

ESE535: Electronic Design Automation

Day 9: February 16, 2009
Partitioning
(Intro, KLFM)



Penn ESE535 Spring 2009 -- DeHon

Today

- Partitioning
 - why important
 - practical attack
 - variations and issues

2

Penn ESE535 Spring 2009 -- DeHon

Motivation (1)

- Divide-and-conquer
 - trivial case: decomposition
 - smaller problems easier to solve
 - net win, if super linear
 - $\text{Part}(n) + 2 \times T(n/2) < T(n)$
 - problems with sparse connections or interactions
 - Exploit structure
 - limited cutsizes is a common structural property
 - random graphs would **not** have as small cuts

3

Penn ESE535 Spring 2009 -- DeHon

Motivation (2)


- Cut size (bandwidth) can determine area
- Minimizing cuts
 - minimize interconnect requirements
 - increases signal locality
- Chip (board) partitioning
 - minimize IO
- Direct basis for placement

4

Penn ESE535 Spring 2009 -- DeHon

Bisection Bandwidth

- Partition design into two equal size halves
- Minimize wires (nets) with ends in both halves
- Number of wires crossing is **bisection bandwidth**
- lower bw = more locality




5

Penn ESE535 Spring 2009 -- DeHon

Interconnect Area

- Bisection is lower-bound on IC width
 - When wire dominated, may be tight bound
- (recursively)



6

Penn ESE535 Spring 2009 -- DeHon

Classic Partitioning Problem

- **Given:** netlist of interconnect cells
- Partition into two (roughly) equal halves (A,B)
- minimize the number of nets shared by halves
- “Roughly Equal”
 - balance condition: $(0.5-\delta)N \leq |A| \leq (0.5+\delta)N$

Penn ESE535 Spring 2009 -- DeHon

7

Balanced Partitioning

- NP-complete for general graphs
 - [ND17: Minimum Cut into Bounded Sets, Garey and Johnson]
 - Reduce SIMPLE MAX CUT
 - Reduce MAXIMUM 2-SAT to SMC
 - Unbalanced partitioning poly time
- Many heuristics/attacks

Penn ESE535 Spring 2009 -- DeHon

8

KL FM Partitioning Heuristic

- Greedy, iterative
 - pick cell that decreases cut and move it
 - repeat
- small amount of non-greediness:
 - look past moves that make locally worse
 - randomization

Penn ESE535 Spring 2009 -- DeHon

9

Fiduccia-Mattheyses (Kernighan-Lin refinement)

- Start with two halves (random split?)
- Repeat until no updates
 - Start with all cells free
 - Repeat until no cells free
 - Move cell with largest gain (balance allows)
 - Update costs of neighbors
 - Lock cell in place (record current cost)
 - Pick least cost point in previous sequence and use as next starting position
- Repeat for different random starting points

Penn ESE535 Spring 2009 -- DeHon

Efficiency

Tricks to make efficient:


- Expend *little* work picking move candidate
 - Constant work $\equiv O(1)$
 - Means amount of work not dependent on problem size
- Update costs on move cheaply $[O(1)]$
- Efficient data structure
 - update costs cheap
 - cheap to find next move

Penn ESE535 Spring 2009 -- DeHon

11

Ordering and Cheap Update

- Keep track of Net gain on node == delta net crossings to move a node
 - cut cost after move = cost - gain
- Calculate node gain as Σ net gains for all nets at that node
 - Each node involved in several nets
- Sort nodes by gain




Penn ESE535 Spring 2009 -- DeHon

12

FM Cell Gains

Gain = Delta in number of nets crossing between partitions
= Sum of net deltas for nets on the node



Penn ESE535 Spring 2009 -- DeHon

13


After move node?

