

Algorithms - 2020

Flows (cont)



Recap

- Finish ch 10 today,
start 11 Wednesday
- HW is due Friday
- Reading as usual

Last time:

Thm: Max-flow = Min cut

Algorithm : Ford-Fulkerson (1956)

MAX FLOW (G):

Let $f(e) = 0$ initially $\forall e$

Compute $G_f \leftarrow$ residual graph

While there is s-t path in G_f :


Let P be a simple s-t path

$f' \leftarrow \text{augment}(f, P)$

$f \leftarrow f'$

update G_f

return f



WFS
+ updating edges
push as much flow on that path as we can

Runtime:


each time get ≥ 1 more flow

so repeats $O(f^*)$ max flow

Total: $O((V+E)f^*)$

Next:

Flow decomposes into
paths.



[use F-F alg. to see
these paths

Faster versions (?)

This is an active area of research!

We'll see two faster examples, both (relatively) simple variations on the Ford-Fulkerson algorithm:

① Edmonds - Karp:
choose largest bottleneck
edge


↳ $O(E^2 \log E \log |F^*|)$

② Shortest ^{# of edges} augmenting path ← BFS tree
↳ $O(VE^2)$

Edmonds-Karp:

Largest bottleneck:
how?

Take G_f :



Grow a tree from s , adding
largest edge out, each
time

Runtime: use heap

$E \log E$

E-K

~~MAXFLOW~~ (G):

Let $f(e) = 0$ initially $\forall e$
Construct G_f

While there is s-t path in G_f :

~~let P be a simple s-t path~~

$f' \leftarrow \text{augment}(f, P)$

$f \leftarrow f'$

update G_f

$O(E+V)$

return f

Let P be the largest
bottleneck path
in G_f

$O(E \log E)$
using
heap

of repetitions in loop?

In FF, went down by ≥ 1 .


Here? Hopefully better!

Key: look at current flow

loop repetitions :

Consider current flow:

In middle of while



Residual graph : Also a flow graph!

Let f' be max flow in G_f .

G_f has f' flow & at most E paths from s to t .

\Rightarrow One of them is "big" *decompose into paths*

$P_1, P_2, P_3, \dots, P_E$

Sum = f'

if all $< \frac{E}{f'}$, then sum would be too small.

$$\boxed{\geq \frac{E}{f'}}$$

So: path must add $\geq \frac{f'}{E}$ to flow
 push more flow along path
 \Rightarrow next residual graph
 will have $\leq (1 - \frac{1}{E})f$ in it.

$$f' - \frac{f'}{E}$$

push at least enough
to be below this

$\hookrightarrow (1 - \frac{1}{E})^l f^*$ repetitions flow
 after l repetitions

What is l ?

rounds

Well, $l = E \ln f^*$ repetitions,

$$(1 - \frac{1}{E})^l f^* < 1$$

math!

Wait: < 1 ??

Integral capacities!

so if < 1 , must $= 0$

\Rightarrow resid. graph is now disconnected

So: total runtime:

$EK(G)$:

Let $f(e)=0$ initially $\forall e$
Construct G_f

While there is s-t path in G_f :

~~let P be a simple s-t path~~

$f' \leftarrow \text{augment}(f, P)$

$f \leftarrow f'$

update G_f

return f

Let P be the largest bottleneck path

Time: $E \log E + O(V+E)$

repetitions: $E \ln f^*$

Total:

$E^2 \log E \ln f^*$

Shortest paths

MAX FLOW (G):

Let $f(e) = 0$ initially $\forall e$
Construct G_f

While there is s-t path in G_f

~~let P be a simple s-t path~~

$f' \leftarrow \text{augment}(f, P)$

$f \leftarrow f'$
update G_f

return f

$O(V+E)$

Strange - no mention
of amt of flow

Let $P =$ shortest s-t path
(# edges, not capacities)

