Warm-starting Push-Relabel

Sami Davies (UC Berkeley), Sergei Vassilvitskii (Google Research—New York),
Yuyan Wang (Google Research—New York)

Network flow

- Transportation
- Intrusion detection
- Network connectivity

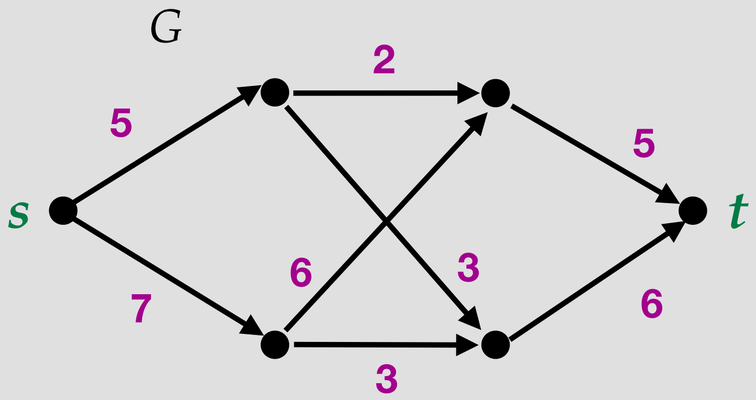

Flows satisfy capacity and conservation constraints

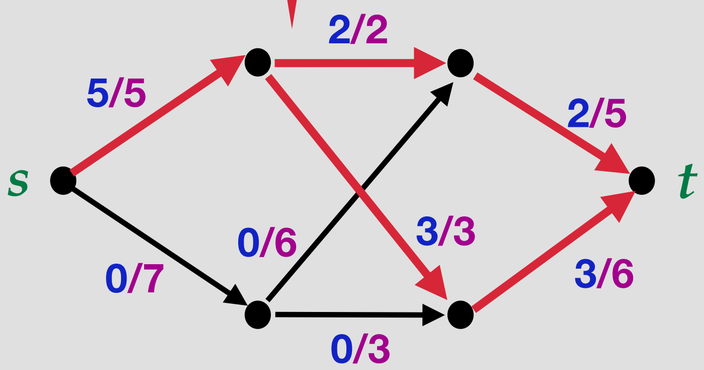

Push-Relabel

Height of a node u estimates length of shortest path between u and t in the residual graph

Push-Relabel (high level)

- Keeps track of a **pre-flow** on edges and **heights** for the nodes
- Algorithm is iterative–while there is a u node with **excess** flow
    1. Push flow from u with height h to a node with height h-1 (send flow "downhill"/ closer to t) or
    2. Increase height of u

