

FIRST PART: CABLE NETWORKS

1



THE ROUTING PROBLEM I.E. THE SHORTEST PATH PROBLEM AND THE LEAST COST PATH PROBLEM

Prof. Tiziana Calamoneri
Network Algorithms
A.A. 2021/22



THE PROBLEM

(already studied!)

3

THE ROUTING PROBLEM (1)

Given a network:

- When **packets** are sent **from a computer to another one** through the network, each computer has to route data on a path passing through intermediate computers.
- This is the very general routing problem.

4

THE ROUTING PROBLEM (2)

- Case 1. Not adaptive routing

A routing algorithm could try to send packets through a network so that the **length** of the used path is **minimized**. Such length can be measured in terms of number of hops between pairs of computers.

If the network is modeled as a graph

(nodes = computers and edges = links), the problem reduces to the **shortest path problem** between two nodes.

5

THE ROUTING PROBLEM (3)

- Case 2. Adaptive Routing

It **takes into account the traffic conditions**: in order to decide next step, the traffic is estimated, so the packet is sent toward the zones of the network not affected by traffic.

If the network is modeled by an edge-weighted graph (nodes = computers, edges = links and weights = dynamic values proportional to the traffic on the connection), the problem reduces to the **(dynamic) least cost path problem**.

6

THE ROUTING PROBLEM (4)

cases 1 and 2. Adaptive and non adaptive routings (cntd)

- Non adaptive routing:
 - Good results with consistent topology and traffic
 - Poor performance if traffic volume or topologies change over time
 - Information about the entire network has to be available
 - Each packet is routed through an outgoing edge in a fixed way
 - **Routing tables are used**

7

THE ROUTING PROBLEM (5)

cases 1 and 2. Adaptive and non adaptive routings (cntd)

- Adaptive routing:
 - Decisions are based on current network state
 - Packets follow dynamically computed routes
 - Routers are able to communicate
 - Rather often re-calculations are necessary
 - Each router **creates** its own routing table

8

THE ROUTING PROBLEM (6)

▪ Case 3. Routing with faults

When the network is modeled as a graph, the length (edge-weight) of an edge may also represent the **probability of its failing** (used, for instance, in networks of telephone lines, or broadcasting systems in computer networks or in transportation routes). In all these cases, one is looking for the route having the highest probability not to fail.

More precisely...

9

THE ROUTING PROBLEM (7)

case 3. Routing with faults (cntd)

- Let $p(e)$ be the probability that edge e does not fail. Under the -not always realistic- assumption that failings of edges occur independently of each other, $p(e_1) \cdot p(e_2) \cdot \dots \cdot p(e_k)$ gives the probability that the path $P=(e_1, e_2, \dots, e_k)$ can be used without any faults.
- We want to maximize this probability over all possible paths with starting point a and arrival point b ...

10

THE ROUTING PROBLEM (8)

case 3. Routing with faults (cntd)

- **Note.** Since function \log is monotonic increasing, the maximum of the product $p(e_1) \cdot p(e_2) \cdot \dots \cdot p(e_k)$ is reached iff the logarithm of the product is maximum, i.e. iff:

$\log p(e_1) + \log p(e_2) + \dots + \log p(e_k)$ is maximum.

- $\log p(e) \leq 0$ for each e because $p(e) \leq 1$.
- Define $w(e) = -\log p(e)$, then $w(e) \geq 0$ for all e ; furthermore, we have to find a path from a to b for which

$w(e_1) + w(e_2) + \dots + w(e_k)$ becomes minimum.

- Thus, our problem is reduced again to the **least cost path problem**.

11

THE ROUTING PROBLEM (9)

Why does routing matter?

- End-to-end performance
 - Quality of the path affects **user performance** (propagation delay, throughput, packet loss)
- Use of network resources
 - **Balance** of the traffic over the routers and the links
 - **Avoidance of congestion** by directing traffic to lightly-loaded links
- Transition
 - The periodical changes of the routing tables reduce the incidence of faults and increase load balancing
 - Limiting packet loss and delay during changes

12



SHORTEST PATHS (1)

- Let $G=(V,E)$ be a graph; let $w(e)$ be the length of each edge e .
- Many versions of the **shortest path problem**:
 - All to all
 - One to one
 - One to all
 - All to one
- Lengths can be:
 - All equal (unit length)
 - Non negative
 - Possibly negative but without negative cycles
 - Creating possible negative cycles

14

SHORTEST PATHS (2)

- Algorithm designed by Moore ['59] for the **one-to-all shortest path problem and unit lengths**:
... Breadth First Search (BFS) ...
- TH.** G is connected iff at the end of the BFS starting from node a , $\text{dist}(a,b) < \infty$ for each node b , where dist is the distance in terms of number of edges.
- Note.** This claim is false if G is a digraph (indeed the notion of connected graph does not exist on digraphs: strong and weak connection...)

15

LEAST COST PATHS (1)

Let $G=(V,E)$ be a graph or a digraph and let $w: E \rightarrow \mathbb{R}$ be an edge-weight function.

- (G,w) is called **network**.
- $w(e)$ is called **length** (though including meanings such as cost, capacity, weight, probability, ...)
- For each path $P=(e_1, e_2, \dots, e_k)$ (if G is a digraph, P is a dipath), the **length** of P is defined as $w(P)=w(e_1)+w(e_2)+ \dots +w(e_k)$.
- Note.** If $w(e)=1$ for each edge, the least cost path problem reduces to the shortest path problem.

