Network Flow

Max-Flow: Edmonds-Karp Algorithm, Dinitz’s Algorithm
Max Bipartite Matching: Hopcroft–Karp–Karzanov algorithm

Residual Network G^f

Given $G = (V, E)$, c, and a flow f

$G^f = (V, E^f)$ and the associated capacity $c^f: E^f \to \mathbb{R}^+$ are defined as follows:

- $(u, v) \in E^f$ if one of the followings holds
  - $(u, v) \in E$ and $f(u, v) < c(u, v)$: in this case, $c^f(u, v) = c(u, v) - f(u, v)$
  - $(v, u) \in E$ and $f(v, u) > 0$: in this case, $c^f(u, v) = f(v, u)$

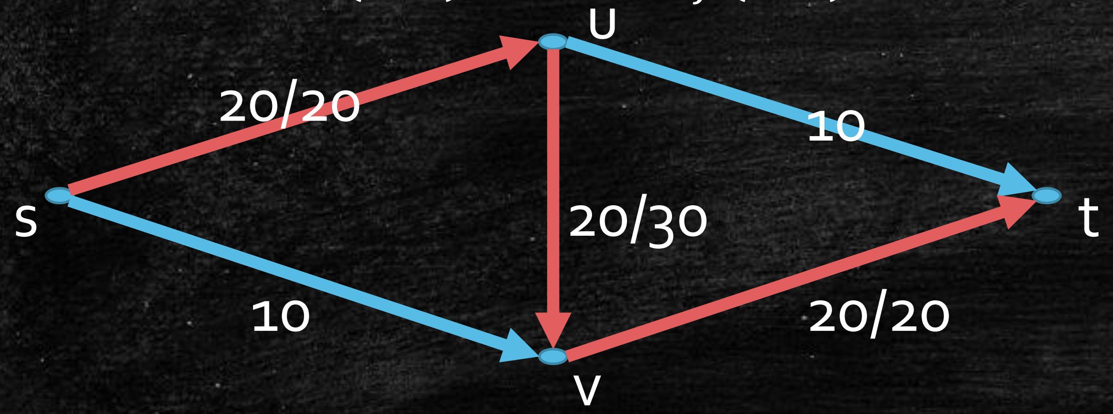

G

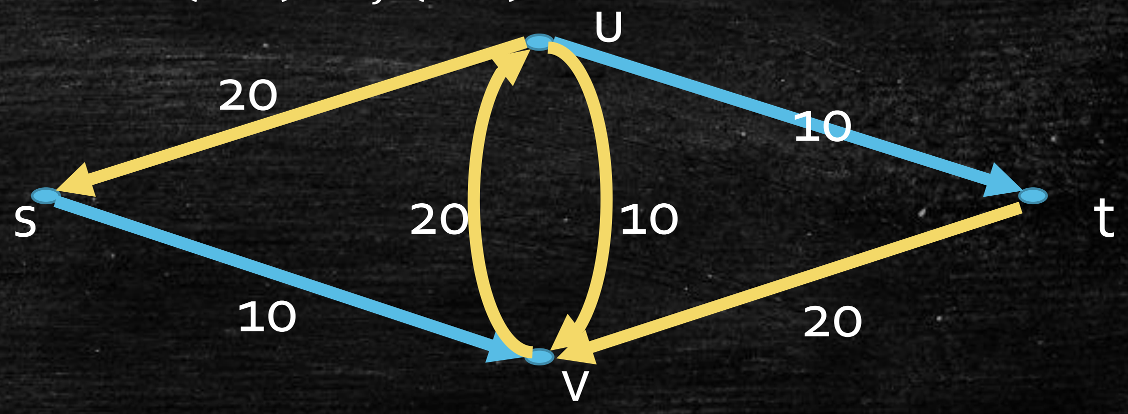

G^f

Last Lecture – Ford-Fulkerson Method

- Always terminates for integer/rational capacities
- Not guaranteed to terminate for irrational capacities
- Time complexity for integer capacities: $O(|E| \cdot v(f)_{\max})$
  - not a polynomial time

Edmonds-Karp Algorithm

Edmonds-Karp Algorithm

EdmondsKarp$(G = (V, E), s, t, c)$:

1. initialize f such that $\forall e \in E: f(e) = 0$; initialize $G^f \leftarrow G$;
2. while there is an s-t path on G^f:
3. **find such a path p by BFS;**
4. find an edge $e \in p$ with minimum capacity b;
5. update f that pushes b units of flow along p;
6. update G^f;
7. endwhile
8. return f

Why BFS?

- BFS maintains the distances
  - distance: num of edges, not weighted distance

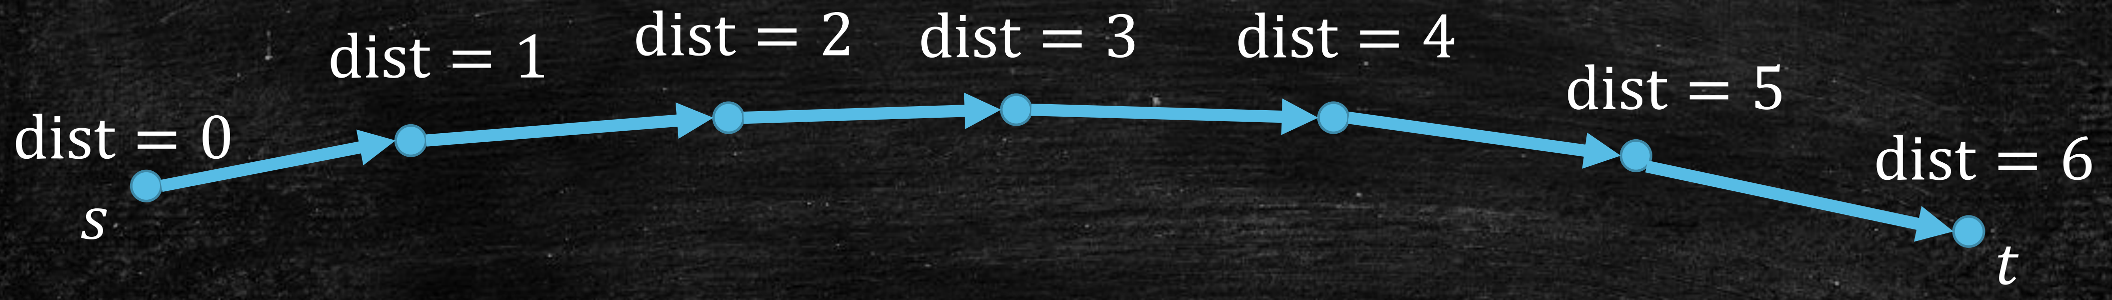

A path found by an iteration of Edmonds-Karp Algorithm

Why BFS?

- In the residual network G^f, a new appeared edge can only goes from a vertex at distance $t+1$ to a vertex at distance t.
- Addition of such edges does not decrease the distance between s and u for every $u \in V$.
- **[Key Observation in this lecture!]** Thus, $\mathrm{dist}(u)$ is non-decreasing throughout the algorithm for every $u \in V$.

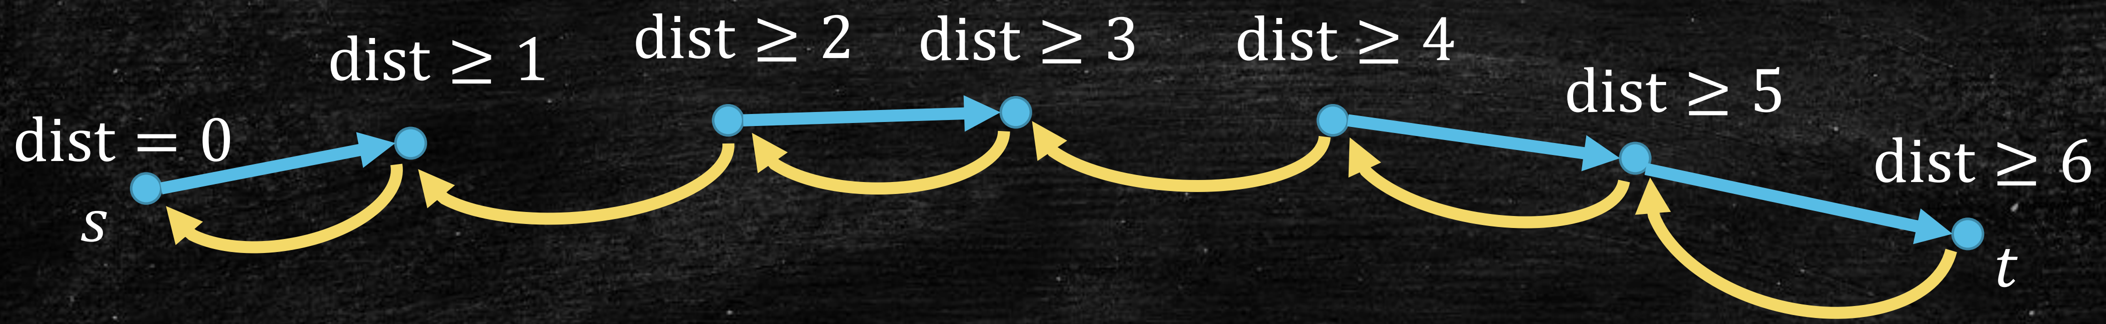

The updates to the edges in G^f

Weak Monotonicity to Strong Monotonicity

- $\mathrm{dist}(u)$ can only be one of $0, 1, 2, \ldots, |V|, \infty$
  - It can only be increased for $|V| + 1$ times!
- It's great that BFS buys us distance monotonicity!
- However, weak monotonicity is not enough.
- To make a progress, we need $\mathrm{dist}(u)$ strictly increases for some $u \in V$, so that we can upper bound the number of iterations.

Counterexample: dist for all vertices remain unchanged after an iteration.