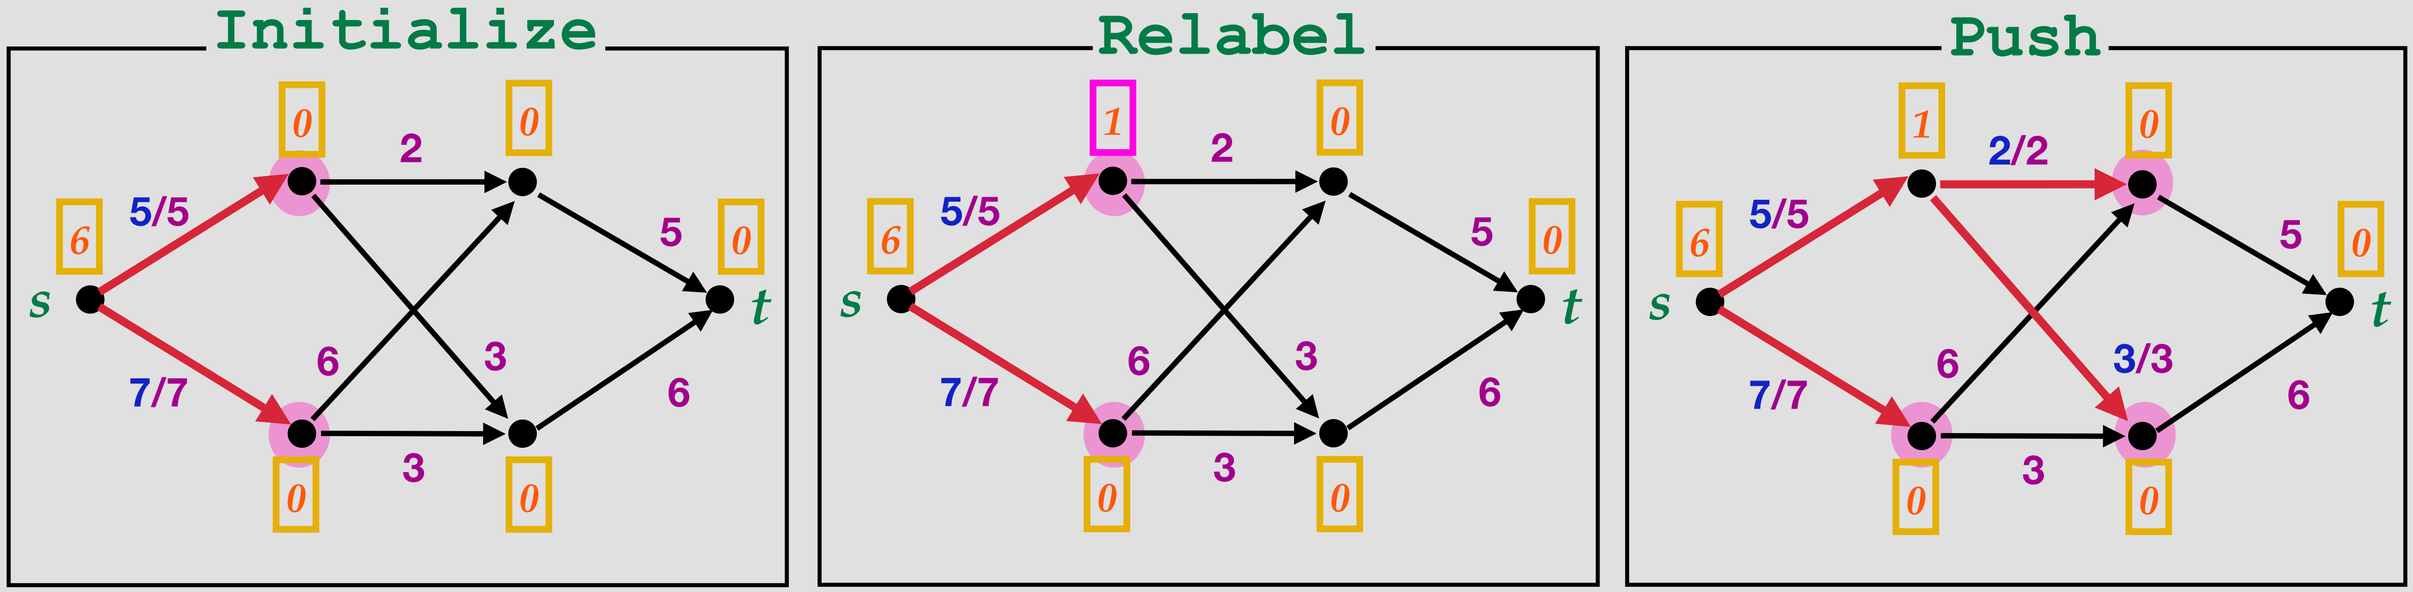

Push-Relabel

Push-Relabel (high level)

- Keeps track of a **pre-flow** on edges and **heights** for the nodes
- Algorithm is iterative—while there is a u node with **excess** flow
    1. Push flow from u with height h to a node with height h-1 (send flow "downhill"/ closer to t) or
    2. Increase height of u

$O(|V|^2|E|)$ running time [Goldberg, Tarjan '86] (or faster $O(|V|^2\sqrt{|E|})$ with heuristics)

Considered the benchmark for max flow algorithms in practice

Why is Push-Relabel so much better in practice than its theoretical guarantees?

Learning-augmented algorithms

Goal: find book by author Lewis Carroll, n books total

Vanilla binary search

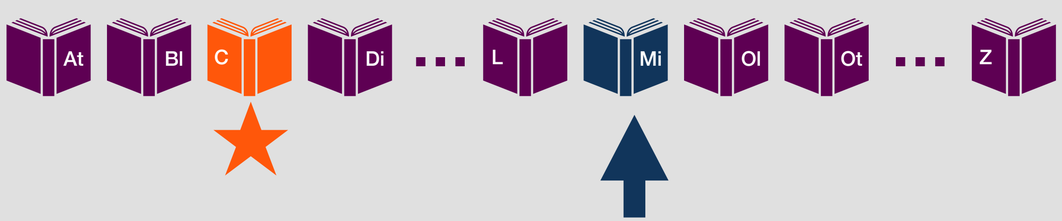

Start in the middle

- Run-time $O(\log n)$

Learning-augmented binary search

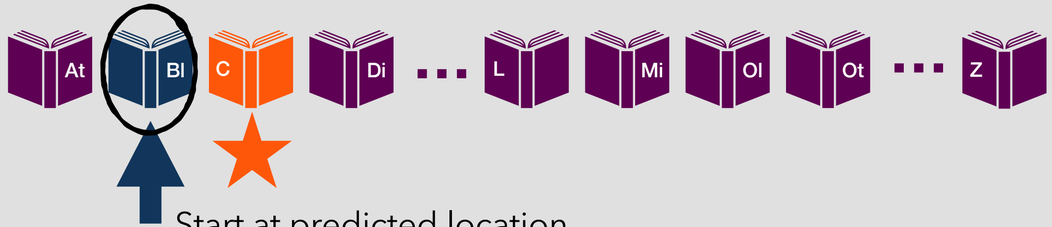

Start at predicted location

- Run-time $O(\log err)$

Algorithm has access to a learned **prediction**
Prediction can guide the algorithm's decisions

Learning-augmented algorithms

Goal: find book by author Lewis Carroll, n books total

Vanilla binary search

Learning-augmented binary search

>200 papers in past 5 years!

https://algorithms-with-predictions.github.io/

- Learned indexes [Kraska et al PODS18]
- Caching [Lykouris and Vassilvitskii JACM21]
- Social welfare maximization [Banerjee et al SODA22]

Learning-augmented Push-Relabel

First proof that the *gap-relabeling heuristic* (popular in practice) can improve the performance of PR!

Main results:

- Can improve running-time of Push-Relabel with good predicted flows

  With predicted flow $\hat{f}$, find optimal $f^\star$ in time $O(|V|^2 \cdot min\{ ||\hat{f}\text{-}f^\star||_1, |E| \})$.

- If optimal $f^\star$ has value $\leq \eta$, can find an optimal flow in time $O(\eta \cdot |V|^2)$.

- Apply our algorithm to image segmantation instances. Empirically obtain speed up using predicted flows, theory is predictive of practice!