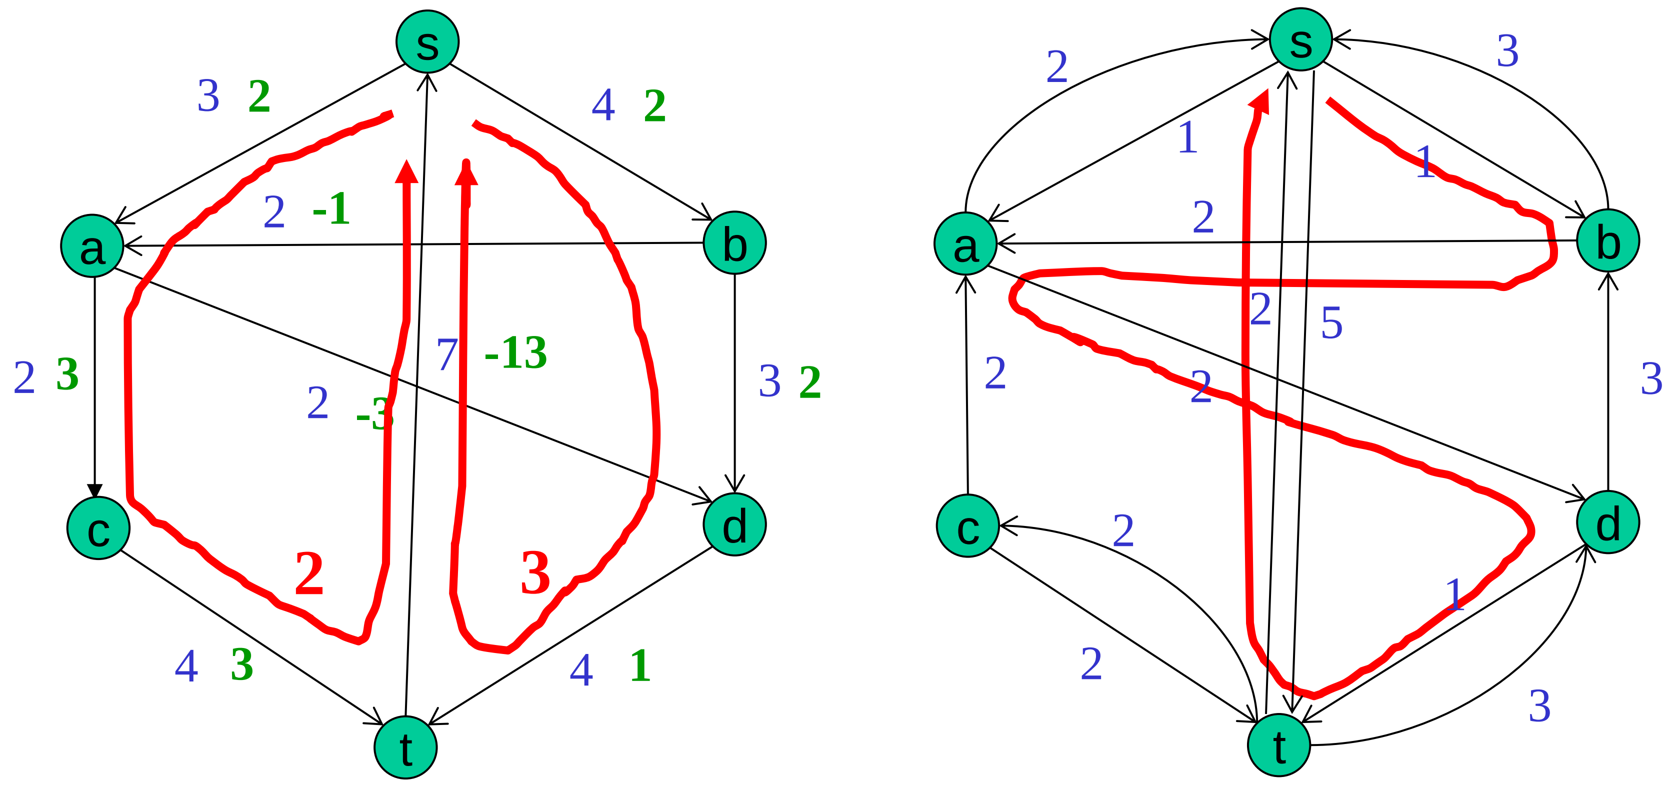

Optimality Criteria 1

f is MCC ⇔ The residual network does not have a negative cost cycle

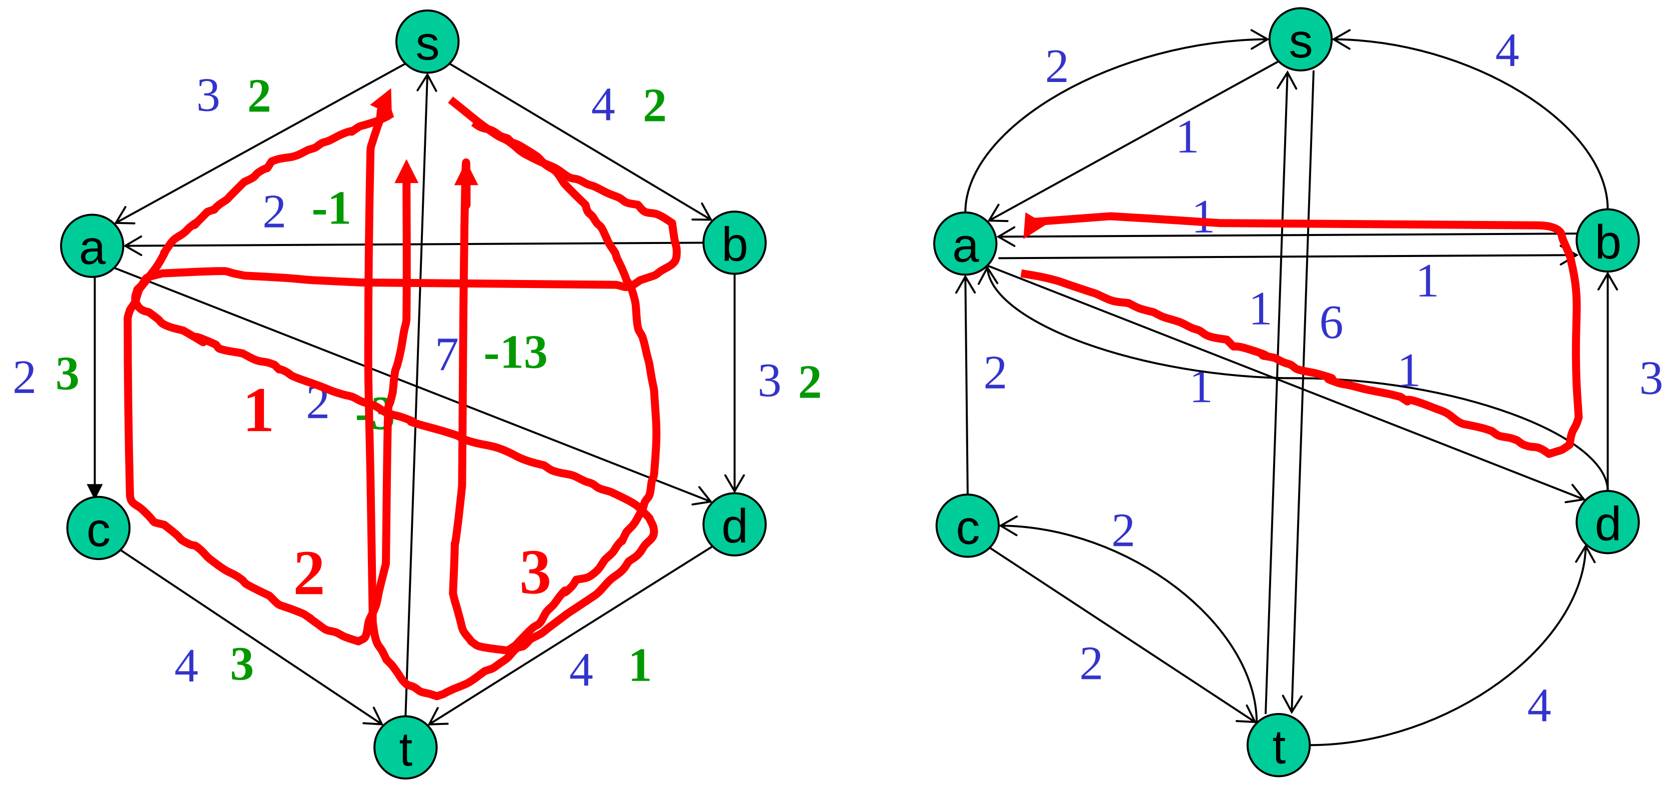

Optimality Criteria 1

f is MCC ⇔ The residual network does not have a negative cost cycle

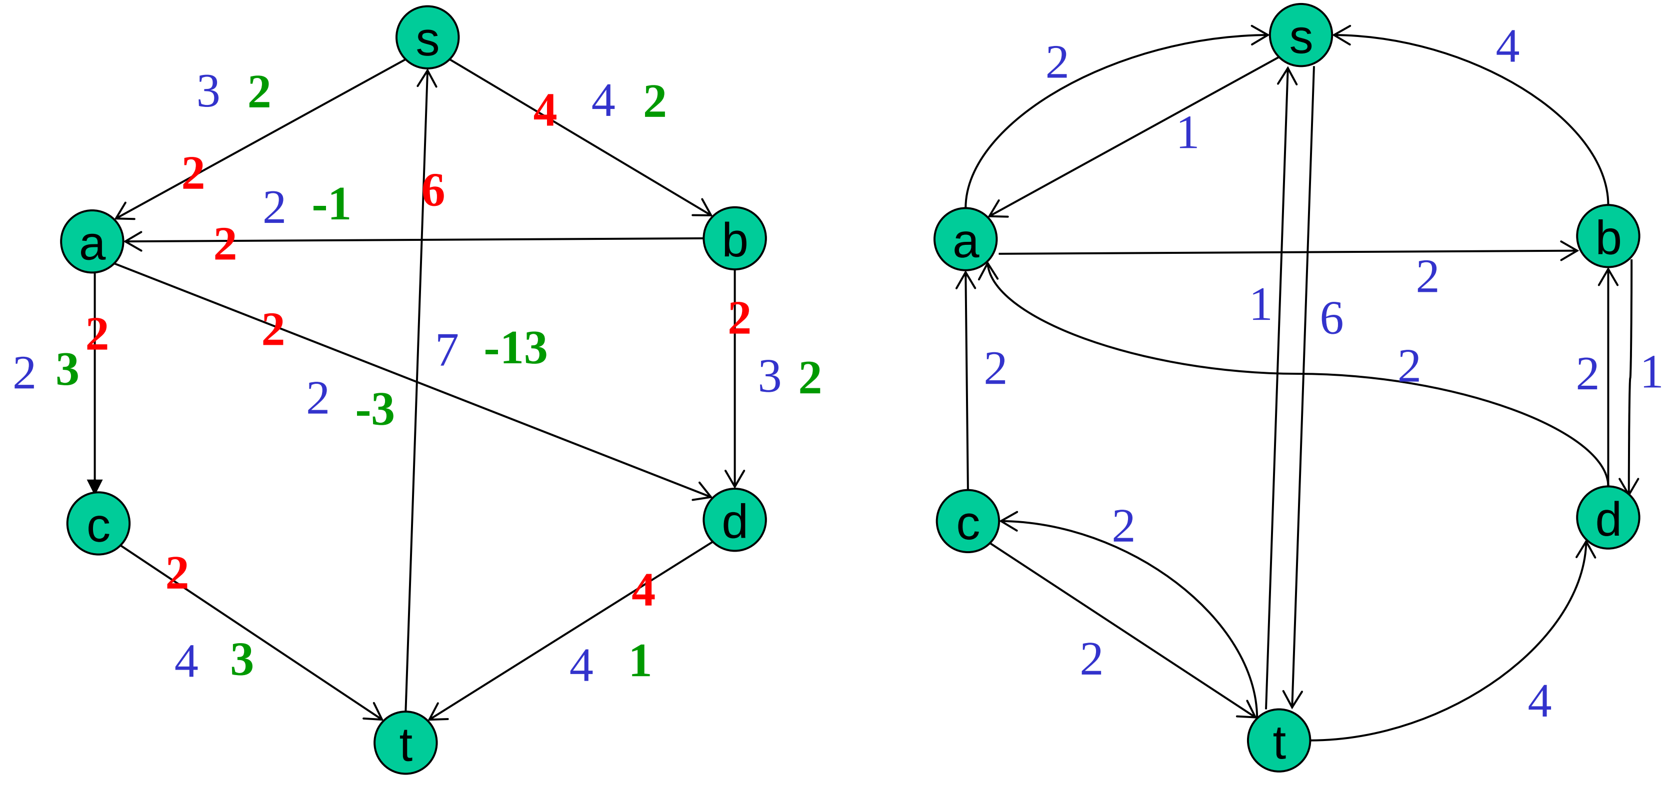

Cost = 4*2 + 2*2 – 2 + 2*2 - 2*3 + 2*3 + 2*3 + 4 - 13*6 = -54

Cycle Canceling Algorithm

Find a negative cost residual cycle

Saturate the cycle

How to find a negative cost residual cycle ?

Bellman-Ford shortest path algorithm ➔ O(mn) time

Cycle Canceling Algorithm (Analysis)

Assume costs and capacities are integers:

$U(e) \le U$, $|c(e)| \le C$

How many iterations can we get ?

mUC , since the cost cannot be smaller than -mUC

➔ $O(nm^2UC)$

Cycle Canceling Algorithm

Find a negative cost residual cycle

Saturate the cycle

If we start with integral capacities then residual capacities remain integral

In particular the minimum cost circulation is integral (we actually know that from the equivalence to min cost max flow)

Node Potentials

- Analogous to the distance labels, d(v) we used for maximum flow

Distance labels (reminder)

- $d(t) = 0$, $d(s) = n$
- $d(v) \le d(w) + 1$

$\Leftrightarrow 1 + d(w) - d(v) \ge 0$

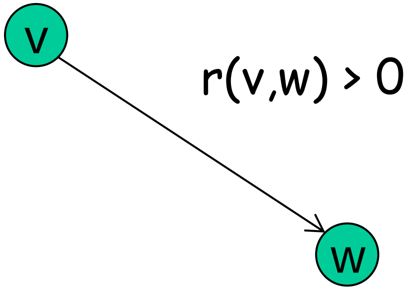

Node Potentials

- $\pi(v)$, mapping nodes to reals
- Intuition: Analogously to distance labels we would like to have: $c(e) + \pi(w) - \pi(v) \geq 0$ for residual arcs

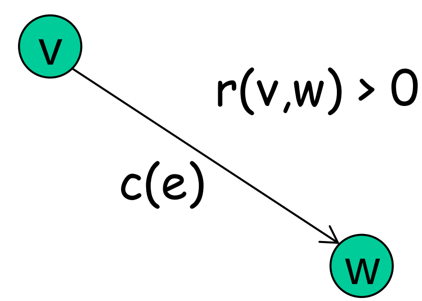

Reduce Costs

We define $c^\pi(v,w) = c(v,w) + \pi(w) - \pi(v)$

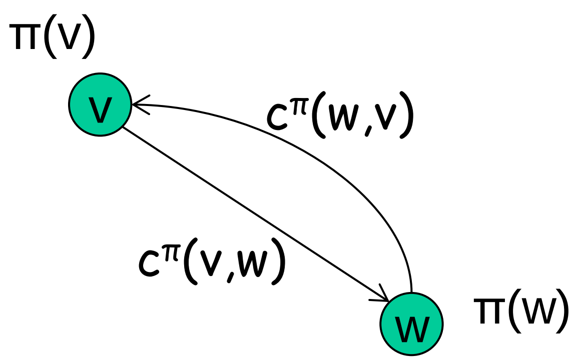

Note that $c^\pi(w,v) = -c^\pi(v,w)$

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

Proof. ➔ Pick an arbitrary vertex t and let $\delta(v,t)$ be the length of the shortest path from v to t in the residual of f.

