MAX FLOW

David Kauchak
CS 140 – Spring 2020

1

Admin

Assignment 7 graded

Assignment 9

2

Checkpoint 3

Greedy algorithms -> all pairs shortest paths

Could ask to decide greedy vs. DP, but no DP solutions/algorithms

(will not include network flow)

2 pages of notes

3

Student networking

You decide to create your own computer network:

- You get three of your friends and string some network cables
- Because of capacity (due to cable type, distance, computer, etc) you can only send a certain amount of data to each person
- If edges denote capacity, what is the maximum throughput you can you send from S to T?

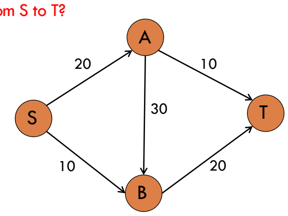

4

Student networking

You decide to create your own campus network:

- You get three of your friends and string some network cables
- Because of capacity (due to cable type, distance, computer, etc) you can only send a certain amount of data to each person
- If edges denote capacity, what is the maximum throughput you can you send from S to T?

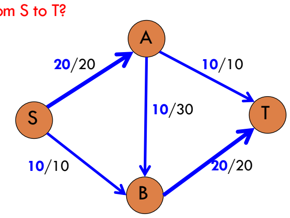

30 units

5

Another flow problem

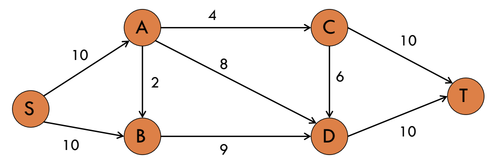

How much water flow can we continually send from s to t?

6

Another flow problem

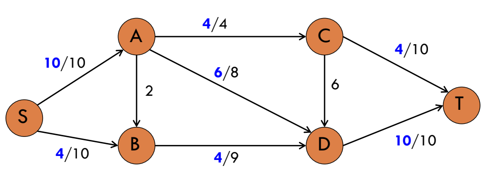

14 units

7

Flow graph/networks

Flow network

- directed, weighted graph (V, E)
- positive edge weights indicating the "capacity" (generally, assume integers)
- contains a single source $s \in V$ with no incoming edges
- contains a single sink/target $t \in V$ with no outgoing edges
- every vertex is on a path from s to t

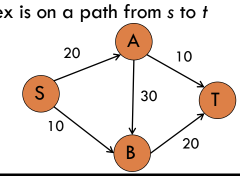

8

Flow

What are the constraints on flow in a network?

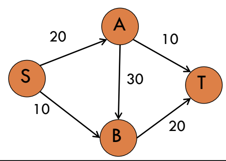

9

Flow

in-flow = out-flow for every vertex (except s, t)

flow along an edge cannot exceed the edge capacity

flows are positive

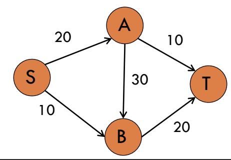

10

Max flow problem

Given a flow network: *what is the maximum flow we can send from s to t that meet the flow constraints?*

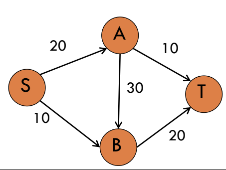

11

Applications?

network flow
- water, electricity, sewage, cellular…
- traffic/transportation capacity

bipartite matching

sports elimination

…

12

Max flow origins

Rail networks of the Soviet Union in the 1950's

The US wanted to know how quickly the Soviet Union could get supplies through its rail network to its satellite states in Eastern Europe.

In addition, the US wanted to know which rails it could destroy most easily to cut off the satellite states from the rest of the Soviet Union.

These two problems are closely related: solving the **max flow problem** also solves the **min cut problem** of figuring out the cheapest way to cut off the Soviet Union from its satellites.