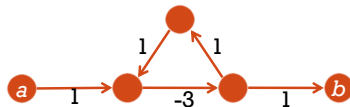
16

LEAST COST PATHS (2)

Given two nodes a and b , the distance $d(a,b)$ is defined as the minimum, over all the paths P connecting a and b , of $w(P)$.

Two problems arise:

- **PR.1:** b could be unreachable from a
- **SOL.:** define $d(a,b)=\infty$ if b is unreachable from a
- **PR.2:** the minimum could not exist (cycles of negative length)
- **SOL.:** only networks without cycles of negative length are feasible



17

LEAST COST PATHS (4)

- In general, when w represents a gain, it seems natural to replace $w(e)$ by $-w(e)$ and look for **least cost paths**, but this could introduce cycles of negative weight.
- There exist good algorithms that find minimum weight paths even when G contains cycles of negative weight.

19

LEAST COST PATHS (3)

Negative lengths may occur!

Example:

- A ship travels from port a to port b , where the route (and possible intermediary ports) may be chosen freely.
- The length $w(x,y)$ signifies the profit gained by going from x to y .
- For some edges, the ship might have to travel empty so that $w(e)$ is negative for these edges: the profit is actually a loss.
- Replacing $w(e)$ by $-w(e)$ for all e in this network, the shortest path represents the route which yields the largest possible profit.

18

LEAST COST PATHS (5)

OBSERVATIONS CONCERNING THE SOLUTION

In any solution:

- **Cycles having negative length** cannot exist (we avoided them by hypothesis)
- **Cycles having positive length** cannot exist (by contradiction: if one of them is in the solution, the new solution without it has a lower cost)
- **Cycles having null length** do not exist without loss of generality: if one of them is in the solution, the new solution without it has the same cost and so is feasible, too
- So: **our solution does not contain any cycles** and hence it passes through at most $n-1$ edges.

20

LEAST COST PATHS (6)

OBSERVATIONS CONCERNING THE SOLUTION (CNT.D)

- In order to univocally determine a path from a to b it is enough, for each node in such path, starting from b and coming back, to store its predecessor on the path.
- To do it: for each node v in G define a pointer $pt(v)$, initially equal to NULL; at the end, it points at the predecessor of v on the path.



BELLMAN-FORD ALGORITHM ['58]

- $G=(V,E)$ directed with edge-weights possibly negative
- It solves the problem of the **shortest path from single source**, hence it outputs the distances from the (single) source to each node
- It assumes that G does not contain any cycles of negative length
- It is based on the principle of relaxation
- Time complexity: $O(nm)$

23

SOME CLASSICAL ALGORITHMS

(already studied!)



22

DIJKSTRA ALGORITHM ['59]

- $G=(V,E)$ directed with non negative edge-weights
- It solves the problem of the **shortest path from single source**, hence it outputs the distances from the (single) source to each node
- It is based on the principle of relaxation
- Time complexity: either $O(n^2)$ or $O(m \log n)$
- The time complexity of the Dijkstra Algorithm is better than the time complexity of the Bellman-Ford Algorithm, but it is less versatile, as it requires not negative edge weight edges.

24

FLOYD-WARSHALL ALGORITHM [62]

- $G=(V,E)$ directed with edge-weights possibly negative
- It solves the problem of the **all pairs shortest path**, hence it outputs a matrix with the distances from each node to each other node
- Repeatedly applying the algs treated before, varying the source over all nodes in V :
 - Bellman-Ford: $n O(nm)=O(n^2m)$
 - Dijkstra: $n O(n^2) =O(n^3)$ or $n O(m \log n)=O(mn \log n)$
- Time complexity: $O(n^3)$ and negative edge weights are allowed

25

THE RELAXATION

- For each node v , let $d(v)$ be a function representing an **estimate of the weight of the shortest path** from s to v .
- At the beginning $d(v)=\infty$ for each v
- One relaxation step is performed as follows:
 - Given an edge (u,v)
 - If $d(u)+w(u,v)<d(v)$

$$d(v)=d(u)+w(u,v)$$

$$pt(v)=u$$
- Time complexity of one relaxation step: $O(1)$

26

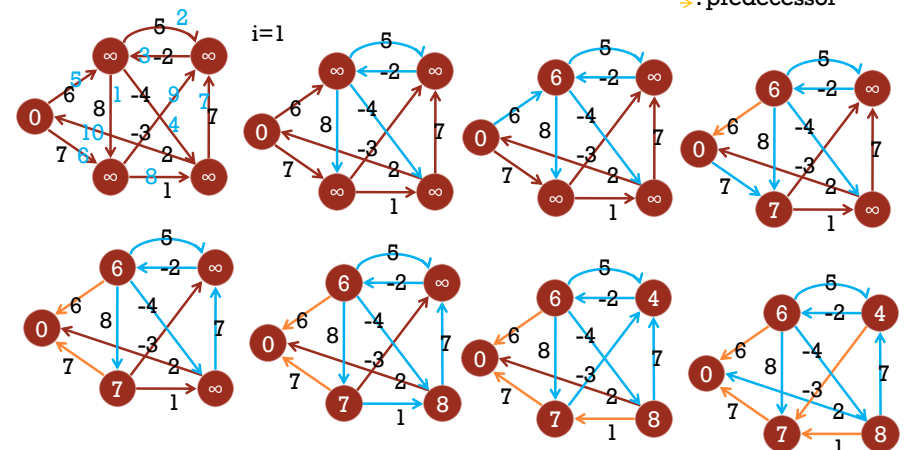
BELLMAN-FORD ALGORITHM (1)