These distances are well defined since we may assume that the residual is strongly connected and by optimality criteria 1 there are no negative cycles

Set $\pi(v) = \delta(v,t)$

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

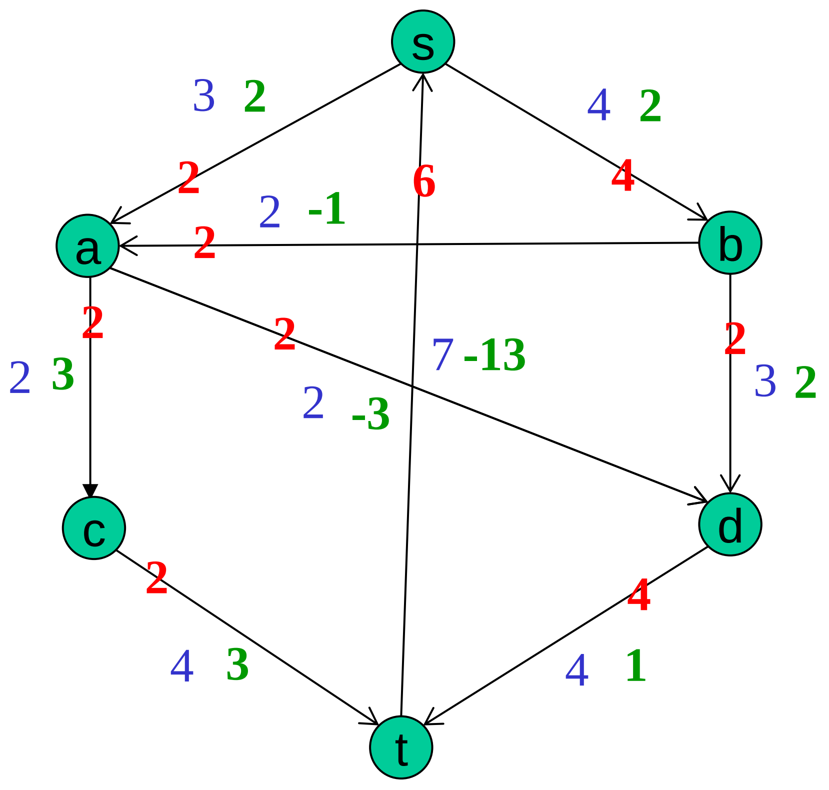

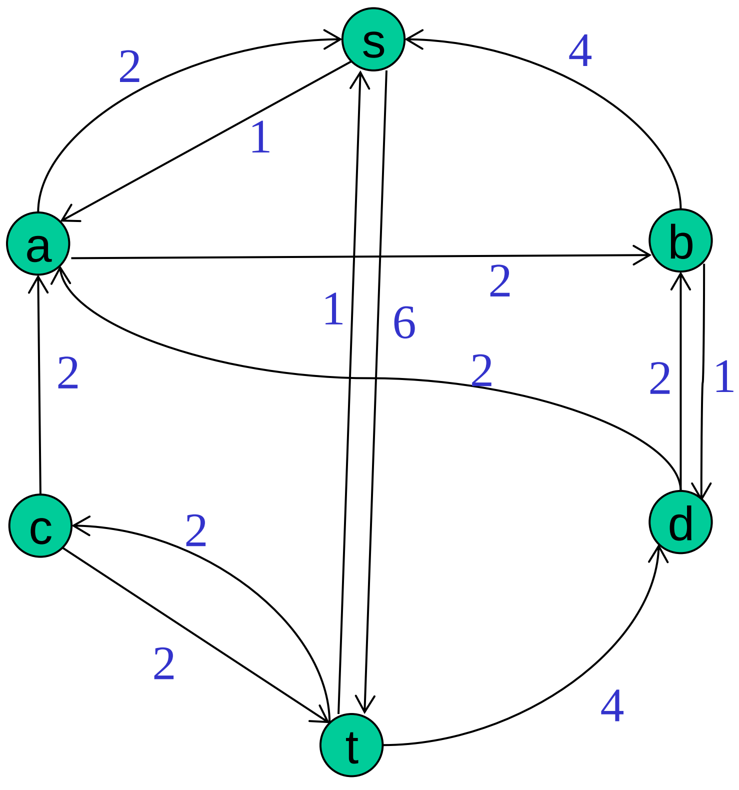

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

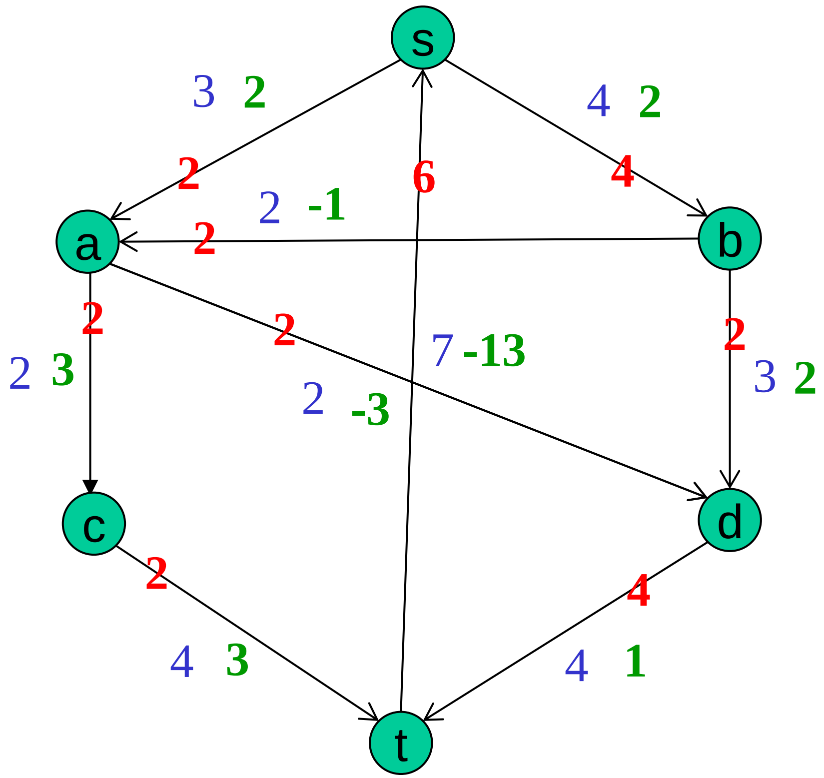

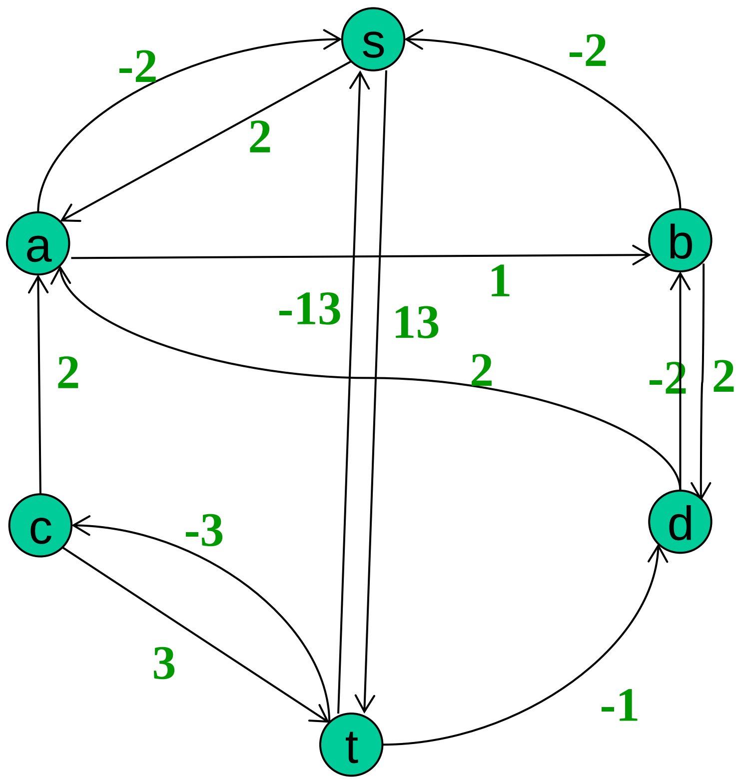

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

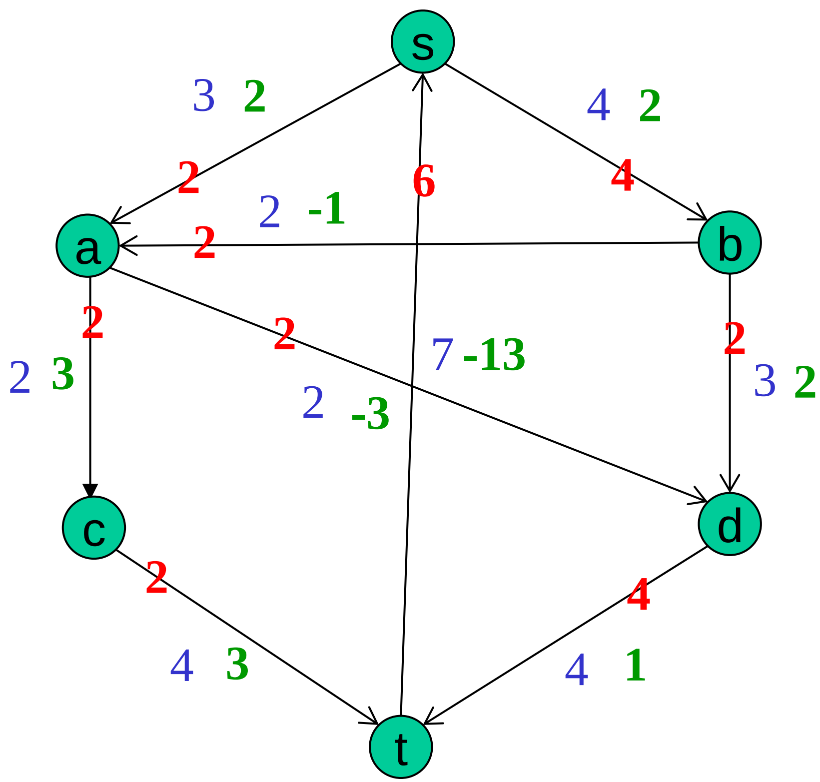

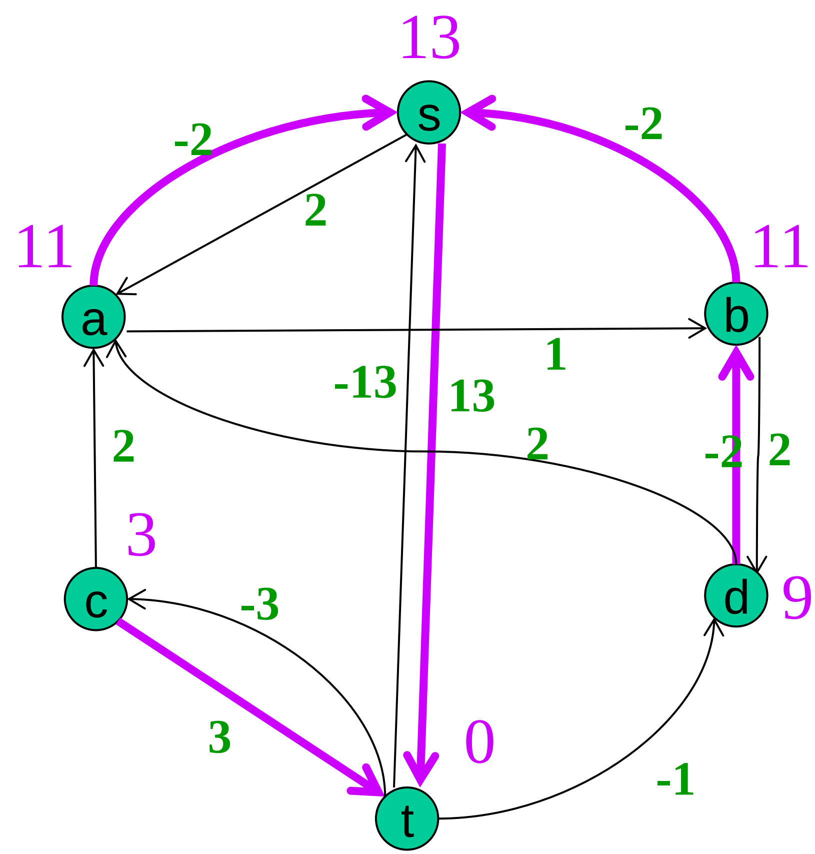

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

Proof. ← Assume there exists potential π such that $c^{\pi}(e) \geq 0$ for every residual arc e

Then the reduced cost of every residual cycle is ≥ 0

The reduced cost of a cycle is the same as its cost and therefore the claim follows from optimality criteria 1

Optimality Criteria 2

f is MCC ⇔ There exists a potential function π such that $c^{\pi}(e) \geq 0$ for every residual arc e

$$c^{\pi}(a,b) = c(a,b) + \pi(b) - \pi(a)$$

$$+\quad c^{\pi}(d,a) = c(d,a) + \pi(a) - \pi(d)$$

$$c^{\pi}(b,d) = c(b,d) + \pi(d) - \pi(b)$$

$$c^{\pi}(C) = c(C) \geq 0$$

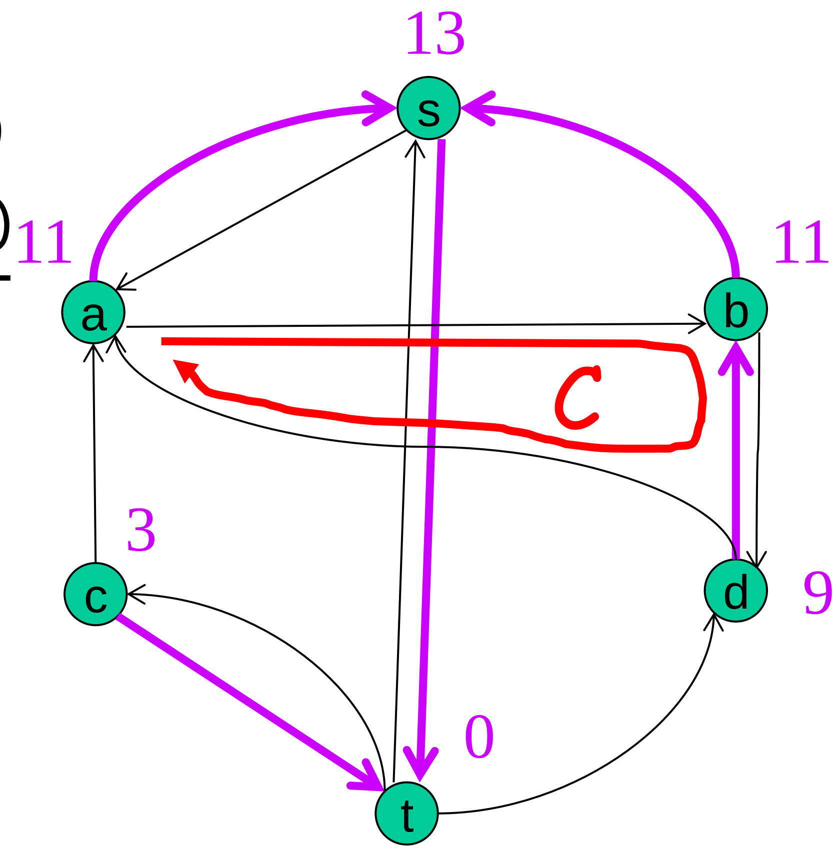

□

The approach

Start with potential $\pi(v) = 0$

Saturate each edge with negative reduced cost

Now you have a pseudo-flow: Some nodes u have excess $e_f(u) > 0$, some deficit $e_f(u) < 0$

