




THE INTERCONNECTION TOPOLOGY LAYOUT PROBLEM i.e. THE ORTHOGONAL GRID GRAPH DRAWING PROBLEM

Prof. Tiziana Calamoneri

Network Algorithms

A.y. 2025/26



THE THOMPSON'S MODEL

2

THOMPSON'S MODEL (1)

- The **interconnection topology layout problem** arises from the problem of producing efficient VLSI (Very Large Scale Integration) layouts on a silicon board.
- It was born in the '40s, but it got a significant interest only relatively recently when the technology has allowed to layout circuits in two and three dimensions at reasonably low prices.



THOMPSON'S MODEL (2)

Exemples of VLSI circuits:

Intel 2004

Intel 2013

Intel 2023



THOMPSON'S MODEL (3)

Model the circuit as a graph (nodes = ports, switches, etc. and edges = wires).

There is a tight relation between the VLSI layout and the **graph drawing**.

Drawing Γ of a graph G : it is a function mapping each node v in a distinct point $\Gamma(v)$, and each edge (u,v) in an open Jordan curve $\Gamma(u,v)$ not crossing any point that is the mapping of a node, starting in $\Gamma(u)$ and arriving in $\Gamma(v)$.


5

THOMPSON'S MODEL (4)

The VLSI technology production imposes many constraints:

- the device pressing the connections can only approximate slanting lines by tiny horizontal and vertical segments ($\Rightarrow$ **orthogonal drawing**);


Orthogonal drawing: drawing of a graph where edges are represented as broken lines whose segments are horizontal or vertical (parallel to the coordinate axes)



6

THOMPSON'S MODEL (5)

- In order to avoid interference, it is necessary to keep wires far enough ($\Rightarrow$ grid drawing);



Grid drawing: drawing of a graph so that all nodes, crosses and bends of the edges are put on grid points (scaling property - resolution)

THOMPSON'S MODEL (6)

- Wires cannot cross; in order to avoid crossings, it is possible to route the crossing wires on the two separate sides of the board, introducing small "holes" trespassing the board from one side to the other; the number of such holes must be small, as their realization is rather expensive ($\Rightarrow$ crossing number minimization)
- ...

THOMPSON'S MODEL (7)

- The silicon is very expensive; so the layout must have a small area ($\Rightarrow$ area minimization).
- Wires should not be too long, as the propagation delay is proportional to their length; in the case of layered topology, wires in the same layer should have (approximately) the same length, so to avoid synchronization problems
($\Rightarrow$ edge length minimization).



THOMPSON'S MODEL (8)

In 1980 Thompson introduced a model that is consistent with all the mentioned constraints:


the layout of a topology G is a plane representation on a bunch of unit distance horizontal and vertical traces that maps:

- nodes of G in the intersection points of the traces,
- ...



THOMPSON'S MODEL (9)

- **edges** of G in disjoint paths constituted by horizontal and vertical segments on traces; such paths cannot cross nodes that are not their extremes and they can cross each other only in correspondence of trace intersection points;
- Overlappings (edge-edge) are not allowed
- Node-edge crosses are not allowed
- "knock-knees" are not allowed



11



12

ORTOGHONAL GRAPH DRAWING

ORTHOGONAL GRID DRAWING (1)

Def. An **orthogonal grid drawing** of a graph $G=(V,E)$ is a bijection mapping:

- nodes $v \in V$ on plane points $\Gamma(v)$ at integer coordinates
- edges $(v,w) \in E$ on not overlapping paths so that the images of their extremes $\Gamma(v)$ and $\Gamma(w)$ are connected by the corresponding paths.
- These paths are constituted by horizontal and vertical segments; the possible bends have integer coordinates

Obs. Only graphs with degree ≤ 4 can be correctly drawn.

13

ORTHOGONAL GRID DRAWING (2)

So, the interconnection topology layout is an **orthogonal grid drawing** of the corresponding graph with the aim of **minimizing the area, the number of crossings and the wire length**.