- Assume that G does not contain any cycles of negative length
 - For each v initialize $d(v)$ and $p(v)$ $O(n)$
 - For $i=1$ to $n-1$ do $|n-1$ times
 - For each (u,v) relax v w.r.t. (u,v) $| mO(1)$
- Time Complexity: $O(nm)$

27

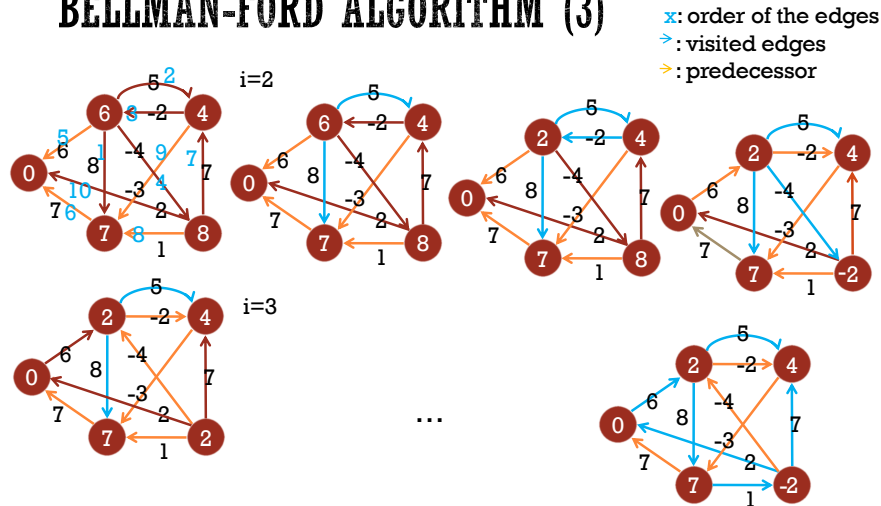
BELLMAN-FORD ALGORITHM (2)

x: order of the edges
 >: visited edges
 >: predecessor



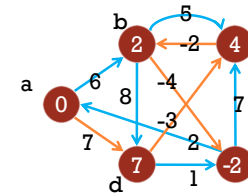
28

BELLMAN-FORD ALGORITHM (3)



29

BELLMAN-FORD ALGORITHM (4)



(Let us consider predecessors in the same direction than original edges)

Forwarding table of a:

- b (a,d) as the shortest path is a-d-c-b
- c (a,d) as the shortest path is a-d-c
- d (a,d)
- e (a,d) as the shortest path is a-d-c-b-e

30

BELLMAN-FORD ALGORITHM (5)

Bellman-Ford Algorithm is used for *Distance Vector routing*, an iterative, asynchronous and distributed protocol:

- $c(x,v)$ =cost for (directed) link from x to v
- $D_x(y)$ =estimate of least cost from x to y; x maintains **distance vector** $D_x=[D_x(y)]$ for each y neighbor of x; for each neighbor y, x maintains also D_y
- Each node x periodically sends D_x to its neighbors
- Neighbors update their own distance vector:
 $D_x(y)=\min\{D_x(y), c(x,v)+D_v(y)\}$
- x notifies neighbors when its distance vector changes
- Over the time, D_x converges

An exemple: Routing Information Protocol (RIP)

31

DIJKSTRA ALGORITHM (1)

- $G=(V,E)$ directed with non negative edge-weights
- It partitions the nodes of G: nodes whose shortest path from s has already been found (S) and all the other nodes (V-S)
- Greedy algorithm
- At each step, let u be the node in V-S with minimum value of d; add u to S and relax all the edges outcoming from u
- Keep V-S in a priority queue (e.g. min heap)

32

DIJKSTRA ALGORITHM (2)

- For each v initialize $d(v)$ and $pt(v)$
- S =empty set
- $Q=V$
- While Q is not empty
 - $u=ExtractMin(Q)$
 - $S=S \cup \{u\}$
 - For each edge (u,v) outcoming from u relax v w.r.t. (u,v)
 - Update Q

The time complexity depends on the data structure used to implement Q :

Queue: $O(n^2)$
 Heap: $O(m \log n)$
 Fibonacci Heap: $O(m+n \log n)$

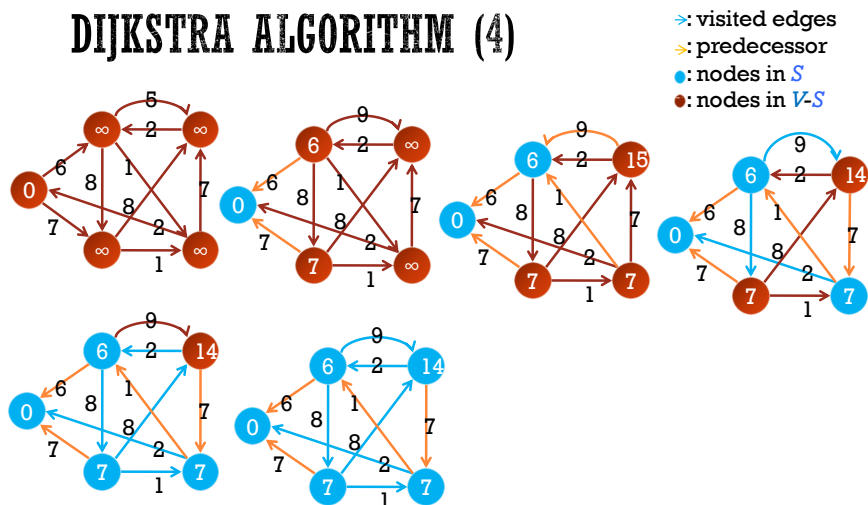
33

DIJKSTRA ALGORITHM (3)