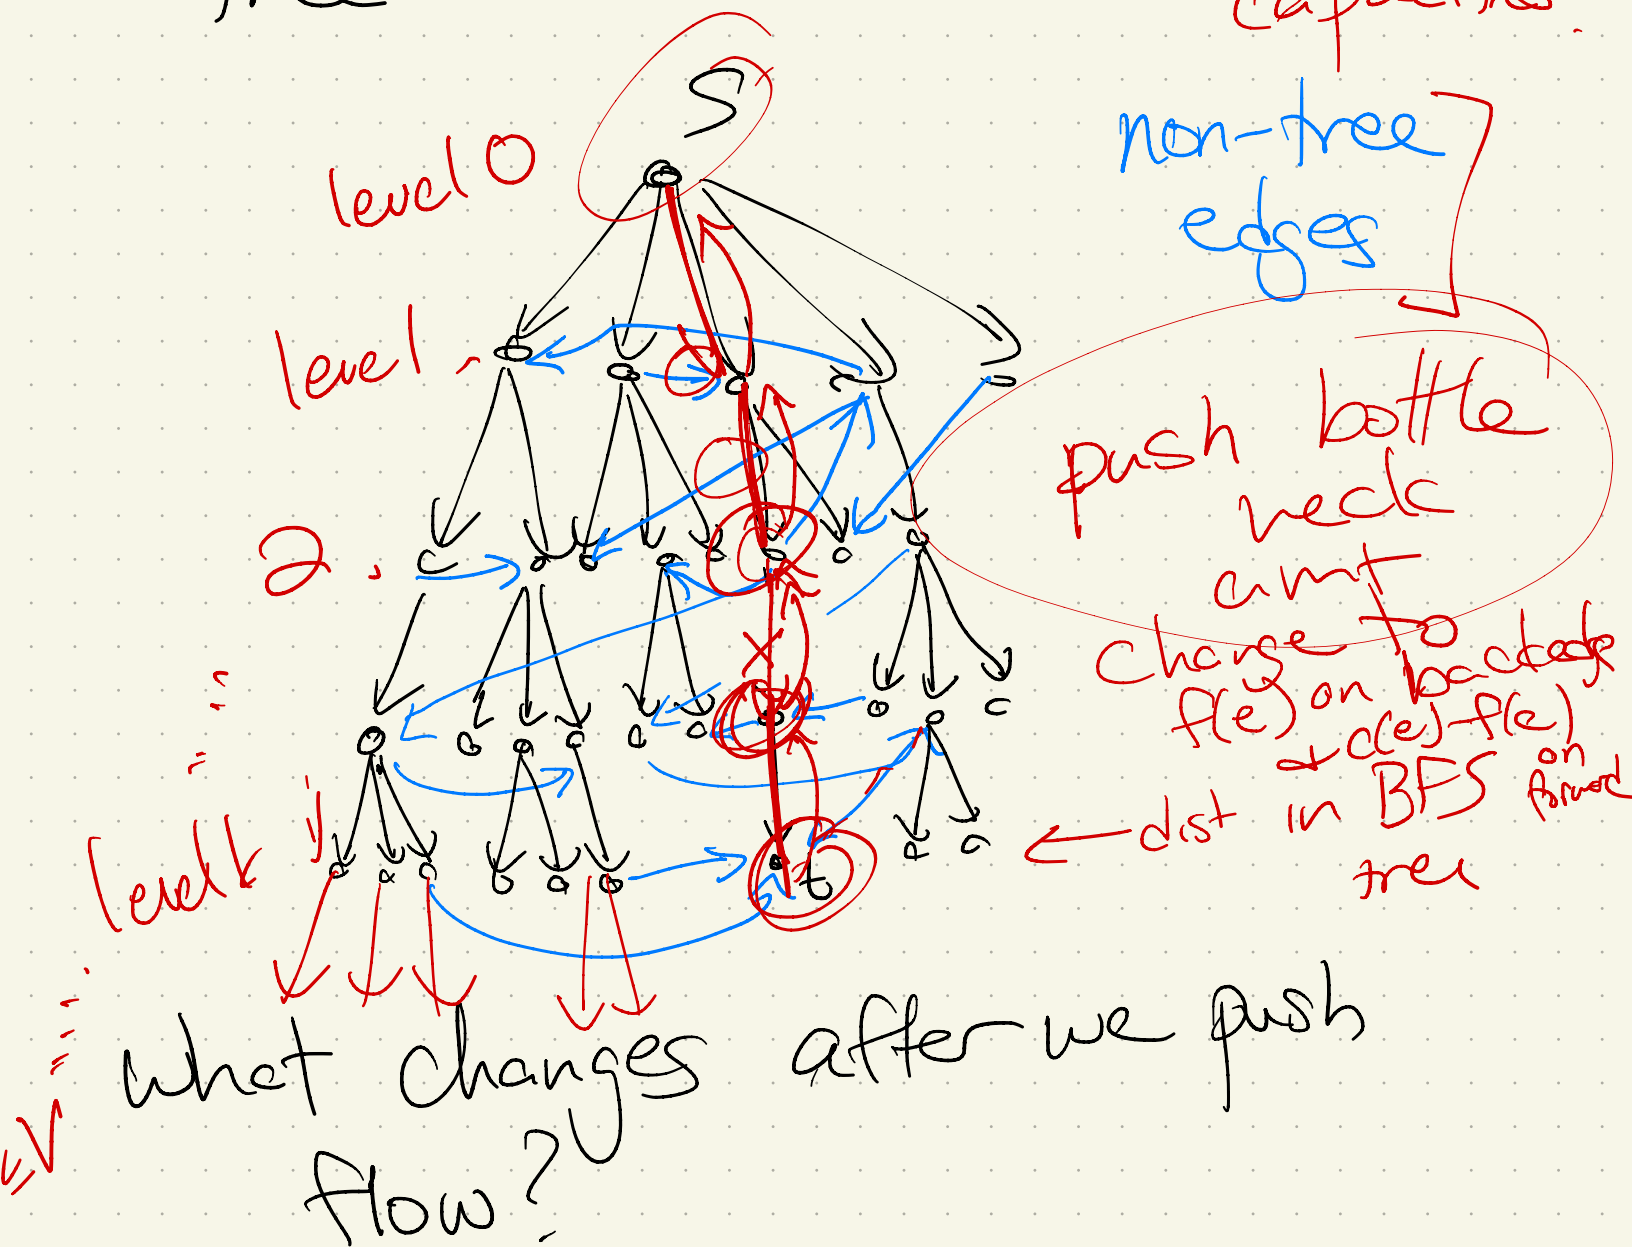
Which traversal?