There is a huge literature in the GRAPH DRAWING area...

Shall we use the known algorithms for orthogonal grid drawing in order to solve the layout problem?

14

ORTHOGONAL GRID DRAWING (3)

No:


these algorithms guarantee some bounds on the optimization functions that hold FOR EACH input graph having the required input hypotheses. Instead, interconnection topologies are very structured graphs (usually regular, symmetric, recursively built, ...), and, by exploiting these properties, it is possible to get better results.

15

ORTHOGONAL GRID DRAWING (4)

- Graph drawing algorithms get a graph in input and draw it on the plane.
- Layout algorithms are designed for a single special interconnection topology and so they get only its dimension in input.
- **Obs.** Improving an optimization function by "only" a constant factor is an important issue (especially the area): if a layout occupies $\frac{1}{2}$ of the area of another one, it will cost the half!

16



COMPLETE BINARY TREE LAYOUT

17

H-TREES (1)



- Result independently achieved by Leiserson ['80] and Valiant ['81].
- H-tree: grid orthogonal rectilinear plane representation of an n -node complete binary tree in area $O(n)$.
- The result is optimum (the area must be $\Omega(n)$)
- Brent e Kung ['80] show that $\Omega(n \log n)$ area is necessary if the leaves must lay on the border.
- Note: each binary tree is contained in a complete binary tree, so we can exploit this approach to represent every binary tree, but the area could be not optimum anymore!

18

H-TREES (2)

Inductive construction:

- $h=0$ 
- inductive step:





H-TREES (3)


An example:

$h=0$ 

$h=1$ 

$h=2$ 


$h=4$ 

$h=3$ 

H-TREES (4)

Th. The area occupied by an n -node H-tree is $2(n+1)+o(n)$.


Proof. Base cases: $l_0=w_0=0$;
 $l_1=0$; $w_1=2$;
 $l_2=2$; $w_2=2$;




General cases:

h odd: $l_h=l_{h-1}$; $w_h=2 w_{h-1}+2$;

h even: $l_h=2 l_{h-1}+2$; $w_h=w_{h-1}$;




h even



h odd

Solving the recurrency equation....



COLLINEAR LAYOUT (1)

- The Thompson model ['79] requires that the wires coming out of each processing element are at most 4 (6 in 3D).

Problem: What if the degree is higher? (end of the '90s)

Sol.: Non-constant node degree model:

- a node of degree d occupies a square of side $\Theta(d)$ (here the degree is $n-1$)
- the wires can run either horizontally or vertically along grid lines.

23

COLLINEAR LAYOUT (2)


- Layout proposed by Yeh and Parami ['98]
- Collinear layout with area $n^4/4$ - optimal
- In a **collinear layout**, all nodes are placed on the same line. Instead of computing its area, it is usual to count the number of necessary **tracks**.

24

COLLINEAR LAYOUT (3)

To obtain the collinear layout of the complete graph:



- let a link be type- i if it connects two nodes whose labels differ by i ; so, the $n(n-1)/2$ links can be classified into types 1, 2, ..., $n-1$, and there are $n-i$ type- i links.
- ...



25

COLLINEAR LAYOUT (4)

- ...
- place the n nodes, labeled 1 through n , along a row;
- place the type-1 links in one track,
- place the type-2 links in two tracks, where links connecting odd nodes are put in one track and links connecting even nodes are put in the other one
- place the type- i links in $\min(i, n-i)$ tracks




26

COLLINEAR LAYOUT (5)

Total number of tracks in this layout:

substituting
 $j=n-i$

$$\begin{aligned} \sum_{i=1}^{n-1} \min(i, n-i) &= \sum_{i=1}^{n/2} i + \sum_{i=\frac{n}{2}+1}^{n-1} (n-i) = \\ &= \sum_{i=1}^{n/2} i + \sum_{j=\frac{n}{2}-1}^{\frac{n}{2}-1} j = \frac{1}{2} \left[\frac{n}{2} \left(\frac{n}{2} + 1 \right) + \frac{n}{2} \left(\frac{n}{2} - 1 \right) \right] = \frac{n^2}{4} \end{aligned}$$



27

COLLINEAR LAYOUT (6)

Def. The **bisection width** of a network is the minimum number of edges one has to cut to disconnect the network into two **equally** sized sub-networks.