Source: lbackstrom, The Importance of Algorithms, at www.topcoder.com

13

Algorithm idea

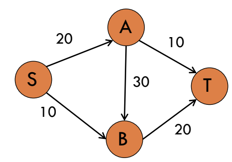

15

Algorithm idea

send some flow down a path

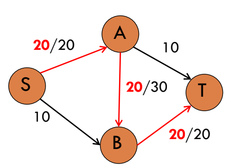

16

Algorithm idea

send some flow down a path

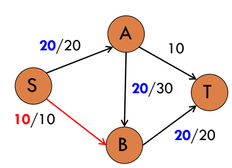

Now what?

17

Algorithm idea

reroute some of the flow

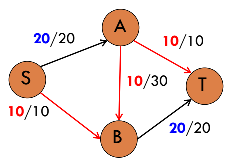

Total flow?

18

Algorithm idea

reroute some of the flow

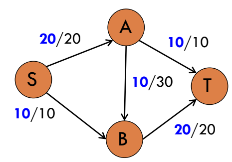

30

19

Algorithm idea

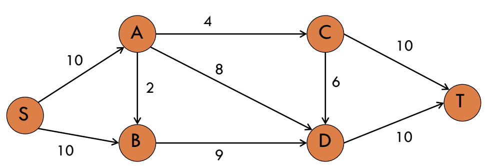

20

Algorithm idea

send some flow down a path

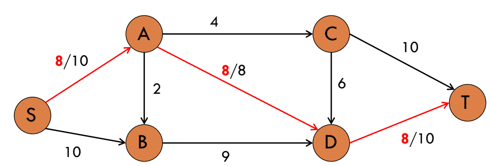

21

Algorithm idea

send some flow down a path

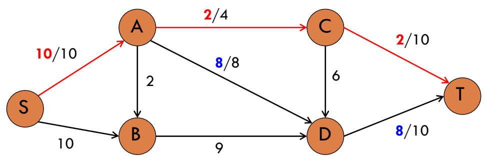

22

Algorithm idea

send some flow down a path

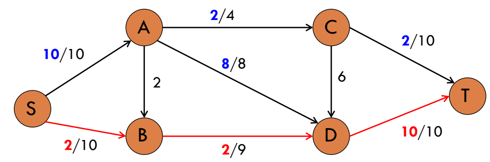

23

Algorithm idea

reroute some of the flow

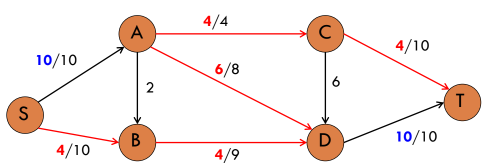

24

Algorithm idea

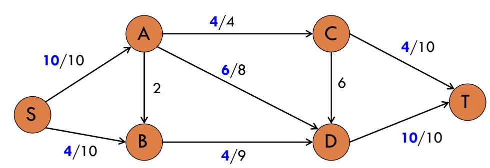

Are we done?
Is this the best we can do?

25

Cuts

A cut is a partitioning of the vertices into two sets S_s and $S_t = V\text{-}S_s$

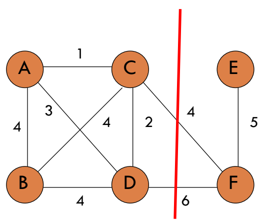

26

Flow across cuts

In flow graphs, we're interested in cuts that separate s from t, that is $s \in S_s$ and $t \in S_t$

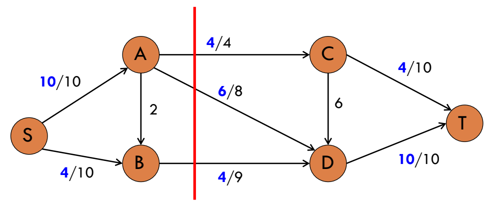

27

Flow across cuts

The flow "across" a cut is the total flow from nodes in S_s to nodes in S_t ***minus*** the total from nodes in S_t to S_s

What is the flow across this cut?