Ship the flow from nodes with excess to nodes with deficit

Example

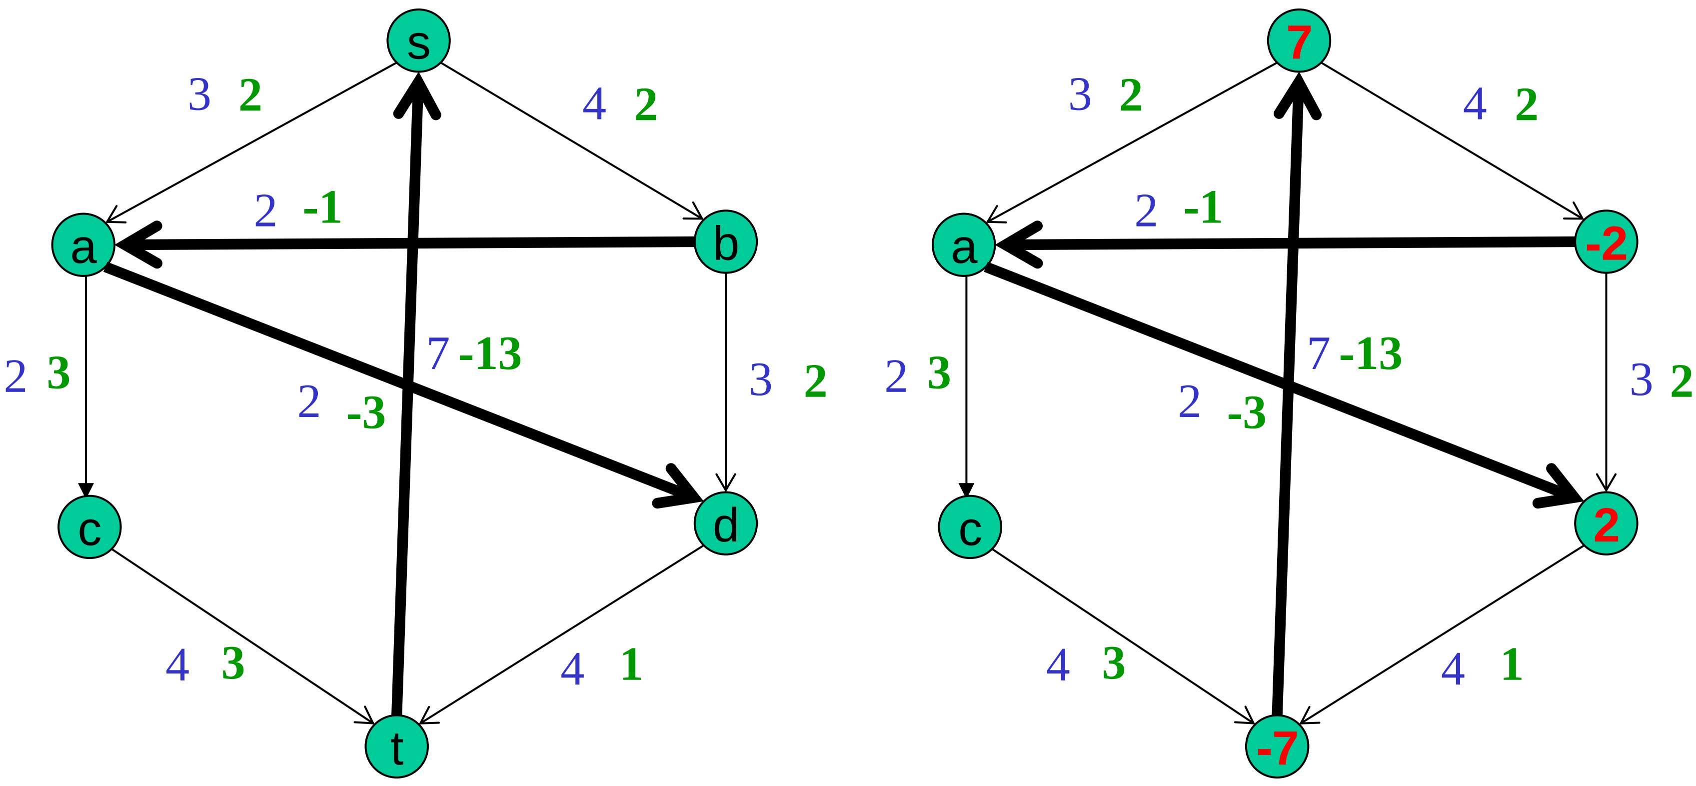

Example

Ship the flow from nodes with excess to node with deficit

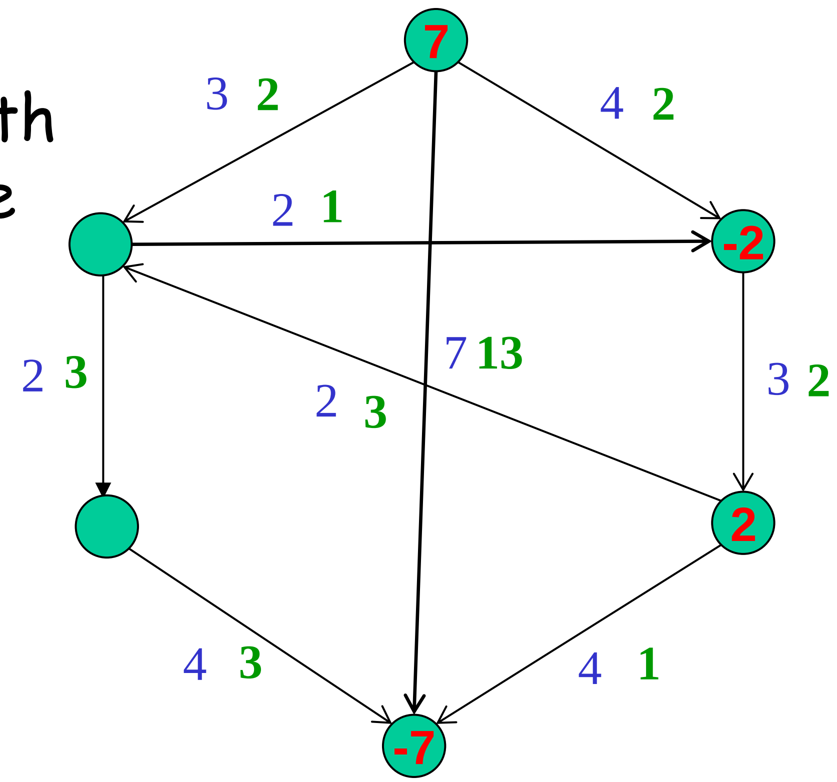

But if we push flow on an edge (v,w) with $c^{\pi}(v,w) > 0$, (w,v) become residual and $c^{\pi}(w,v) < 0$

Example

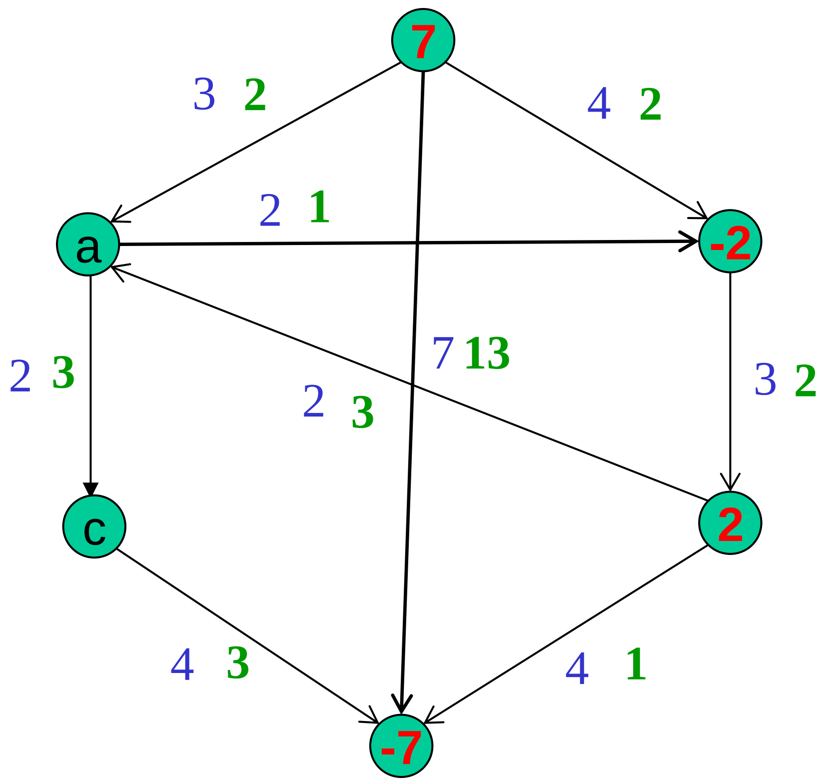

Push flow on an edge (v,w) with $c^{\pi}(v,w) = 0$?

Approximate Optimality

f is ε-optimal ⇔ There exists a potential function π such that $c^{\pi}(e) \geq -\varepsilon$ for every residual arc e

We use both approx. opt flow and approx. opt pseudo flows

Approximate Optimality and Optimality

If all costs are integers then f is ε-optimal for $\varepsilon < 1/n \Leftrightarrow$ f is MCC

Proof:

The Successive Approximation Alg

$\varepsilon \leftarrow C$

$\forall v, \pi(v) \leftarrow 0$

$\forall e, f(e) \leftarrow 0$

while ($\varepsilon \geq 1/n$)
  $\varepsilon \leftarrow \varepsilon/2$
  refine (f, π)

return(f)

Admissible arcs

An arc (v,w) is admissible if it is residual $u_f(v,w) > 0$ and its reduced cost $c_\pi(v,w) < 0$

Refine

$\forall (v,w) \in E$ do if $c^{\pi}(v,w) < 0$ then $f(v,w) \leftarrow u(v,w)$

/* admissible subgraph is empty, there are excesses $A = \{v \mid e_f(u) > 0\}$ and deficits $D = \{v \mid e_f(u) < 0\}$ */

While f is not a circulation do {

  $S \leftarrow$ Collect all vertices v reachable from A in the admissible subgraph