- Update cost
 - Newcost=cost-gain
- Also need to update gains
 - on all nets attached to moved node
 - but moves are nodes, so push to
 - all nodes affected by those nets

Penn ESE535 Spring 2009 -- DeHon

14

Composability of Net Gains




Penn ESE535 Spring 2009 -- DeHon

15

FM Recompute Cell Gain

- For each net, keep track of number of cells in each partition [F(net), T(net)]
- Move update:(for each net on moved cell)
 - if T(net)==0, increment gain on F side of net
 - (think -1 => 0)




Penn ESE535 Spring 2009 -- DeHon

16

FM Recompute Cell Gain

- For each net, keep track of number of cells in each partition [F(net), T(net)]
- Move update:(for each net on moved cell)
 - if T(net)==0, increment gain on F side of net
 - (think -1 => 0)
 - if T(net)==1, decrement gain on T side of net
 - (think 1=>0)




Penn ESE535 Spring 2009 -- DeHon

17

FM Recompute Cell Gain

- Move update:(for each net on moved cell)
 - if T(net)==0, increment gain on F side of net
 - if T(net)==1, decrement gain on T side of net
 - decrement F(net), increment T(net)




Penn ESE535 Spring 2009 -- DeHon

18

FM Recompute Cell Gain

- Move update:(for each net on moved cell)
 - if $T(\text{net})=0$, increment gain on F side of net
 - if $T(\text{net})=1$, decrement gain on T side of net
 - decrement $F(\text{net})$, increment $T(\text{net})$
 - if $F(\text{net})=1$, increment gain on F cell




Penn ESE535 Spring 2009 -- DeHon

19

FM Recompute Cell Gain

- Move update:(for each net on moved cell)
 - if $T(\text{net})=0$, increment gain on F side of net
 - if $T(\text{net})=1$, decrement gain on T side of net
 - decrement $F(\text{net})$, increment $T(\text{net})$
 - if $F(\text{net})=1$, increment gain on F cell
 - if $F(\text{net})=0$, decrement gain on all cells (T)



Penn ESE535 Spring 2009 -- DeHon

20

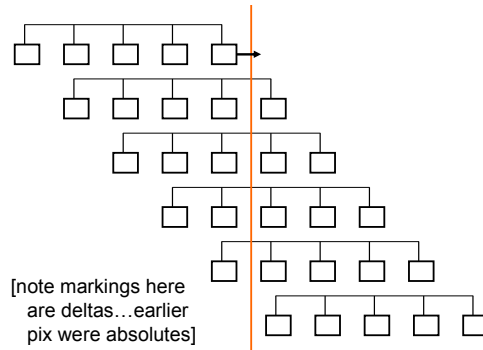
FM Recompute Cell Gain

- For each net, keep track of number of cells in each partition [$F(\text{net})$, $T(\text{net})$]
- Move update:(for each net on moved cell)
 - if $T(\text{net})=0$, increment gain on F side of net
 - (think $-1 \Rightarrow 0$)
 - if $T(\text{net})=1$, decrement gain on T side of net
 - (think $1 \Rightarrow 0$)
 - decrement $F(\text{net})$, increment $T(\text{net})$
 - if $F(\text{net})=1$, increment gain on F cell
 - if $F(\text{net})=0$, decrement gain on all cells (T)

Penn ESE535 Spring 2009 -- DeHon

21

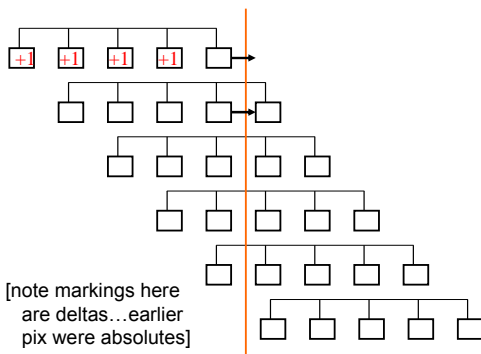
FM Recompute (example)



Penn ESE535 Spring 2009 -- DeHon

22

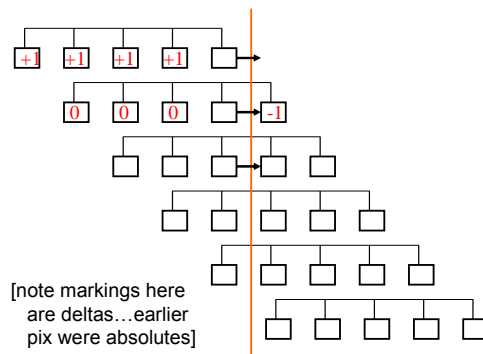
FM Recompute (example)