- For each v initialize $d(v)$ and $pt(v)$ $\theta(n)$
 - S =empty set $\theta(1)$
 - $Q=V$ $\theta(n)$
 - While Q is not empty
 - $u=ExtractMin(Q)$ n times $| O(\log n)$
 - $S=S \cup \{u\}$ $| \theta(1)$
 - For each edge (u,v) outcoming from u relax v w.r.t. (u,v) $| \theta(deg u)$ times $| \theta(1)$
 - Update Q $| \theta(\log n)$
- tot. $O(n \log n + m \log n) = O(m \log n)$

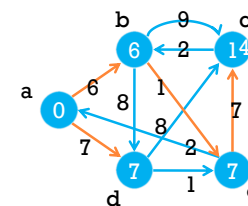
34

DIJKSTRA ALGORITHM (4)



35

DIJKSTRA ALGORITHM (5)



(Let us consider drawings of predecessors in the same direction than original edges)

Forwarding table of a:

- b (a,b)
- c (a,b) as the shortest path is a-b-e-c
- d (a,d)
- e (a,b) as the shortest path is a-b-e

36

DIJKSTRA ALGORITHM (6)

Dijkstra Algorithm is used for dynamic routing protocols.

- Each router:
 - Keeps trace of its incident links
 - Whether the link is up or down
 - The cost on the link (varying in time)
 - Broadcasts the link state (flooding)
 - So, every router has a complete view of the graph
 - Runs Dijkstra Algorithm
 - To compute the shortest paths...
 - ...and construct the forwarding table
- An example: Open Shortest Path First (OSPF) used in the networks with Internet Protocol (IP)

37

FLOYD-WARSHALL ALGORITHM (1)

Sometimes, it is not enough to calculate the distances w.r.t. a certain node s : we need to know the distances between all pairs of nodes.

- $G=(V,E)$ directed with any edge-weight.
- Algorithm for the **all-to-all shortest path problem**
- **Trick 1**: all edges are in; the non-existing ones have $w=\infty$
- **Trick 2**: in order to go from i to j you can either go directly or passing through a third node k
- dynamic programming

38

FLOYD-WARSHALL ALGORITHM (2)

Algorithm:

- For each node i , initialize $dist(i,i)=0$ $\theta(n)$
- For each edge (i,j) initialize $dist(i,j)=w(i,j)$ $\theta(n^2)$
- For each node k n times
 - For each node i $|n$ times
 - For each node j $|| n$ times
 - $dist(i,j)=\min\{dist(i,j), dist(i,k)+dist(k,j)\}$ $||| \theta(1)$
- Time Complexity $\theta(n^3)$

39

ANOTHER APPLICATION (1)

Difference Constraints:

- Let be given some **tasks** with precedence constraints and running lengths, and an unlimited (or limited by n =number of tasks) number of processors:
 - Each task i has:
 - **Starting time** s_i
 - **Time to complete** $b_i > 0$
 - **Constraint** $s_j + b_j \leq s_i$ if task i can be started after that task j has been completed
 - First task can start at time 0
 - When can we finish last task?

40

ANOTHER APPLICATION (2)

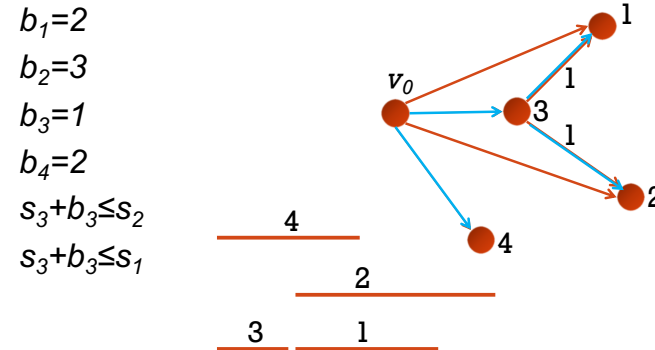
This is the Shortest Path Problem on directed acyclic graphs:

- Define a graph having a node for each task
- Insert a dummy node v_0 (that models time 0)
- Insert an arc (v_0, i) for each task node i and let 0 be its weight
- For each precedence constraint $s_j + b_j \leq s_i$ insert an arc (j, i) with weight b_i
- Optimal Solution:** start each task i at time equal to the length L_i of the longest path from v_0 to i . All the tasks are surely completed within time $\max_i (L_i + b_i)$

41


ANOTHER APPLICATION (3)

To transform to the shortest path problem, multiply all lengths and times by -1:



42

PACKET-ROUTING ALGORITHMS ON INTERCONNECTION TOPOLOGIES



43

PACKET-ROUTING ON INTERCONNECTION TOPOLOGIES (1)

- Up to now, in the routing problem we have considered the network as a **graph unknown to the nodes and variable in time** (faults, varying traffic, etc.)
- Nevertheless, when the network is an **interconnection topology** (and connects, for example, processors), it is **known and fixed in time**. Furthermore, **efficiency** is a primary issue.
- Solutions having stronger properties than the simple shortest path algorithms are required.

44

PACKET-ROUTING ON INTERCONNECTION TOPOLOGIES (2)

Many different types of routing models.

Here, we will focus on the **store-and-forward** model (also known as the **packet-switching** model):

- Each packet is maintained as an entity that is passed from node to node as it moves through the network
 - A single packet can cross each edge during each step of the routing
 - Depending on the algorithm, we may or may not allow packets to pile up in queues located at each node.
- When queues are allowed: effort to keep them short.