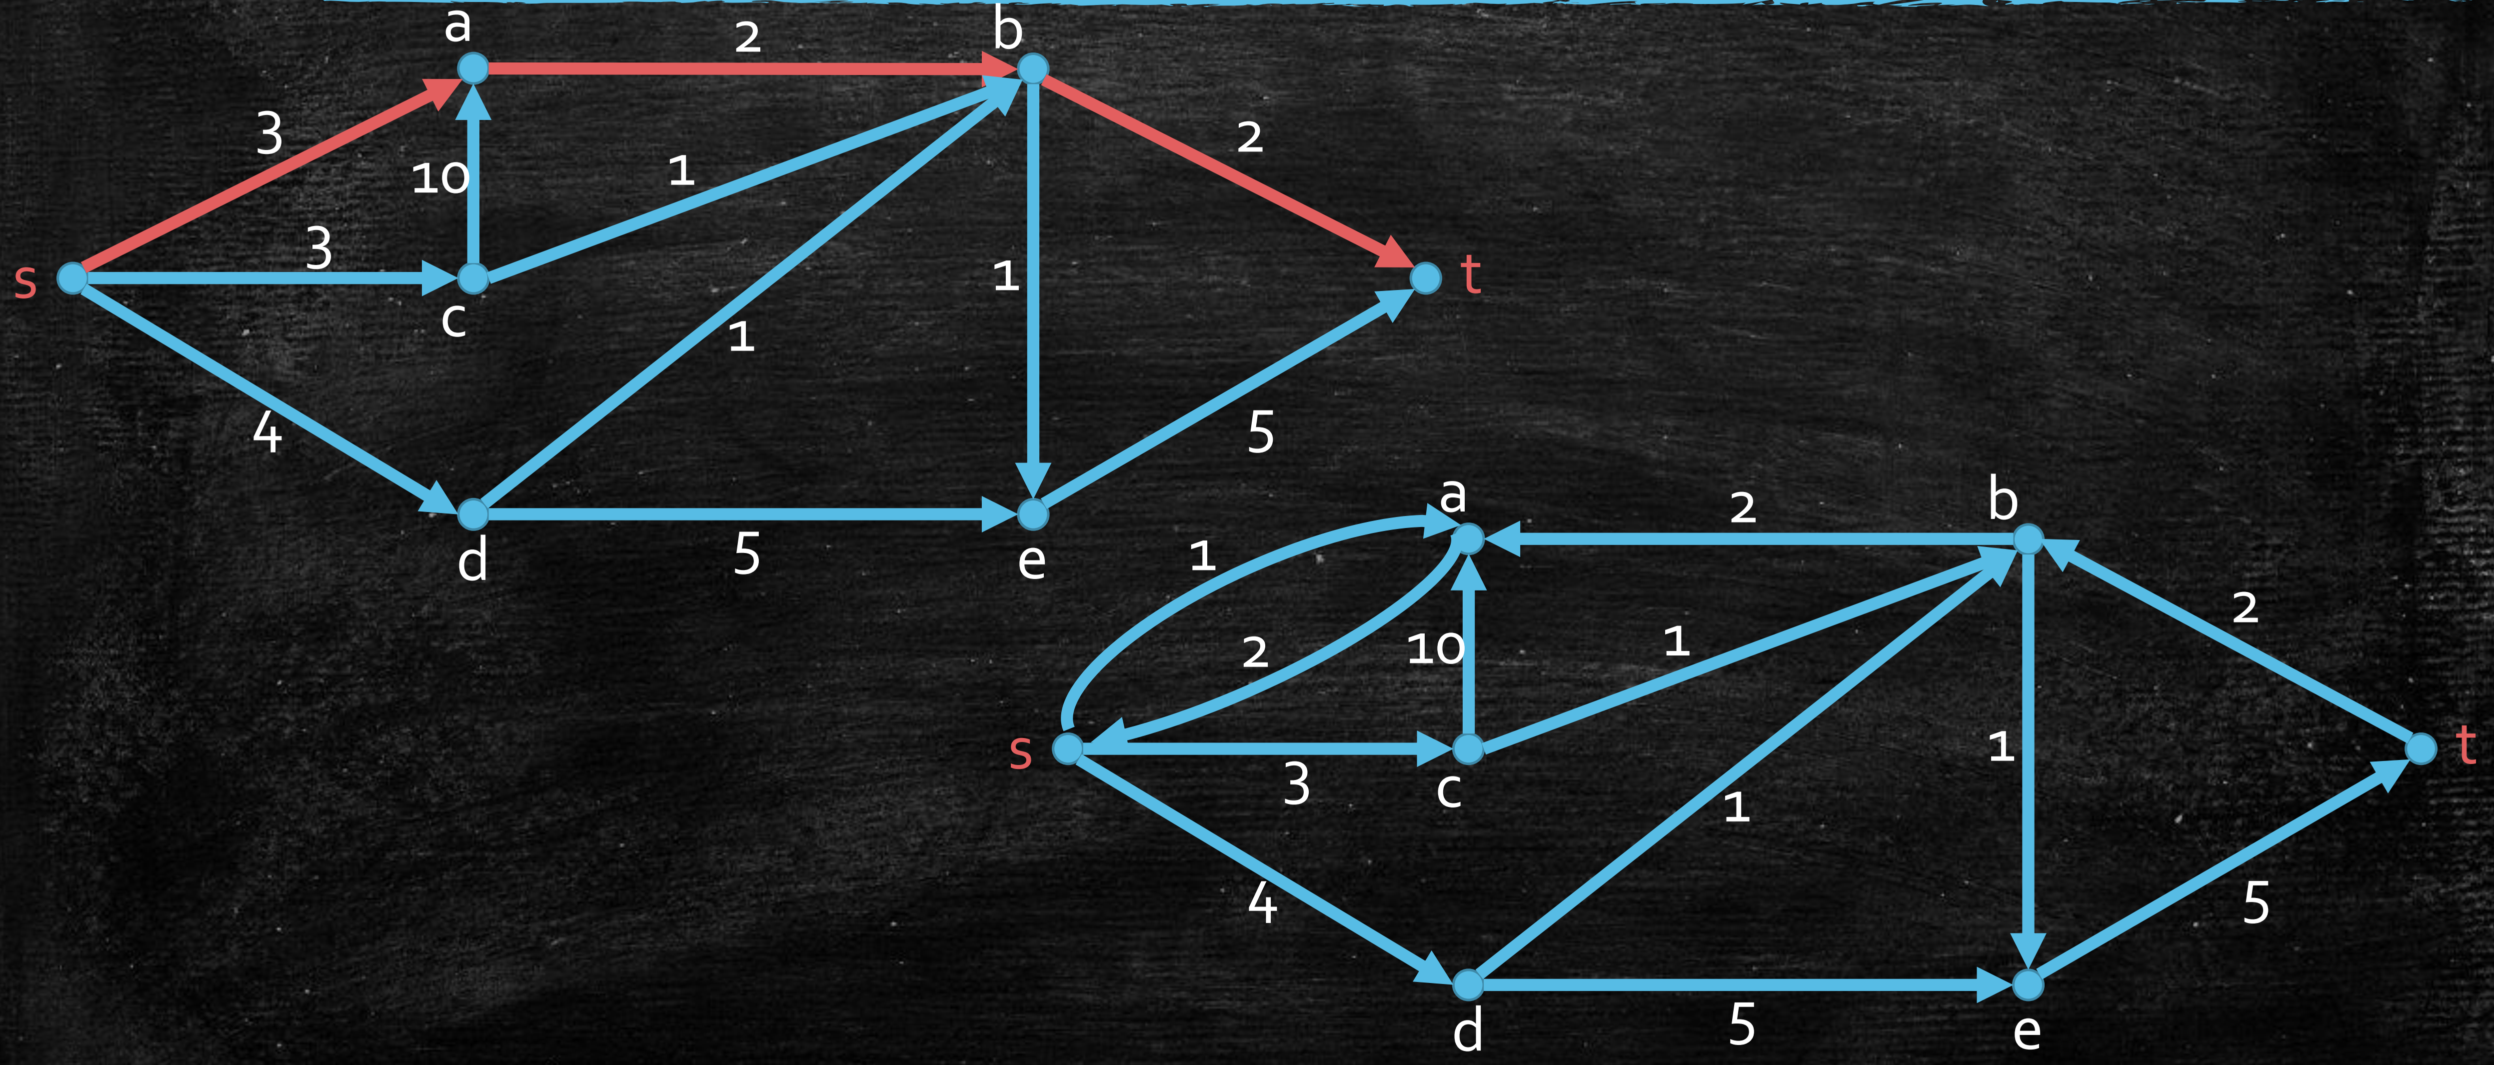

Towards Strong Monotonicity...

- Observation: At least one edge (u, v) on p is "saturated", and this edge will be deleted in the next iteration.
- Each iteration will remove an edge from a vertex at distance i to a vertex at distance $i + 1$.
- Intuitively, we cannot keep removing such edges while keeping the distances of all vertices unchanged.

Towards Strong Monotonicity

- Suppose we are at the $(i+1)$-th iteration. f_i is the current flow, and p is the path found in G^{f_i} at the $(i+1)$-th iteration.
- We say that an edge (u,v) is critical if the amount of flow pushed along p is $c^{f_i}(u,v)$.
- A critical edge disappears in $G^{f_{i+1}}$, but it may reappear in the future…
- We will try to bound the number of times (u,v) becomes critical.

Between two “critical”

A flow along p in G^{f_i} where (u, v) becomes critical

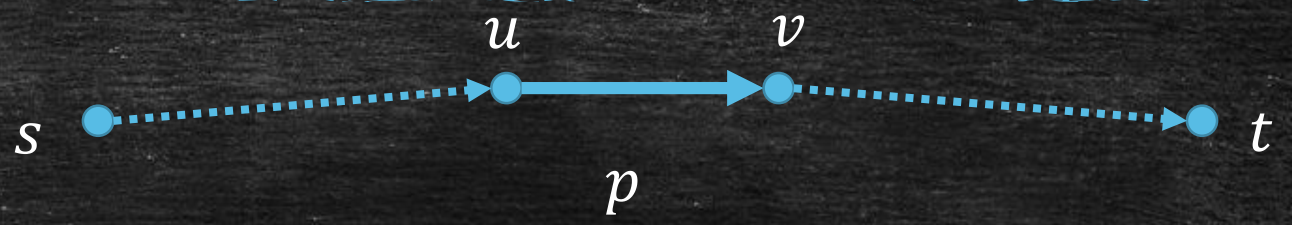

In $G^{f_{i+1}}$, (u, v) disappears, and (v, u) appears.

Before the next time (u, v) becomes critical again, (u, v) must first reappear!

Before (u, v) reappears, the algorithm must have found p going through (v, u).

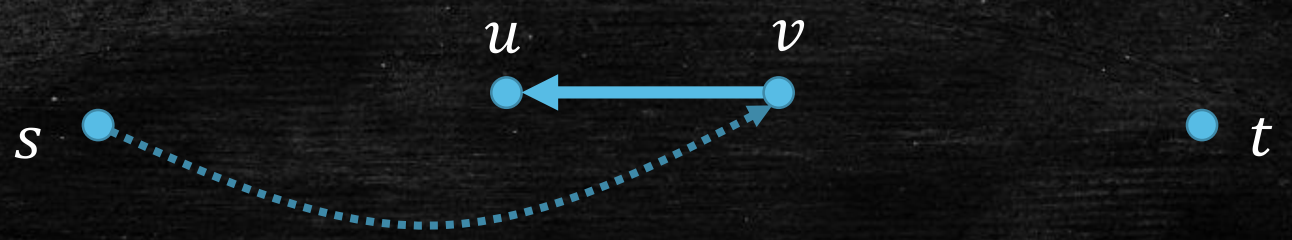

Between two "critical"

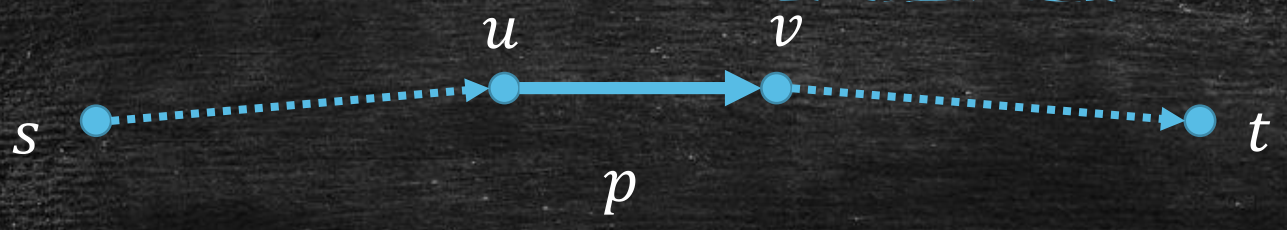

$\text{dist}^i(v) = \text{dist}^i(u) + 1$

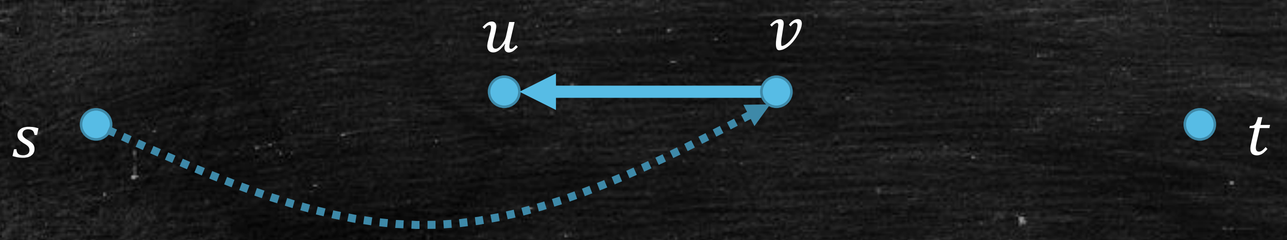

$\text{dist}^{i+j}(u) = \text{dist}^{i+j}(v) + 1$

- Distance monotonicity: $\text{dist}^{i+j}(v) \geq \text{dist}^i(v)$.
- Thus, $\text{dist}^{i+j}(u) = \text{dist}^{i+j}(v) + 1 \geq \text{dist}^i(v) + 1 \geq \text{dist}^i(u) + 2$.
- The distance of u from s increases by 2 between two "critical".

Putting Together

- The distance of u from s increases by 2 between two “critical”.
- Distance takes value from $\{0, 1, \ldots, |V|, \infty\}$, and never decrease.
- Thus, each edge can only be critical for $O(|V|)$ times.
- At least 1 edge become critical in one iteration.
- Total number of iterations is $O(|V| \cdot |E|)$.
- Each iteration takes $O(|E|)$ time.
- Overall time complexity for Edmonds-Karp: $O(|V| \cdot |E|^2)$.
- It can handle the issue with irrational numbers!

Can we improve?

- Each iteration takes $O(|E|)$ time to find a shortest s-t path by BFS.
- However, each shortest s-t path has length at most $|V|$.
- Idea: push flow from multiple shortest s-t paths in one iteration!

Dinic's Algorithm (Dinitz's Algorithm)

- Proposed by Yefim Dinitz
- Updated by Shimon Even and Alon Itai
- Time complexity: $O(|V|^2 \cdot |E|)$

Dinic's Algorithm – high-level ideas

- Build a level graph:
  - Vertices at Level i are at distance i.
  - Only edges go from a level to the next level are kept.
  - Can be done in $O(|E|)$ time using a similar idea to BFS.

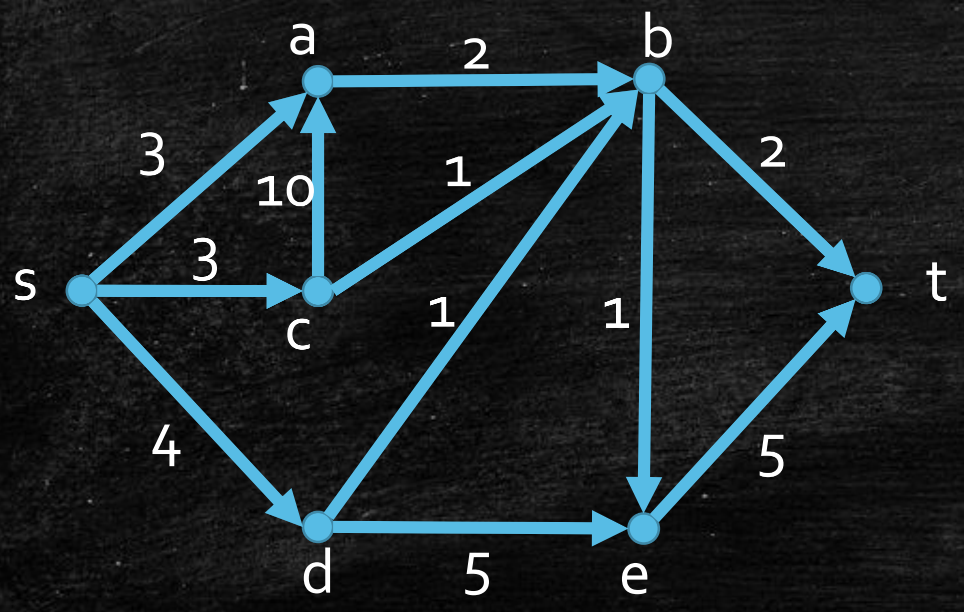

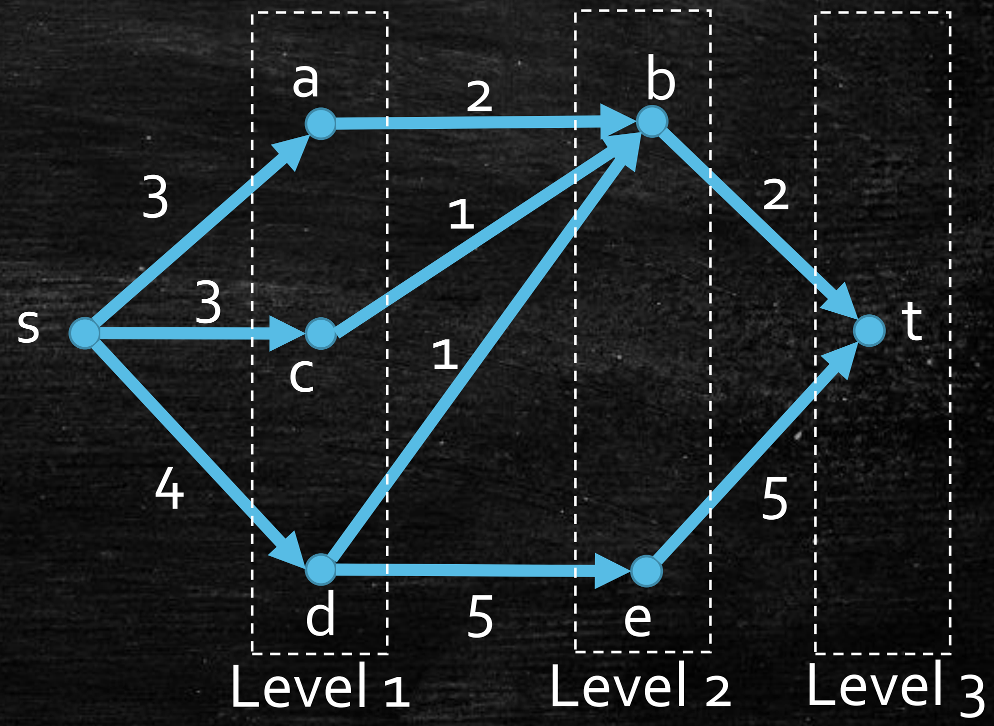

Dinic's Algorithm – high-level ideas

- Find a blocking flow on the level graph:
  - Push flow on multiple s-t paths.
  - Each s-t path must contain a critical edge!

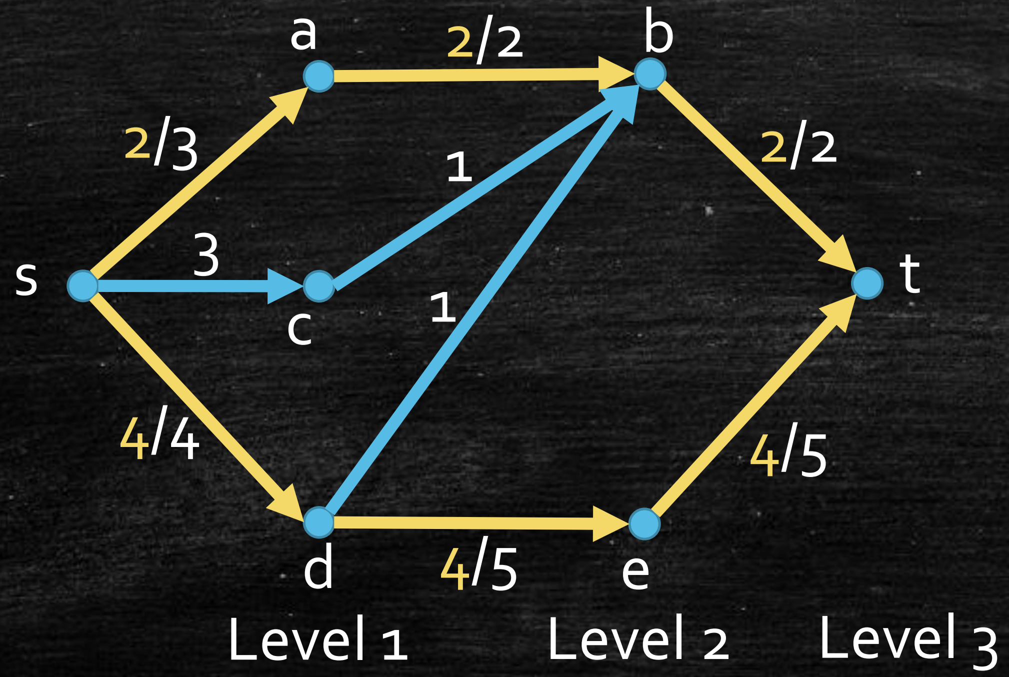

Dinic's Algorithm – high-level ideas

- Find a blocking flow on the level graph:
  - Push flow on multiple s-t paths.
  - Each s-t path must contain a critical edge!