BFS: $O(E+V)$

Q: loop??

Q: How many times do we need a path?
(ie: how many repetitions of the while loop?)

Think of G_f , + the BFS tree rooted at s : forget capacities!



Let $G_0 = \text{initial } G_f$

& G_i = residual graph after i repetitions of the loop.

G_i has a BFS tree, so
let $\text{level}_i(v) = \text{depth in tree}$
 $0 \rightarrow \leq V-1$

(Note: once s can't reach t ,
then $\text{level}(t) = \infty$.)

Claim: levels only get
get bigger (\geq) in each round.

proof: induction on level. (fix i)

base case: s (level 0), always
 $k=0$ stays ≥ 0 ($=0$)


G_i



← level k

IH: consider levels $\leq k$
 in G_i . IH: for any such
 u on level $\leq k$,
 $\text{level}_{i-1}(u) \leq \text{level}_i(u)$
 $\text{in } G_{i-1}$ $\text{in } G_i$

IS: now take v on level
 $k+1$ of G_i :
 Must be a path $S \rightsquigarrow v$

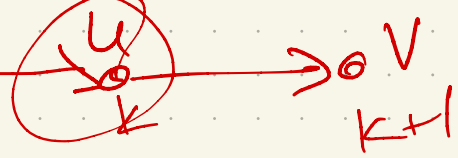


take u just
 before v on path
 \Rightarrow by IH,

$$\text{level}_{i-1}(u) \leq \text{level}_i(u)$$

Now: how did we
 get this path
 in G_i ?

HTS: holds for v

Cont: 

→ $u \rightarrow v$ is an edge in G_{i-1}
 $\text{level}_{i-1}(v) \leq \text{level}_{i-1}(u) + 1$

$\text{level}_i(u) \neq \text{level}_{i-1}(u)$
 $\rightarrow v$ satisfies property

→ might have come from pushing
 in G_{i-1} :
 here, it came from pushing
 a shortest path in the
 last round, so $v \rightarrow u$ was
 used last round.