  $\forall v \in S$, $\pi(v) \leftarrow \pi(v) + \varepsilon$

  Find a blocking flow b from A to D

  $f \leftarrow f + b$

}

Find a blocking flow b from A to D

By finding a blocking flow from s to t in the following network

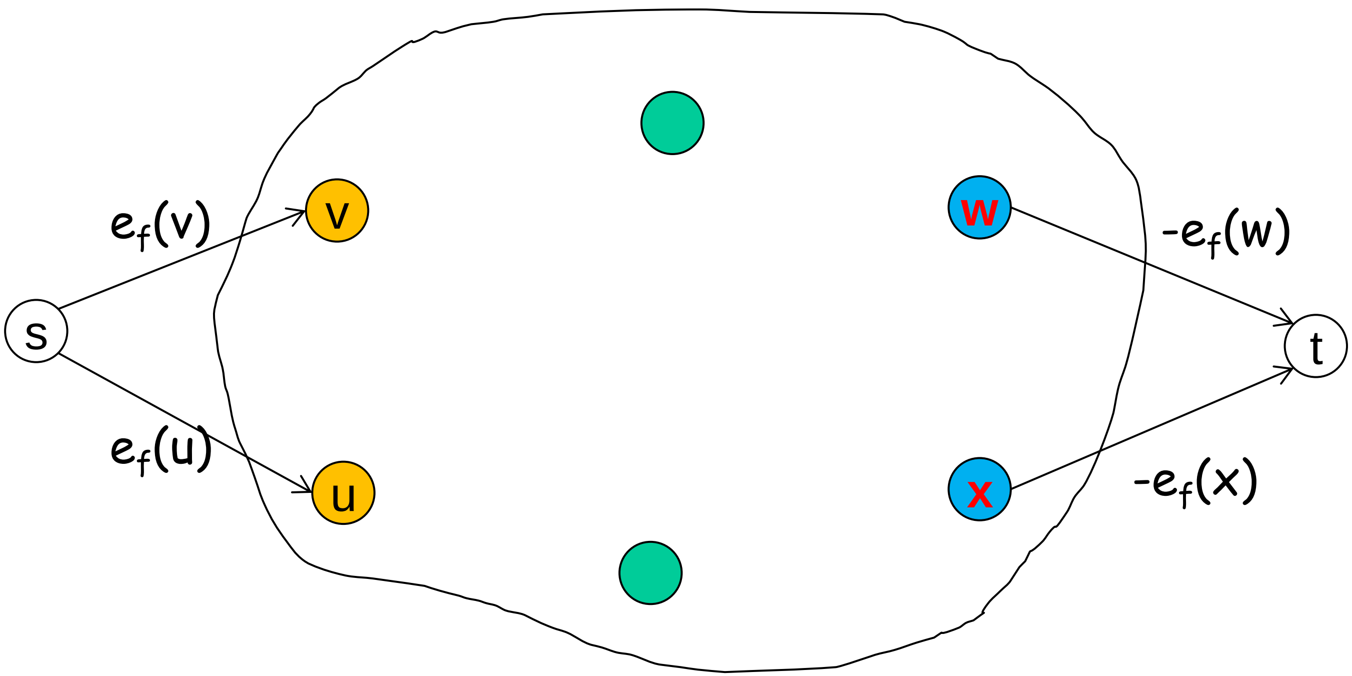

Example

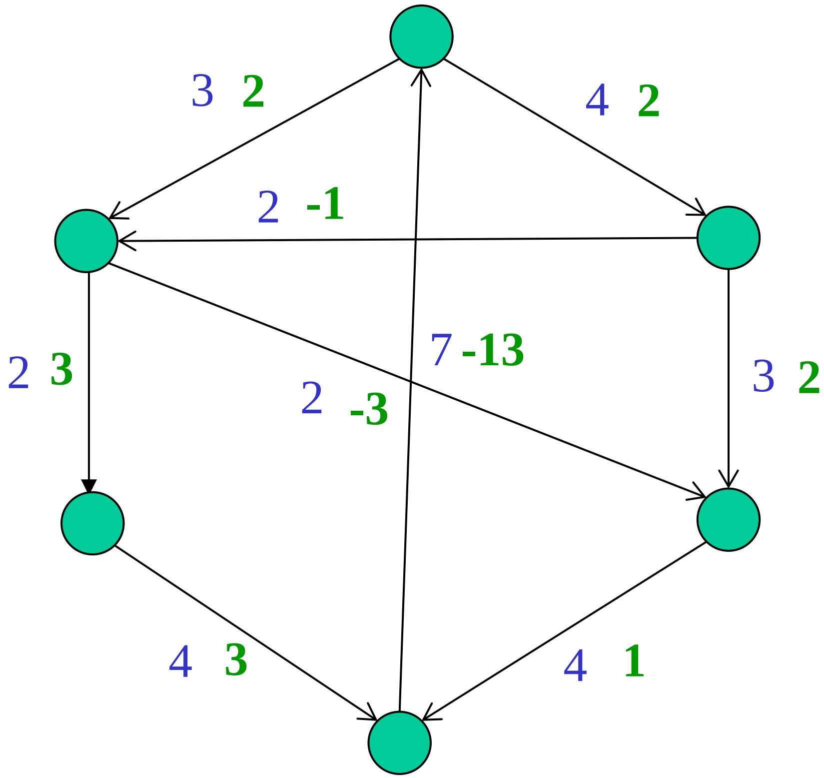

Initialize:
$\varepsilon \leftarrow C = 13$

The 0 circulation and potentials are ε optimal

First round, $\varepsilon \leftarrow \varepsilon/2 = 6.5$

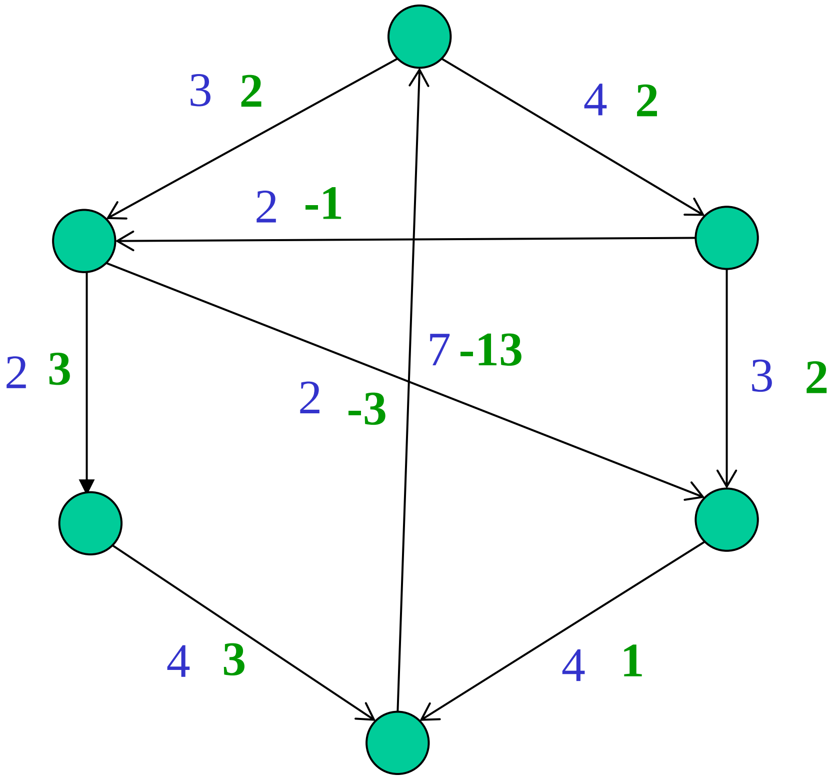

Saturate arcs with negative reduced cost

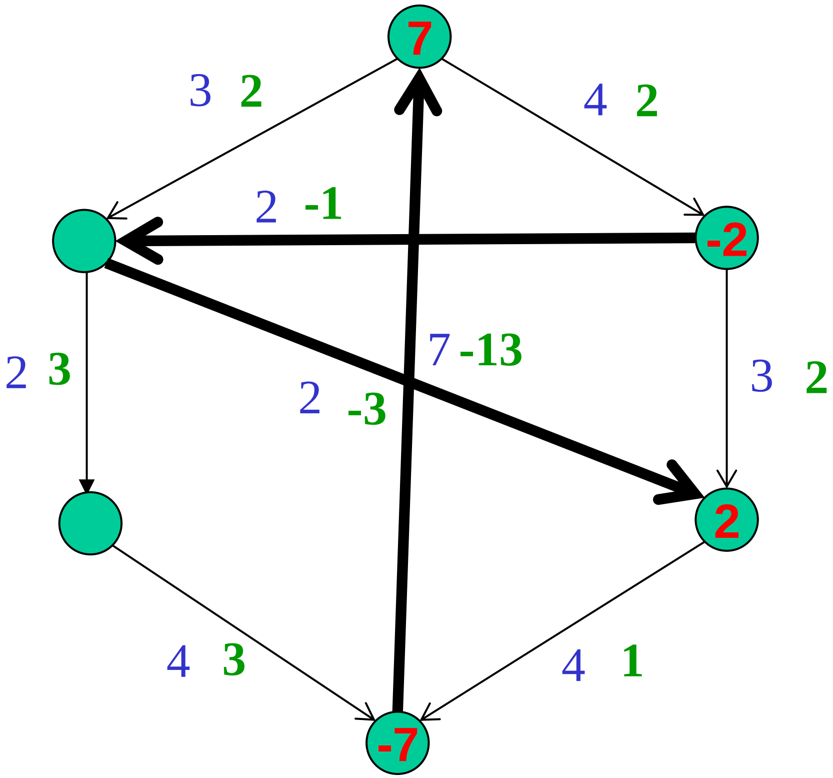

Ship the excesses into deficits

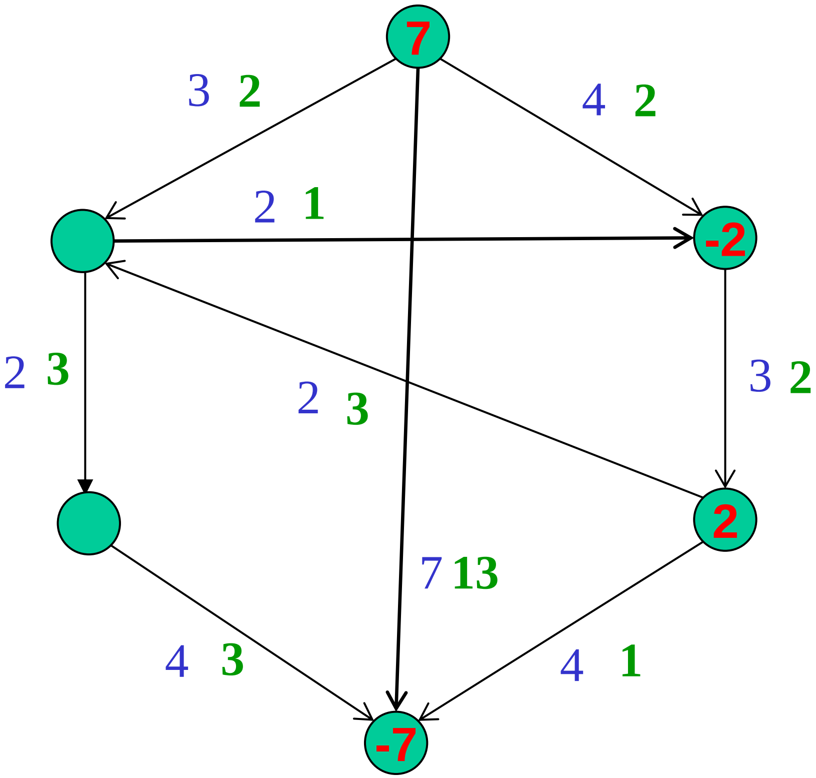

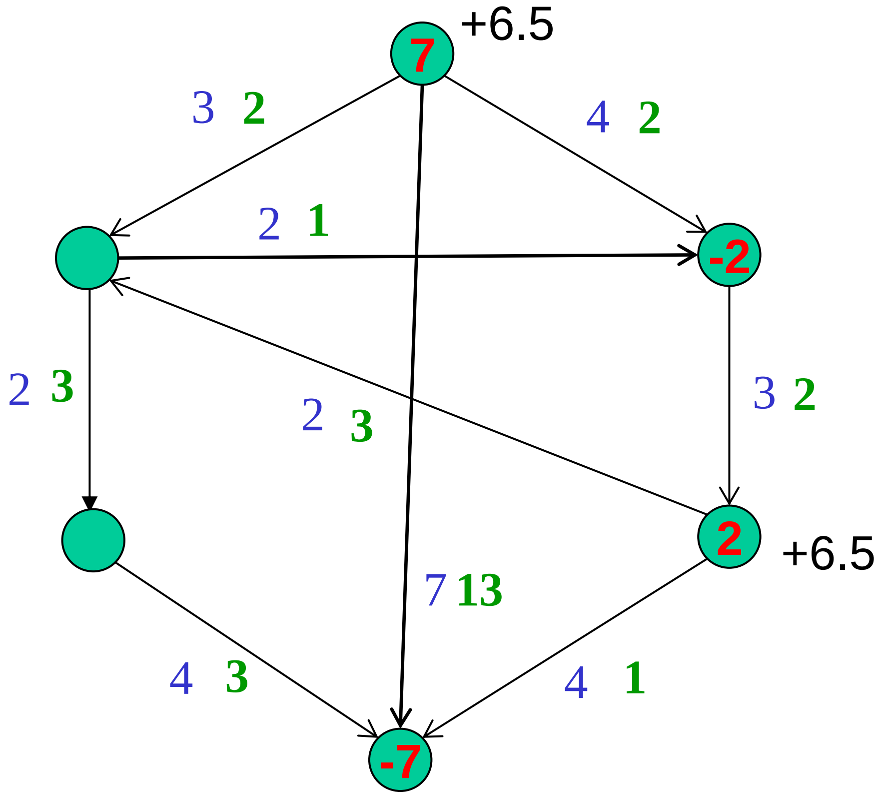

+6.5
7
3 2
4 2
2 1
-2
2 3
2 3
3 2
2
+6.5
7 13
4 3
4 1
-7

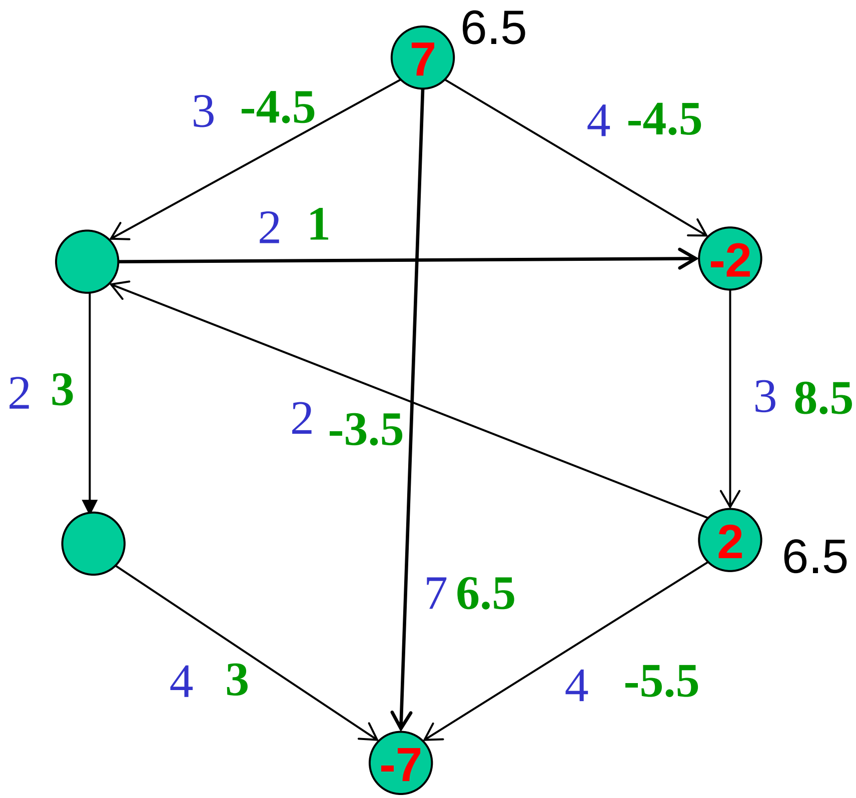
6.5
7
3 -4.5
4 -4.5
2 1
-2
2 3
2 -3.5
3 8.5
2
6.5
7 6.5
4 3
4 -5.5
-7

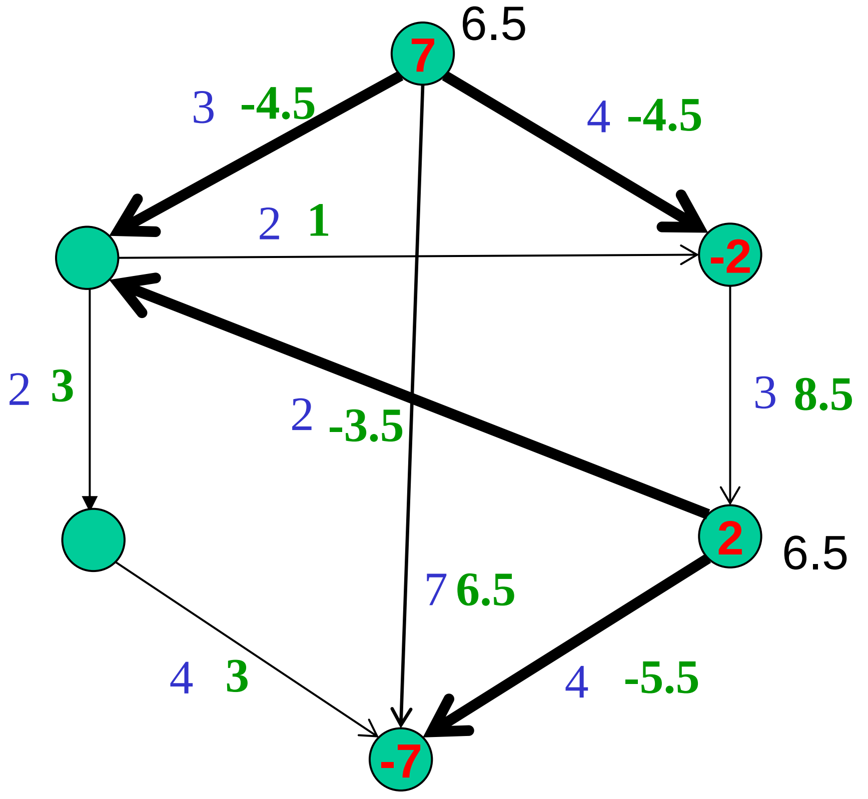

6.5
7
3 -4.5
4 -4.5
2 1
-2
2 3
3 8.5
2 -3.5
2
6.5
7 6.5
4 3
4 -5.5
-7

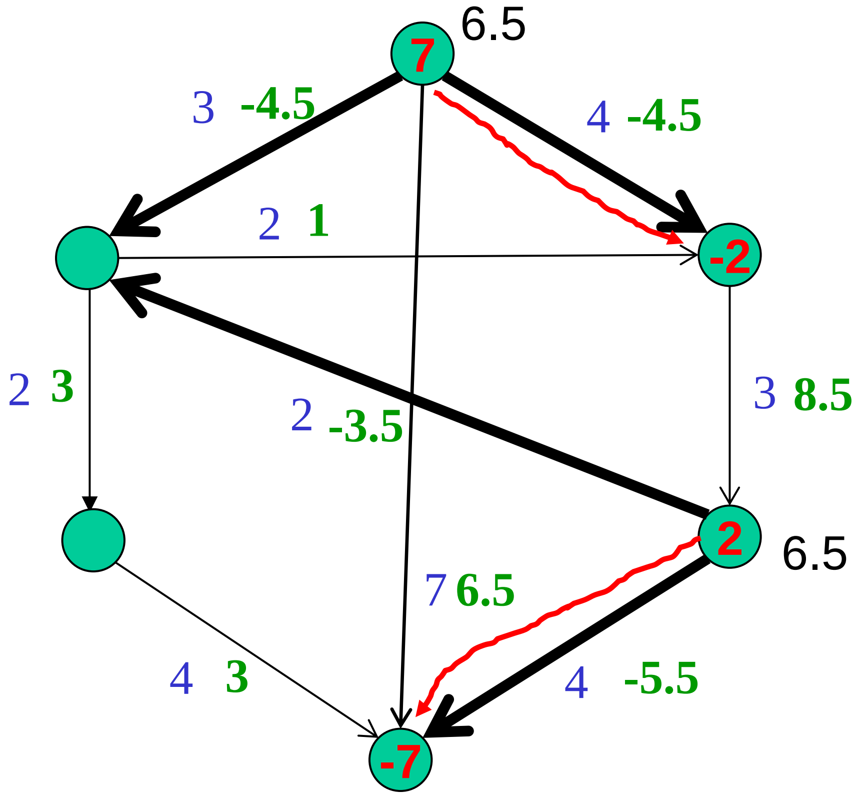

6.5
7
3 -4.5
4 -4.5
2 1
-2
2 3
3 8.5
2 -3.5
2
6.5
7 6.5
4 3
4 -5.5
-7

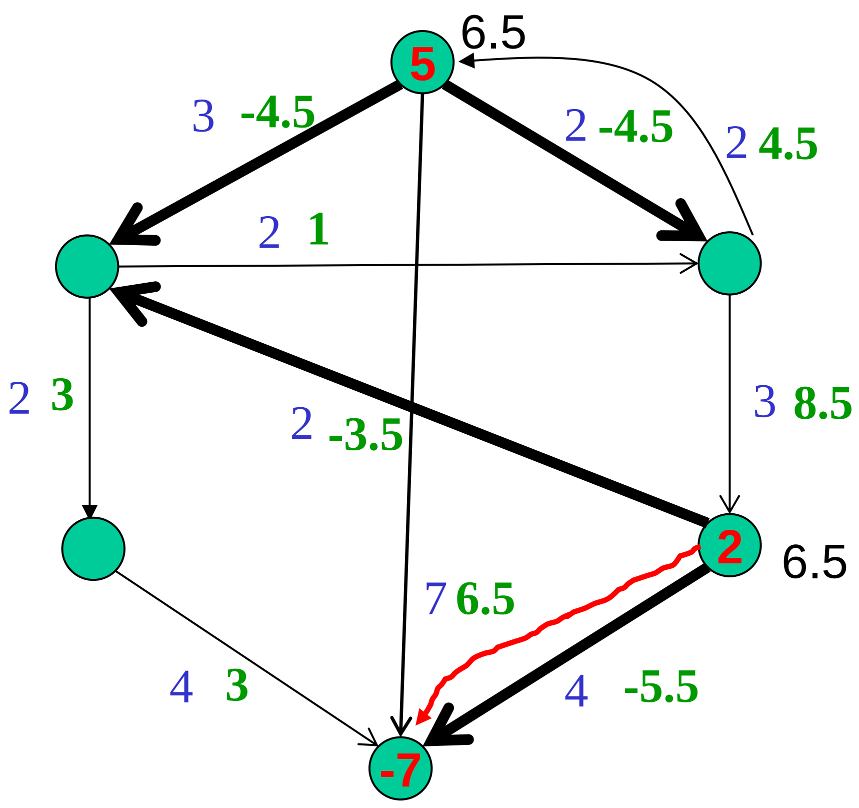

5
6.5
3 -4.5
2 -4.5
2 4.5
2 1
2 3
2 -3.5
3 8.5
2
6.5
7 6.5
4 3
4 -5.5
-7

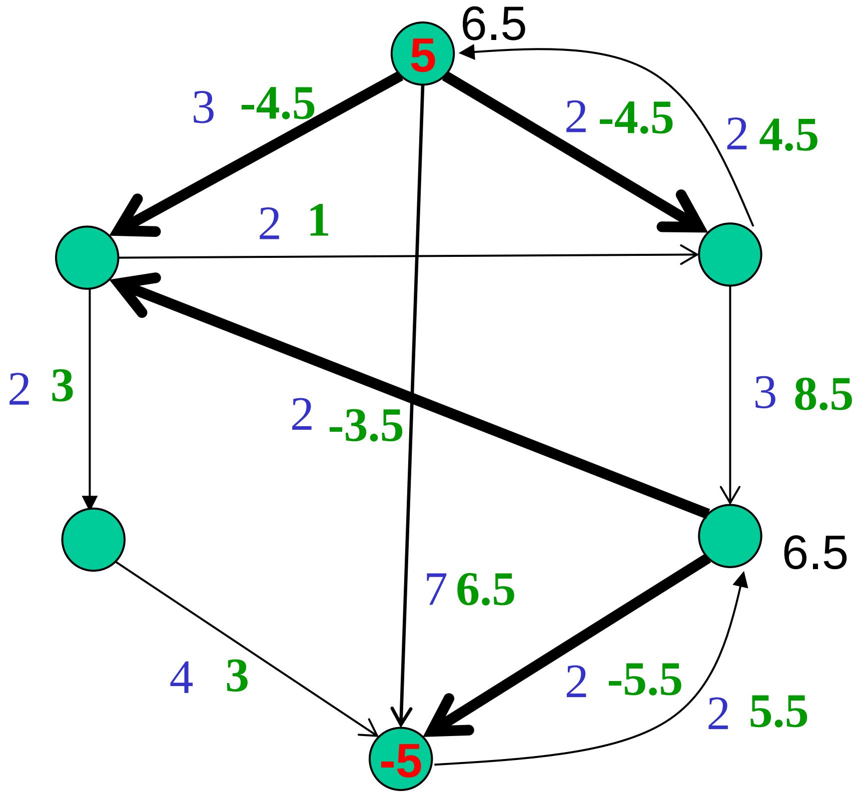

5
6.5
3 -4.5
2 -4.5
2 4.5
2 1
2 3
2 -3.5
3 8.5
6.5
7 6.5
4 3
2 -5.5
2 5.5
-5

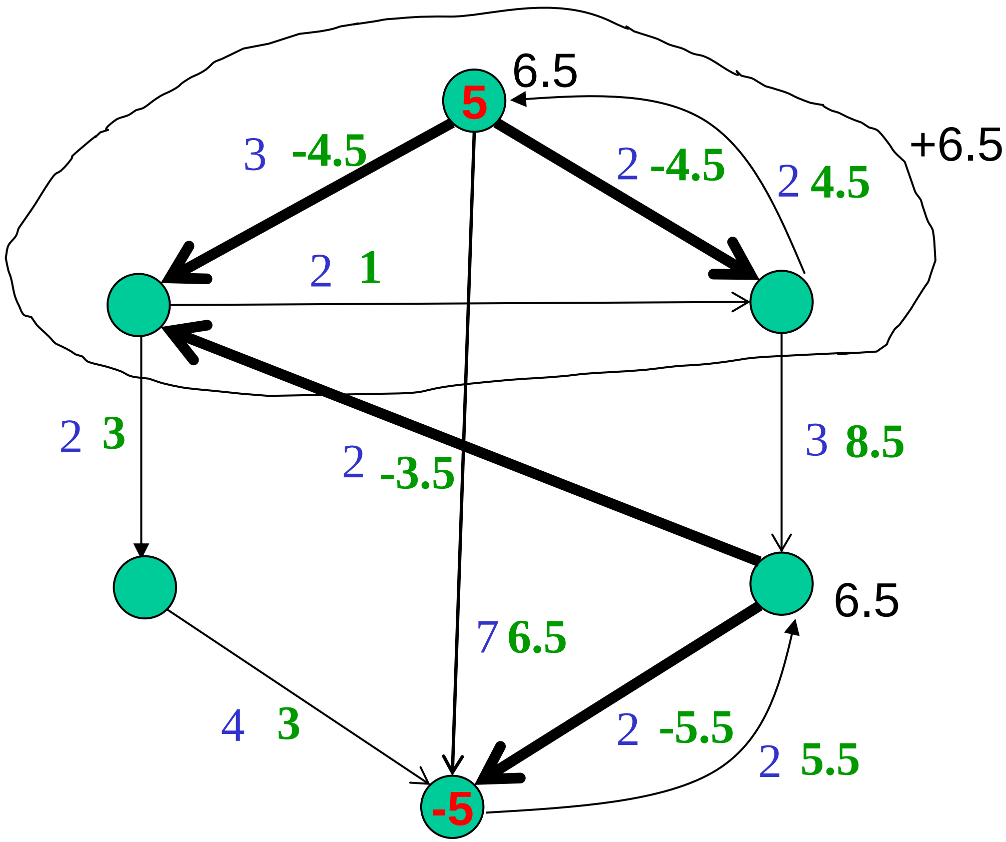
5
6.5
3 -4.5
2 -4.5
2 4.5
+6.5
2 1
2 3
2 -3.5
3 8.5
6.5
7 6.5
4 3
2 -5.5
2 5.5
-5

