

ESE535: Electronic Design Automation


Day 14: March 16, 2015
Routing 2
(Pathfinder)



Penn ESE 535 Spring 2015 -- DeHon

Today

- Routing
 - Pathfinder
 - graph based
 - global routing
 - simultaneous global/detail




Penn ESE 535 Spring 2015 -- DeHon

2

Global Routing

- **Problem:** Find sequence of channels for all routes
 - minimizing channel sizes
 - minimize max channel size
 - meeting channel capacity limits




Penn ESE 535 Spring 2015 -- DeHon

3

Global → Graph

- Graph Problem on routes through regions




Penn ESE 535 Spring 2015 -- DeHon

4


Global/Detail

- With limited switching (e.g. FPGA)
 - can represent routing graph exactly



Penn ESE 535 Spring 2015 -- DeHon

5




Penn ESE 535 Spring 2015 -- DeHon

6

Routing in Graph

- Find {shortest,available} path between source and sink
 - search problem (e.g. BFS, A*)



Penn ESE 535 Spring 2015 -- DeHon

7


Breadth First Search (BFS)

- Start at source src
- Put src node in priority queue with cost 0
 - Priority queue orders by cost
- While (not found sink)
 - Pop least cost node from queue
 - Get: current_node, current_cost
 - Is this sink? → found
 - For each outgoing edge from current_node
 - Push destination onto queue
 - with cost current_cost+edge_cost

Penn ESE 535 Spring 2015 -- DeHon

8


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

9

Search Animation




Not possible to use minimum Manhattan distance route.
Why?

Penn ESE 535 Spring 2015 -- DeHon

10


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

11

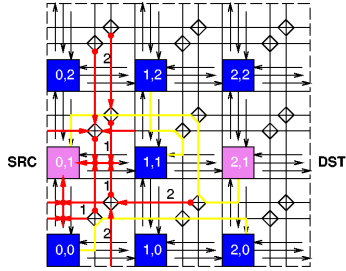
Search Animation



Penn ESE 535 Spring 2015 -- DeHon

12


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

13


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

14


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

15


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

16


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

17


Search Animation



Penn ESE 535 Spring 2015 -- DeHon

18

Search Animation



Penn ESE 535 Spring 2015 -- DeHon

19


Easy?

- Finding a path is moderately easy
- What's hard?
- Can I just iterate and pick paths?
 - Does greedy selection work?

Penn ESE 535 Spring 2015 -- DeHon

20


Example



Penn ESE 535 Spring 2015 -- DeHon

21

Challenge



- Satisfy *all* routes simultaneously
- Routes share potential resources
- Greedy/iterative
 - not know who will need which resources
 - E.g. consider routing s3->d3 then s2->d2 then s1->d1
 - *i.e.* resource/path choice looks arbitrary
 - ...but earlier decisions limit flexibility for later
 - like scheduling
 - order effects result

Penn ESE 535 Spring 2015 -- DeHon

22

Negotiated Congestion

- Idea:
 - try once
 - see where we run into problems
 - undo problematic/blocking allocation
 - rip-up
 - use that information to redirect/update costs on subsequent trials
 - retry

Penn ESE 535 Spring 2015 -- DeHon

23

Negotiated Congestion

- Here
 - route signals
 - allow overuse
 - identify overuse and encourage signals to avoid
 - reroute signals based on overuse/past congestion

Penn ESE 535 Spring 2015 -- DeHon

24

Basic Algorithm


- Route signals along minimum cost path
- If congestion/overuse
 - assign higher cost to congested resources
 - Makes problem a shortest path search
 - Allows us to adapt costs/search to problem
- Repeat until done

Key Idea

- Congested paths/resources become expensive
- When there is freedom
 - future routes with freedom to avoid congestion will avoid the congestion
- When there is less freedom
 - must take congested routes
- Routes that must use congested will, others will chose uncongested paths


Cost Function (1)

- $\text{PathCost} = \sum (\text{link costs})$
- $\text{LinkCost} = \text{base} \times f(\#\text{routes using, time})$
- Base cost of resource
 - E.g. delay of resource
 - Encourage minimum resource usage
 - (minimum length path, if possible)
 - minimizing delay = minimizing resources
- Congestion
 - penalizes (over) sharing
 - increase sharing penalty over time




Example (first order congestion)

FUTURE: maybe
So they can callou




What is preferred path for $s1 \rightarrow d1$, $s2 \rightarrow d2$, $s3 \rightarrow d3$?

Example (first order congestion)



All, individual routes prefer middle; create congestion.