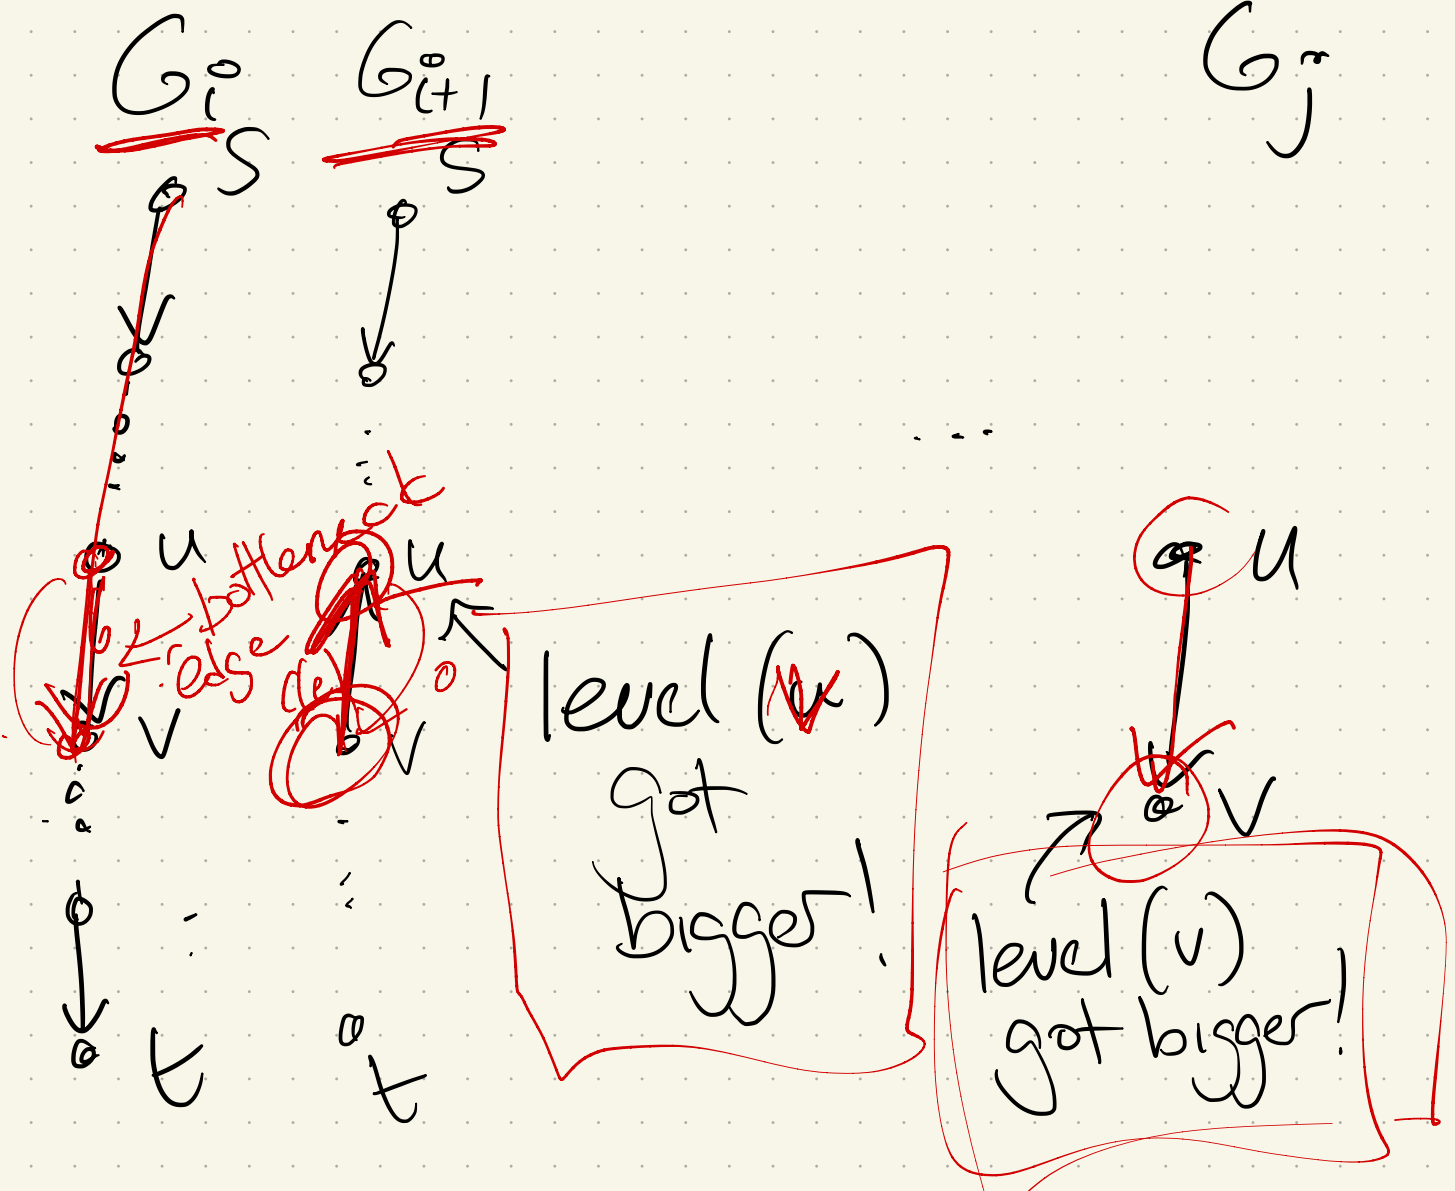
$i-1$

$$\text{level}_i(u) = \text{level}_{i-1}(v) + 1$$

$$\leq \text{level}_{i-1}(u) + 1$$



Then: Edges can disappear & reappear.
How many times?



If level increases by $+2$:
after $\frac{V}{2}$ changes, must
be ∞

In each iteration of the loop — some edge disappears!

$\Rightarrow \boxed{\frac{V}{2} \cdot E}$ repetitions

Total $\frac{V}{2} \cdot E (V+E) \in (E^2 V)$

MAX FLOW (G):

Let $f(e) = 0$ initially $\forall e$
Construct G_f

While there is s - t path in G_f

~~Let P be a simple s - t path~~

$f' \leftarrow \text{augment}(f, P)$

$f \leftarrow f'$

update G_f

return f

Let $P =$ shortest s - t path
(# edges, not capacities)

And... not done!

Technique	Direct	With dynamic trees	Source(s)
Blocking flow	$O(V^2E)$	$O(VE \log V)$	[Dinitz; Karzanov; Even and Itai; Sleator and Tarjan]
Network simplex	$O(V^2E)$	$O(VE \log V)$	[Dantzig; Goldfarb and Hao; Goldberg, Grigoriadis, and Tarjan]
Push-relabel (generic)	$O(V^2E)$	—	[Goldberg and Tarjan]
Push-relabel (FIFO)	$O(V^3)$	$O(VE \log(V^2/E))$	[Goldberg and Tarjan]
Push-relabel (highest label)	$O(V^2\sqrt{E})$	—	[Cheriyani and Maheshwari; Tunçel]
Push-relabel-add games	—	$O(VE \log_{E/(V \log V)} V)$	[Cheriyani and Hagerup; King, Rao, and Tarjan]
Pseudoflow	$O(V^2E)$	$O(VE \log V)$	[Hochbaum]
Pseudoflow (highest label)	$O(V^3)$	$O(VE \log(V^2/E))$	[Hochbaum and Orlin]