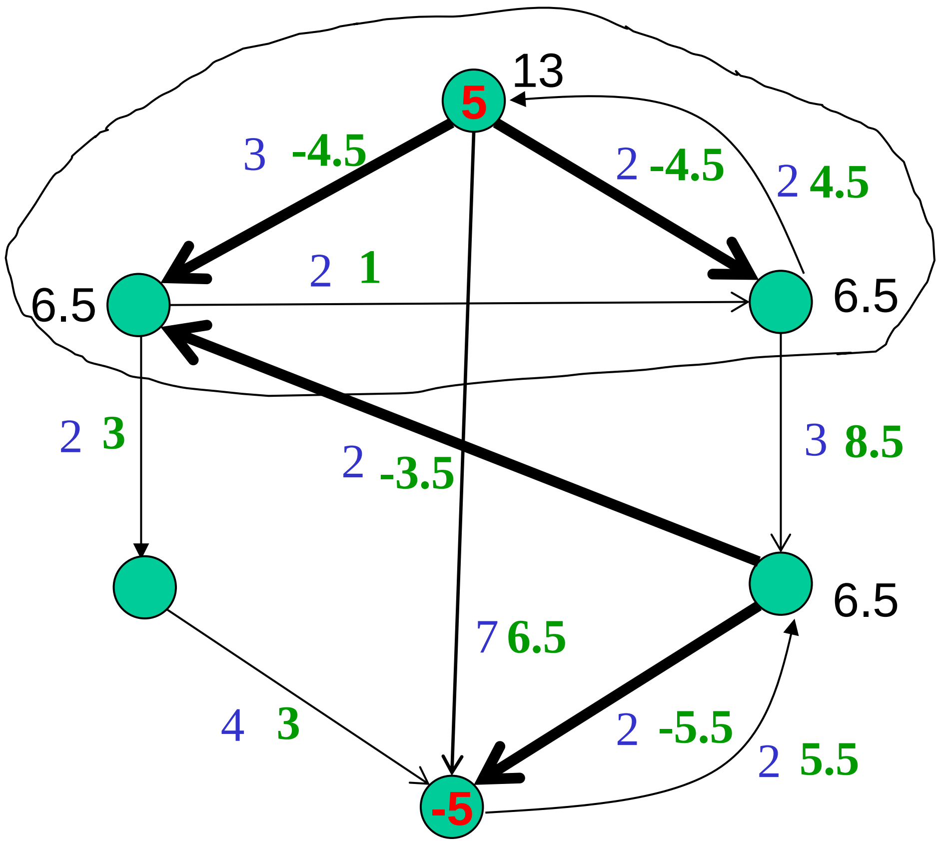
5
13
3 -4.5
2 -4.5
2 4.5
6.5
2 1
6.5
2 3
2 -3.5
3 8.5
6.5
7 6.5
4 3
2 -5.5
2 5.5
-5

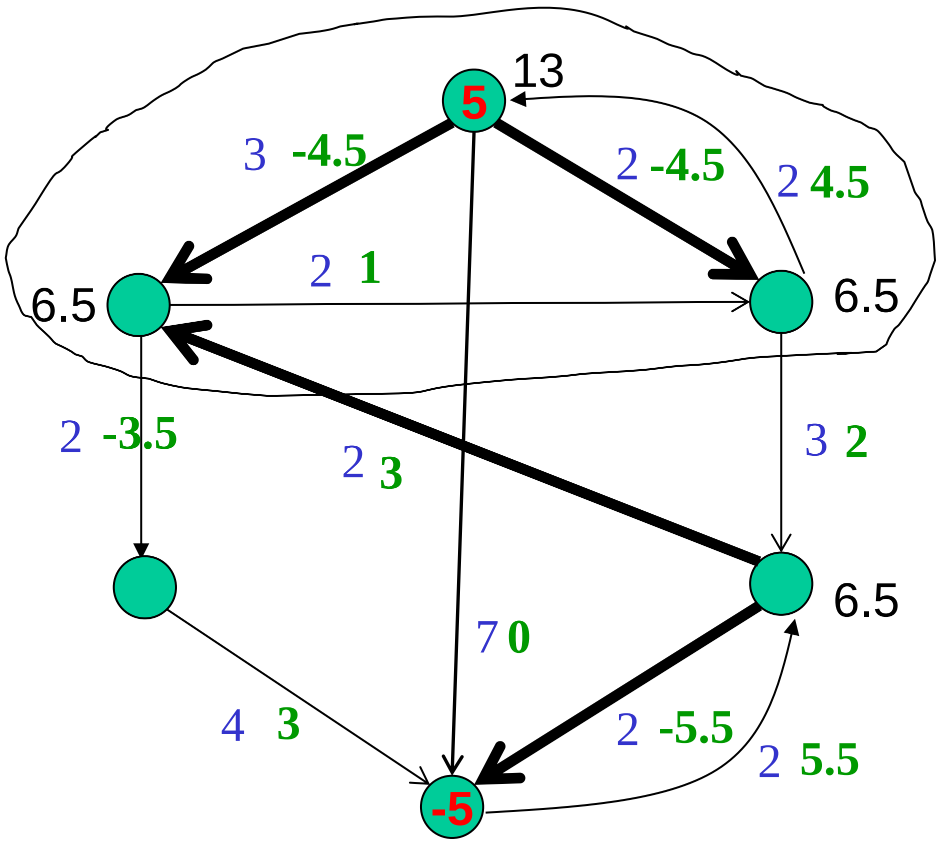

13
5
3 -4.5
2 -4.5
2 4.5
6.5
2 1
6.5
2 -3.5
2 3
3 2
6.5
7 0
4 3
2 -5.5
2 5.5
-5

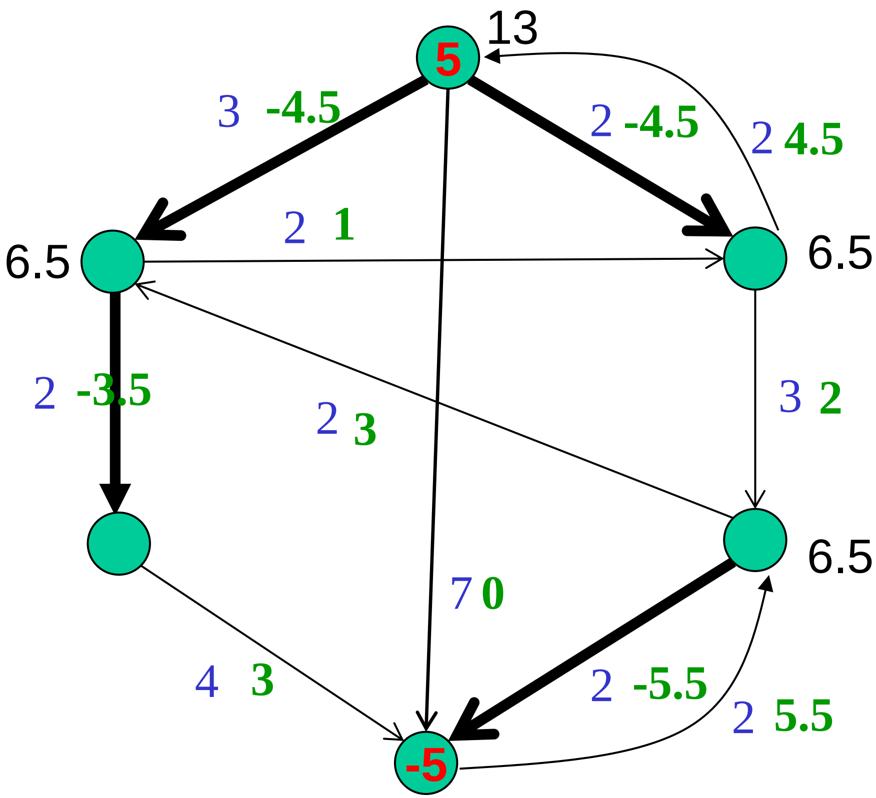

13
5
3 -4.5
2 -4.5
2 4.5
6.5
2 1
6.5
2 -3.5
2 3
3 2
6.5
7 0
4 3
2 -5.5
2 5.5
-5

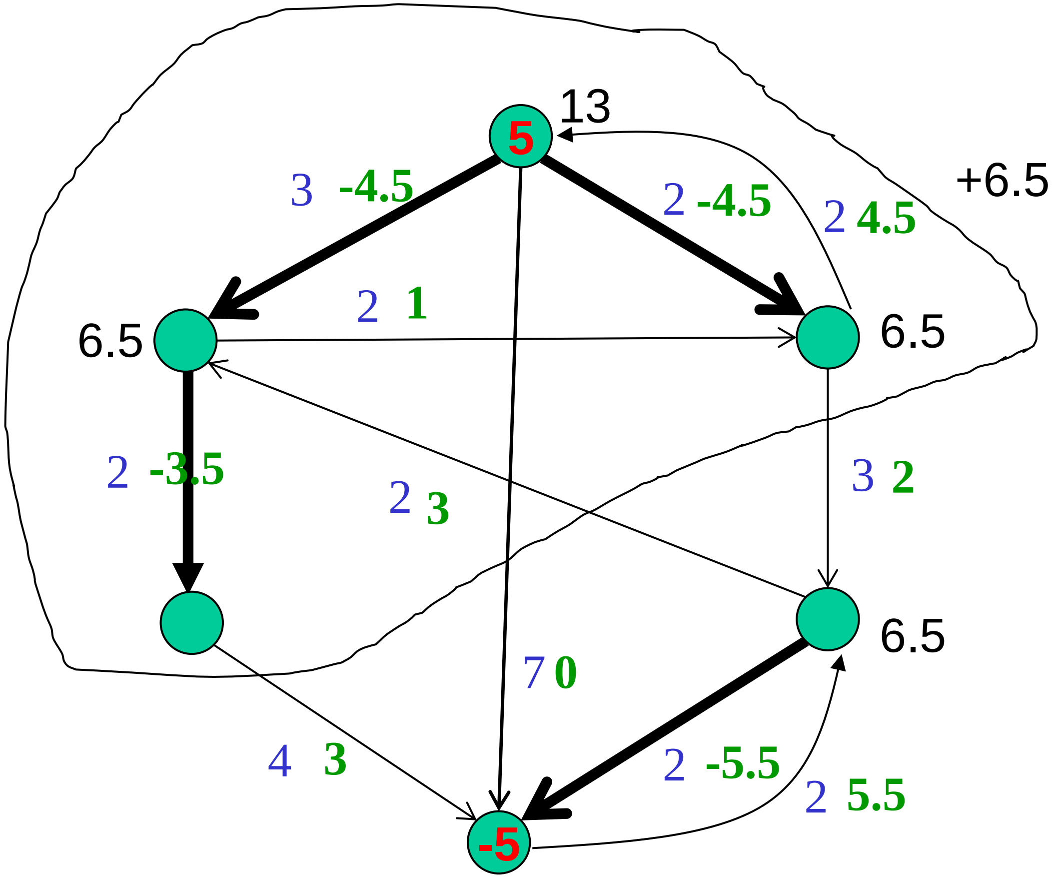

13
5
3 -4.5
2 -4.5
2 4.5
+6.5
6.5
2 1
6.5
2 -3.5
2 3
3 2
6.5
7 0
4 3
2 -5.5
2 5.5
-5

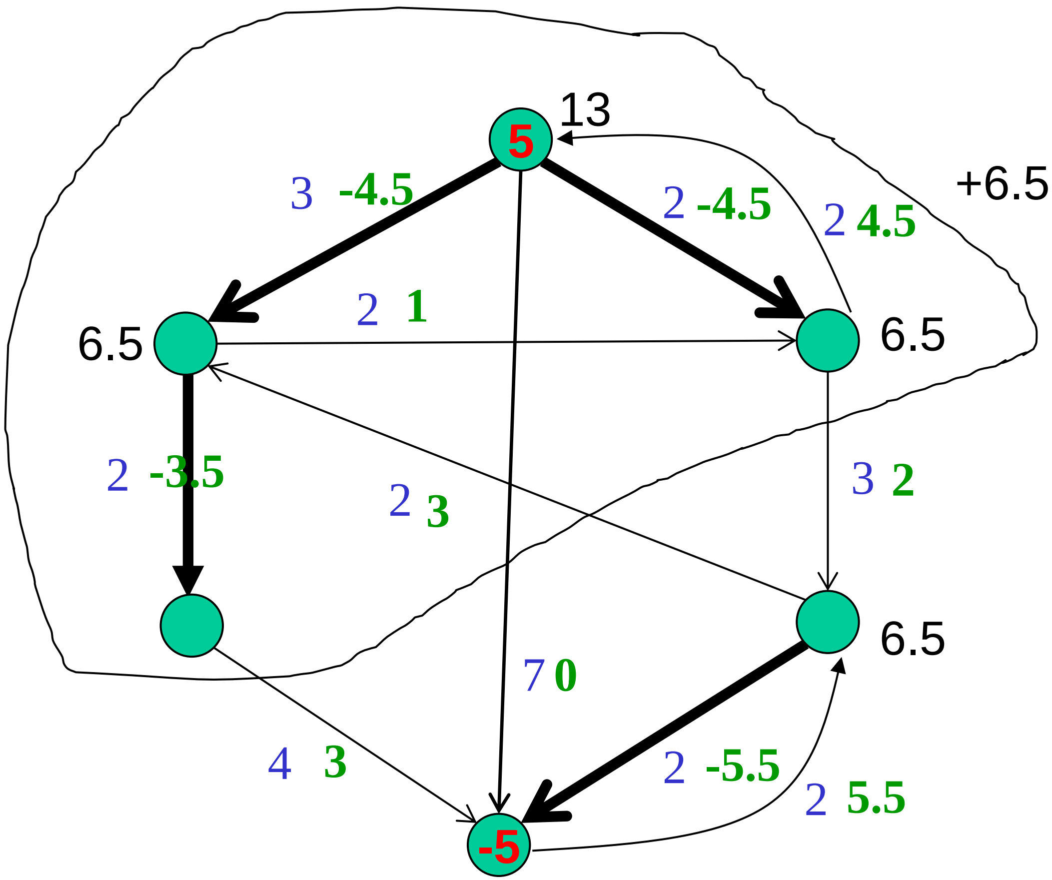
13
5
+6.5
3 -4.5
2 -4.5
2 4.5
2 1
6.5
6.5
2 -3.5
2 3
3 2
6.5
7 0
4 3
2 -5.5
2 5.5
-5