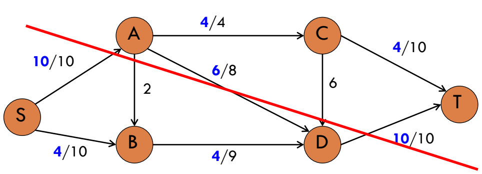

28

Flow across cuts

The flow "across" a cut is the total flow from nodes in S_s to nodes in S_t ***minus*** the total from nodes in S_t to S_s

10+10-6 = 14

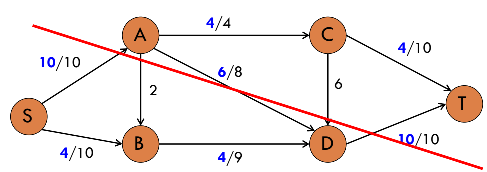

29

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

What do we know about the flow across the any such cut?

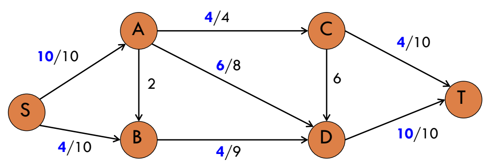

30

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

The flow across ANY such cut is the same and is the current flow in the network

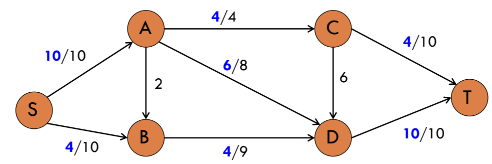

31

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

4+10 = 14

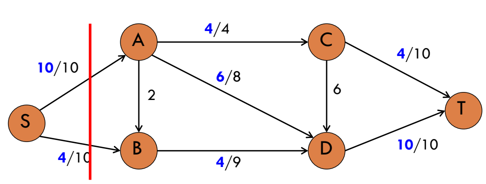

32

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

4+6+4 = 14

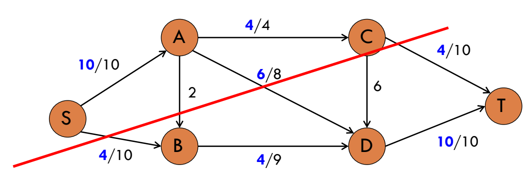

33

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

10+10-6 = 14

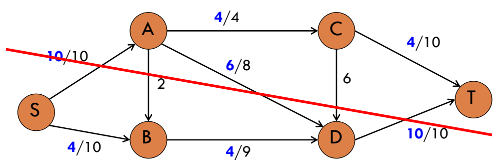

34

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

The flow across ANY such cut is the same and is the current flow in the network

Why? Can you prove it?

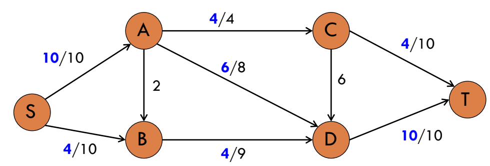

35

Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Inductively?

- every vertex is on a path from *s* to *t*
- in-flow = out-flow for every vertex (except s, t)
- flow along an edge cannot exceed the edge capacity
- flows are positive

36

Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Base case: S_s = s

- Flow is total from from s to t: therefore the total flow out of s should be the flow
- All flow from s gets to t
  - every vertex is on a path from s to t
  - in-flow = out-flow

37

Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Inductive case: Consider moving a node x from S_t to S_s

Is the flow across the different partitions the same?

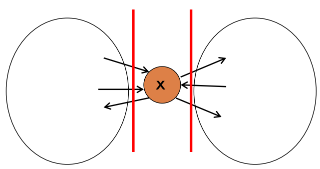

38

Flow across cuts

Inductive case: Consider moving a node x from S_t to S_s

cut = left-inflow(x) – left-outflow(x)

cut = right-outflow(x) – right-inflow(x)

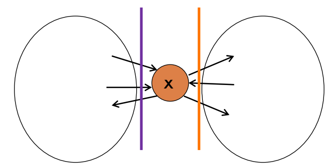

left-inflow(x) + right-inflow(x) = left-outflow(x) + right-outflow(x)

in-flow = out-flow

left-inflow(x) - left-outflow(x) = right-outflow(x) – right-inflow(x)

39

Flow across cuts

Consider any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

The flow across ANY such cut is the same and is the current flow in the network

40

Capacity of a cut

The "**capacity of a cut**" is the maximum flow that we *could* send from nodes in S_s to nodes in S_t (i.e. across the cut)

How do we calculate the capacity?