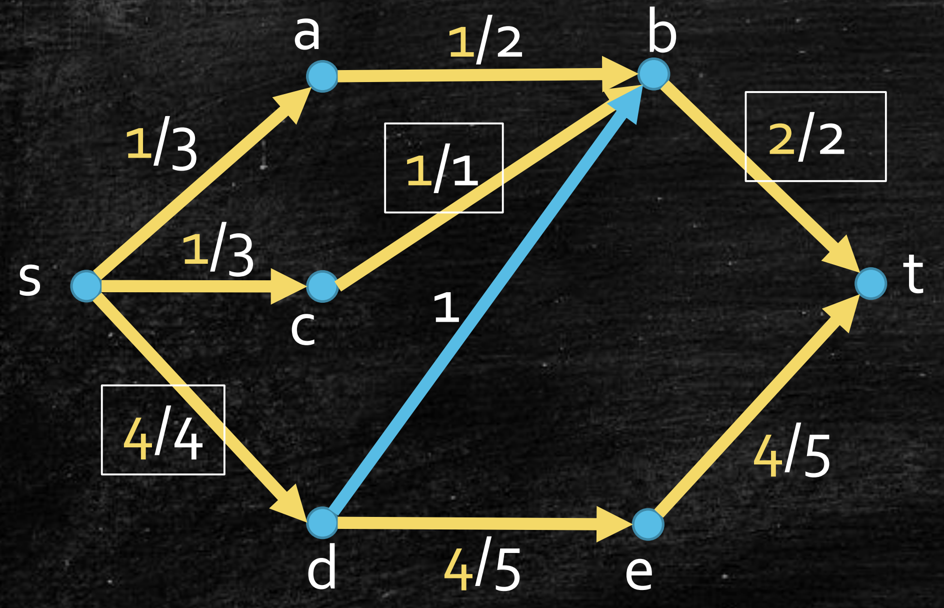

a blocking flow

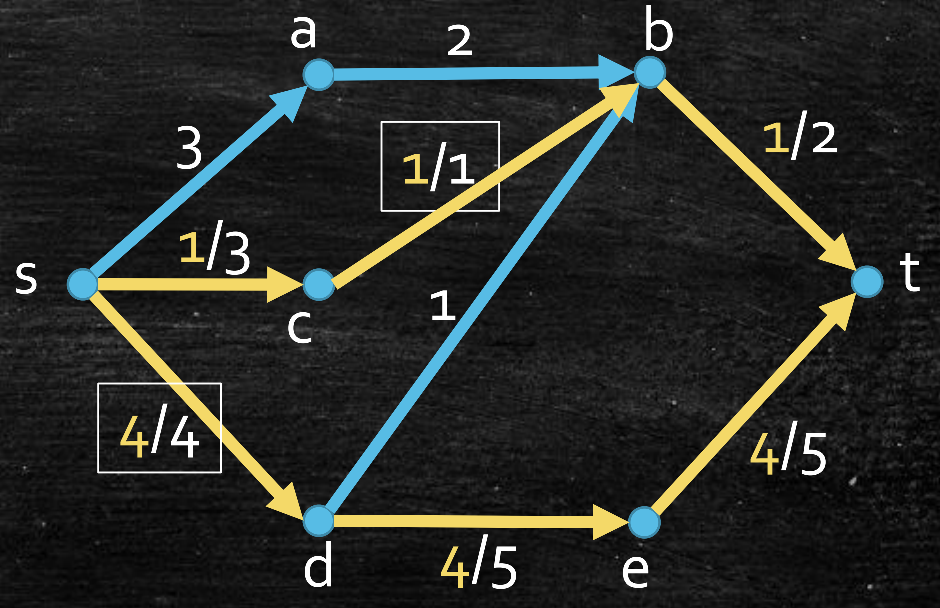

not a blocking flow: path s-a-b-t contains no critical edge

Dinic's Algorithm – Overview

- Initialize f to be the empty flow and $G^f = G$.
- Iteratively do the followings until $\mathrm{dist}(t) = \infty$:
  - Construct the level graph G_L^f for G^f.
  - Find a blocking flow on G_L^f.
  - Update f and G^f.

Questions Remain

1. How many iterations do we need before termination?
2. How do we find a blocking flow?

Questions Remain

1. How many iterations do we need before termination?
2. How do we find a blocking flow?

Simple yet important observations

- In the level graph $G_L^{f_i}$ at every iteration i, every s-t path has length $\mathrm{dist}^i(t)$.
- Every shortest s-t path in G^{f_i} also appears in $G_L^{f_i}$.

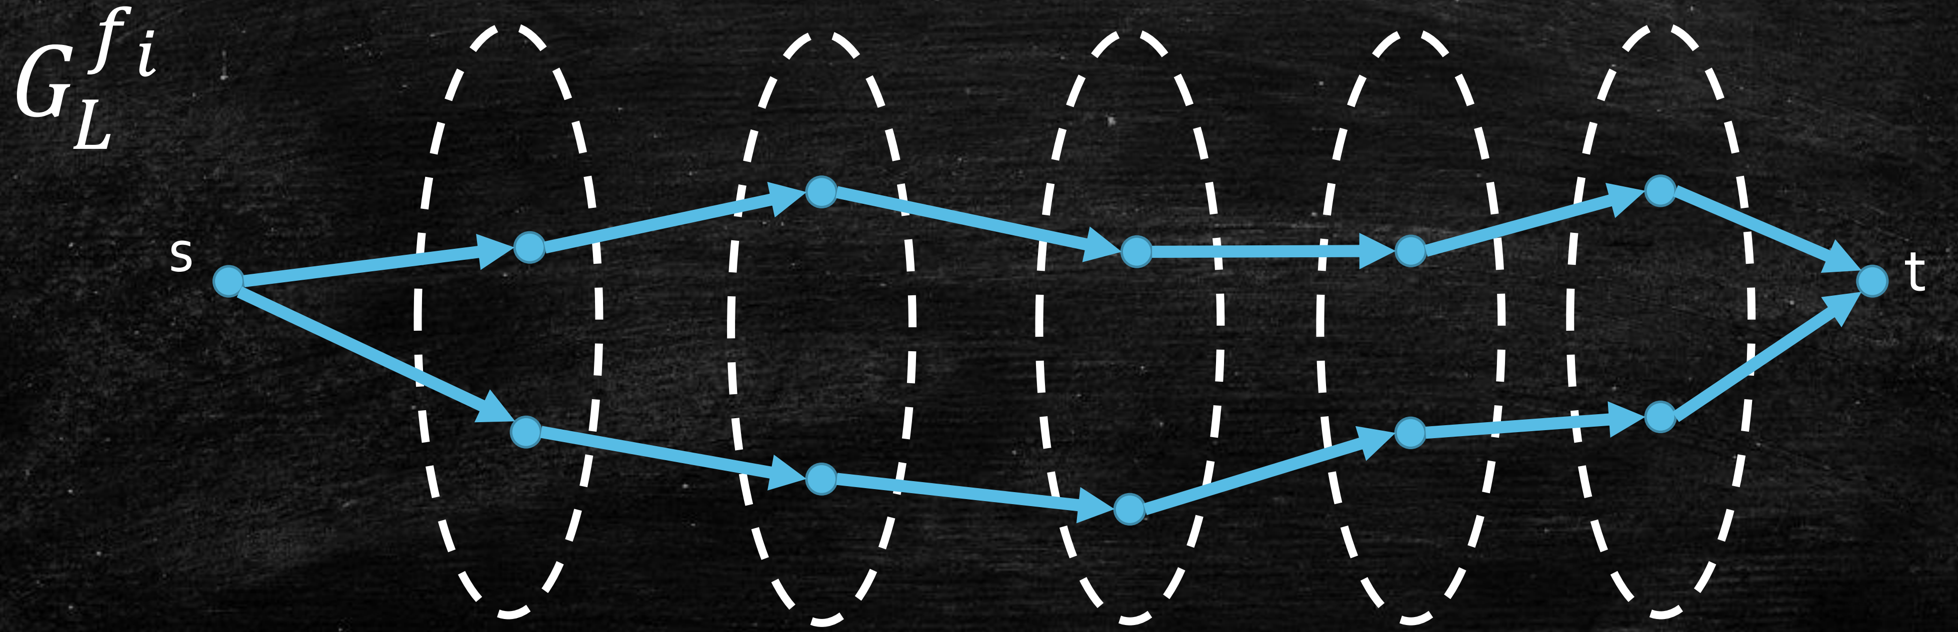

Distance Monotonicity

- After one iteration, a new edge (u, v) appearing in $G^{f_{i+1}}$ (but not in G^{f_i}) must be "backward": $\text{dist}^i(u) = \text{dist}^i(v) + 1$.

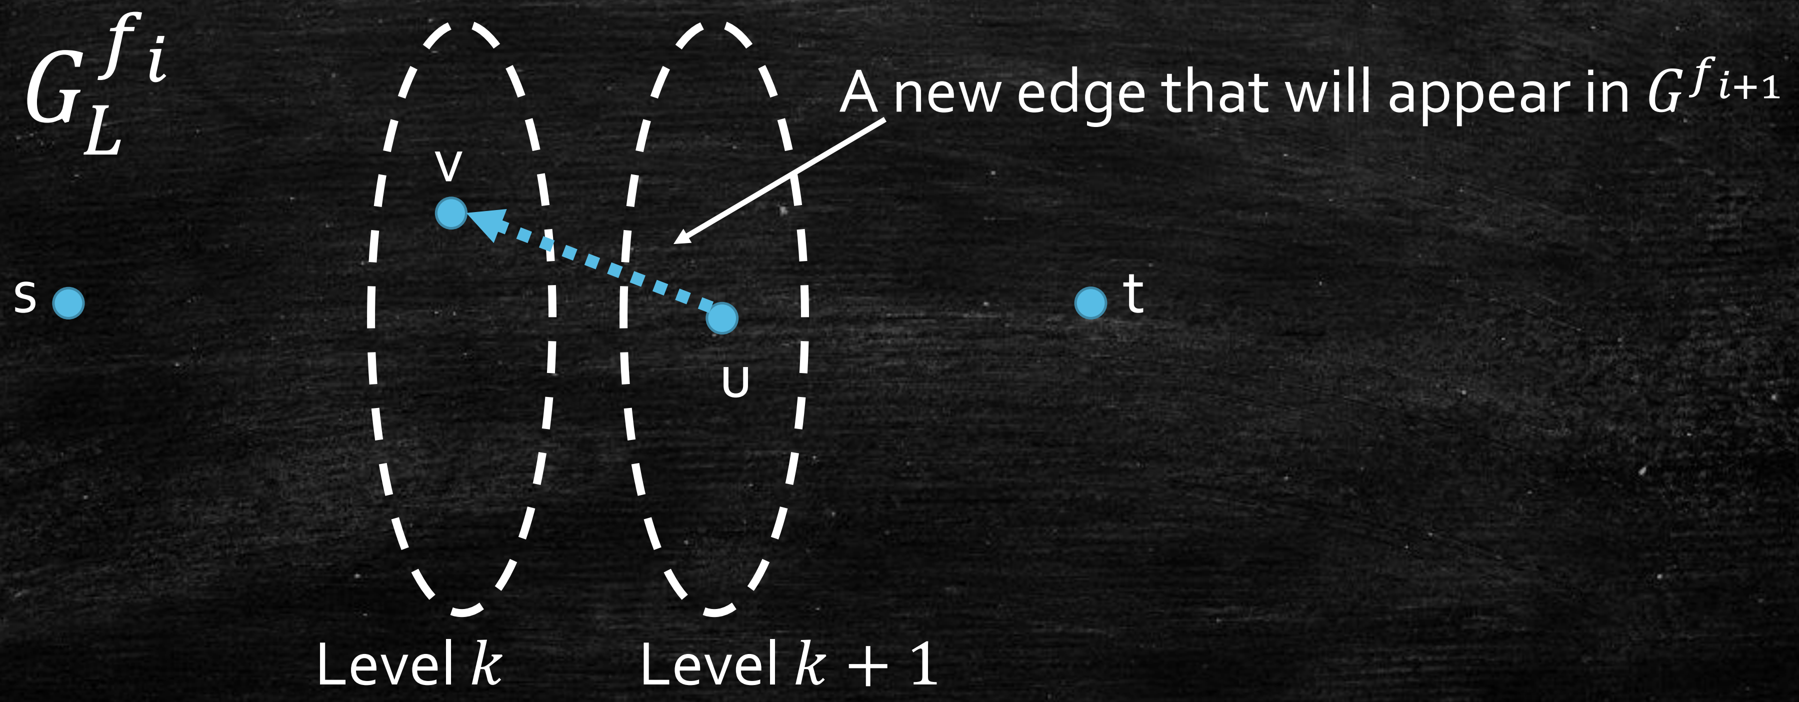

Distance Monotonicity

- After one iteration, a new edge (u, v) appearing in $G^{f_{i+1}}$ (but not in G^{f_i}) must satisfy $\mathrm{dist}^i(u) = \mathrm{dist}^i(v) + 1$.
- Such additions of edges cannot reduce the distance for any vertex!
- We again have that $\mathrm{dist}(u)$ is non-decreasing!
- Can we have strong monotonicity?

Taking a closer look...

- All the paths in $G_L^{f_i}$ with length $\mathrm{dist}^i(t)$ are “blocked” after the i-th iteration.
- Thus, a path in the $(i+1)$-th iteration must use some edges that are not in $G_L^{f_i}$.

Taking a closer look…

- This new edge may be a “backward” edge whose reverse was a critical edge in the previous iteration.
- In this case, $\mathrm{dist}(t)$ is increased by at least 2.

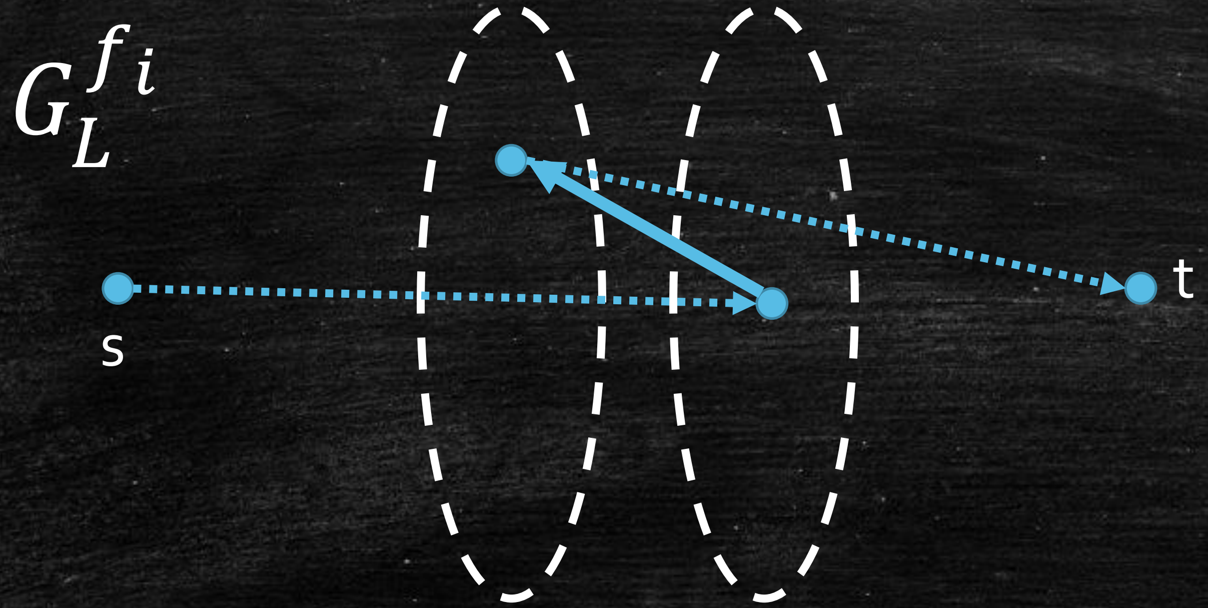

Taking a closer look...