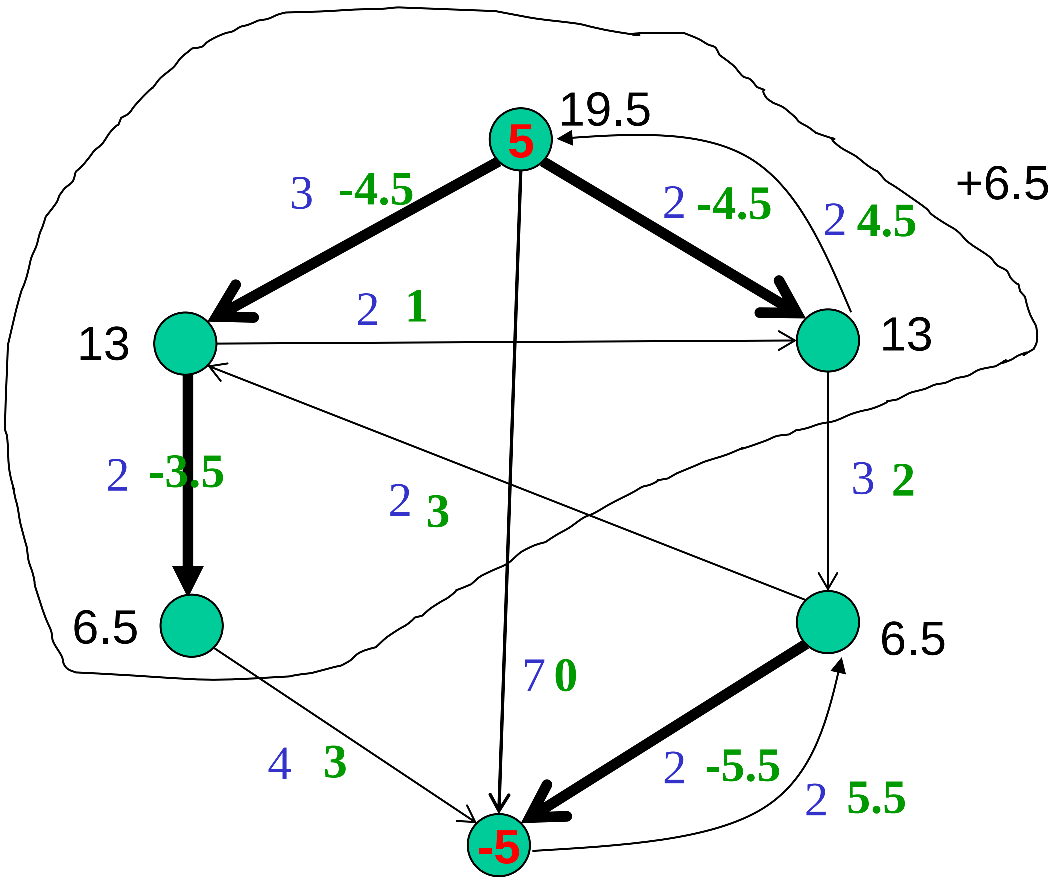

19.5
5
+6.5
3 -4.5
2 -4.5
2 4.5
2 1
13
13
2 -3.5
2 3
3 2
6.5
6.5
7 0
4 3
2 -5.5
2 5.5
-5

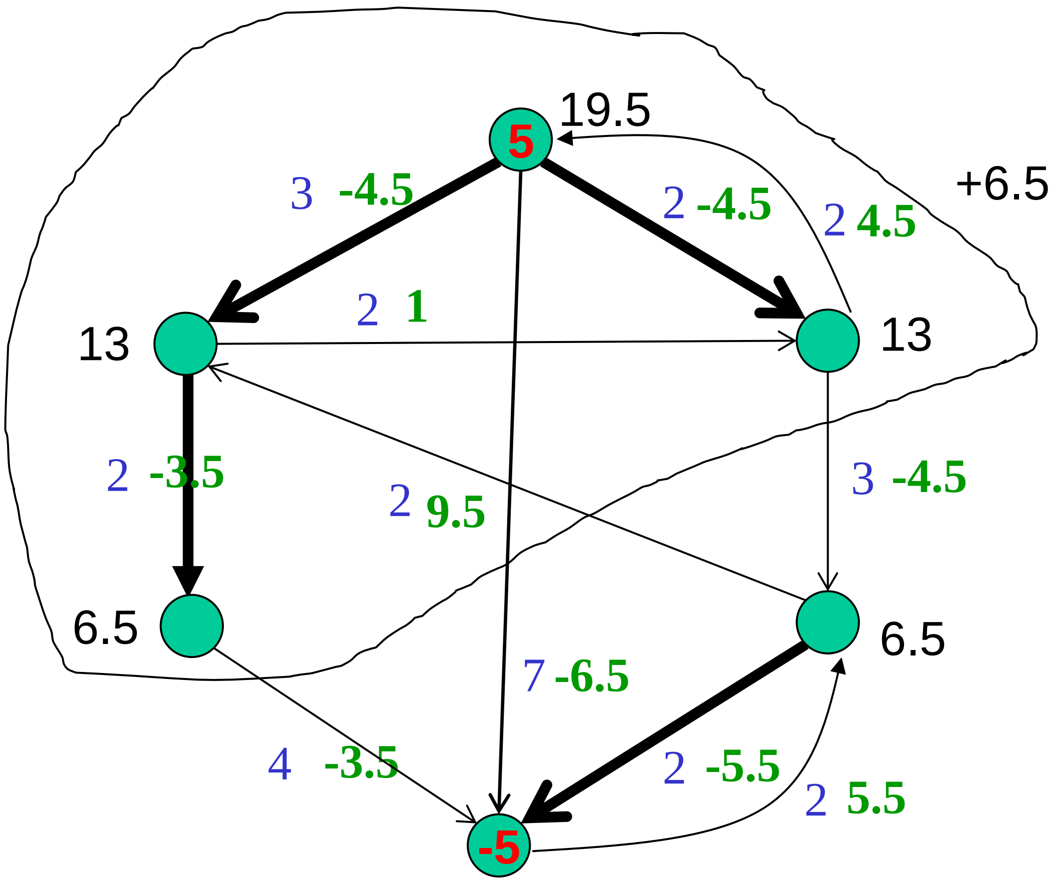

19.5
5
+6.5
3 -4.5
2 -4.5
2 4.5
13
2 1
13
2 -3.5
2 9.5
3 -4.5
6.5
6.5
7 -6.5
4 -3.5
2 -5.5
2 5.5
-5

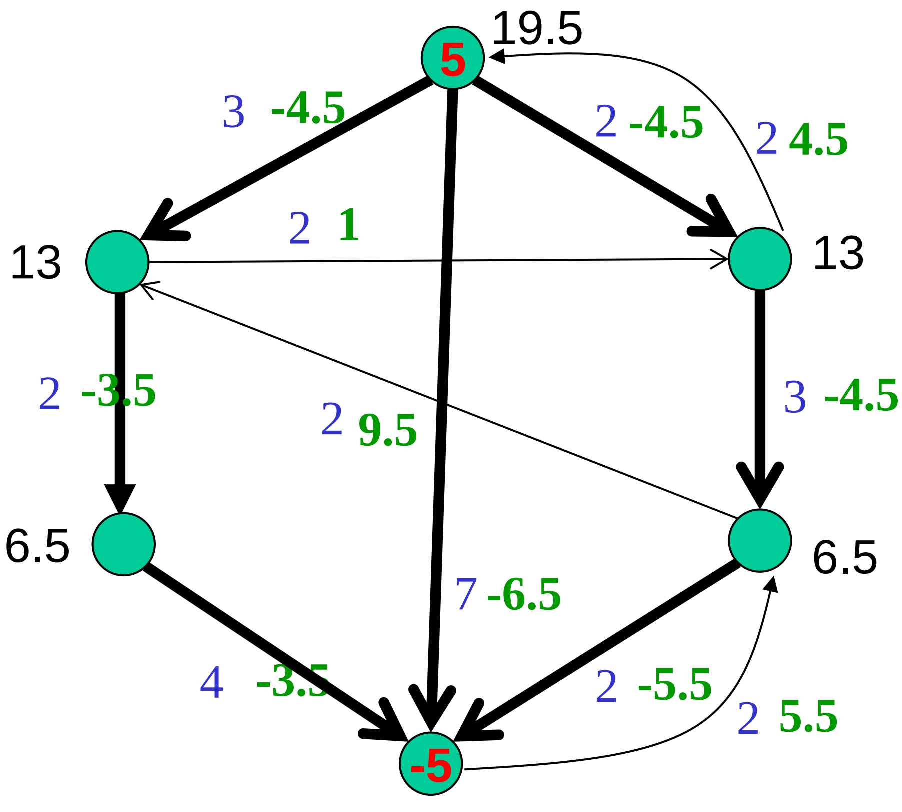

19.5
5
3 -4.5
2 -4.5
2 4.5
13
2 1
13
2 -3.5
2 9.5
3 -4.5
6.5
6.5
7 -6.5
4 -3.5
2 -5.5
2 5.5
-5

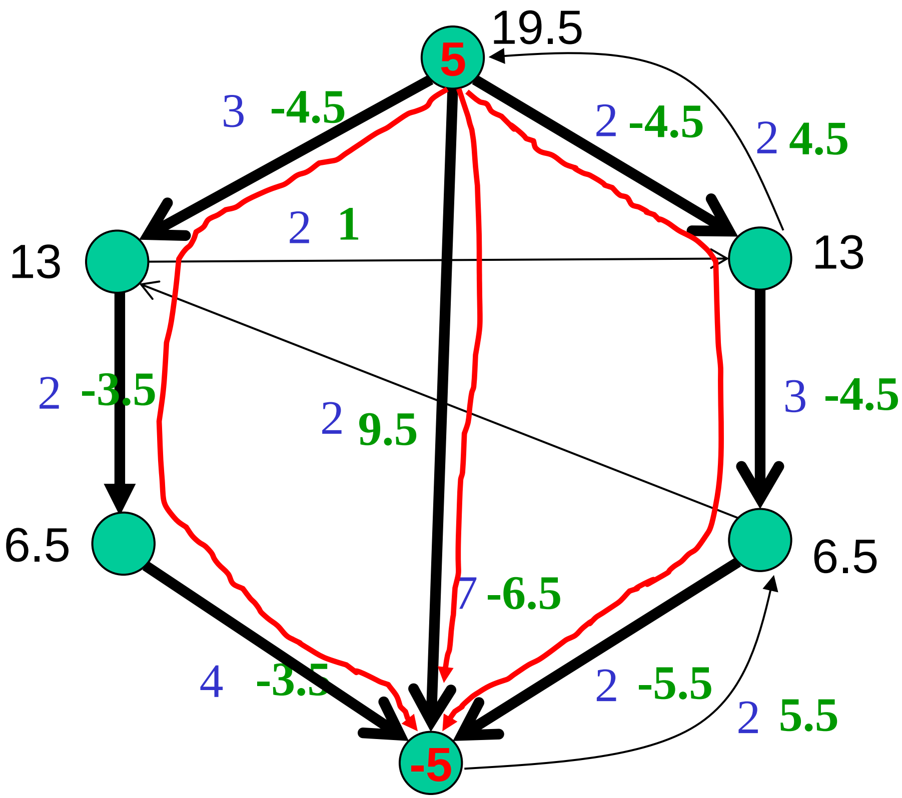

19.5
5
3 -4.5
2 -4.5
2 4.5
13
2 1
13
2 -3.5
2 9.5
3 -4.5
6.5
6.5
7 -6.5
4 -3.5
2 -5.5
2 5.5
-5

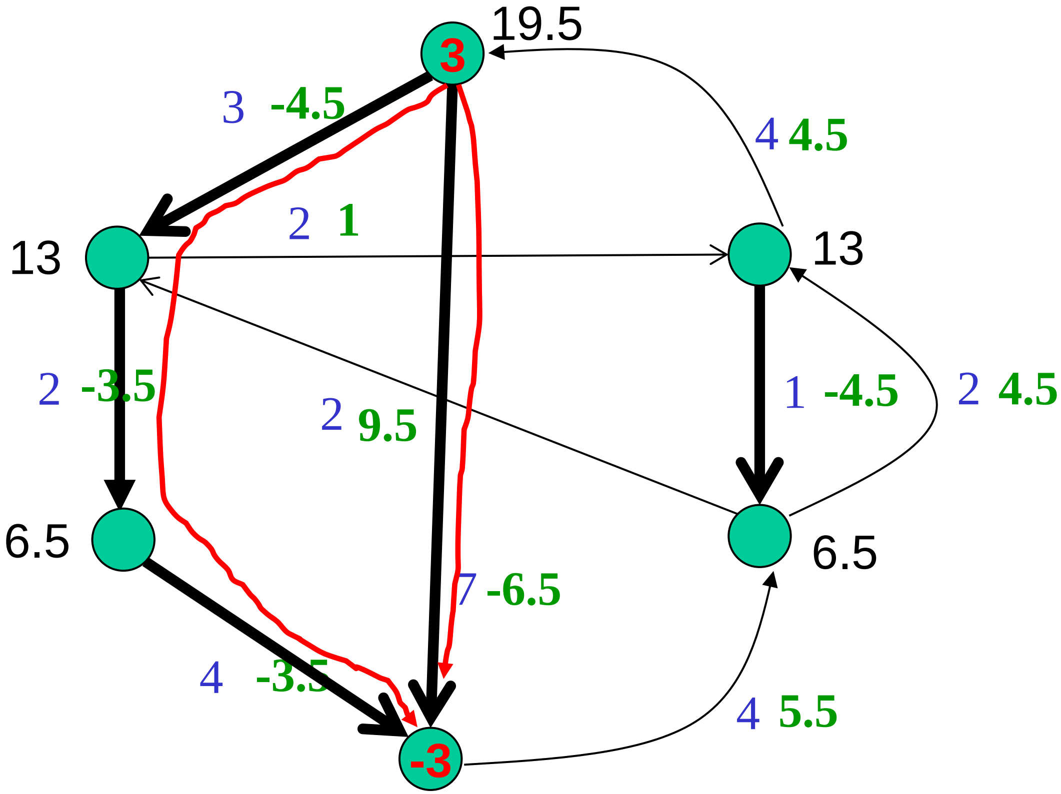

19.5
3
3 -4.5
4 4.5
13
2 1
13
2 -3.5
2 9.5
1 -4.5
2 4.5
6.5
6.5
7 -6.5
4 -3.5
4 5.5
-3

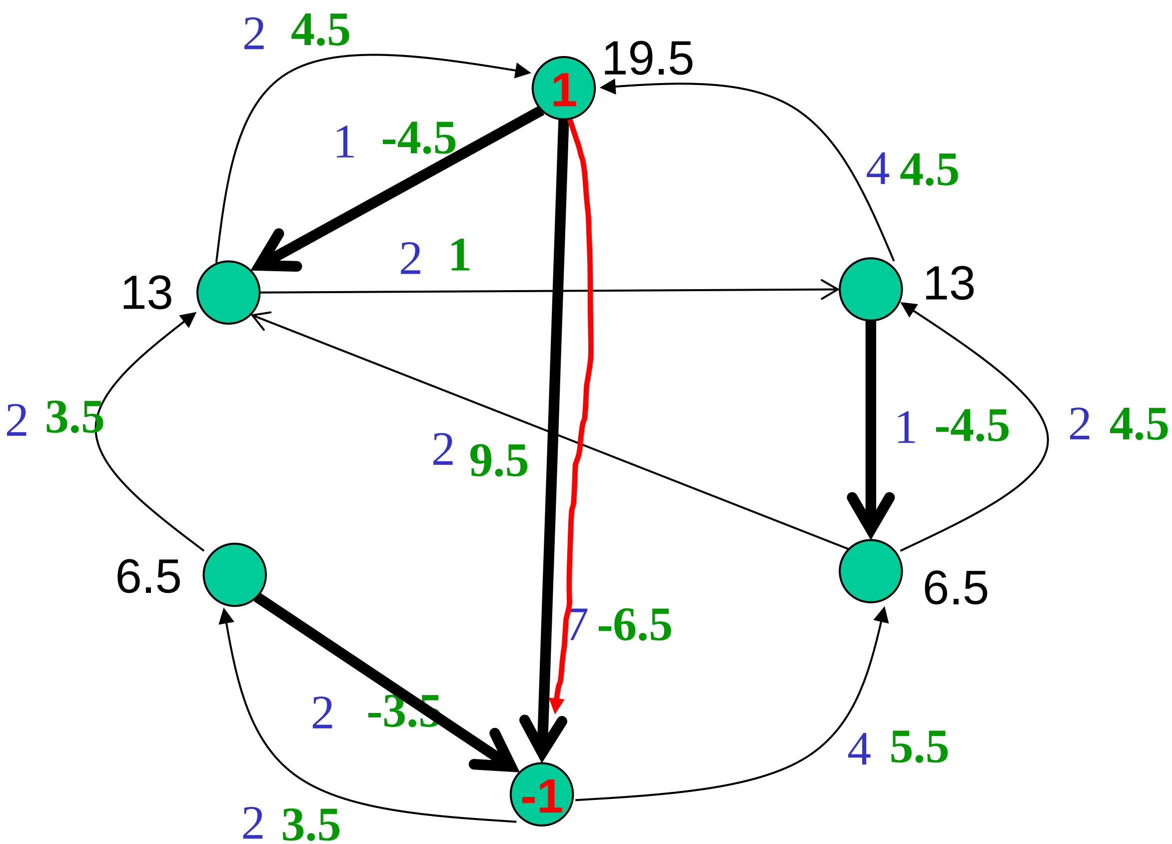

2 4.5
19.5
1
1 -4.5
4 4.5
2 1
13
13
2 3.5
1 -4.5
2 4.5
2 9.5
6.5
6.5
7 -6.5
2 -3.5
4 5.5
-1
2 3.5