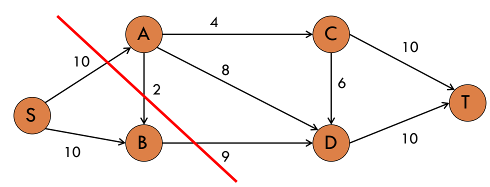

41

Capacity of a cut

The "**capacity of a cut**" is the maximum flow that we *could* send from nodes in S_s to nodes in S_t (i.e. across the cut)

Capacity is the sum of the edges from S_s to S_t

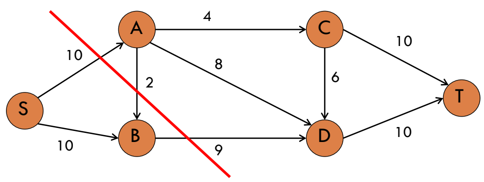

10 + 9 = 19

42

Capacity of a cut

The "**capacity of a cut**" is the maximum flow that we *could* send from nodes in S_s to nodes in S_t (i.e. across the cut)

Capacity is the sum of the edges from S_s to S_t

Why?

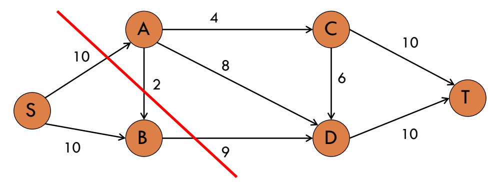

43

Capacity of a cut

The "**capacity of a cut**" is the maximum flow that we *could* send from nodes in S_s to nodes in S_t (i.e. across the cut)

Capacity is the sum of the edges from S_s to S_t

- Any more and we would violate the edge capacity constraint
- Any less and it would not be maximal, since we could simply increase the flow

44

Max Power

https://www.youtube.com/watch?v=BSVms6cT9nk

45

Quick recap

A cut is a partitioning of the vertices into two sets S_s and $S_t = V - S_s$

For any cut where $s \in S_s$ and $t \in S_t$, i.e. the cut partitions the source from the sink

- the flow across any such cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from S_s to S_t

46

Maximum flow

For any cut where $s \in S_s$ and $t \in S_t$

- the flow across the cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from S_s to S_t

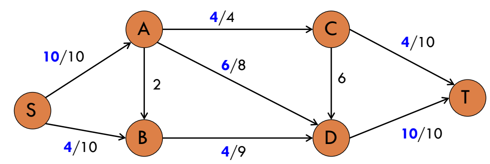

Are we done?
Is this the best we can do?

47

Maximum flow

For any cut where $s \in S_s$ and $t \in S_t$

- the flow across the cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from S_s to S_t

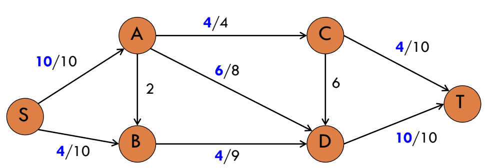

We can do no better than the minimum capacity cut!

48

Maximum flow

What is the minimum capacity cut for this graph?

Capacity = 10 + 4

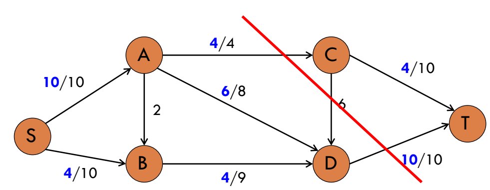

Is this the best we can do?

49

Maximum flow

What is the minimum capacity cut for this graph?