Property. The bisection width of the complete graph is $n^2/4 + o(n^2)$.

Th. A lower bound on the number of tracks in the collinear layout of a network is its bisection width (to be proved later).


Cor. A lower bound on the number of tracks in the collinear layout of the complete graph is $n^2/4 + o(n^2)$.

28

ORTHOGONAL LAYOUT (1)

Note. The area of the collinear layout is:

$$\frac{n^4}{4}$$



29


ORTHOGONAL LAYOUT (2)

- Although the collinear layout leads to the smallest possible number of tracks, layouts with smaller area can be obtained.
- An area-efficient layout for complete graphs is based on the previous collinear layout.
- W.l.o.g. $n = m_1 \times m_2$, where m_1 and m_2 are $\Theta(\sqrt{n})$
- Each node can be labeled (i, j) with $i = 1, \dots, m_1$ and $j = 1, \dots, m_2$.

30

ORTHOGONAL LAYOUT (3)

- Put node (i,j) at coordinates (i,j) on an $m_1 \times m_2$ grid.
- Without entering into details:



$$Area = \frac{n^4}{16} + o(n^4)$$

31

ORTHOGONAL LAYOUT (4)

Th. A lower bound on the layout area of a network is the square of its bisection width.

Reminder. The bisection width of the complete graph is $n^2/4 + o(n^2)$.

Cor. A lower bound on the layout area of the complete graph is $n^4/16 + o(n^4)$.


32

ORTHOGONAL LAYOUT (5)

Let us prove the theorem:

Th. [Thompson'79] A lower bound on the layout area of a network is the square of its bisection width.


Proof. Suppose that the bisection width of a network G can be counted when partitioning its nodes in two sets of k and $n-k$ nodes, respectively.



width at least as large as the bisection width...
the same holds for the height...



BUTTERFLY LAYOUT



BUTTERFLY NETWORK


(MEMORANDUM)

Def. (reminder) Let $N=2^n$ (and $n=\log N$); an **n-dim. Butterfly** is a layered graph having N ($n+1$) nodes ($n+1$ layers, with 2^n nodes each) and $2Nn$ edges.

The nodes are labeled with a pair (w, i) , where i is the layer of the node and w is an n -bit binary number indicating the row of the node.

Two nodes (w, i) and (w', i') are adjacent iff $i'=i+1$ and:


- $w=w'$ (straight edge) or
- w e w' differ in exactly the i -th bit (cross edge).



35

WISE LAYOUT (1)

Layout proposed by D.S.Wise ['81]




He writes:

"This paper offers a result that can be described as a picture. [...] The perceptive reader may stop here, since the remainder of this paper only describes it."

36

WISE LAYOUT (2)



Properties:


1. (very important in a layered topology)

All the wires in the same layer are of equal length.

Nevertheless, this length grows exponentially with the layer.

37

WISE LAYOUT (3)




2. The **longest** path length from any input to any output is linear in N (namely, $2(N-1)$).

Indeed:

- All the paths have the same length.
- For the sake of simplicity, consider the path from the upper-left node to the lower-right node: its length coincides with the diagonal of the square having side $\sqrt{2}(N-1)$, so it is $2(N-1)$.

38


WISE LAYOUT (4)



3. The layout is performed on the two sides of the silicon board, so it can be considered a 2-layer layout; one layer is composed of all diagonal wires running from lower-left to upper-right -red lines- and the other layer is composed of all diagonal wires from lower-right to upper-left -black lines.