- Or, it may be an edge in G^{f_i}, but not in $G_L^{f_i}$.
- In this case, $\mathrm{dist}(t)$ is increased by at least 1.

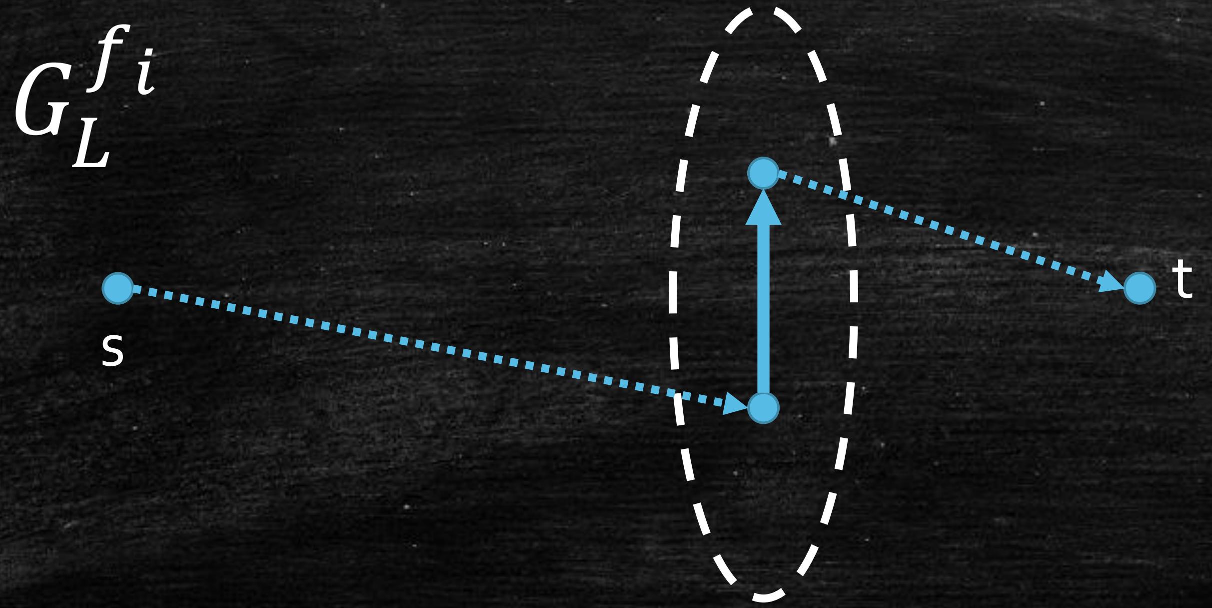

Taking a closer look…

- In both cases: $\mathrm{dist}^{i+1}(t) > \mathrm{dist}^{i}(t)$
- Let's prove it rigorously then…

Proving $\mathrm{dist}^{i+1}(t) > \mathrm{dist}^{i}(t)$

- Consider an arbitrary s-t path p in $G_L^{f_{i+1}}$ with length $\mathrm{dist}^{i+1}(t)$.
- We have $\mathrm{dist}^{i+1}(t) \geq \mathrm{dist}^{i}(t)$ by monotonicity.
- Suppose for the sake of contraction that $\mathrm{dist}^{i+1}(t) = \mathrm{dist}^{i}(t)$.
- Case 1: all edges in p also appear in $G_L^{f_i}$
- Then p is a shortest path containing no critical edges in $G_L^{f_i}$
- Contracting to the definition of blocking flow!

Proving $\mathrm{dist}^{i+1}(t) > \mathrm{dist}^{i}(t)$

- Case 2: p contains an edge (u, v) that is not in $G_L^{f_i}$
- If (u, v) was not in G^{f_i}, then (v, u) was critical in the last iteration. We have $\mathrm{dist}^{i}(u) = \mathrm{dist}^{i}(v) + 1$.
- If (u, v) was in G^{f_i} but not $G_L^{f_i}$, by the definition of level graph, we have $\mathrm{dist}^{i}(u) \geq \mathrm{dist}^{i}(v)$.
- In both cases above, $\mathrm{dist}^{i}(u) \geq \mathrm{dist}^{i}(v)$.
- We have $\mathrm{dist}^{i+1}(u) \geq \mathrm{dist}^{i}(u)$ by monotonicity,
- and we have $\mathrm{dist}^{i+1}(v, t) \geq \mathrm{dist}^{i}(v, t)$. (why?)

Proving $\mathrm{dist}^{i+1}(t) > \mathrm{dist}^{i}(t)$

- Case 2: p contains an edge (u, v) that is not in $G_L^{f_i}$
- Fact i: $\mathrm{dist}^{i}(u) \geq \mathrm{dist}^{i}(v)$.
- Fact ii: $\mathrm{dist}^{i+1}(u) \geq \mathrm{dist}^{i}(u)$.
- Fact iii: $\mathrm{dist}^{i+1}(v, t) \geq \mathrm{dist}^{i}(v, t)$.

Putting together:

$$\begin{aligned}
\mathrm{dist}^{i+1}(t) &= \mathrm{dist}^{i+1}(u) + 1 + \mathrm{dist}^{i+1}(v, t) \\
&\geq \mathrm{dist}^{i}(u) + 1 + \mathrm{dist}^{i}(v, t) && \text{(Fact ii and iii)} \\
&\geq \mathrm{dist}^{i}(v) + 1 + \mathrm{dist}^{i}(v, t) && \text{(Fact i)} \\
&\geq \mathrm{dist}^{i}(t) + 1 && \text{(triangle inequality)}
\end{aligned}$$

Putting Together…

- $\mathrm{dist}(t)$ is increased by at least 1 after each iteration.
- $\mathrm{dist}(t)$ takes value from $\{0, 1, \ldots, |V|, \infty\}$, so it can be increased for at most $O(|V|)$ times.
- Total number of iterations is at most $O(|V|)$.

Questions Remain

1. How many iterations do we need before termination?
   - $O(|V|)$
2. How do we find a blocking flow?

Finding a blocking flow in a level graph…

Iterative do the followings, until no path from s to t:

- Find an arbitrary maximal path in G_L^f starting from s:
  - At every vertex, find an arbitrary edge in G_L^f and append it to the path.
- Two possibilities:
  - End up at t: in this case, we update f (by pushing flow along the path) and remove the critical edge
  - End up at a dead-end, a vertex v with no out-going edges in G_L^f: in this case, we remove all the incoming edges of v

Finding a blocking flow in a level graph…

- At least one edge is removed after each search.
  - Total number of searches: $O(|E|)$
- Each search takes at most $|V|$ steps.
- Time complexity for each iteration of Dinic's algorithm: $O(|V| \cdot |E|)$.

Overall Time Complexity for Dinic's Algorithm

- Each iteration: $O(|V| \cdot |E|)$.
- We need at most $O(|V|)$ iterations.
- Overall time complexity for Dinic's algorithm: $O(|V|^2 \cdot |E|)$.

Other Algorithms for Max-Flow

- Improvements to Dinic's algorithm:
  - [Malhotra, Kumar & Maheshwari, 1978]: $O(|V|^3)$
  - Dynamic tree: $O(|V| \cdot |E| \cdot \log|V|)$
- Push-relabel algorithm [Goldberg & Tarjan, 1988]
  - $O(|V|^2|E|)$, later improved to $O(|V|^3)$, $O\left(|V|^2\sqrt{|E|}\right)$, $O\left(|V||E|\log\frac{|V|^2}{|E|}\right)$
- [King, Rao & Tarjan, 1994] and [Orlin, 2013]: $O(|V| \cdot |E|)$
- Interior-point-method-based algorithms:
  - [Kathuria, Liu & Sidford, 2020] $|E|^{\frac{4}{3}+o(1)}U^{\frac{1}{3}}$
  - [BLNPSSSW, 2020] [BLLSSSW, 2021] $\tilde{O}\left(\left(|E| + |V|^{\frac{3}{2}}\right)\log U\right)$
  - [Gao, Liu & Peng, 2021] $\tilde{O}\left(|E|^{\frac{3}{2}-\frac{1}{328}}\log U\right)$

Hopcroft–Karp–Karzanov algorithm

- Find a maximum bipartite matching in $O\left(|E| \cdot \sqrt{|V|}\right)$ time.
- Proposed independently by Hopcroft-Karp and Karzanov.
- Can be viewed as a special case of Dinic's algorithm.

Conversion to Max-Flow Problem

Set the capacity to 1 for all edges.

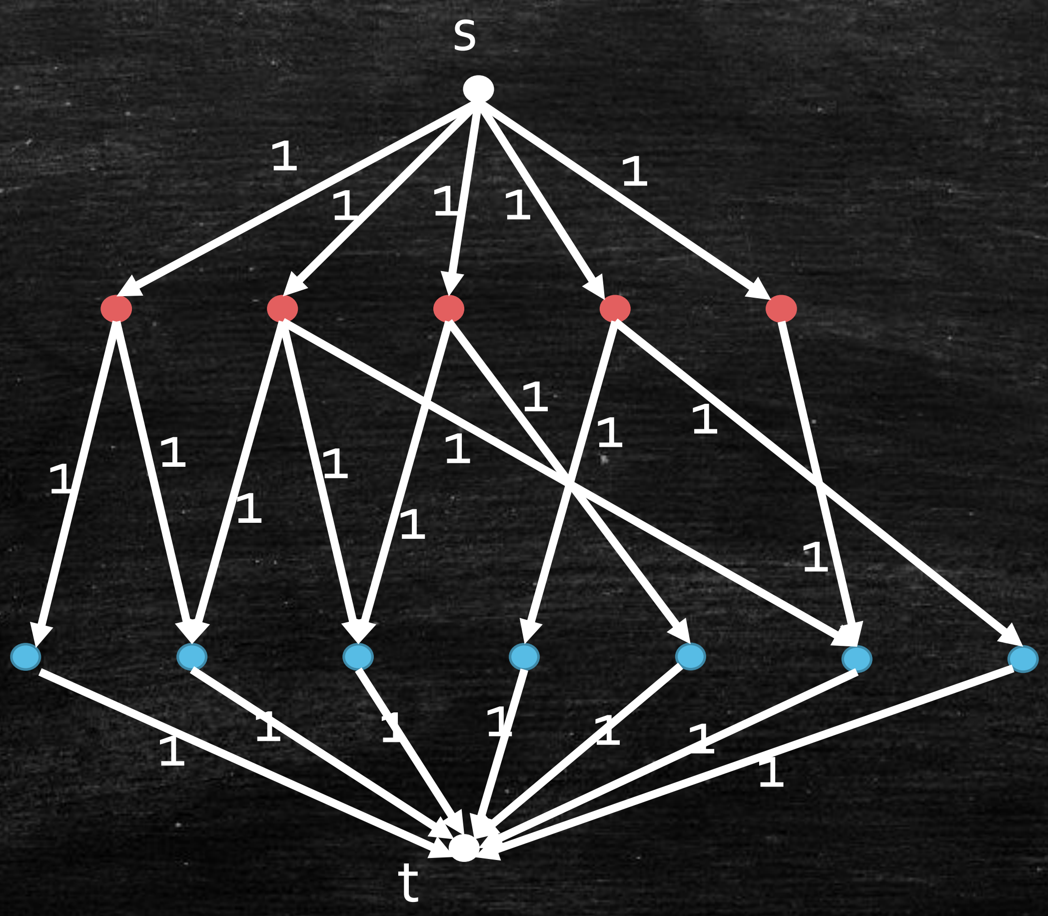

Conversion to Max-Flow Problem

Dinic's algorithm runs in $O\left(|E| \cdot \sqrt{|V|}\right)$ time for this special case.

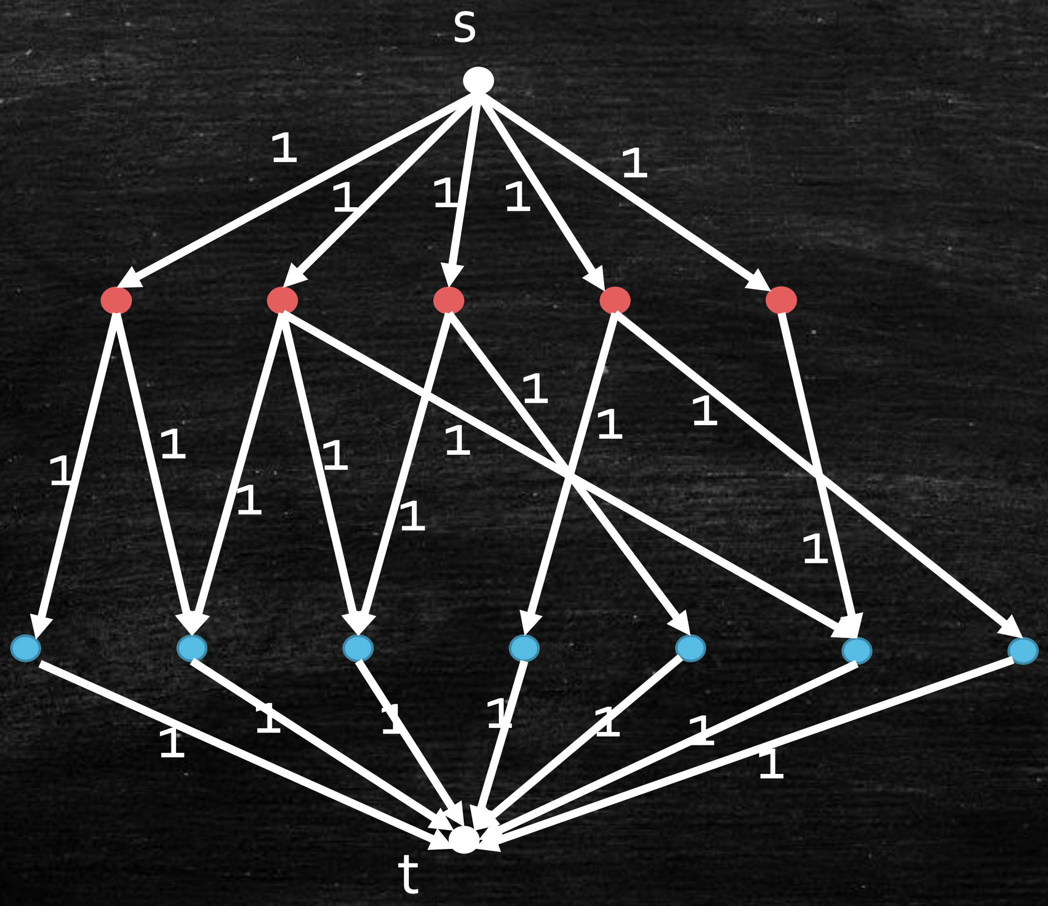

Conversion to Max-Flow Problem

- Integrality theorem also holds for Dinic's algorithm:
  - The flow output by Dinic's algorithm in our case is integral.
- We aim to show Dinic's algorithm runs in $O\left(|E| \cdot \sqrt{|V|}\right)$ time.
- Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.
- Step 2: Number of iterations is at most $2\sqrt{|V|}$.

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Iteratively do the followings, until no path from s to t:

- Perform DFS from s
- If we reach t, delete all edges on the s-t path (why can we do this?) and start over from s.
- If we ever go backward, delete the edge just travelled. (why can we do this?)