45

PACKET-ROUTING ON INTERCONNECTION TOPOLOGIES (3)

- Global controller to precompute routing paths not allowed
- Problem handled using only local control
- A routing problem is called **one-to-one** if at most one packet must be addressed to every node and each packet has a different destination.
- In contrast, **one-to-many** and **many-to-one**

46

BUTTERFLY NETWORK (1)

Def. Let $N=2^n$ (hence $n=\log N$); the **n -dimensional Butterfly** is a layered graph with:

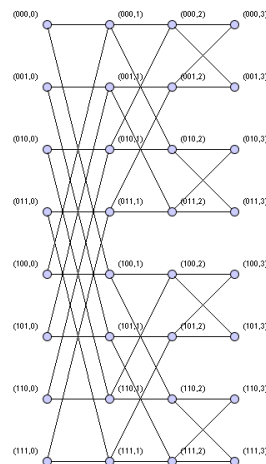
- $N(n+1)$ nodes ($n+1$ layers with 2^n nodes each) and
- $2Nn$ edges.

Nodes:

nodes correspond to pairs (w, i) , where:

- i is the *layer* of the node
- w is an n -bit *binary number* that denotes the *row* of the node.

...



BUTTERFLY NETWORK (2)

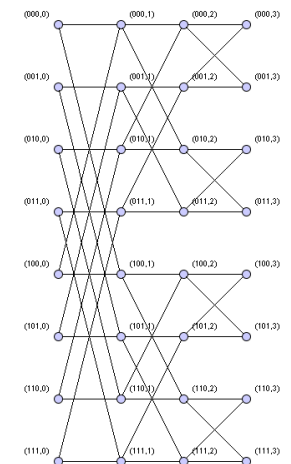
def. of n -dimensional butterfly (cntd)

...

Edges:

Two nodes (w, i) e (w', i') are linked by an edge iff $i'=i+1$ and either:

- $w=w'$ (*straight edge*) or
- w and w' differ in precisely the i -th bit (*cross edge*)



BUTTERFLY NETWORK (3)

- The nodes of the Butterfly are *crossbar switches*, i.e. switches with two input and two output values and can assume two states, *cross* and *bar*.
- Hence, the butterfly can be seen as a switching network connecting $2N$ ($N=2^n$) *input* units to $2N$ *output* units through a $\log N + 1$ layered network, having N nodes each.
- Input and output devices are usually processors and are often omitted in the graphical representations for the sake of simplicity.

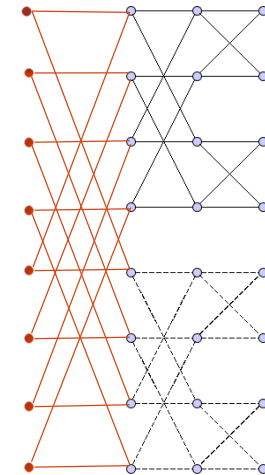


49

BUTTERFLY NETWORK (4)

The butterfly has a simple recursive structure:

one n -dim. butterfly contains two $(n-1)$ -dim. butterflies as subgraphs (just remove either the layer 0 nodes or the layer n nodes of the n -dim. butterfly to get two $(n-1)$ -dimensional butterflies).

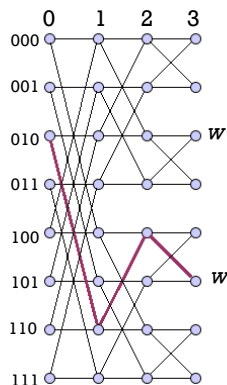


50

BUTTERFLY NETWORK (5)

For each pair of rows w and w' , there exists a unique path of length n (known as *greedy path*) from $(w, 0)$ to (w', n) ;

this path passes through each layer exactly once, using a cross-edge from layer i to layer $i+1$ ($i=0, \dots, n$) iff w and w' differ in the i -th bit and using a straight-edge otherwise.



51

ROUTING ON THE BUTTERFLY (1)

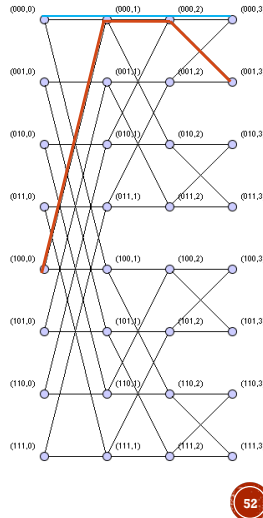
Problem of routing N packets from layer 0 to layer n in an n -dimensional butterfly:

- Each node $(u, 0)$ on layer 0 of the butterfly contains a packet that is destined for node $(\pi(u), n)$ on layer n , where $\pi: [1, N] \rightarrow [1, N]$ is a permutation.
- In the *greedy routing algorithm*, each packet is constrained to follow its *greedy path*.
- When there is only one packet to route, the greedy algorithm performs very well.
- Trouble can arise when many packets have to be routed in parallel...

52

ROUTING ON THE BUTTERFLY (2)

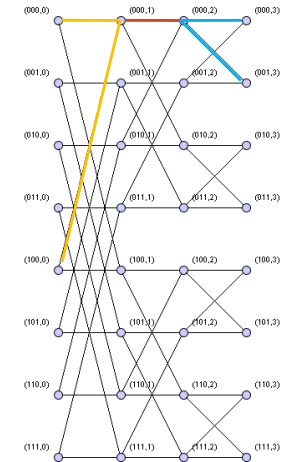
- Many greedy paths might pass through a single node or edge.
- Since only one of these packets can use the edge at a time, one of them must be delayed before crossing the edge.
- The butterfly is not able to route each permutation without delays, i.e. is a **blocking network**.
- The congestion problem arising in this example is not overly serious. When N is larger, however, the problem can be much serious. In fact...