Capacity = 10 + 4

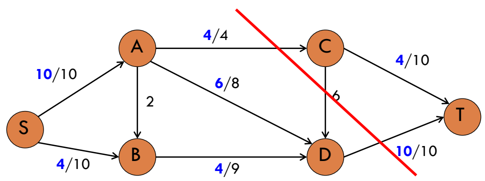

flow = minimum capacity, so we can do no better

50

Algorithm idea

send some flow down a path

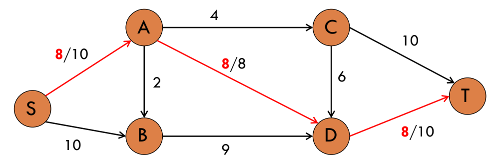

How do we determine the path to send flow down?

51

Algorithm idea

send some flow down a path

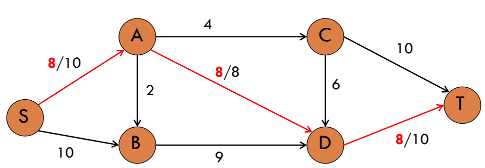

Search for a path with remaining capacity from s to t

52

Algorithm idea

reroute some of the flow

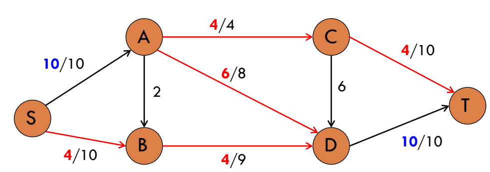

How do we handle "rerouting" flow?

53

Algorithm idea

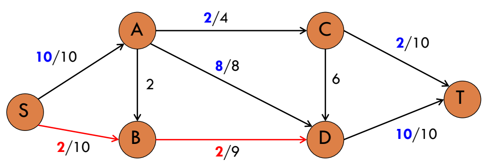

During the search, if an edge has some flow, we consider "reversing" some of that flow

54

Algorithm idea

reroute some of the flow

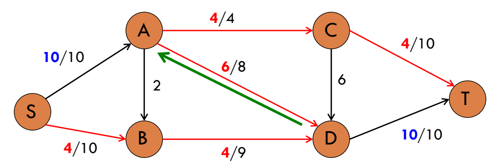

During the search, if an edge has some flow, we consider "reversing" some of that flow

55

The residual graph

The *residual graph* G_f is constructed from G

For each edge e in the original graph (G):

- if *flow*(e) < capacity(e)
  - introduce an edge in G_f with capacity = capacity(e)-flow(e)
  - this represents the remaining flow we can still push
- if flow(e) > 0
  - introduce an edge in G_f in the *opposite direction* with capacity = flow(e)
  - this represents the flow that we can reroute/reverse

56

Algorithm idea

G

G_f

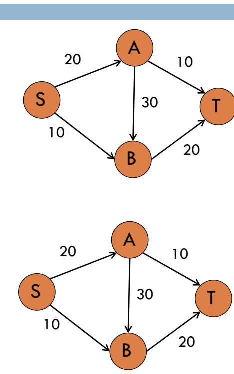

Find a path from s to t in G_f

57

Algorithm idea

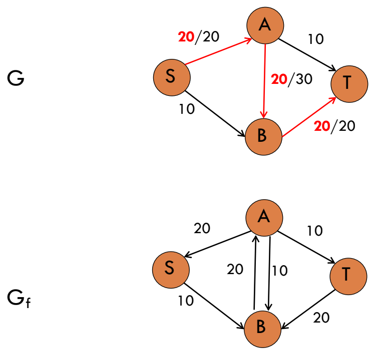

Find a path from s to t in G_f

58

Algorithm idea

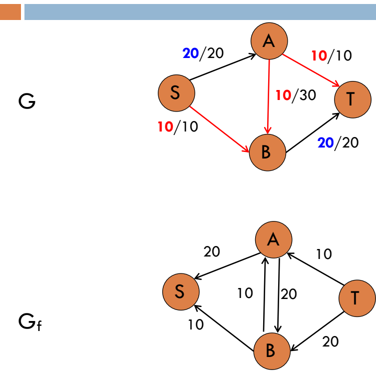

Find a path from s to t in G_f

59

Algorithm idea

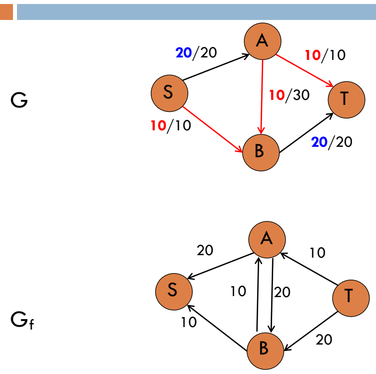

None exist… done!