39

WISE LAYOUT (5)




PROs:

- Good area : $\sqrt{2}(N-1) \times \sqrt{2}(N-1) = 2N^2 + o(N^2)$.
- Same wire length on each layer; this is not true in every layout: in the classical drawing of the butterfly network, for example, straight-edges on the last layer have unit length while the cross-edges on the same layer have a linear length in the input size N ; this is extremely bad, because synchronization of the information flow goes lost;
- The input and output nodes lie on the boundary of the layout, and this can be required by some applications.

40

WISE LAYOUT (6)



CONs:


- “slanted” lines, so that the area of the layout is measured by a rectangle whose sides are not parallel to coordinate axes but lie at 45° ; if we follow the standard definition of layout area, it becomes $2(N-1) \times 2(N-1) = 4N^2 + o(N^2)$; indeed, the circumscribed square with sides parallel to the coordinate axes has side equal to the length of the path from the upper-left node to the lower-right node, that is $2(N-1)$;
- ...

41

WISE LAYOUT (7)

CONs (cntd):

- it is a ‘cheating’ layout, indeed the “knock-knees” are not avoided but arranged in the layout thanks to some devices that have no null area and so enlarge the layout area.
- The Wise layout “looks like” the usual representation. Nevertheless, in order to get the Wise layout from the usual representation, nodes must be permuted:



42

EVEN AND EVEN LAYOUT (1)

- Layout presented by G. Even and S. Even [’00], and based on the notion of Layered Cross Product (LCP):
- Def.** A **layered graph** of $l+1$ layers $G=(V_0, V_1, \dots, V_l, E)$ consists of $l+1$ layers of nodes; V_i is the (non-empty) set of nodes in layer i ; E is a set of directed edges: edge (u,v) connects two nodes of two adjacent layers, that is, if u lies on layer i then v lies on layer $i+1$.


43

EVEN AND EVEN LAYOUT (2)

Def. [Even & Litman ’92] The **LCP** of two layered graphs of $l+1$ layers each:

$G^1=(V_0^1, V_1^1, \dots, V_l^1, E^1)$ and $G^2=(V_0^2, V_1^2, \dots, V_l^2, E^2)$, is a layered graph of $l+1$ layers, $G=(V_0, V_1, \dots, V_l, E)$, where:


- For every $i=0, \dots, l$, $V_i = V_i^1 \times V_i^2$ (i.e. each layer is the cartesian product of the corresponding layers in G^1 and G^2);
- There is an edge (u,v) in G connecting nodes (u^1, u^2) and (v^1, v^2) iff (u^1, v^1) and (u^2, v^2) are edges in G^1 and G^2 , respectively.



44

EVEN AND EVEN LAYOUT (3)

Exemples of LCP



Even and Litman proved that many well-known topologies are the LCP of simple structures (e.g. trees).


Specifically, the butterfly network is the LCP of two complete binary trees, an upward one and a downward one.

45

EVEN AND EVEN LAYOUT (4)

The Projection Methodology (PM):

Let G^1 and G^2 two layered graphs of $l+1$ layers each and let G denote their LCP. A layout of G is obtained with the PM as follows:



46


EVEN AND EVEN LAYOUT (5)

The Projection Methodology (cntd)

Consider a cube and draw the graph G^1 on the xy face so that:

- (a) the y -coordinate of every node $u \in V_i^1$ equals i
- (b) the x -coordinate of every node is an integer.

Similarly, draw the graph G^2 on the yz face




EVEN AND EVEN LAYOUT (6)

The Projection Methodology (cntd)

A three-dimensional drawing of the LCP G is constructed in the cube as follows:

- if $u \in V_i^1$ is drawn in coordinates $(x_u, i, 0)$ and $v \in V_i^2$ is drawn in coordinates $(0, i, z_v)$, then the coordinates of node $(u, v) \in V_i$ are (x_u, i, z_v) .