52

ROUTING ON THE BUTTERFLY (3)

- Assume for simplicity that n is odd (but similar results hold when n is even), and consider edge $e = ((00\dots 0, (n-1)/2), (00\dots 0, (n+1)/2))$
- Node $(00\dots 0, (n-1)/2)$ is the root of a complete binary tree extending to the left having $2^{(n-1)/2}$ leaves
- Analogously to the right
- ...



ROUTING ON THE BUTTERFLY (4)

- ...
- The permutation can be such that each greedy path from a leaf of the left tree arrives to a leaf of the right tree traverses e
- There are $2^{(n-1)/2} = \sqrt{N}/2$ possible such paths, and thus $2^{(n-1)/2} = \sqrt{N}/2$ packets may traverse e . So at least one of them may be delayed by $\sqrt{N}/2 - 1$ steps.
- It takes at least $n = \log N$ steps to traverse the whole networks and to route a packet to its destination.
- In this case, the greedy algorithm can take $\sqrt{N}/2 + \log N - 1$ steps to route a permutation.
- In general...

55

ROUTING ON THE BUTTERFLY (5)

Th. Given any routing problem on an n -dimensional butterfly for which at most one packet starts at each layer 0 node and at most one packet is destined for each layer n node, the greedy algorithm will route all the packets to their destinations in $O(\sqrt{N})$ steps.

Proof. For simplicity, assume that n is odd (but the case n even is similar).

- Let e be any edge in layer i , $0 < i \leq n$, and define n_i to be the number of greedy paths that traverse e
- $n_i \leq 2^{i-1}$ (left tree) and, similarly, $n_i \leq 2^{n-i}$ (right tree) so $n_i \leq \min\{2^{i-1}, 2^{n-i}\}$
- ...

56

ROUTING ON THE BUTTERFLY (6)

- ...
- Any packet crossing e can be delayed by at most the other $n_i - 1$ packets that want to cross the edge.
- As this packet traverses layers 1, 2, ..., n , the total delay encountered can be at most:

$$\sum_{i=1}^n (n_i - 1) = \sum_{i=1}^{(n+1)/2} (n_i - 1) + \sum_{i=(n+3)/2}^n (n_i - 1) \leq \sum_{i=1}^{(n+1)/2} (2^{i-1} - 1) + \sum_{i=(n+3)/2}^n (2^{n-i} - 1) \leq$$

recalling that $\sum_{j=0}^k 2^j = 2^{k+1} - 1$

$$\leq 2^{(n+1)/2} + 2^{(n-1)/2} - n = O(\sqrt{N}) - n = O(\sqrt{N}) \quad \blacksquare$$

Annotations in the original image:

- $\sum_{i=1}^{(n+1)/2} (n_i - 1) = (n+1)/2 + 1$
- $\sum_{i=1}^{(n+1)/2} (2^{i-1} - 1) = \sum_{j=0}^{(n+1)/2-1} (2^j - 1)$
- $\sum_{i=(n+3)/2}^n (2^{n-i} - 1) = \sum_{j=0}^{(n-3)/2} (2^j - 1)$

57

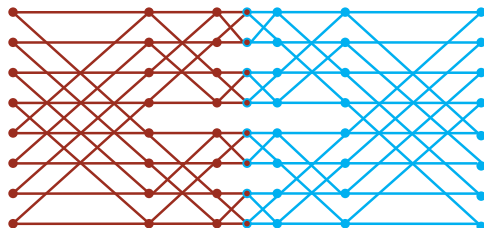
ROUTING ON THE BUTTERFLY (7)

- Despite the fact that the greedy routing algorithm performs poorly in the worst case, the greedy algorithm is very useful in practice.
- For many useful classes of permutations, the greedy algorithm runs in n steps, which is optimal and, for most permutations, the greedy algorithm runs in $n + o(n)$ steps.
- As a consequence, the greedy algorithm is widely used in practice.

58

BENEŠ NETWORK (1)

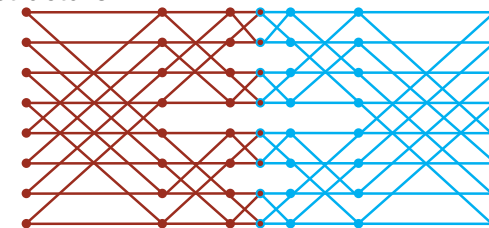
- A possibility to avoid a routing with delays is providing a **non blocking topology**.
- Beneš network has this property
- It consists of two back-to-back butterflies



59

BENEŠ NETWORK (2)

- The n -dimensional Beneš network has $2n+1$ layers, each with 2^n nodes.
- The first and last $n+1$ layers in the network form an n -dimensional Butterfly (the middle layer is shared).
- Not surprisingly, the Beneš network is very similar to the Butterfly, in terms of both its computational power and its network structure.



60

BENEŠ NETWORK (3)

- The reason for defining the Beneš network is that it is an excellent example of a **rearrangeable network**.
- **Def.** A network with N inputs and N outputs is said to be **rearrangeable** if for any one-to-one mapping π of the inputs to the outputs (i.e. for any permutation), we can construct edge-disjoint paths in the network linking the i -th input to the $\pi(i)$ -th output for $1 \leq i \leq N$.
- In the case of the n -dimensional Beneš network, we can have *two* inputs for each node at layer 0 and *two* outputs for each node at layer $2n$, and still connect every permutation of inputs to outputs with edge-disjoint paths.
- Hence, in this case, # of inputs = 2^{n+1} .