60

Algorithm idea

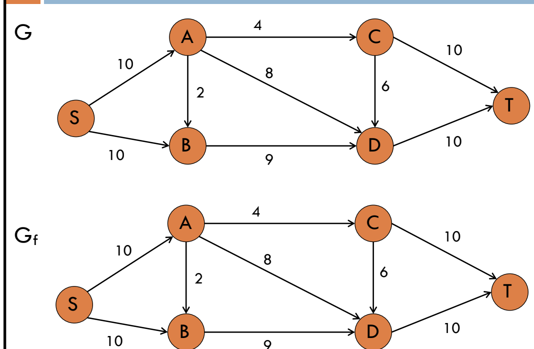

61

Algorithm idea

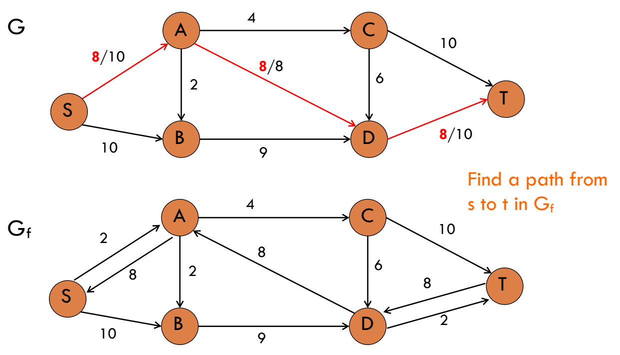

62

Algorithm idea

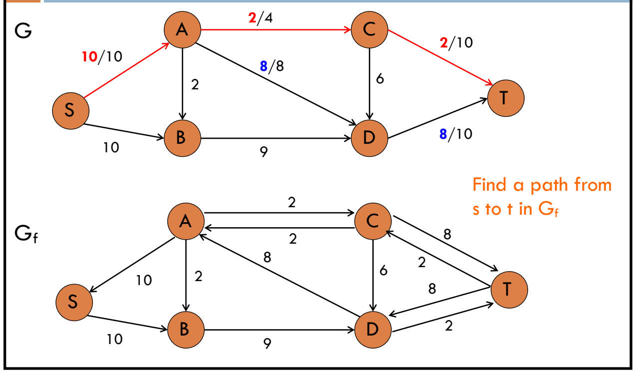

63

Algorithm idea

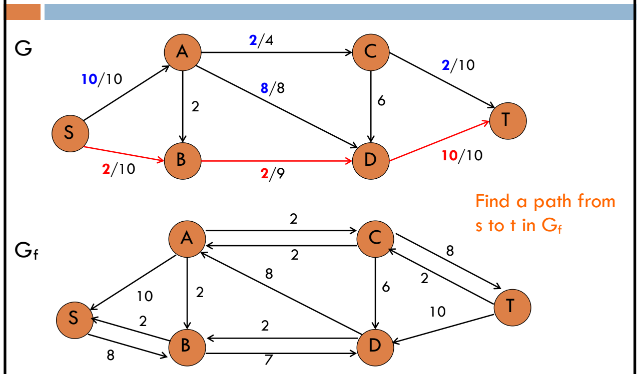

64

Algorithm idea

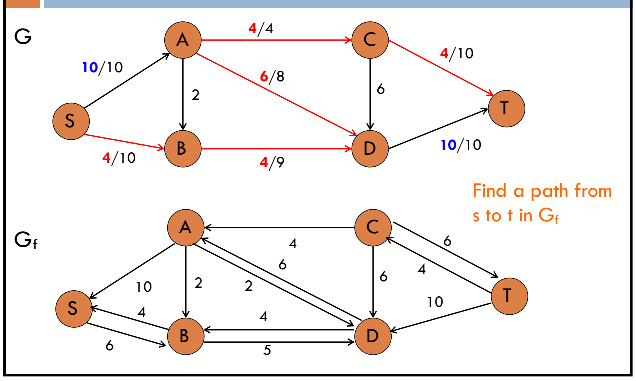

65

Algorithm idea

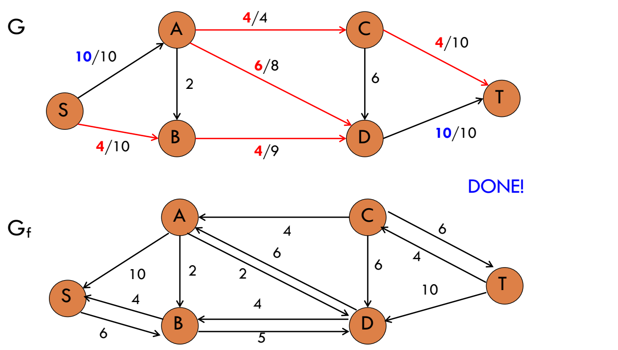

66

Ford-Fulkerson

Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  G_f = residualGraph(G)
  while a simple path exists from s to t in G_f
    send as much flow along the path as possible
    G_f = residualGraph(G)
  return flow

a simple path contains no repeated vertices

67

Ford-Fulkerson: is it correct?

Does the function terminate?

Every iteration increases the flow from s to t
- Every path must start with s
- The path has positive flow (or it wouldn't exist)
- The path is a simple path (so it cannot revisit s)
- conservation of flow

Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  G_f = residualGraph(G)
  while a simple path exists from s to t in G_f
    send as much flow along path as possible
    G_f = residualGraph(G)
  return flow

68

Ford-Fulkerson: is it correct?

Does the function terminate?
- Every iteration increases the flow from s to t
- the flow is bounded by the min-cut

Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  G_f = residualGraph(G)
  while a simple path exists from s to t in G_f
    send as much flow along path as possible
    G_f = residualGraph(G)
  return flow

69

Ford-Fulkerson: is it correct?

When it terminates is it the maximum flow?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

70

Ford-Fulkerson: is it correct?

When it terminates is it the maximum flow?

Assume it didn't

- We know then that the flow < min-cut
- therefore, the flow < capacity across EVERY cut
- therefore, across each cut there must be a forward edge in G_f
- thus, there must exist a path from s to t in G_f
  - start at s (and A = s)
  - repeat until t is found
    - pick one node across the cut with a forward edge
    - add this to the path
    - add the node to A (for argument sake)
- However, the algorithm would not have terminated… a contradiction

71

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

72

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

- traverse the graph
- at most add 2 edges for original edge
- $\boxed{\theta(V + E)}$

Can we simplify this expression?

73

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

- traverse the graph
- at most add 2 edges for original edge
- $\theta(V + E) = \theta(E)$
- (all nodes exists on paths from s to t)

74

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

- BFS or DFS
- $O(V + E) = O(E)$

75

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

- max-flow!
- increases ever iteration
- integer capacities, so integer increases

Can we bound the number of times the loop will execute?