In other words, the nodes of G are the intersections between the lines orthogonal to plane xy and passing through nodes of G^1 and the lines orthogonal to plane yz and passing through nodes of G^2 .



EVEN AND EVEN LAYOUT (7)


The Projection Methodology (cntd)

A 2D drawing of G is obtained by projecting the 3D drawing to the xz plane.




EVEN AND EVEN LAYOUT (8)

Obs. It is possible to avoid constructing the 3D representation by immediately using the prolongations on plane xz of the projections of nodes in layer i of G^1 on the x -axis and of nodes in layer i of G^2 on the z -axis, $i=0,\dots,l$



EVEN AND EVEN LAYOUT (9)




The PM may produce layouts that do not satisfy the constraints required by the Thompson model.

For example, the drawing above is a grid drawing but it is not an orthogonal drawing.

51

EVEN AND EVEN LAYOUT (10)



We now describe how rectilinear layouts of G can be obtained via the PM. First, we formalize necessary and sufficient conditions:


- for the edges of the xz projection of G to be along grid paths,
- for nodes to be mapped to different grid points,
- for not using any grid edge more than once.

52

EVEN AND EVEN LAYOUT (11)

Four types of edges in the product graph G :

1. The product of two diagonal edges yields a diagonal edge;
2. The product of a vertical edge and a diagonal edge yields a vertical edge;
3. The product of a diagonal edge and a vertical edge yields a horizontal edge;
4. The product of two vertical edges yields a single grid point.



EVEN AND EVEN LAYOUT (12)

In order to get a feasible layout through the PM, we have to impose that the product of either two diagonal edges or two vertical edges never occurs.

More precisely:


1. The PM generates a layout of G in which the edges are grid lines if and only if the drawings of G^1 and G^2 on the faces of the cube satisfy the following condition: For every edge $e \in E$, exactly one of its factor is drawn diagonally.

This claim avoids overlappings of nodes of the same layer, too.

EVEN AND EVEN LAYOUT (13)

We need to impose that nodes in different layers do not overlap:

2. The PM generates a layout of G in which at most one node is mapped to each grid point if and only if the sets $\{(x_u, z_v): u \in V_i^1 \text{ e } v \in V_i^2\}$ are disjoint, for each $i=0, \dots, l$.




55

EVEN AND EVEN LAYOUT (14)

Consider now two diagonal edges (a,b) and (c,d) in G^1 ; the coordinates of nodes a, b, c, d are:

- node a : $(x_a, i, 0)$;
- node b : $(x_b, i+1, 0)$;
- node c : $(x_c, j, 0)$;
- node d : $(x_d, j+1, 0)$.




We say that these two edges are **consistent** if the open intervals (x_a, x_b) and (x_c, x_d) are disjoint.

56


EVEN AND EVEN LAYOUT (15)

3. The PM generates a layout of G in which no grid edge is used twice if and only if for every two inconsistent edges of one of the multiplicands the following condition holds: The two edges are not in the same layer, and on the two layers in which they appear, there are no (straight) edges of the other multiplicand which are collinear.

Inconsistent edges on the same layer



Inconsistent edges on different layers



EVEN AND EVEN LAYOUT (16)

In order to produce a feasible layout, we need to ensure that all three claims are satisfied.

Let us consider the Claims one by one:


EVEN AND EVEN LAYOUT (17)

1. The PM generates a layout of G in which the edges are grid lines if and only if the drawings of G^1 and G^2 on the faces of the cube satisfy the following condition: For every edge $e \in E$, exactly one of its factors is drawn diagonally.

A solution is to double the number of edge levels so that the edges in the drawing of G^1 are diagonal in odd layers and straight in the even layers, while the edges in the drawing of G^2 are straight in the odd layers and diagonal in the even layers.

59

EVEN AND EVEN LAYOUT (18)



The doubling of the number of edge levels is achieved by stretching each edge of the two multiplicands to become a path of two edges. In this way, we simulate the creation of edge bends.