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

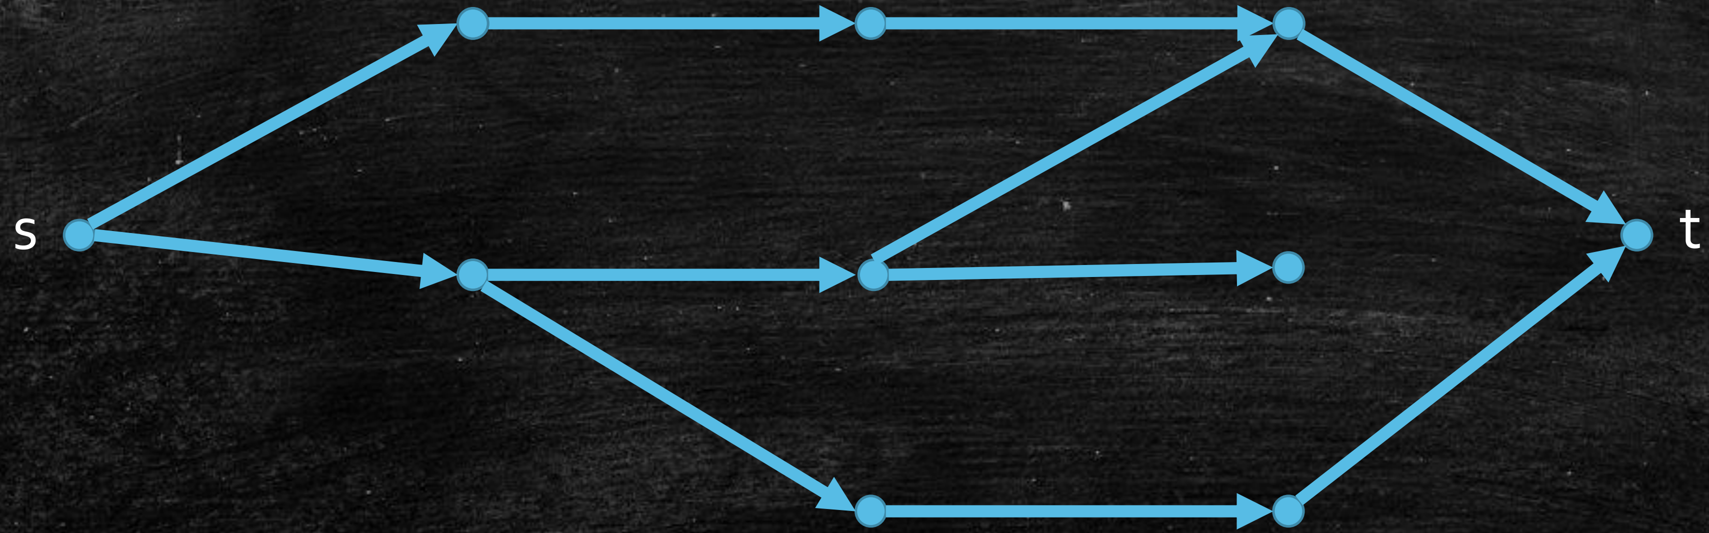

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

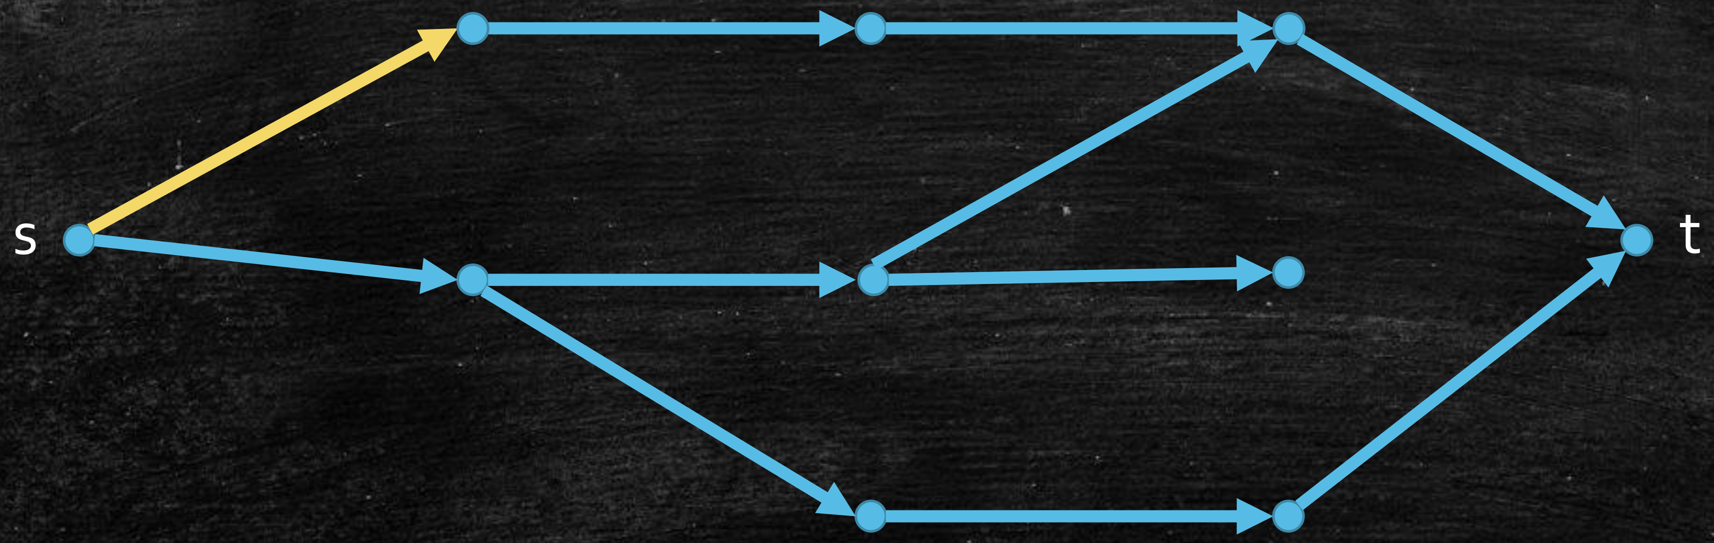

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

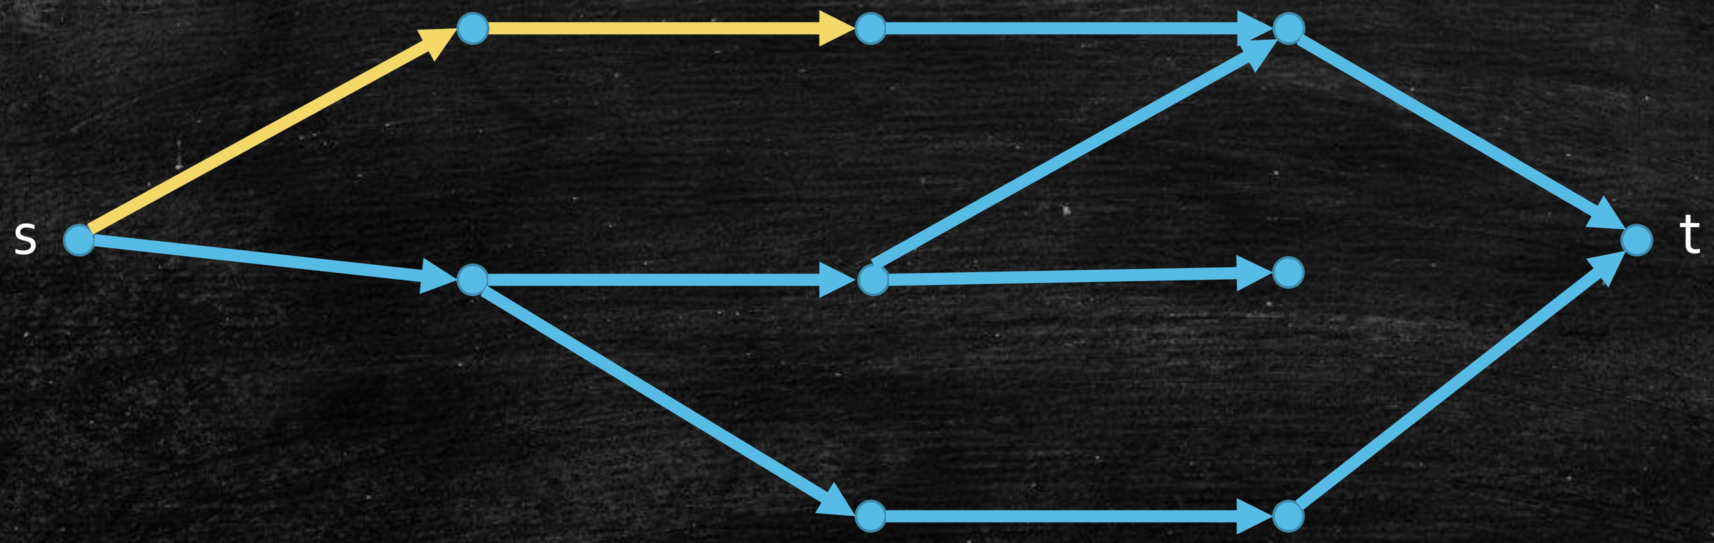

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

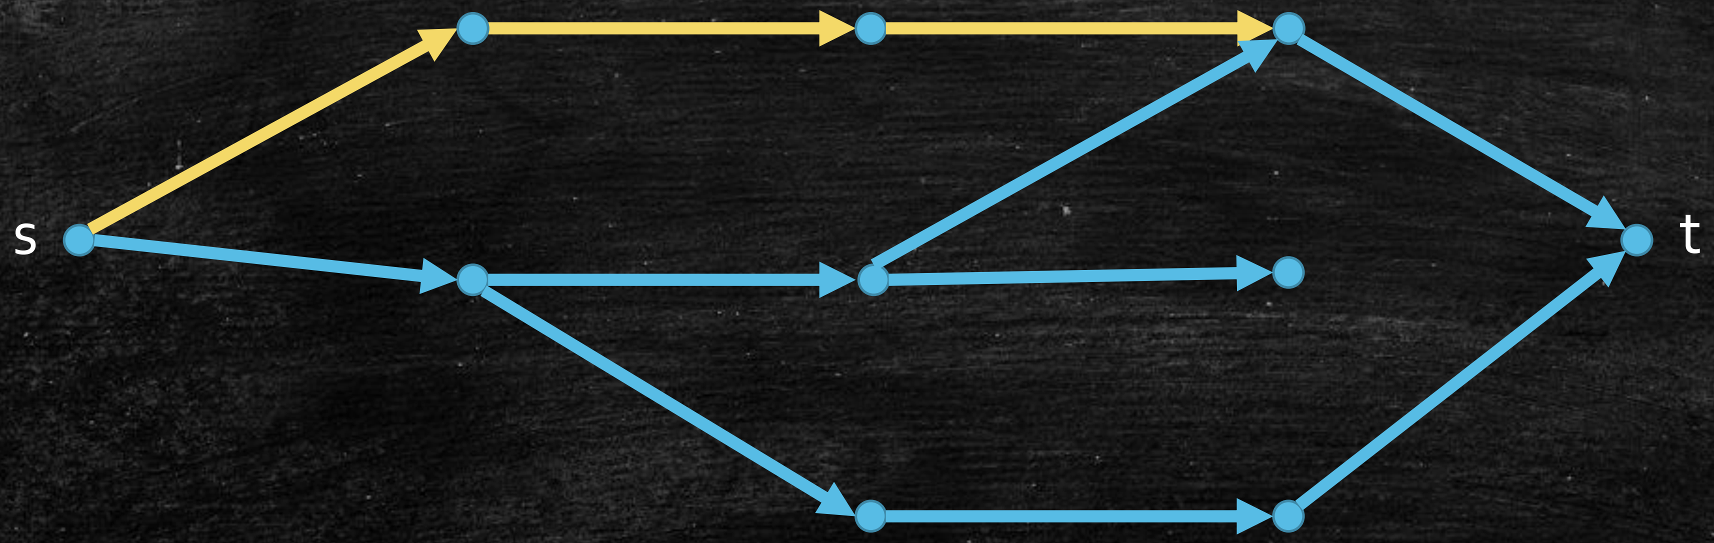

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

An s-t path is found, remove all edges from the path.