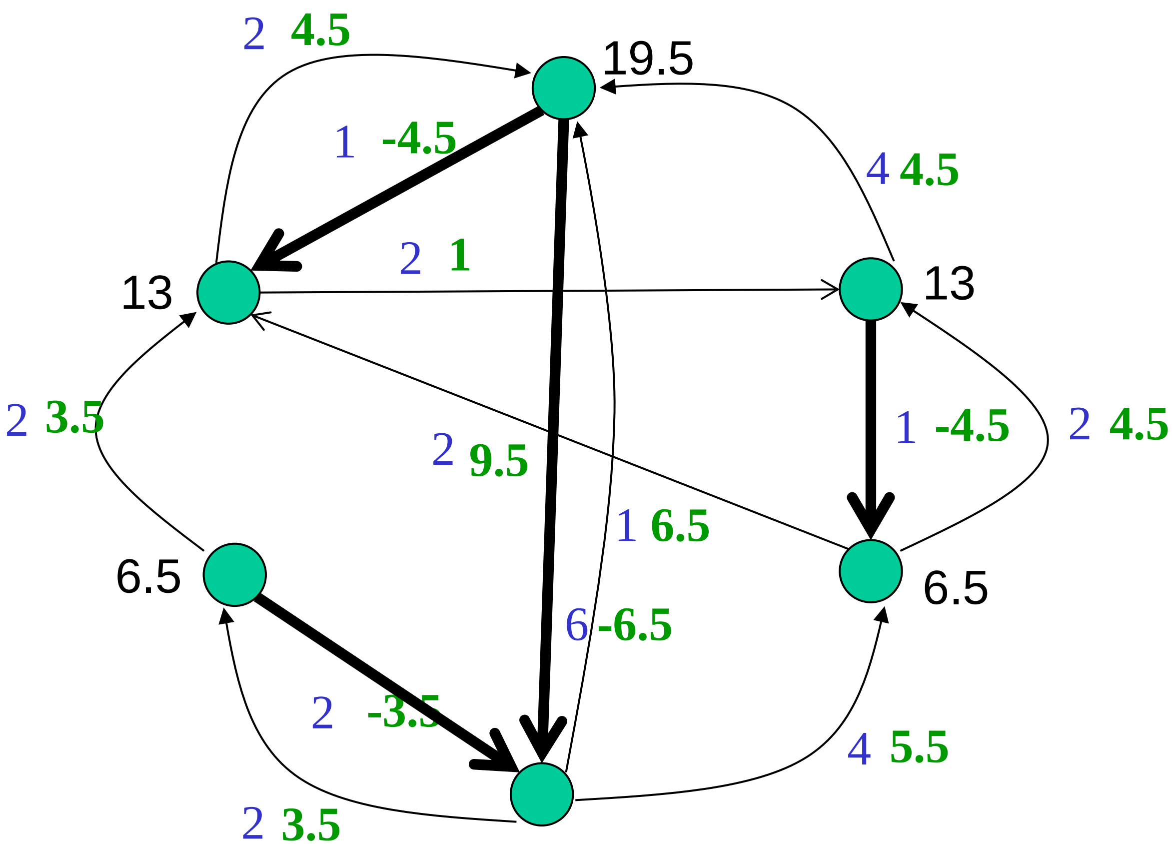

Optimal, but the algorithm does not know about this and continues…

First round, ε ← ε/2 = 3.25

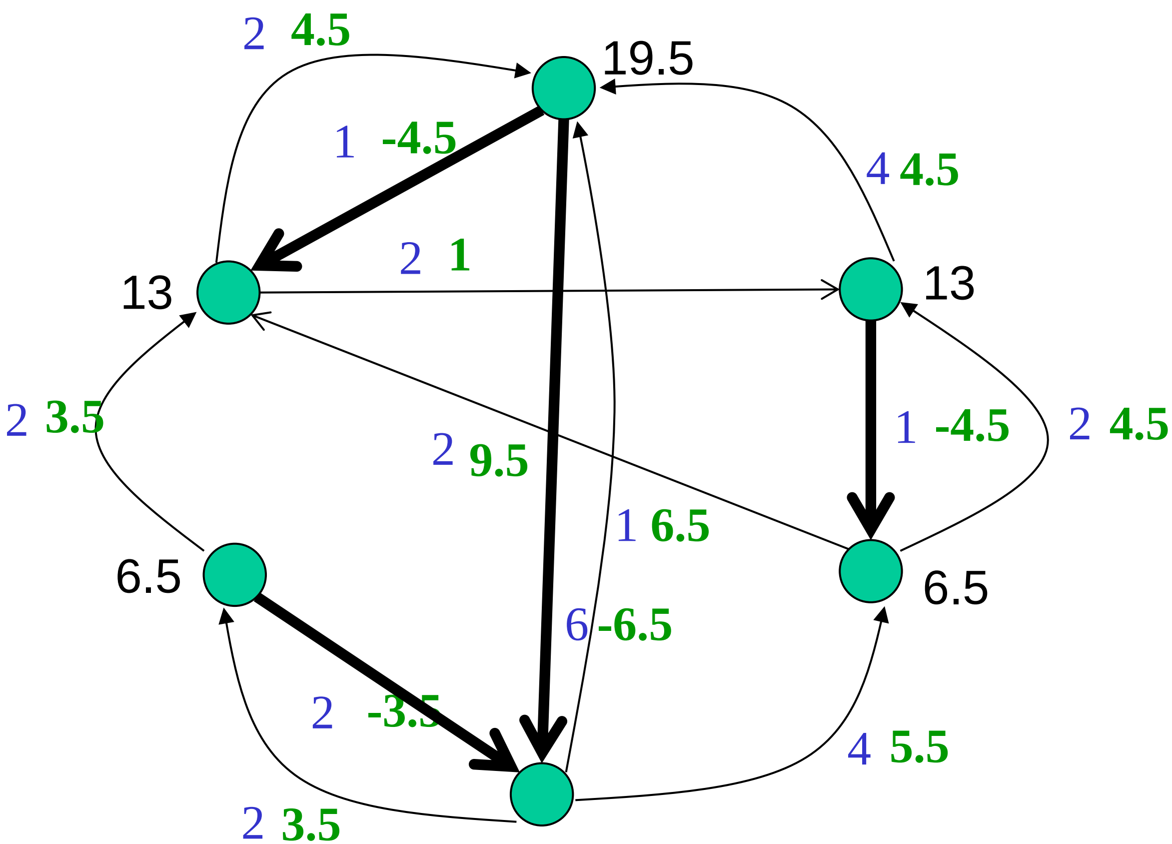

Second round, ε ← ε/2 = 3.25

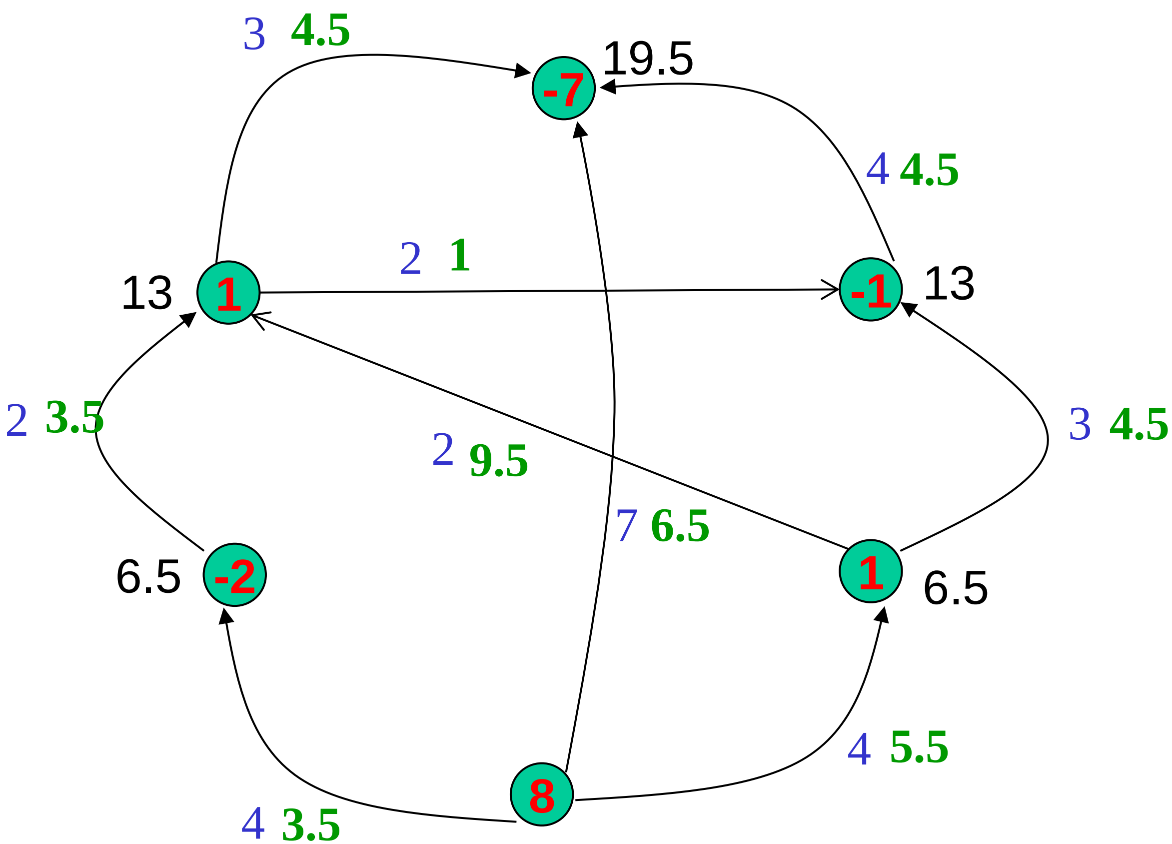

To be continued...

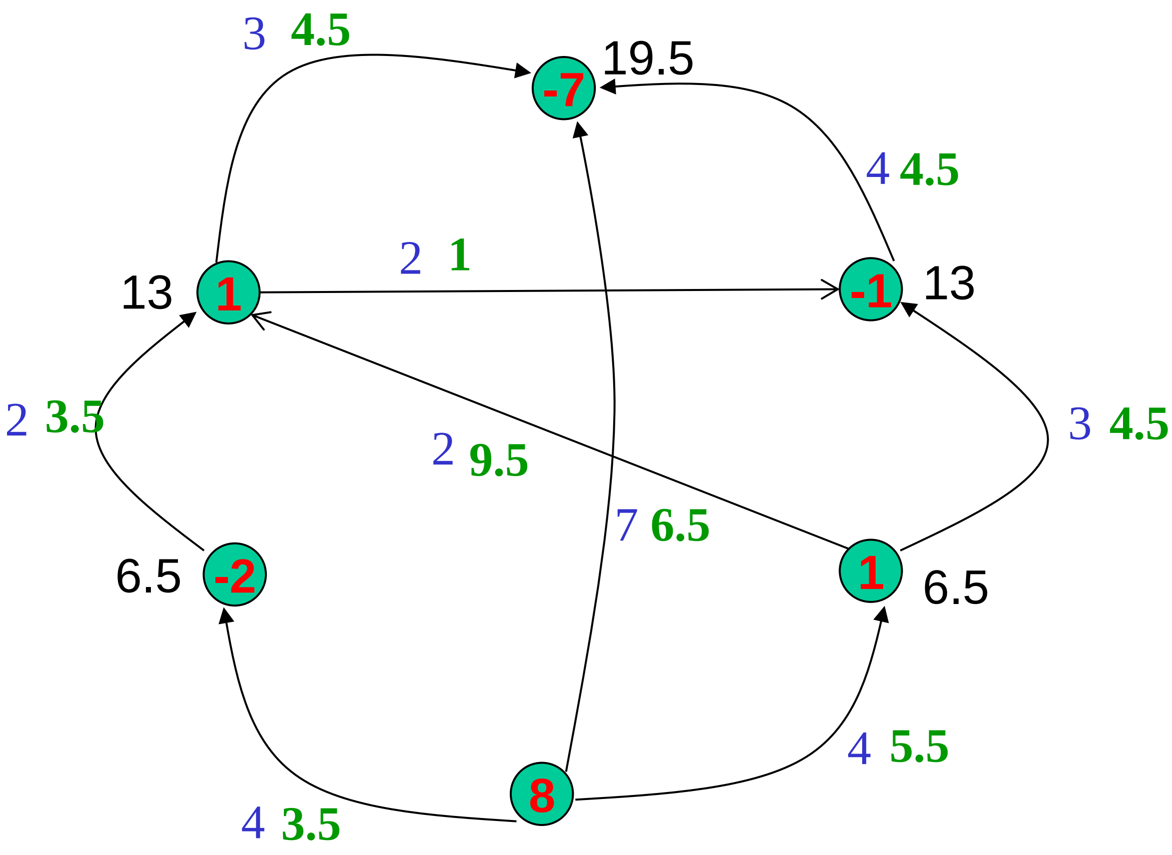

Correctness

Lemma: Refine maintains ε-optimality

Proof: Induction

Clear when refine starts (no negative reduced cost arcs at all)

Raising the potentials of the set S:

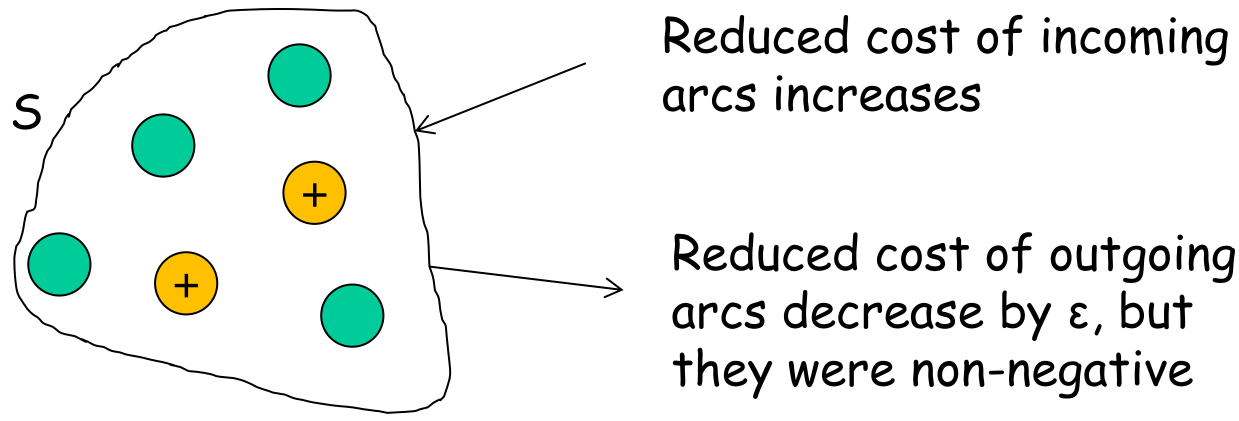

Adding the blocking flow is ok because it uses only admissible arcs

□

Another conclusion from previous proof

- Admissible subgraph is acyclic at all times..

➔ Each blocking flow computation takes $O(m\log(n))$ time !

Running time

The key Lemma

Lemma: Let f be a pseudo flow, f' be a circulation, and v an active vertex with respect to f. Then there is a path from v to a vertex w with $e(w) < 0$, which is residual for f, and its reverse is residual for f'.

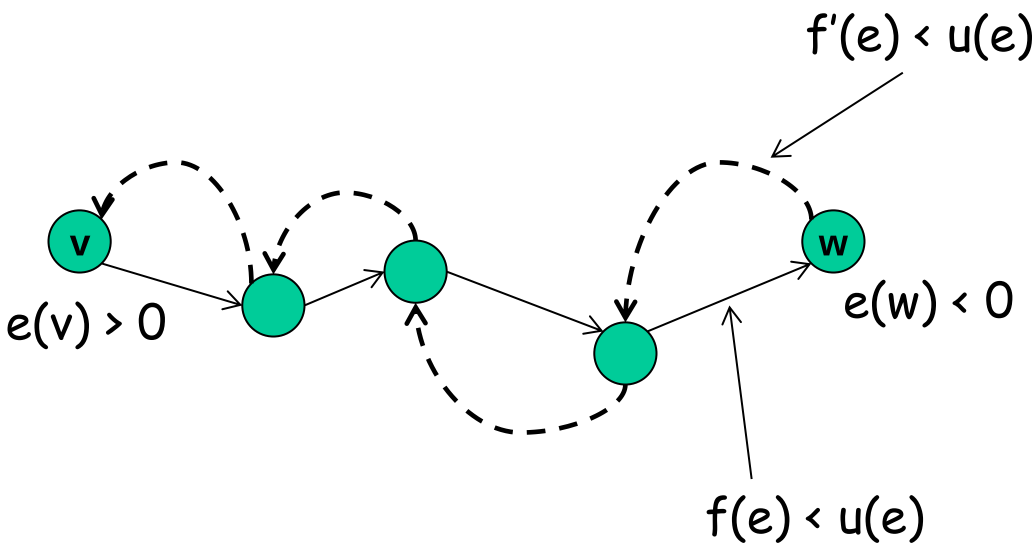

Some intuition

We collect everything reachable in the residual from an active vertex v (e(v) > 0)

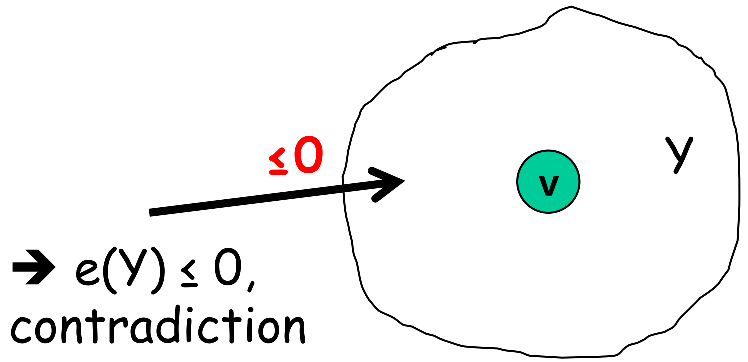

(Recall that the excess of a set is the sum of the excesses of its vertices)