61

BENEŠ NETWORK (5)

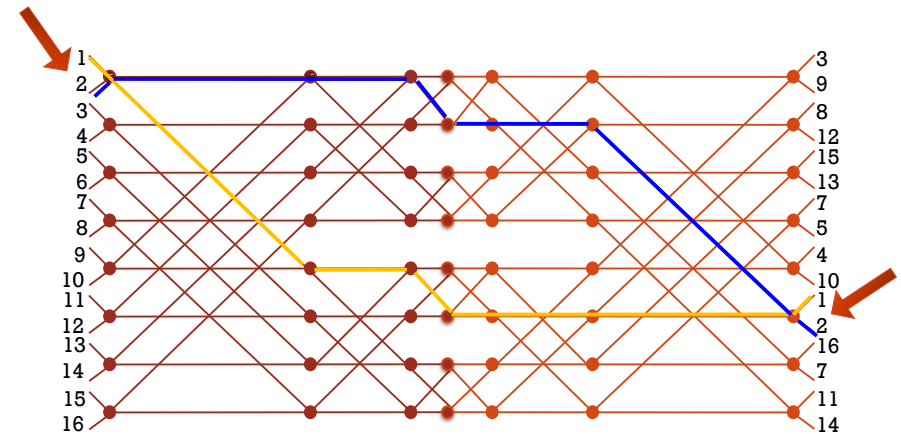
It seems extraordinary that we can find edge-disjoint paths for any permutation. Nevertheless, the result is true, and it is even fairly easy to prove, as we show in the following:

Th. Given any one-to-one mapping π of 2^{n+1} inputs to 2^{n+1} outputs on an n -dimensional Beneš network, there is a set of edge-disjoint paths from the inputs to the outputs connecting input i to output $\pi(i)$ for $1 \leq i \leq 2^{n+1}$.

Proof. ...

63

BENEŠ NETWORK (4)



62

BENEŠ NETWORK (6)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

Proof. By induction on n .

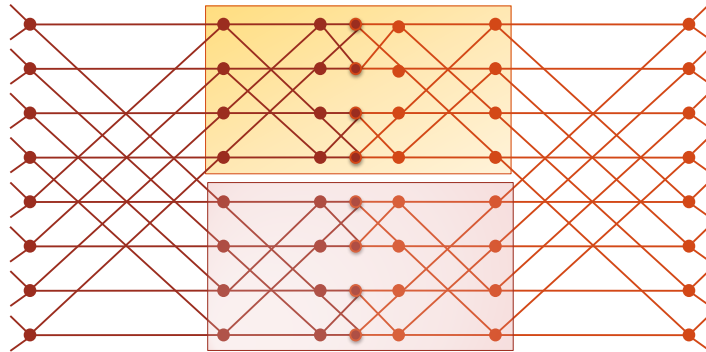
- **Basis:** if $n=0$, the Beneš network consists of a single node (i.e. a single 2×2 switch) and the result is obvious.
- **Induction:** assume that the result is true for an $(n-1)$ -dimensional Beneš network
- **Key observation:** the middle $2n-1$ layers of an n -dimensional Beneš network comprise two $(n-1)$ -dimensional Beneš networks

▪ ...

64

BENEŠ NETWORK (7)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

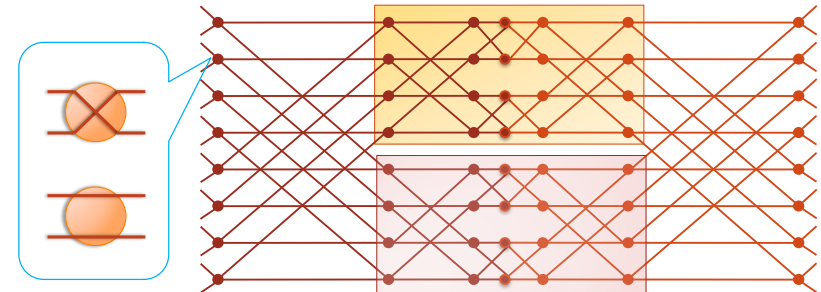


65

BENEŠ NETWORK (8)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

- Hence, for each path, it will be sufficient to decide whether it is to be routed through the upper sub-Beneš network or through the lower sub-Beneš network.



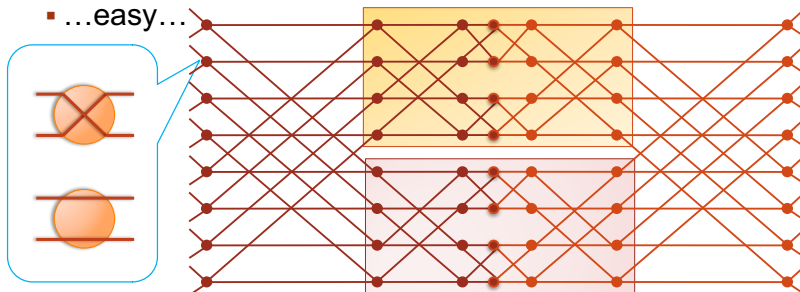
66

BENEŠ NETWORK (9)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

- The only constraints we have to consider to decide whether paths use the upper or lower subnetworks are that paths from inputs $2i-1$ and $2i$ must use different subnetworks for $1 \leq i \leq 2n$, and that paths to outputs $2i-1$ and $2i$ must use different sub-networks.