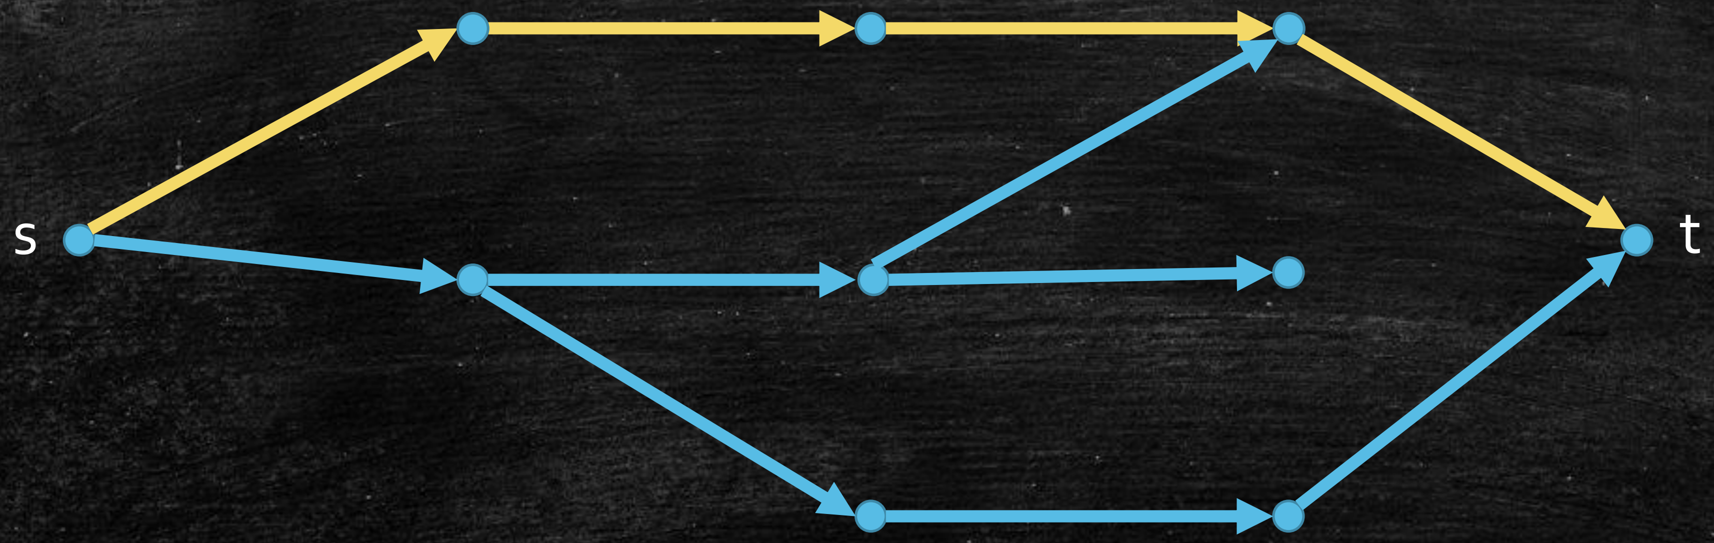

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Start over…

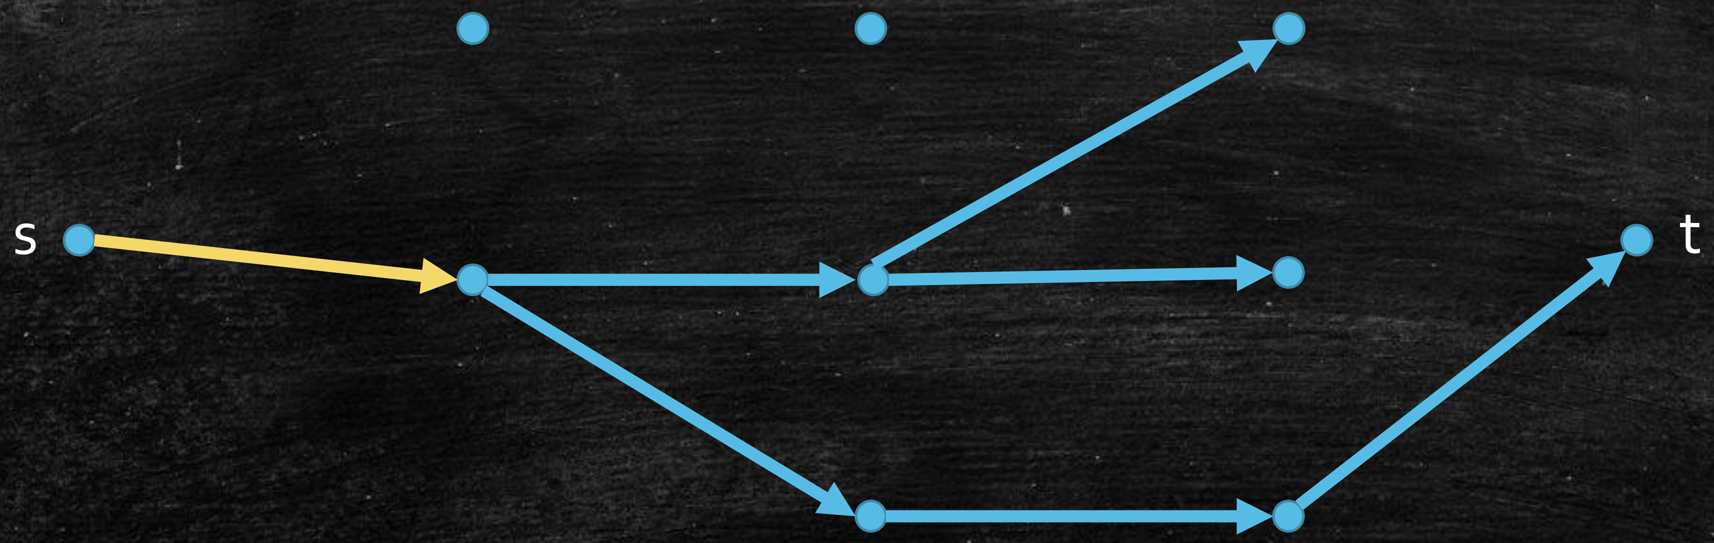

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

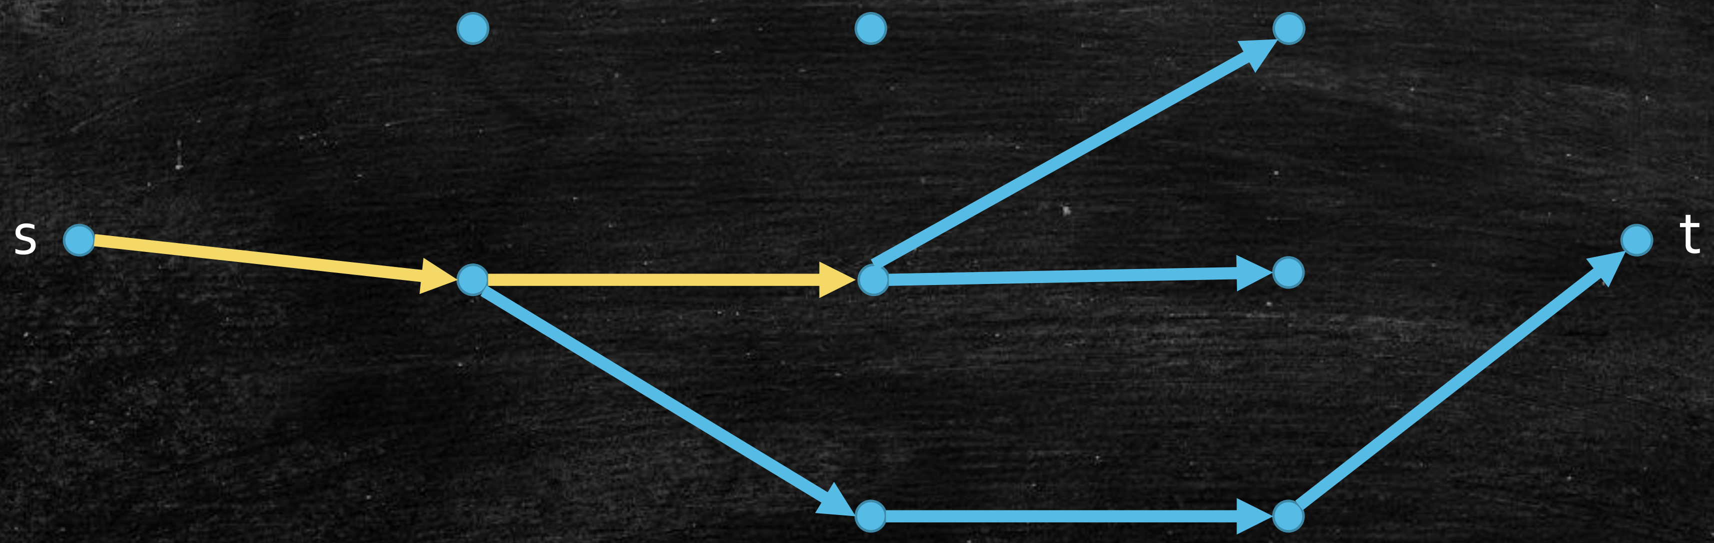

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

We have to go backward now; delete the edge just travelled.

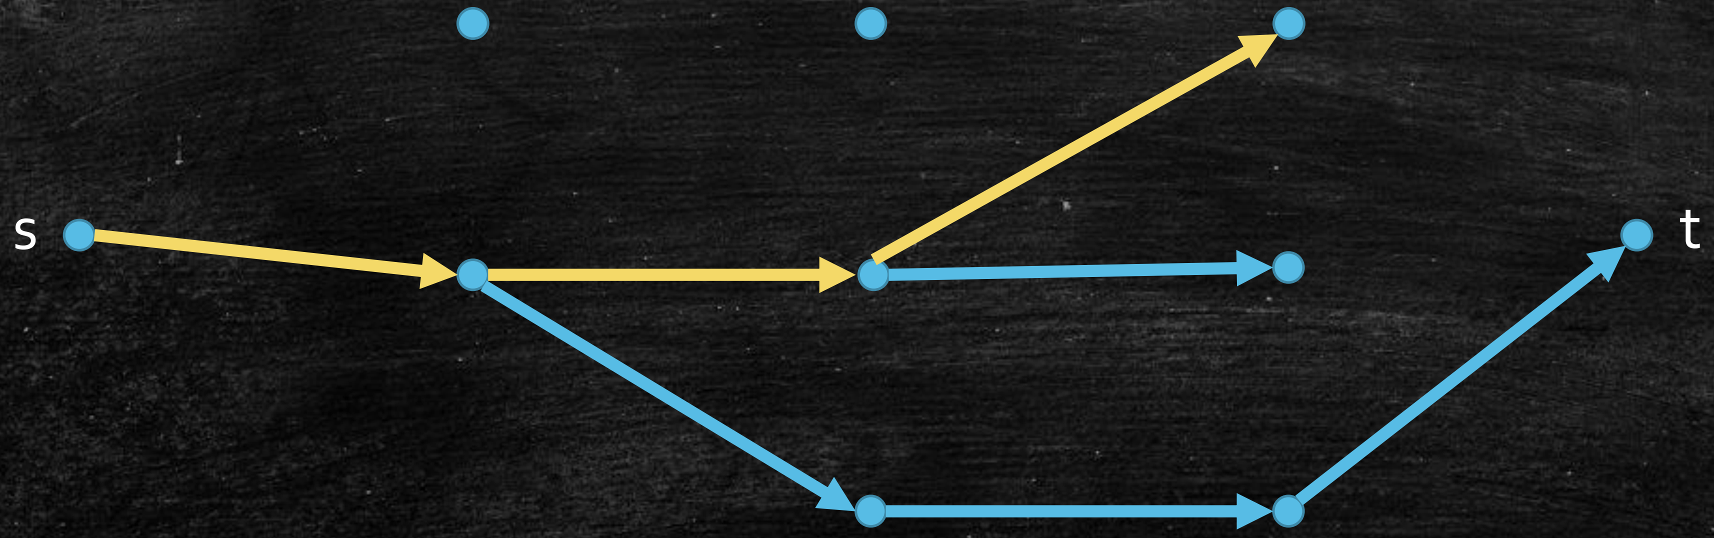

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

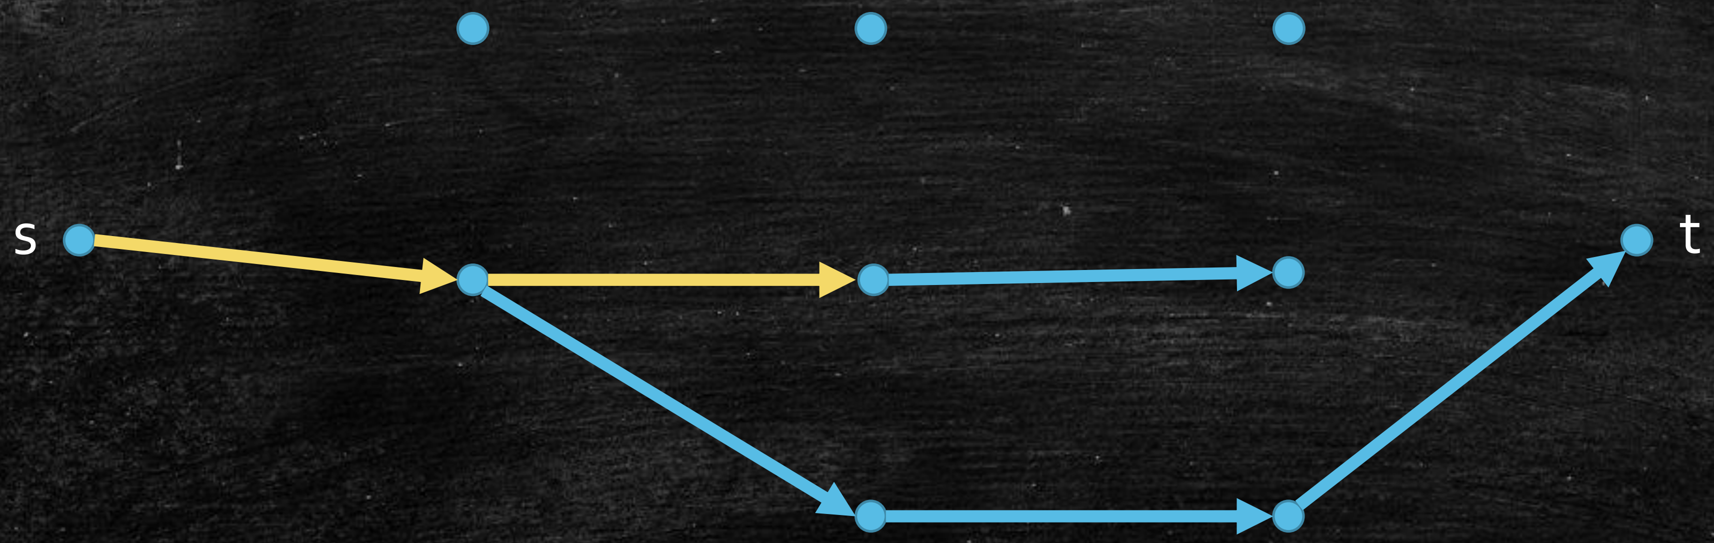

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Again, we have to go backward; delete the edge just travelled.

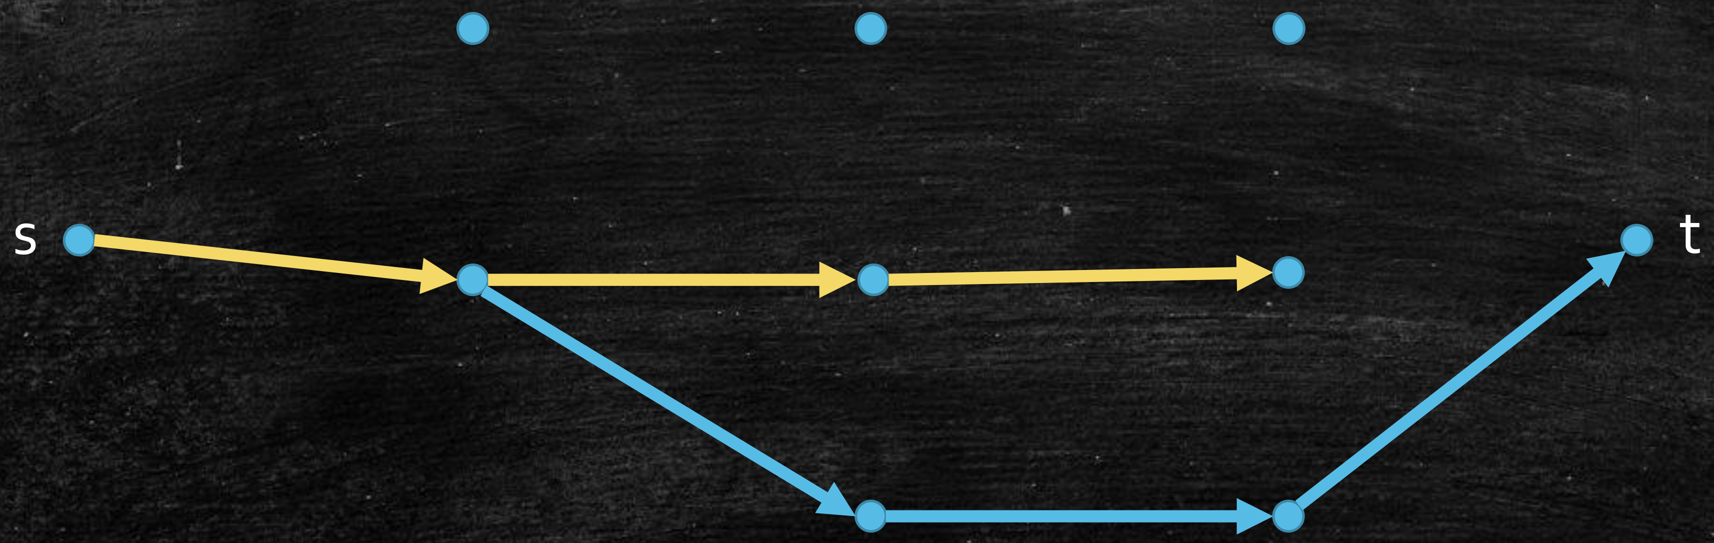

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Again, we have to go backward; delete the edge just travelled.