60

EVEN AND EVEN LAYOUT (19)

2. The PM generates a layout of G in which at most one node is mapped to each grid point if and only if the sets $\{(x_u, z_v): u \in V_i^1 \text{ e } v \in V_i^2\}$ are disjoint, for each $i=0, \dots, l$.

A simple way to guarantee that this condition will hold is to make sure that no two nodes in the drawing of G^1 (G^2), except for the two end-points of the same straight edge, share the x -coordinate (z -coordinate).

This is always possible if we opportunely enlarge the drawings of the two factors.

61

EVEN AND EVEN LAYOUT (20)


3. The PM generates a layout of G in which no grid edge is used twice if and only if for every two inconsistent edges of one of the multiplicands, the following condition holds:

- The two edges are not in the same layer of the multiplicand, and
- on the two layers in which they appear, there are no (straight) edges of the other multiplicand which are collinear.

This condition is harder to enforce and is a severe limitation of this technique. For this reason, we limit to networks, each of which is the LCP of two trees.

62

EVEN AND EVEN LAYOUT (21)



63

EVEN AND EVEN LAYOUT (22)

The butterfly network is the LCP of an upward and a downward binary. (We dedicate a column to each vertex to prevent vertices of the layout from colliding.)

Proceed as follows:


- Draw one tree on the xy plane and the other on the yz plane
- Construct their LCP in 3D inside the cube with the PM
- The projection of this 3D figure on the floor is a planar layout of the butterfly

64

EVEN AND EVEN LAYOUT (23)

This layout has the following properties:

- It's symmetric 👍
- Its height and its width are $H=W=2(N-1)$, so its area is $4N^2+o(N^2)$ 🙄
- Input and output nodes are not on the boundary 🙄
- All the edges on the same layer has the same length 👍



65

COMPARING THE TWO TECHNIQUES

- **WISE – PROs:**
 - relatively small area
 - It “looks like” a butterfly
 - Input/output nodes on the boundary
- **WISE – CONs:**
 - knock-knees
 - “slanted” grid
- **EVEN & EVEN – PROs:**
 - It eliminates all the flaws
- **EVEN & EVEN – CONs:**
 - Larger area
 - input/output nodes inside the layout

66

OTHER RESULTS (1)

With the aim of optimizing the layout area, other layout algorithms have been proposed:

- Dinitz [’98] proves that the area of the Even & Even layout can be decreased by means of some local adjustments, so to achieve area $11/6 N^2 + o(N^2)$
- Later, Avior et al. [’98] prove that any butterfly layout cannot have area smaller than $N^2 + o(N^2)$ if “slanted” drawing is not allowed, and they provide an algorithm producing a layout of optimal area.

67

OTHER RESULTS (2)

Finally, Dinitz et al. [’99] prove that, if a “slanted” drawing is allowed, area $1/2 N^2 + o(N^2)$ is necessary and sufficient.

These works definitively close the optimal area layout problem of the Butterfly network.

68

OPTIMAL AREA LAYOUT OF THE BUTTERFLY NETWORK




69

OPTIMAL AREA LAYOUT - IDEA (1)

The two papers that provide an optimal area layout base their results on the following lemma:

Lemma: For any non-negative integers j, k , $0 \leq j \leq j+k \leq n$, the subgraph of the n -dim. Butterfly induced by the nodes of levels $j, j+1, \dots, j+k$ is the disjoint union of 2^{n-k} copies of k -dimensional butterflies.

In particular, if $j=0$ and $k=n-1$:




70

OPTIMAL AREA LAYOUT - IDEA (2)

Hence, an $(n-1)$ -dimensional butterfly can be built as a pair of $(n-2)$ -dim. butterflies connected by one node layer and one edge layer.