Penn ESE535 Spring 2009 -- DeHon

23

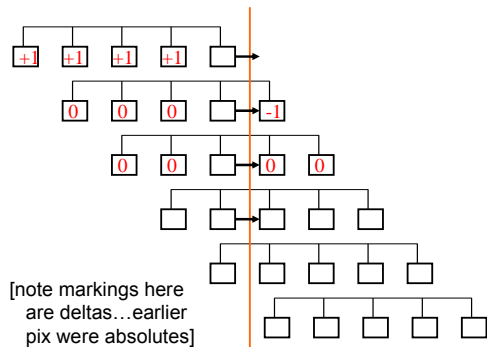
FM Recompute (example)



Penn ESE535 Spring 2009 -- DeHon

24

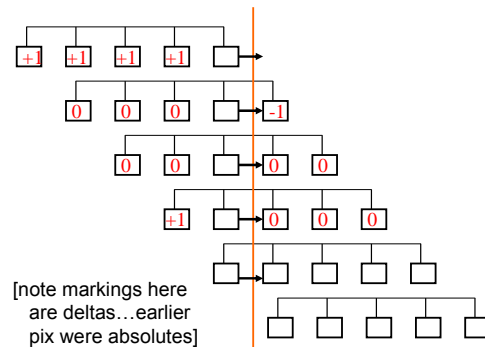
FM Recompute (example)



Penn ESE535 Spring 2009 -- DeHon

25

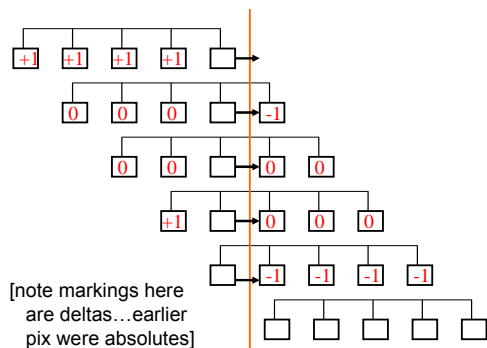
FM Recompute (example)



Penn ESE535 Spring 2009 -- DeHon

26

FM Recompute (example)




Penn ESE535 Spring 2009 -- DeHon

27

FM Data Structures

- Partition Counts A,B
- Two gain arrays
 - One per partition
 - **Key:** constant time cell update
- Cells
 - successors (consumers)
 - inputs
 - locked status



Penn ESE535 Spring 2009 -- DeHon

28

FM Optimization Sequence (ex)

+3	+2	-1
+3	+1	-1
+2	0	-1
+2	-1	0
+1	0	+1
0	-1	-1
+1	-1	-1
0	-2	-2
0	-2	-2
-1	-1	-2
+1	0	-1
-1	-1	-2
-2	-2	-1
-2	-3	-3
-3	-3	-3
-3	-3	-3
+12	+3	0

Penn ESE535 Spring 2009 -- DeHon

29

FM Running Time?

- Randomly partition into two halves
- Repeat until no updates
 - Start with all cells free
 - Repeat until no cells free
 - Move cell with largest gain
 - Update costs of neighbors
 - Lock cell in place (record current cost)
 - Pick least cost point in previous sequence and use as next starting position
- Repeat for different random starting points

Penn ESE535 Spring 2009 -- DeHon

30


FM Running Time

- **Claim:** small number of passes to converge
 - Constant passes?
- Small (constant?) number of random starts
- N cell updates each round (swap)
- **Updates K + fanout work (avg. fanout K)**
 - assume at most K inputs to each node
 - For every net attached (K+1)
 - For every node attached to those nets ($O(K)$)
- **Maintain ordered list $O(1)$ per move**
 - every io move up/down by 1
- Running time: $O(K^2N)$
 - Algorithm significant for its speed
 - (more than quality)

Penn ESE535 Spring 2009 -- DeHon

31

FM Starts?



So, FM gives
a **not bad**
solution
quickly


21K random starts, 3K network -- Alpert/Kahng

Penn ESE535 Spring 2009 -- DeHon

32

Weaknesses?