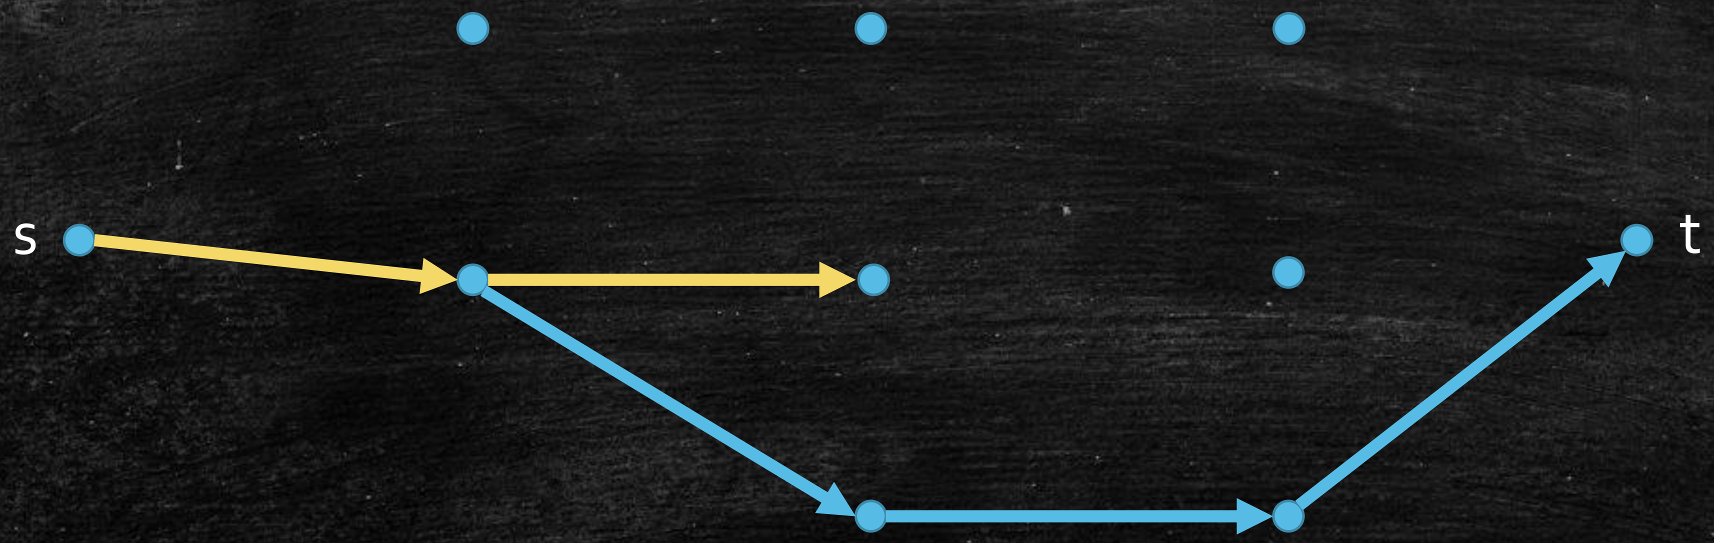

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Again, we have to go backward; delete the edge just travelled.

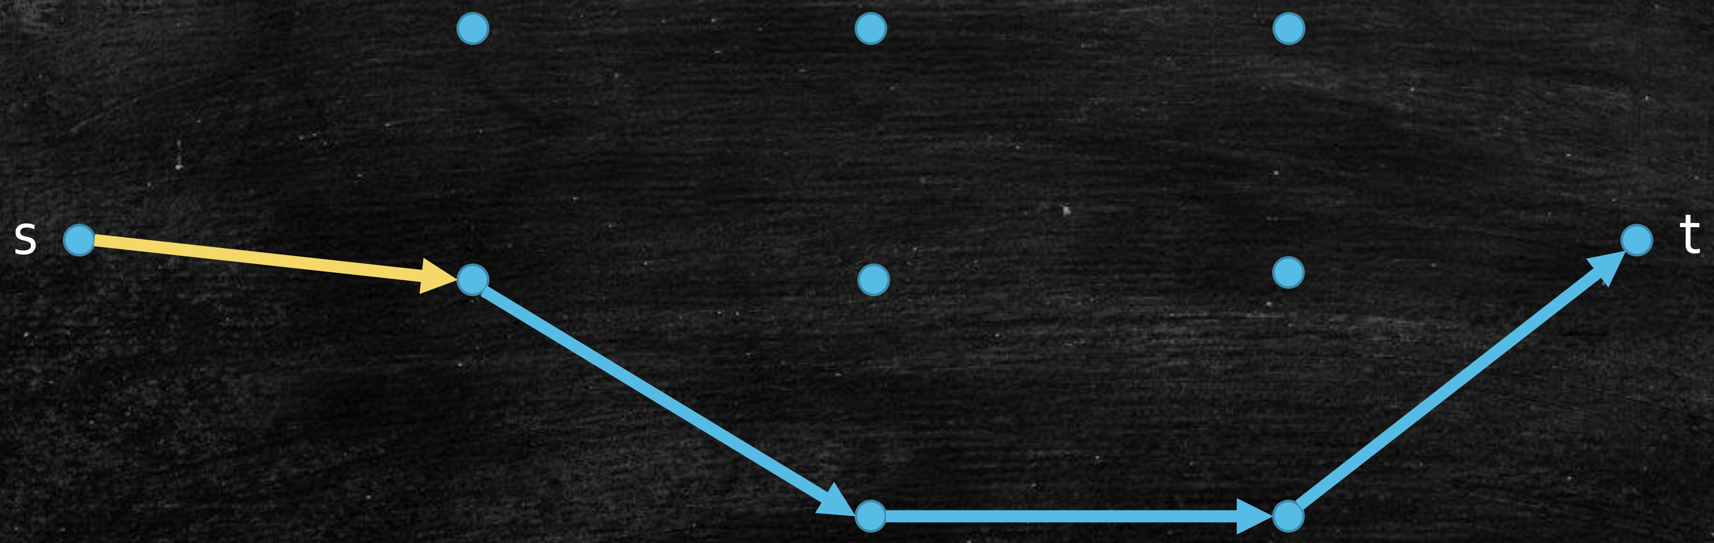

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

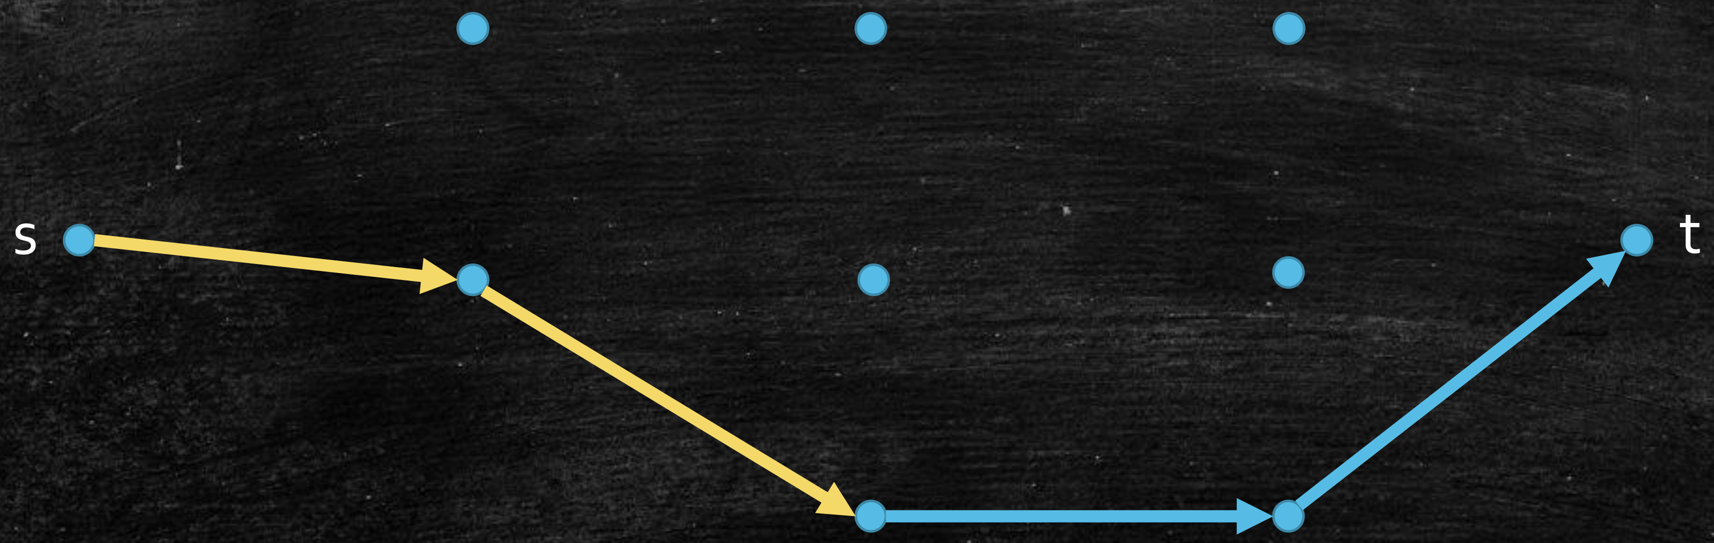

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

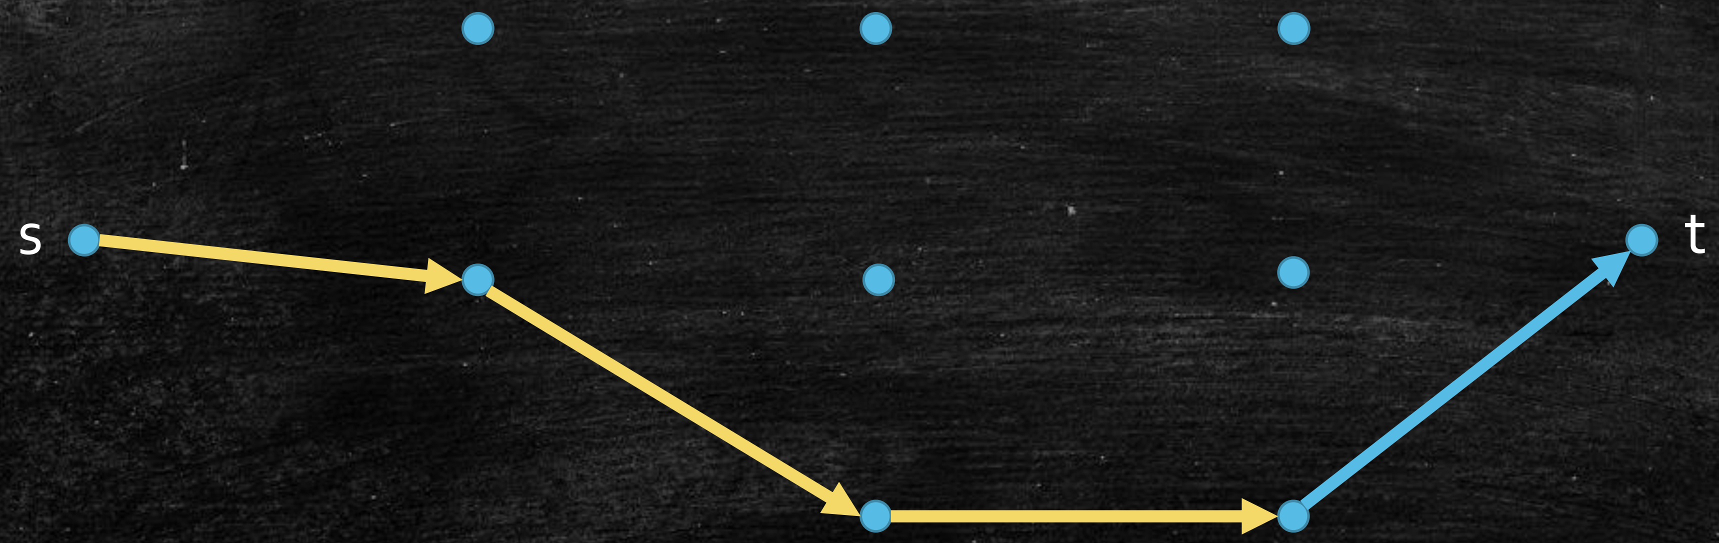

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Find another s-t path; delete all edges on the path

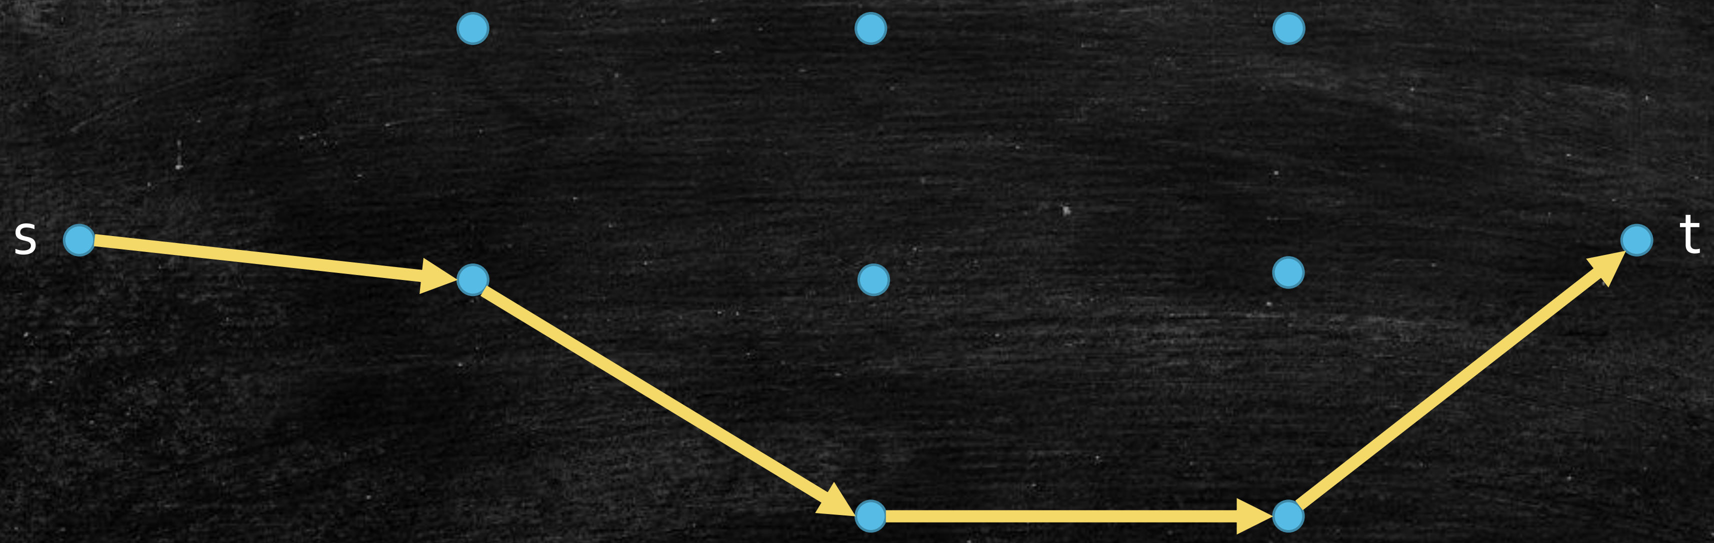

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

We are done!