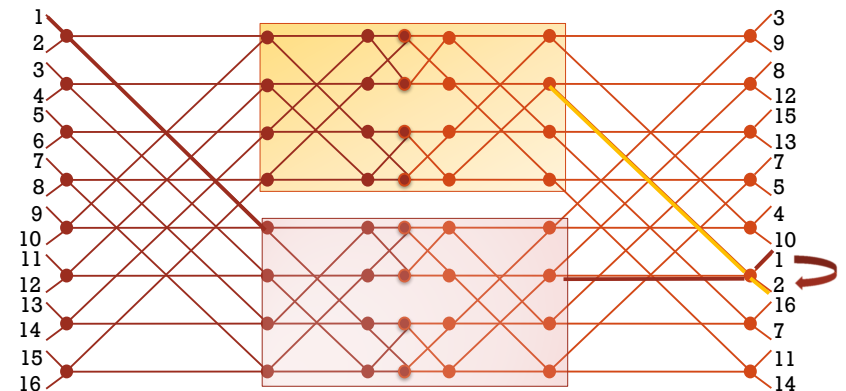
...easy...



67

BENEŠ NETWORK (10)

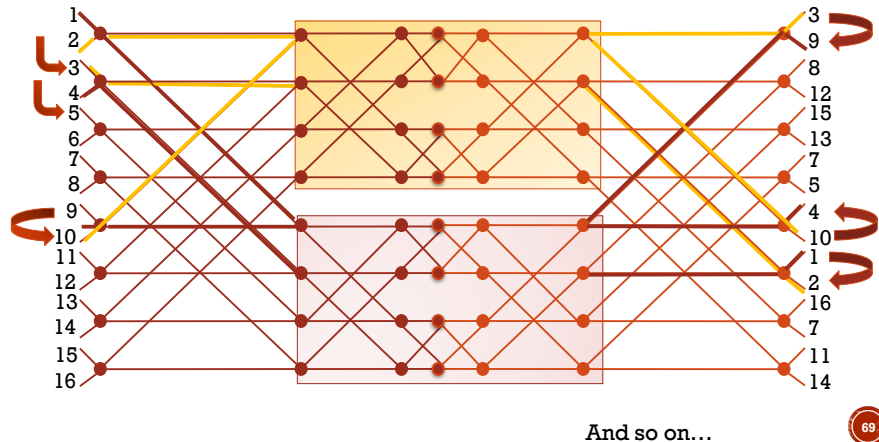
PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)



68

BENEŠ NETWORK (11)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)



BENEŠ NETWORK (13)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

- ...
- Eventually, we will close the loop by routing a path through the lower sub-network (in response to an output constraint) that shares an input switch with the first path that was routed.
- If any additional paths needs to be routed, we continue as before, starting over again with an arbitrary unrouted path.
- In this way, all paths can be assigned to the upper or lower sub-networks without conflict.

71

BENEŠ NETWORK (12)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

Summary of the steps:

- We start by routing the first path through the upper sub-network.
- We next satisfy the constraint generated at the output by routing the corresponding path through the lower sub-network.
- We keep on going back and forth through the network, satisfying constraints at the inputs by routing through the upper sub-network and satisfying constraints at the outputs by routing through the lower sub-network.
- ...

70

BENEŠ NETWORK (14)

PROOF OF THE REARRANGEABILITY OF THE BENEŠ NETWORK (CNTD)

- This algorithm is called **looping algorithm**.
- It is easy to see that all paths can be assigned to the upper or lower sub-networks without conflict:
- By construction, if we start going to the upper sub-network, we will arrive to the corresponding output in the upper sub-network and we will leave it to the lower sub-network, and so on.
- For parity reason, when a loop is close, we will correctly arrive from the right sub-network.
- The remainder of the path routing and switch setting is handled by induction in the sub-networks.

72

BENEŠ NETWORK (15)

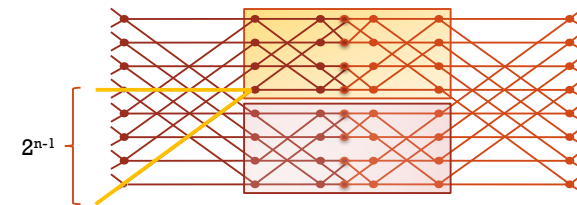
- In the case that each layer 0 node of the n -dimensional Beneš network has just one input and each layer $2n$ node has just one output, then the paths from the inputs to the outputs can be constructed so as to be node-disjoint (instead of only edge-disjoint):

▪ ...

73

BENEŠ NETWORK (16)

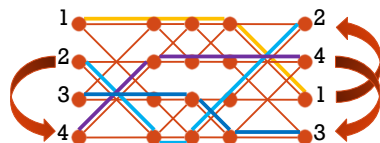
- **Th.** Given any one-to-one mapping of π of 2^n inputs to 2^n outputs in an n -dim. Beneš network, there is a set of node-disjoint paths from the inputs to the outputs connecting input i to output $\pi(i)$ for $1 \leq i \leq 2^n$.
- **Proof.** Identical to the previous one, but the paths needing to use different Beneš networks are now i and $i+2^{n-1}$, $1 \leq i \leq 2^{n-1}$ (and not $2i-1$ and $2i$). ■



74

BENEŠ NETWORK (17)

- Exemple: $n=2$, hence $2^{n-1}=2$



75

BENEŠ NETWORK (18)

Drawbacks of the looping algorithms (both versions):

- we do not know how to set the switches on-line. In other words, each switch needs to be told what to do by a **global control** that has knowledge of the permutation being routed
 - There exist numerous methods for overcoming this difficulty (not studied here).
- every time a new permutation must be routed, $\theta(N \log N)$ time is necessary to re-set switches.

76