76

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  Gf = residualGraph(G)
  while a simple path exists from s to t in Gf
    send as much flow along path as possible
    Gf = residualGraph(G)
  return flow
```

- max-flow!
- increases ever iteration
- integer capacities, so integer increases

Overall runtime? O(max-flow * E)

77

O(max-flow * E)

Can you construct a graph that *could get* this running time?

Hint:

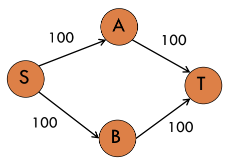

78

O(max-flow * E)

Can you construct a graph that *could get* this running time?

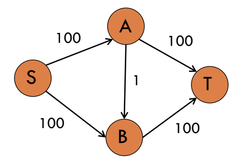

79

O(max-flow * E)

Can you construct a graph that *could get* this running time?

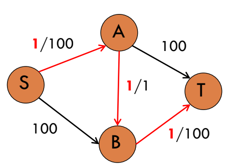

80

O(max-flow * E)

Can you construct a graph that *could get* this running time?

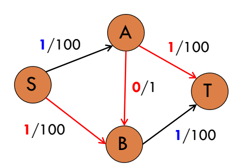

81

O(max-flow * E)

Can you construct a graph that *could get* this running time?

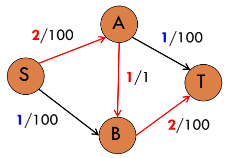

82

O(max-flow * E)

Can you construct a graph that *could get* this running time?

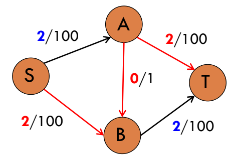

83

O(max-flow * E)

Can you construct a graph that *could get* this running time?

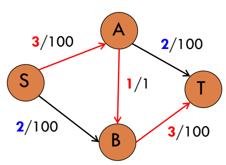

84

O(max-flow * E)

Can you construct a graph that *could get* this running time?

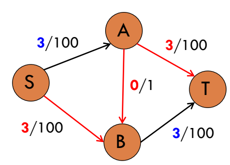

What is the problem here?
Could we do better?

85

Faster variants

Edmunds-Karp

- Select the *shortest path* (in number of edges) from s to t in G_f
  - How can we do this?
    - use BFS for search
- Running time: $O(V E^2)$
  - avoids issues like the one we just saw
  - see the book for the proof
  - or http://www.cs.cornell.edu/courses/CS4820/2011sp/handouts/edmondskarp.pdf

preflow-push (aka push-relabel) algorithms

- $O(V^3)$

86

Other variations…

Method	Complexity
Linear programming	
Ford–Fulkerson algorithm	$O(E \max\| f \|)$
Edmonds–Karp algorithm	$O(VE^2)$
Dinitz blocking flow algorithm	$O(V^2E)$
General push-relabel maximum flow algorithm	$O(V^2E)$
Push-relabel algorithm with *FIFO* vertex selection rule	$O(V^3)$
Dinitz blocking flow algorithm with dynamic trees	$O(VE \log(V))$
Push-relabel algorithm with dynamic trees	$O(VE \log(V^2/E))$
Binary blocking flow algorithm [1]	$O(E\min(V^{2/3},\sqrt{E})\log(V^2/E)\log U)$
MPM (Malhotra, Pramodh-Kumar and Maheshwari) algorithm	$O(V^3)$

http://en.wikipedia.org/wiki/Maximum_flow

TABLE I. POLYNOMIAL-TIME ALGORITHMS FOR THE MAXIMUM FLOW PROBLEM[a]

Algorithm no.	Date	Discoverer	Running time	References
1	1969	Edmonds and Karp	$O(nm^2)$	[5]
2	1970	Dinic	$O(n^2m)$	[4]
3	1974	Karzanov	$O(n^3)$	[18]
4	1977	Cherkasky	$O(n^2m^{1/2})$	[3]
5	1978	Malhotra, Pramodh Kumar, and Maheshwari	$O(n^3)$	[21]
6	1978	Galil	$O(n^{5/3}m^{2/3})$	[11]
7	1978	Galil and Naamad; Shiloach	$O(nm(\log n)^2)$	[12, 25]
8	1980	Sleator and Tarjan	$O(nm \log n)$	[27, 28]
9	1982	Shiloach and Vishkin	$O(n^3)$	[26]
10	1983	Gabow	$O(nm \log U)$	[10]
11	1984	Tarjan	$O(n^3)$	[31]
12	1985	Goldberg	$O(n^3)$	[14]
13	1986	Goldberg and Tarjan	$O(nm \log(n^2/m))$	[16, 15]
14	1986	Ahuja and Orlin	$O(nm + n^2 \log U)$	[1]

[a] Algorithm 13 is presented in this paper.

http://akira.ruc.dk/~keld/teaching/algoritmedesign_f03/Artikler/08/Goldberg88.pdf

87