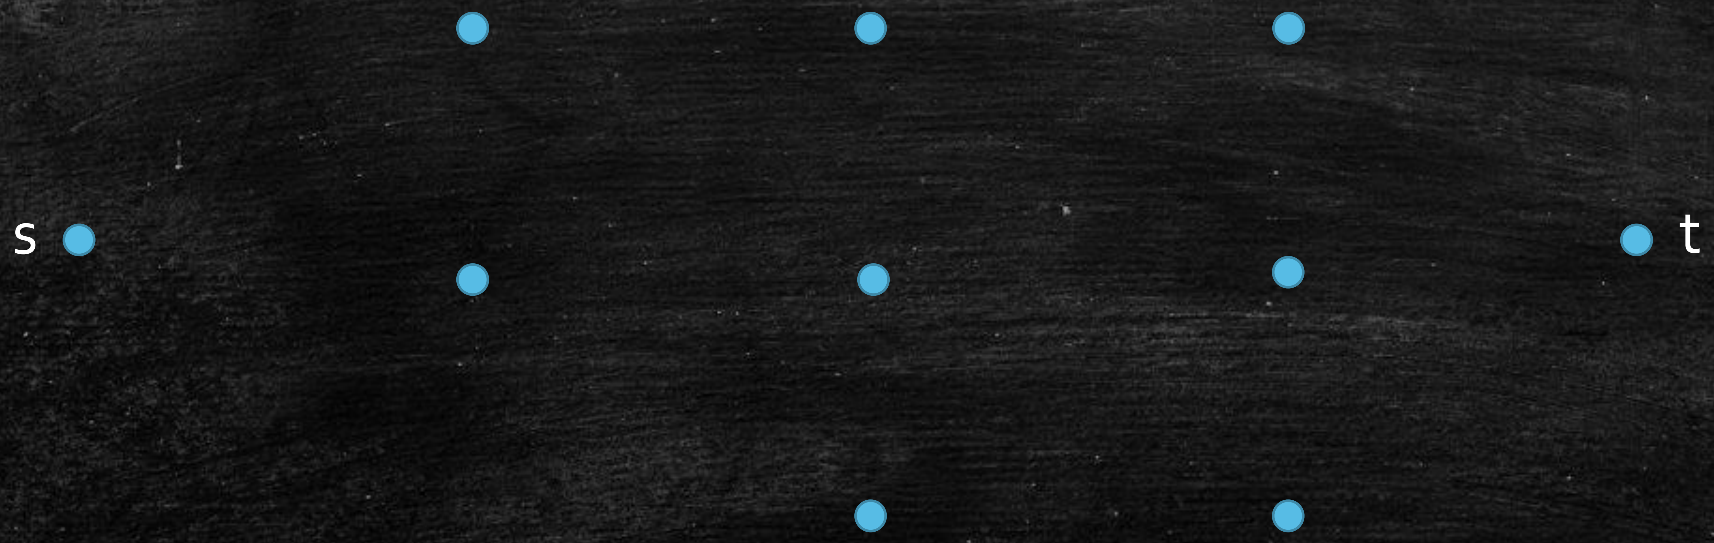

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

We have obtained a blocking flow!

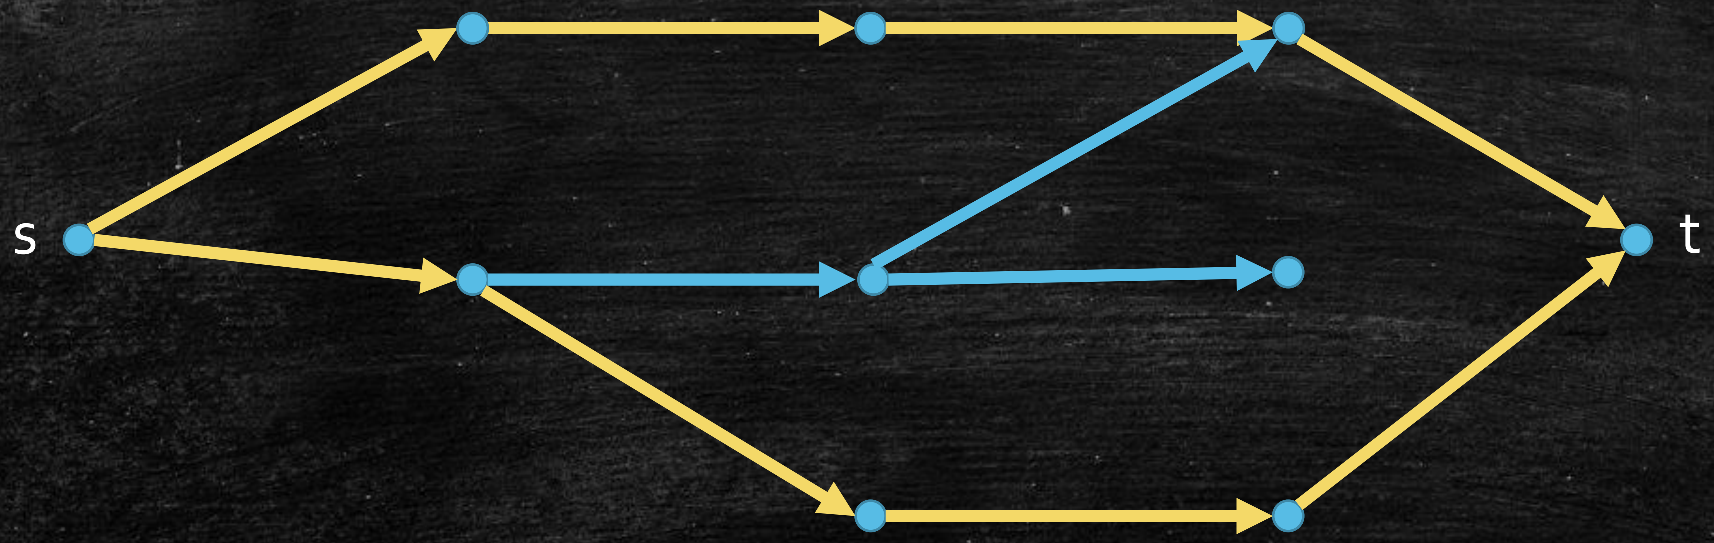

Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.

Time complexity: $O(|E|)$

- Each edge is visited at most once.

Step 2: Number of iterations is at most $2\sqrt{|V|}$.

- If the algorithm terminates within $\sqrt{|V|}$ iterations, we are already done!
- Otherwise, let f be the flow after $\sqrt{|V|}$ iterations.
- Claim: the maximum flow in G^f has value at most $\sqrt{|V|}$.

Observation on G^f

- In each iteration, for each $v \in V \setminus \{s, t\}$, either its in-degree is 1, or its out-degree is 1.
- Proof. By Induction…
- At the beginning, this is clearly true.

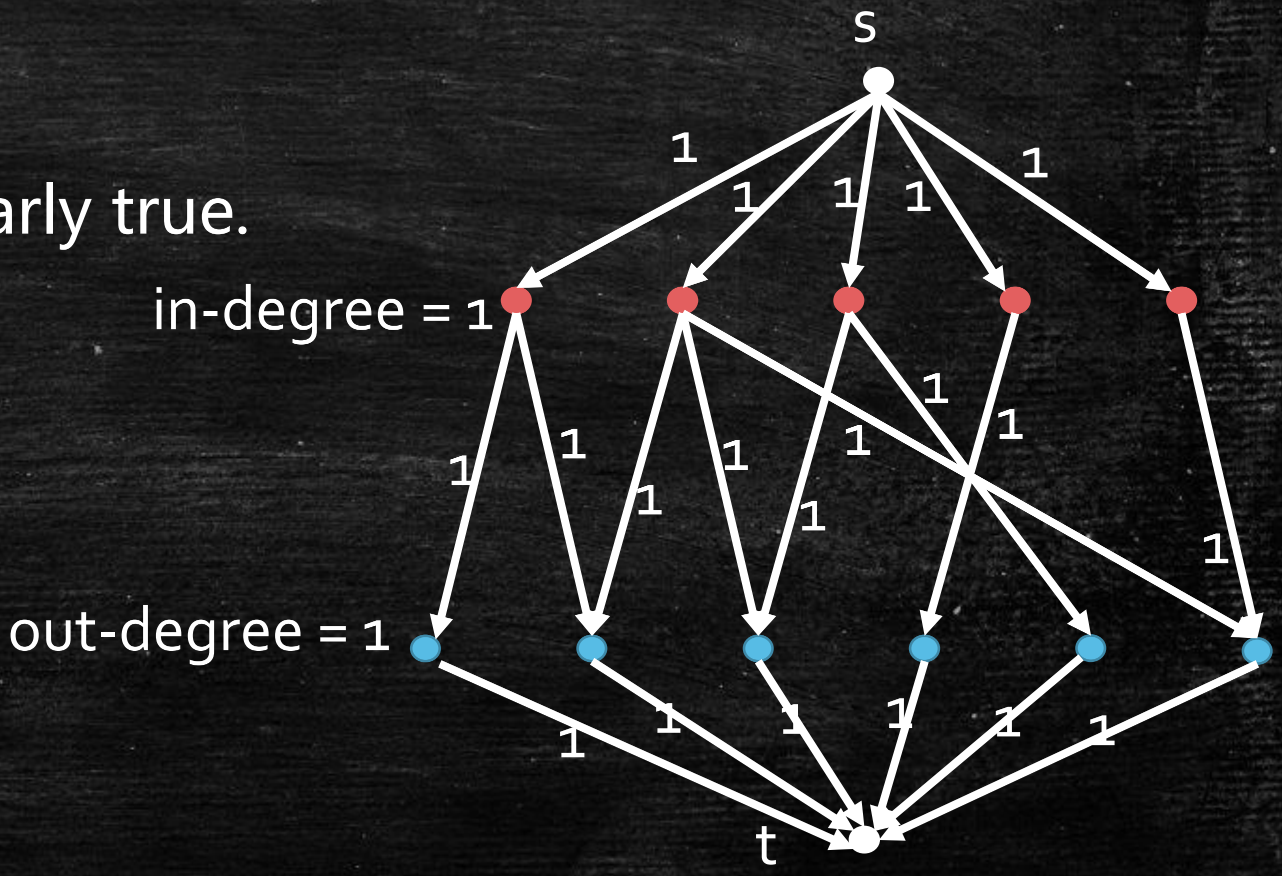

Observation on G^f

- In each iteration, for each $v \in V \setminus \{s, t\}$, either its in-degree is 1, or its out-degree is 1.
- Proof. At the beginning, this is clearly true.
- For each iteration, the amount of flow going through v is either 0 or 1.
- If it is 0, v's in-degree and out-degree are unchanged.
- Otherwise, exactly one in-edge and one out-edge are flipped; the property is still maintained.

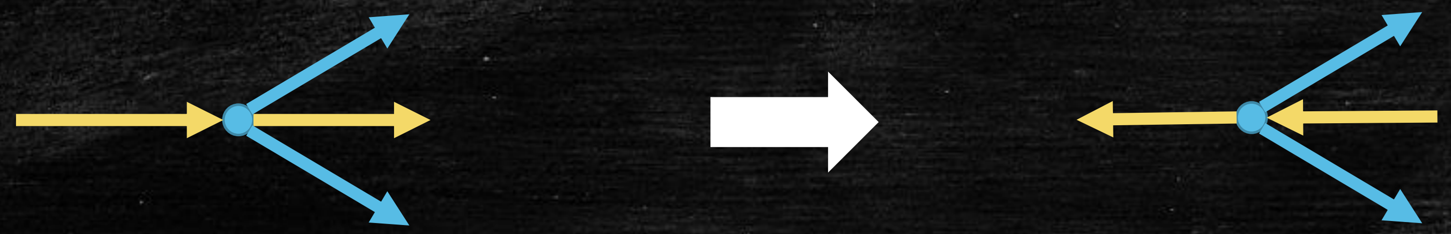

The maximum flow in G^f has value at most $\sqrt{|V|}$

- Integrality Theorem: there exists a maximum integral flow f' in G^f.
- f' consists of edge-disjoint paths.
- Edge-disjointness implies vertex-disjointness by previous observation on G^f.

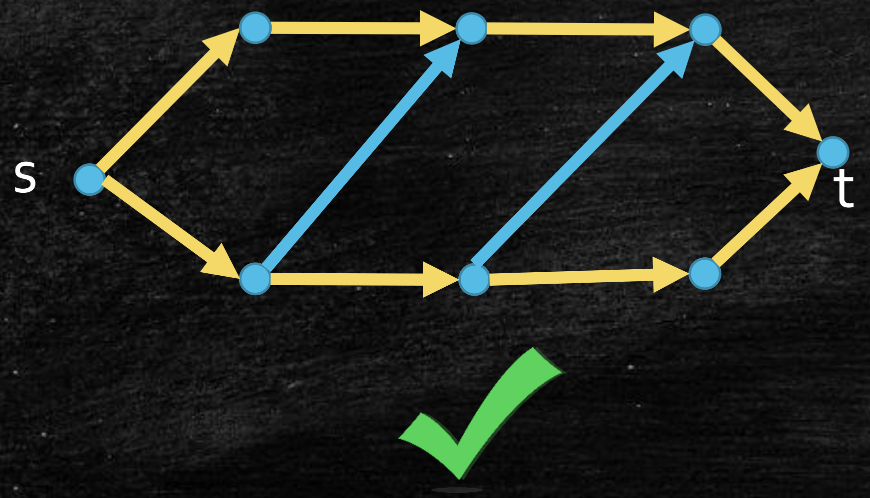

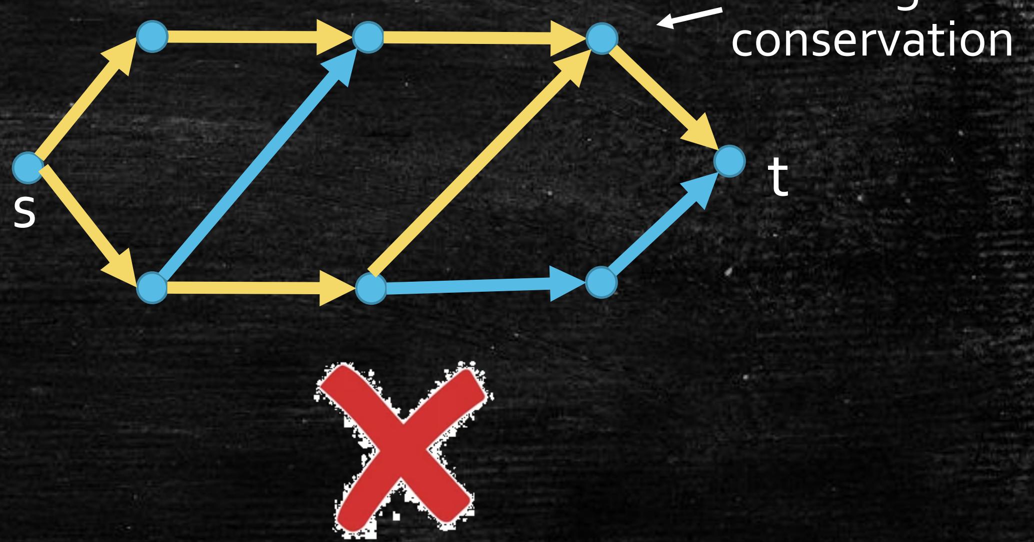

The maximum flow in G^f has value at most $\sqrt{|V|}$

- Max-flow on G^f, f', is integral and consists of edge-disjoint paths.
- By our analysis to Dinic's algorithm, $\mathrm{dist}^{G^f}(s,t) \geq \sqrt{|V|}$.
- Each path in f' has length at least $\sqrt{|V|}$.
- There are at most $\frac{|V|}{\sqrt{|V|}} = \sqrt{|V|}$ paths in f' by vertex-disjointness.
- $v(f') \leq \sqrt{|V|}$

Step 2: Number of iterations is at most $2\sqrt{|V|}$.

- If the algorithm terminates within $\sqrt{|V|}$ iterations, we are already done!
- Otherwise, let f be the flow after $\sqrt{|V|}$ iterations.
- Claim: the maximum flow in G^f has value at most $\sqrt{|V|}$.
- Each iteration increase the value of flow by at least 1.
- Thus, the algorithm will terminates within at most another $\sqrt{|V|}$ iterations.
- Total number of iterations: $2\sqrt{|V|}$.

Putting Together...

- Step 1: Finding a blocking flow in a level graph takes $O(|E|)$ time.
- Step 2: Number of iterations is at most $2\sqrt{|V|}$.
- Overall time complexity: $O\left(|E| \cdot \sqrt{|V|}\right)$

Today's Lecture

Maximum Flow Problem:

- Edmonds-Karp Algorithm
  - Implement Ford-Fulkerson method by BFS
  - $O(|V| \cdot |E|^2)$
- Dinic's Algorithm
  - Push flow on multiple paths at one iteration
  - Level graph and blocking flow
  - $O(|V|^2 \cdot |E|)$

Maximum Bipartite Matching Problem:

- Hopcroft–Karp–Karzanov algorithm
  - Apply Dinic's algorithm
  - $O\left(|E| \cdot \sqrt{|V|}\right)$