➔ We cannot get stuck without reaching a vertex w with e(w) < 0…

Intuition

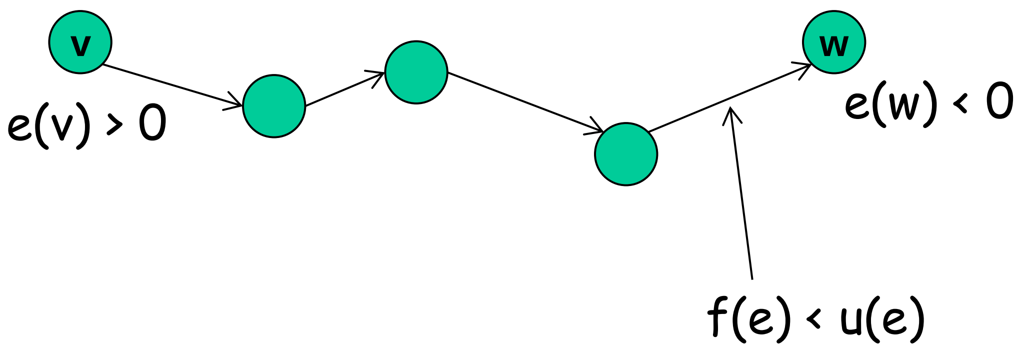

So showing the existence of a path residual for f is straightforward

The proof of the lemma

Collect everything we can reach from v with residual edges e such that $f(e) < f'(e)$, assume (by contradiction) that we get stuck without reaching a vertex w with $e(w) < 0$

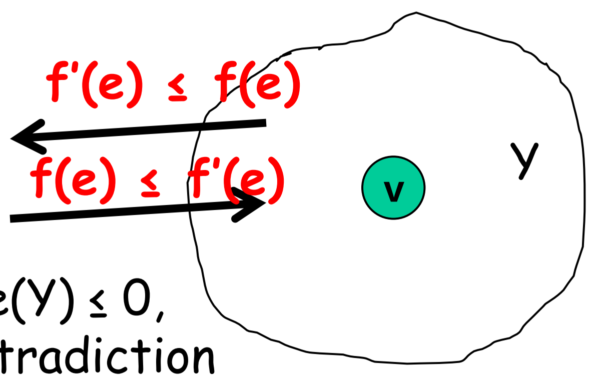

➔ $e(Y) \leq 0$, contradiction

(The sum of all edges incoming Y of $f'(e)$ is 0)

The proof of the lemma

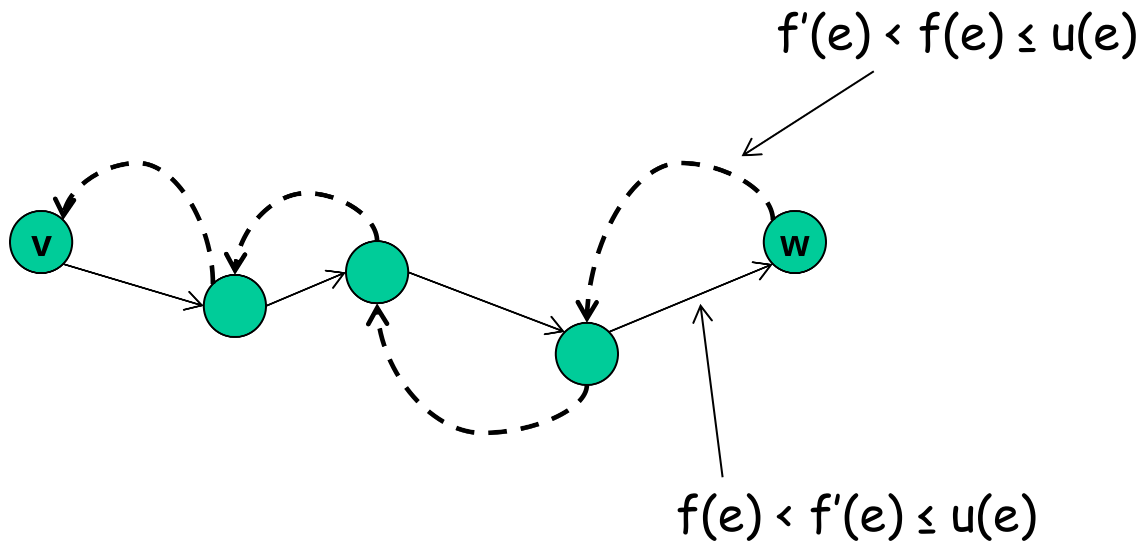

This implies that a residual path whose reverse is residual for f’ from v to w exists □

Bounding the # of relabels

Lemma: The potential of a vertex cannot increase by more than $3n\varepsilon$

Proof:

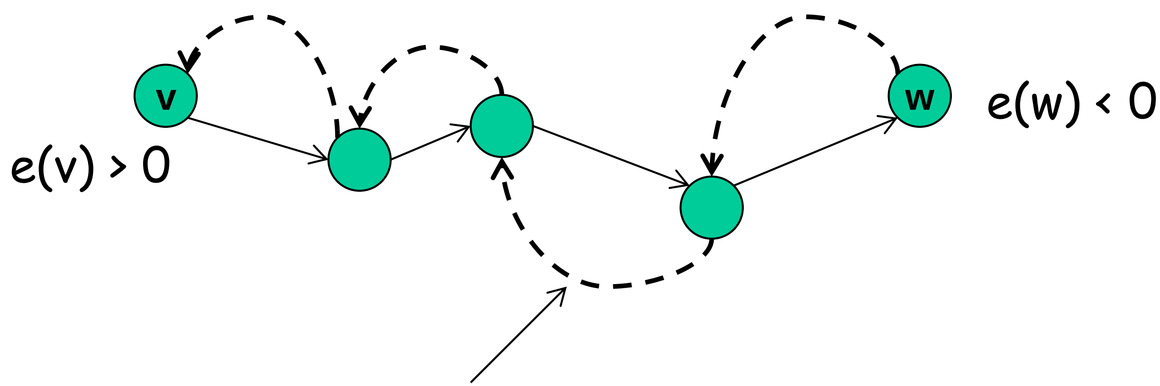

This is residual for the circulation at the beginning of the iteration

Bounding the # of relabels

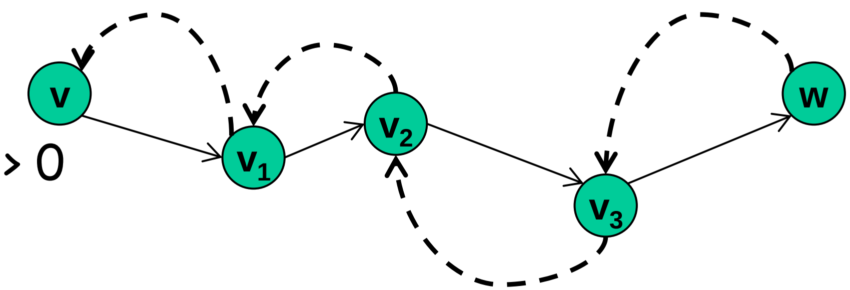

$\pi(v) \le \pi(v_1) + c(v,v_1) + \varepsilon \le \pi(v_2) + c(v,v_1) + c(v_1,v_2) + 2\varepsilon \le \ldots.$

$\pi(w) + c(v,v_1,v_2,v_3,w) + 4\varepsilon = \pi'(w) + c(v,v_1,v_2,v_3,w) + 4\varepsilon \le$

$\pi'(v_3) + c(w,v_3) + 2\varepsilon + c(v,v_1,v_2,v_3,w) + 4\varepsilon \le \ldots..$

$\pi'(v) + c(w,v_3,v_2,v_1,v) + 8\varepsilon + c(v,v_1,v_2,v_3,w) + 4\varepsilon = \pi'(v) + 12\varepsilon$

Bounding the # of relabels

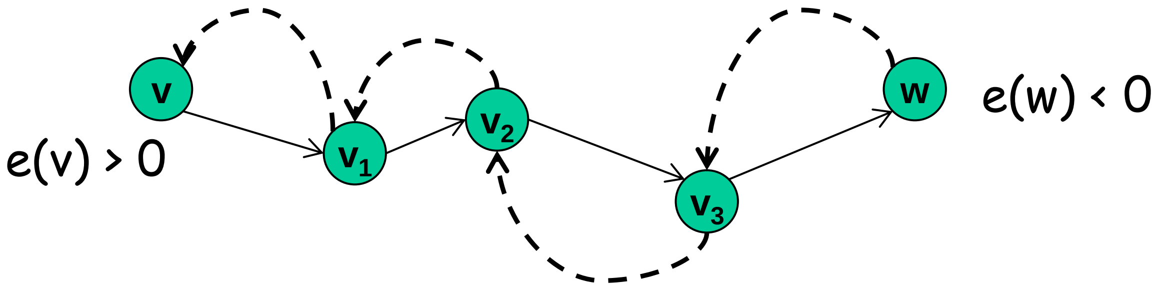

In general we get that

$$\pi(v) \leq \pi'(v) + 3\ell\varepsilon \leq \pi'(v) + 3n\varepsilon$$

➔ Lemma: O(n) blocking flow computations

Summary

We have got an $O(nm\log(n)\log(nC))$ time algorithm.

Can be slightly improved to $O(nm\log(n^2/m)\log(nC))$ by an even fancier blocking flow alg.

Open: Get an $O(m^{3/2}\log(U)\log(nC)\log(n))$ time algorithm

Some recent progress on this using interior point methods and other sophisticated tools

A strongly polynomial algorithm

ε-tight potentials

Def: f is ε-optimal ⇔ There exists a potential function π such that $c^{\pi}(e) \geq -\varepsilon$ for every residual arc e

Given a pseudoflow f, find the smallest ε such that there are potentials for which f is ε optimal

How do we find this ε ?

ε-tight potentials

Suppose f is ε-optimal

➔ The mean cost of a residual cycle Γ is $\geq -\varepsilon$

➔ The minimum mean cost of a res. cycle is $\geq -\varepsilon$

It turns out that the opposite is also true:

Thm: f is ε-optimal $\Leftrightarrow$
The minimum mean cost of a res. cycle is $\geq -\varepsilon$

ε-tight potentials

Suppose the minimum mean cost of a res. cycle is ≥ $-\varepsilon$

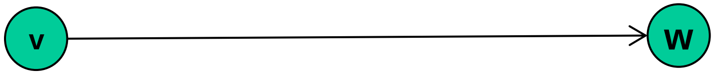

$c^{\varepsilon}(v,w) \leftarrow c(v,w) + \varepsilon$

No negative res. cycles with respect to c^{ε}

Take a shortest path tree in the residual to some vertex with respect to c^{ε}

$c^{\varepsilon}(v,w) + d(w) - d(v) \geq 0$

$c(v,w) + d(w) - d(v) \geq -\varepsilon$

Set $\pi(v) = d(v)$ □

ε-tight potentials

f is ε-optimal $\Leftrightarrow$
The minimum mean cost of a res. cycle is $\geq -\varepsilon$

Corollary:

f is ε-tight $\Leftrightarrow$ The minimum mean cost of a res. cycle $= -\varepsilon$

Can find the value of the minimum mean cost cycle in $O(nm)$ time

The Successive Approximation Strongly Poly. Alg

$\varepsilon \leftarrow C$

$\forall v,\ \pi(v) \leftarrow 0$

$\forall e,\ f(e) \leftarrow 0$

repeat

find η and π_η such that f is η-tight with respect to π_η

if $(\eta=0)$ return

$\varepsilon \leftarrow \eta/2$

refine (f, π_η)

Analysis

Thm. f is ε-optimal with respect to prices π and $|c^{\pi}(v,w)| \geq 2n\varepsilon$. Then $f(v,w) = f'(v,w)$ for any other ε-optimal f'.

Pf. Since f is ε-optimal $c^{\pi}(v,w) \leq -2n\varepsilon$
→ $f(v,w) = u(v,w)$

So if $f'(v,w) \neq f(v,w)$ then $f'(v,w) < f(v,w)$

→ $f'(w,v) > f(w,v)$

f'-f has a cycle containing (w,v),
and this cycle is res. for f

Analysis

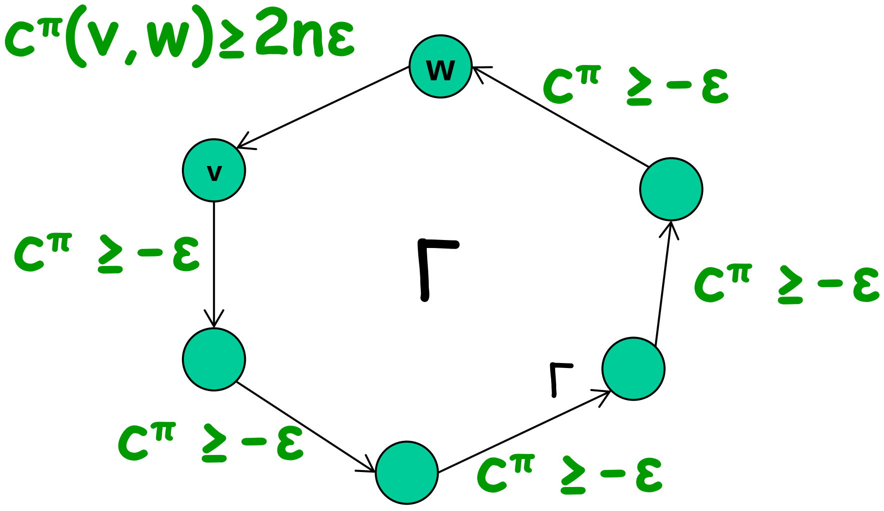

$c^\varepsilon(\Gamma) \geq 2n\varepsilon - (n-1)\varepsilon > n\varepsilon$

rev(Γ) res. for f' but its mean weight is smaller than -ε

➔ f' is not ε-optimal

□

Analysis

Define F_ε to be the set of arcs such that the flow through them is the same for every ε-optimal f

Assume there is an ε-tight f, let $\varepsilon' \leq \varepsilon/(2n)$

Thm. $F_{\varepsilon'}$ properly contains F_ε

Conclusion: Every log(n) iterations we fix the flow on one additional edge

iter. is min{log(nC), mlog(n)}

Run time $O(nm\log(n^2/m) \min\{\log(nC), m\log(n)\})$

Analysis

Thm. $F_{\varepsilon'}$ properly contains F_{ε}

Pf. $F_{\varepsilon'}$ clearly contains F_{ε} since an ε′-optimal flow is an ε-optimal flow.

Consider potentials π such that f is ε-optimal with respect to π. Consider the minimum mean cycle res. Γ.

Arcs of Γ are not in F_{ε}

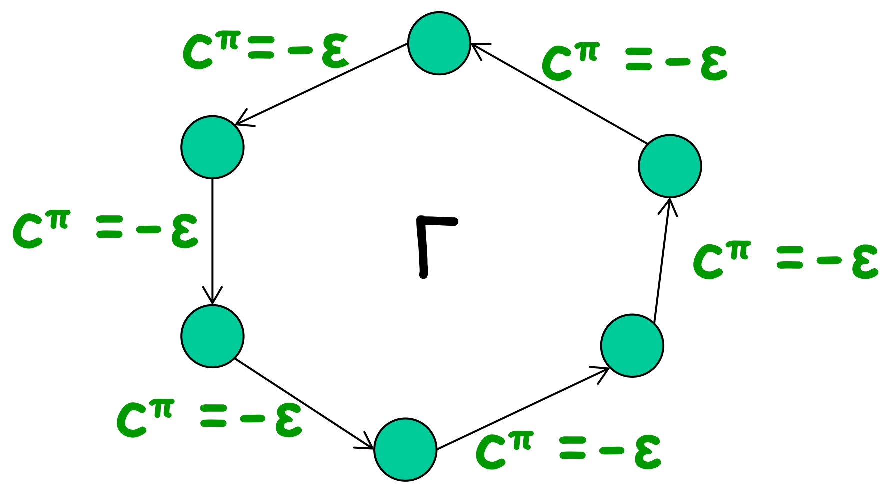

Analysis

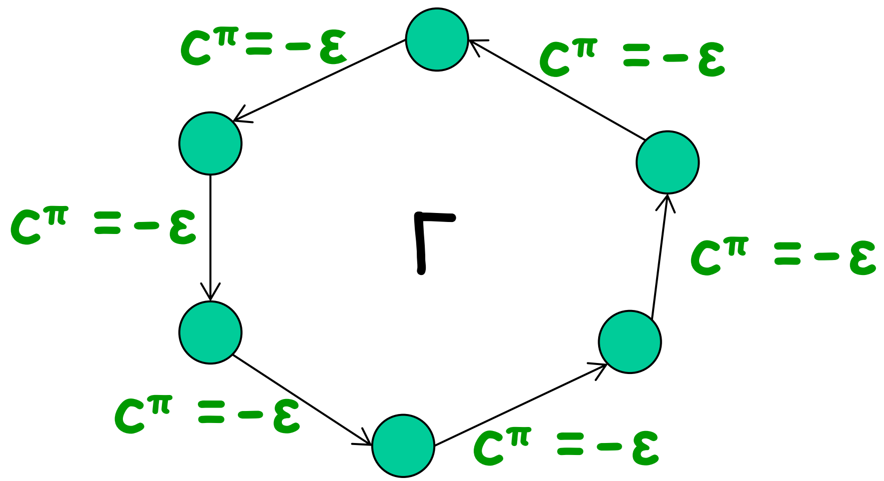

Consider π' and f'

There is an edge $e \in \Gamma$ such that $c^{\pi'}(e) \leq -\varepsilon \leq -2n\varepsilon'$

So $e \in F_{\varepsilon'} \setminus F_{\varepsilon}$