If we cut out the input and output nodes from an n -dim. Butterfly, we get:



71

OPTIMAL AREA LAYOUT - IDEA (3)


Each one of these $(n-2)$ -dim. Butterflies can be, in turn, cut into many smaller butterflies:



72

OPTIMAL AREA LAYOUT - IDEA (4)


The previous layout can be better specified as follows:



73

OPTIMAL AREA LAYOUT - IDEA (5)

Each rectangle contains a Butterfly that can be represented, either horizontally or vertically, layer by layer as follows:




Obs.: this layout is far from being optimal; nevertheless it allows to produce a final optimal layout.

74

OPTIMAL AREA LAYOUT - IDEA (6)


It remains to connect the small rectangular butterflies:



75

OPTIMAL AREA LAYOUT - IDEA (7)

In the case of slanted layout, it can be bent along the line:



76

OPTIMAL AREA LAYOUT - IDEA (8)


It is possible to prove tight lower and upper bounds on the layout area for both the models (usual and slanted).

The interested students can look at:

- A. Avior, T.C., S. Even, A. Litman, A.L. Rosenberg: A Tight Layout of the Butterfly Network. Theory of Computing Systems 31, 1998.
- Y. Dinitz, S. Even, M. Zapolotsky: A Compact Layout of the Butterfly. J. of Interconnection Networks 4, 2003.

77

LAYOUT OF THE HYPERCUBE NETWORK



78

THE HYPERCUBE (1)


Widely used for parallel computation, thanks to its nice properties (high regularity, logarithmic diameter, good fault tolerance, ...).

Def. The n -dimensional **Hypercube**, Q_n , has $N=2^n$ nodes and $\frac{1}{2}n2^n$ edges. Each node is labeled with an n -bit binary string, and two nodes are linked with an edge iff their binary strings differ in precisely one bit.

The edges of the hypercube can be naturally partitioned according to the dimensions that they traverse and Q_n can be seen as $Q_{n-1} \equiv Q_{n-1} \dots$

79

THE HYPERCUBE (2)




Q_n can be built by joining with an edge nodes in two different copies of Q_{n-1} if they have the same label.

Obs.: These edges form a perfect matching.

80

THE HYPERCUBE (3)



Property: Q_n has diameter $\log N$.

Sketch of proof. Any two nodes

$$u = u_1 u_2 \dots u_{\log N} \text{ and } v = v_1 v_2 \dots v_{\log N}$$

are connected by the path:

$$u_1 u_2 \dots u_{\log N} v_1 u_2 \dots u_{\log N} v_1 v_2 \dots u_{\log N} \dots v_1 v_2 \dots v_{\log N}$$

There are pairs of nodes requiring exactly $\log N$ steps. ■


81

THE HYPERCUBE (4)

Reminder: The **bisection width** of a network is the minimum number of edges one has to cut to disconnect the network into two equally sized subnetworks.

Property. $BW(Q_n) = N/2$.

Sketch of proof. the green edges (i.e. in a single dim.) divide the hypercube into two equally sized subnetworks; they are $N/2$ and it is not possible to cut a smaller number of edges to get the same result.



82

THE HYPERCUBE (5)

Th. A lower bound on the layout area of a network is the square of its bisection width (already proved).

Cor. Each layout of Q_n has area at least $N^2/4$.


In the following: layout with area $4/9N^2 + o(N^2)$, that hence is almost optimal (far from the lower bound by a factor of 1.7) [Yeh, Varvarigos, Parhami, '99].

83

COLLINEAR LAYOUT (1)

Reminder: In a **collinear layout** all nodes are placed on the same line. Instead of computing its area, it is usual to count the number of necessary **tracks**.