Incremental BFS	$O(V^2E)$	$O(VE \log(V^2/E))$	[Goldberg, Held, Kaplan, Tarjan, and Werneck]
Compact networks	—	$O(VE)$	[Orlin]

Figure 10.10. Several purely combinatorial maximum-flow algorithms and their running times.

use this

Many use very different techniques

- linear programming
- complex data structures
- not residual graphs

Still active!

8:46 AM Mon Oct 26

arxiv.org

98%

Showing 1–4 of 4 results

Search v0.5.6 released 2020-02-24

Feedback?

Query: order: -announced_date_first; size: 50; classification: Computer Science (cs); include_cross_list: True; terms: AND title=Max flow algorithm

Simple Search

Refine query

New search

50

results per page. Sort results by

Announcement date (newest first)

Go

1. [arXiv:2009.03260](#) [pdf, ps, other] [cs.DS](#) [math.OC](#)

A Potential Reduction Inspired Algorithm for Exact Max Flow in Almost $\tilde{O}(m^{4/3})$ Time

Authors: Tarun Kathuria

Abstract: We present an algorithm for computing s - t maximum flows in directed graphs in $\tilde{O}(m^{4/3+o(1)}U^{1/3})$ time. Our algorithm is inspired by potential reduction interior point methods for linear programming. Instead of using scaled gradient/Newton steps of a potential function, we take the step which maximizes the decrease in the potential value subject to advancing a certain amount α ... [More](#)
Submitted 7 September, 2020; originally announced September 2020.