- Local, incremental moves only
 - hard to move clusters
 - no lookahead
- Looks only at local structure



Penn ESE535 Spring 2009 -- DeHon

33

Improving FM

- Clustering
- Initial partitions
- Runs
- Partition size freedom
- Replication

Following comparisons from Hauck and Boriello '96

Penn ESE535 Spring 2009 -- DeHon

34

Clustering

- Group together several leaf cells into cluster
- Run partition on clusters
- Uncluster (keep partitions)
 - iteratively
- Run partition again
 - using prior result as starting point
 - instead of random start

Penn ESE535 Spring 2009 -- DeHon

35

Clustering Benefits

- Catch local connectivity which FM might miss
 - moving one element at a time, hard to see move whole connected groups across partition
- Faster (smaller N)
 - METIS -- fastest research partitioner exploits heavily
 - FM work better w/ larger nodes (???)

Penn ESE535 Spring 2009 -- DeHon

36

How Cluster?

- Random
 - cheap, some benefits for speed
- Greedy “connectivity”
 - examine in random order
 - cluster to most highly connected
 - 30% better cut, 16% faster than random
- Spectral (next time)
 - look for clusters in placement
 - (ratio-cut like)
- Brute-force connectivity (can be $O(N^2)$)

Penn ESE535 Spring 2009 -- DeHon

37

Initial Partitions?

- Random
- Pick Random node for one side
 - start imbalanced
 - run FM from there
- Pick random node and Breadth-first search to fill one half
- Pick random node and Depth-first search to fill half
- Start with Spectral partition

Penn ESE535 Spring 2009 -- DeHon

38


Initial Partitions

- If run several times
 - pure random tends to win out
 - more freedom / variety of starts
 - more variation from run to run
 - others trapped in local minima

Penn ESE535 Spring 2009 -- DeHon

39

Number of Runs



Penn ESE535 Spring 2009 -- DeHon

40


Number of Runs

- 2 - 10%
- 10 - 18%
- 20 < 20% (2% better than 10)
- 50 (4% better than 10)
- ...but?

Penn ESE535 Spring 2009 -- DeHon

41

FM Starts?




21K random starts, 3K network -- Alpert/Kahng

Penn ESE535 Spring 2009 -- DeHon

42


Unbalanced Cuts

- Increasing slack in partitions
 - may allow lower cut size



Penn ESE535 Spring 2009 – DeHon

Unbalanced Partitions




Small/large is benchmark size not small/large partition IO.

Following comparisons from Hauck and Boriello '96

Penn ESE535 Spring 2009 – DeHon

Replication

- Trade some additional logic area for smaller cut size
 - Net win if wire dominated




Replication data from: Enos, Hauck, Sarrafzadeh '97

45

Penn ESE535 Spring 2009 – DeHon

Replication

- 5% → 38% cut size reduction
- 50% → 50+% cut size reduction



46

Penn ESE535 Spring 2009 – DeHon


What Bisection doesn't tell us

- Bisection bandwidth purely geometrical
- No constraint for delay
 - I.e. a partition may leave critical path weaving between halves

47

Penn ESE535 Spring 2009 – DeHon

Critical Path and Bisection



Minimum cut may cross critical path multiple times.
Minimizing long wires in critical path => increase cut size.

48

Penn ESE535 Spring 2009 – DeHon

So...

- Minimizing bisection
 - good for area
 - oblivious to delay/critical path

Partitioning Summary

- Decompose problem
- Find locality
- NP-complete problem
- linear heuristic (KLFM)
- many ways to tweak
 - Hauck/Boriello, Karypis
- even better with replication
- only address cut size, not critical path delay

Admin

- Reading for Wed. online
- No class next Monday (23rd)
 - Use time to finish Assignment 2B
 - Due 25th

Today's Big Ideas:

- Divide-and-Conquer
- Exploit Structure
 - Look for sparsity/locality of interaction
- Techniques:
 - greedy
 - incremental improvement
 - randomness avoid bad cases, local minima
 - incremental cost updates (time cost)
 - efficient data structures