We start with a 2-dim. Hypercube, and inductively move to hypercubes of higher dimensions:



84

COLLINEAR LAYOUT (2)

- If n odd: Assume that we have a collinear layout for Q_{n-1} that requires $f(n-1)$ tracks:
 Q_n




2 tracks 2 tracks 1 track

Tot. $f(n)=2f(n-1)+1$ tracks

85

COLLINEAR LAYOUT (3)

- If n is even: To obtain the collinear layout of Q_n we start with the layouts of four Q_{n-2} s:



$$f(n+2)=4f(n)+2$$

86

COLLINEAR LAYOUT (4)

- **Th.** The number of tracks required for the collinear layout of Q_n is $2/3N$ (with $N=2^n$ no. of nodes).
- **Proof.** We solve the following recurrence equation:
 - $f(n)=2f(n-1)+1$ if n odd
 - $f(n)=4f(n-2)+2$ if n even
 - $f(2)=2$

Even case:

$$f(n)=4f(n-2)+2=4^2f(n-4)+4 \times 2+2=$$

$$=4^3f(n-6)+2^5+2^3+2=\dots=$$

$$=\dots \text{when } n-2k=2 \text{ iff } k=(n-2)/2 \dots = 4^k f(n-2k) + \sum_{i=0}^{k-1} 2 \cdot 4^i =$$

$$= 4^{\frac{n-2}{2}} f(2) + 2 \sum_{i=0}^{\frac{n-2}{2}-1} 4^i = 2 \sum_{i=0}^{\frac{n-2}{2}} 4^i \cong 2 \cdot \frac{4^{\frac{n-2}{2}+1} - 4}{3} = \frac{2}{3} 2^n = \frac{2}{3} N$$

87

COLLINEAR LAYOUT (5)

(proof cntd)

The odd case is analogous. ■

The area of this layout is $(2/3N+n) \times (nN)$.

Th. Q_n can be laid out in $4/9N^2+o(N^2)$ area.

Sketch of Proof. Let $n=n_1+n_2$.

Let us use 2^{n_1} copies of the collinear layout of Q_{n_2} , each placed along a row.

We connect the 2^{n_1} nodes that belong to the same column vertically according to the collinear layout of a Q_{n_1} .

88

COLLINEAR LAYOUT (6)

(proof cntd)

Reminder: Q_k needs of $2/3 2^k$ tracks

of horiz. tracks (rows):

2^{n_1} copies $\times$ $2/3 2^{n_2}$ tracks


Additive height (nodes):

2^{n_1} copies $\times$ n_2

Total height:

$2^{n_1} \times (2/3 2^{n_2} + n_2) =$

$2/3 2^{n_1+n_2+n_2} 2^{n_1}$



89

COLLINEAR LAYOUT (7)

(proof cntd)

Reminder:

height = $2/3 2^{n_1+n_2+n_2} 2^{n_1}$

The **width** is computed analogously, switching the roles of n_1 and n_2 .

Total width:


$2/3 2^{n_1+n_2+n_1} 2^{n_2}$

Area = $(2/3 N + n_1 2^{n_2})(2/3 N + n_2 2^{n_1}) =$


$= 4/9 N^2 + o(N^2)$ if n_1 and n_2 are $o(N)$, e.g. if

$n_1 = \Theta(n_2) \cong n/2$.

👉 possible students' lessons



90




3D LAYOUT

91

3D LAYOUT PROBLEM (1)


- The diffusion of the 3D layout has increased in the last thirty years.
- The topology is laid out on a series of slices.
- Further optimization of the wire length and number of bends
- Less silicon used.

3D Structure





Via Hole

2D Structure



Wiring



92

3D LAYOUT PROBLEM (2)

Def. A 3D layout of a topology G is a 1-1 function between G and the 3D grid such that:

- the nodes are mapped into grid points;
it is better if they lie on the external slice to minimize: energy consuming, production of heat and difficulty of connection with other devices
- the wires are mapped on grid paths so that:
 - these paths are edge-disjoint;
 - there are no "knock-knees"
 - these paths do not cross any mapping of a node that is not an extreme of the corresponding wire.

Aim: minimizing volume and keeping wires short.

👉 possible students' lessons