2. [arXiv:1910.04848](#) [pdf, other] [cs.DS](#) [cs.CC](#)

A Fast Max Flow Algorithm

Authors: James B. Orlin, Xiao-Yue Gong

Abstract: In 2013, Orlin proved that the max flow problem could be solved in $O(nm)$ time. His algorithm ran in $O(nm + m^{1.94})$ time, which was the fastest for graphs with fewer than $n^{1.06}$ arcs. If the graph was not sufficiently sparse, the fastest running time was an algorithm due to King, Rao, and Tarjan. We describe a new variant of the excess scaling algorithm for the max flow problem whose runn... [More](#)
Submitted 10 October, 2019; originally announced October 2019.
Comments: 35 pages

3. [arXiv:1901.01412](#) [pdf, other] [cs.DS](#)

New Algorithms and Lower Bounds for All-Pairs Max-Flow in Undirected Graphs

Authors: Amir Abboud, Robert Krauthgamer, Ohad Trabelsi

Abstract: We investigate the time-complexity of the All-Pairs Max-Flow problem: Given a graph with n nodes and m edges, compute for all pairs of nodes the maximum-flow value between them. If Max-Flow (the version with a given source-sink pair s, t) can be solved in time $T(m)$, then an $O(n^2) \cdot T(m)$ is a trivial upper bound. But can we do better? For directed graphs, recent results in fine-grai... [More](#)
Submitted 9 July, 2019; v1 submitted 5 January, 2019; originally announced January 2019.

4. [arXiv:1304.2338](#) [pdf, other] [cs.DS](#)

An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Multicommodity Generalizations

Authors: Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, Aaron Sidford

Abstract: In this paper, we introduce a new framework for approximately solving flow problems in capacitated, undirected graphs and apply it to provide asymptotically faster algorithms for the maximum s - t flow and maximum concurrent multicommodity flow problems. For graphs with n vertices and m edges, it allows us to find an ϵ -approximate maximum s - t flow in time $O(m^{1+o(1)}\epsilon^{-2})$, improvi... [More](#)
Submitted 23 September, 2013; v1 submitted 8 April, 2013; originally announced April 2013.

About

Help

Copyright

Privacy Policy

Contact

Subscribe

Web Accessibility Assistance

arXiv Operational Status

Get status notifications via [email](#) or [slack](#)

Even more known about
"special" graphs:

Computer Science > Data Structures and Algorithms

Minimum Cuts in Surface Graphs

Erin W. Chambers, Jeff Erickson, Kyle Fox, Amir Nayyeri

(Submitted on 9 Oct 2019)

We describe algorithms to efficiently compute minimum (s, t) -cuts and global minimum cuts of undirected surface-embedded graphs. Given an edge-weighted undirected graph G with n vertices embedded on an orientable surface of genus g , our algorithms can solve either problem in $g^{O(g)}n \log \log n$ or $2^{O(g)}n \log n$ time, whichever is better. When g is a constant, our $g^{O(g)}n \log \log n$ time algorithms match the best running times known for computing minimum cuts in planar graphs.

Our algorithms for minimum cuts rely on reductions to the problem of finding a minimum-weight subgraph in a given \mathbb{Z}_2 -homology class, and we give efficient algorithms for this latter problem as well. If G is embedded on a surface with b boundary components, these algorithms run in $(g + b)^{O(g+b)}n \log \log n$ and $2^{O(g+b)}n \log n$ time. We also prove that finding a minimum-weight subgraph homologous to a single input cycle is NP-hard, showing it is likely impossible to improve upon the exponential dependencies on g for this latter problem.

A note from Ch 11

Using flows:

A few common themes:

① Matching:

Identify a way to
pair up items.

More pairs \Rightarrow larger flow

② Disjoint paths:

Find paths that avoid
each other.

③ "Tuple" Selection



Example:

Sham-Poobanana University has hired you to write an algorithm to schedule their final exams. There are n different classes, each of which needs to schedule a final exam in one of r rooms during one of t different time slots. At most one class's final exam can be scheduled in each room during each time slot; conversely, classes cannot be split into multiple rooms or multiple times. Moreover, each exam must be overseen by one of p proctors.⁴ Each proctor can oversee at most one exam at a time; each proctor is available for only certain time slots; and no proctor is allowed oversee more than 5 exams total. The input to the scheduling problem consists of three arrays:

- An integer array $E[1..n]$ where $E[i]$ is the number of students enrolled in the i th class.
- An integer array $S[1..r]$, where $S[j]$ is the number of seats in the j th room. The i th class's final exam *can* be held in the j th room if and only if $E[i] \leq S[j]$.
- A boolean array $A[1..t, 1..p]$ where $A[k, \ell] = \text{TRUE}$ if and only if the ℓ th proctor is available during the k th time slot.⁵

Idea: Build a graph!

any valid assignment \Rightarrow path in G