Example (first order congestion)



If make congestion expensive:
e.g. $\text{cost}(\text{congest}) = 2^{\text{users}}$


What happens when
reroute $s1 \rightarrow d1$?

Example (first order congestion)




Reroute, avoid congestion.

Example (need for history)

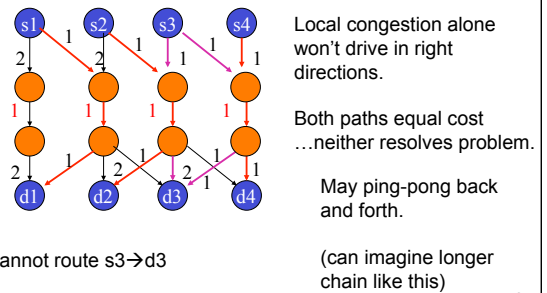


Need to redirect uncongested paths.

Example (need for history)



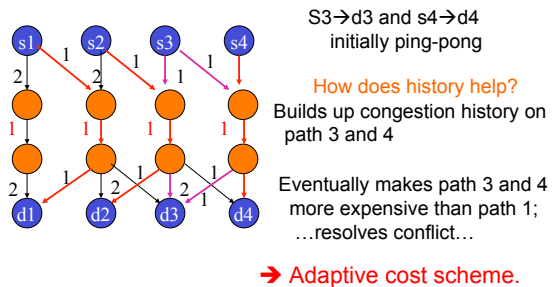
Example (need for history)



Cost Function (2)

- Cost = (base + history)*f(#resources,time)
- History
 - avoid resources with history of congestion
 - E.g. add 1 to history every time resource is congested

Example (need for history)



Delay

What about delay?

- Existing formulation uses delay to reduce resources, but doesn't directly treat
- How do we want to optimize delay?
- Want:
 - prioritize critical path elements for shorter delay
 - allow nodes with slack to take longer paths

Integrate Delay into Cost Function

- Cost=
 - $(1-W(\text{edge})) \times \text{delay} + W(\text{edge}) \times \text{congest}$
 - congest as before
 - $(\text{base} + \text{history}) \times f(\#\text{signals}, \text{time})$
- $W(\text{edge}) = \text{Slack}(\text{edge}) / D_{\max}$
 - 0 for edge on critical path
 - >0 for paths with slack
- Use $W(\text{edge})$ to order routes
- Update critical path and W each round

Cost Function (Delay)

- Cost=
 - $(1-W(\text{edge})) \times \text{delay} + W(\text{edge}) \times \text{congest}$
 - congest as before
 - $(\text{base} + \text{history}) \times f(\#\text{signals}, \text{time})$
- $W(\text{edge}) = \text{Slack}(\text{edge}) / D_{\max}$
- What happens if multiple slack 0 nets contend for edge?
- $W(\text{edge}) = \text{Max}(\min W, \text{Slack}(\text{edge}) / D_{\max})$
 - $\min W > 0$

Problem


- Are nanoseconds and congestion comparable?
- How normalize/weight so can add together?

VPR

- If doesn't uncongect, weight congestion more
- Cost=
 - $(1-W(e)) \times \text{delay} + W(e) \times \text{PF}^{(\text{iter})} \times \text{congest}$
 - PF=Pressure Factor Multiplier
- Eventually congest dominates delay
- What might go wrong?

VPR Pressure Factor

- Converges quickly
- But may “freeze” with higher delay than necessary
- Netlist Shuffle experiment




Penn ESE 535 Spring 2015 -- DeHon

[Rubin / FPGA 2011]

43

VPR Pressure Factor Tuning



$$(1-W(e)) \times \text{delay} + W(e) \times \text{PF}^{(\text{iter})} \times \text{congest}$$

Penn ESE 535 Spring 2015 -- DeHon

[Raphael Rubin 2010]

44

Alternate Delay Approach


- Believe Pathfinder can resolve congestion
- Pathfinder has trouble mixing delay and congestion
- **Idea:** Turn delay problem into congestion problem
 - Reject paths that are too long
 - All signals compete only for resources that will allow them to meet their timing goals

Penn ESE 535 Spring 2015 -- DeHon

45

Outlaw Long Paths

- **Issue:** Critical path may go through multiple gates
 - Contain more than one gate → gate path
 - How allocate slack among paths?




Penn ESE 535 Spring 2015 -- DeHon

46

Outlaw Long Paths

- **Issue:** Critical path may go through multiple gates
 - Contain more than one gate → gate path
 - How allocate slack among paths?




Penn ESE 535 Spring 2015 -- DeHon

47

Outlaw Long Paths

- **Issue:** Critical path may go through multiple gates
 - Contain more than one gate → gate path
 - How allocate slack among paths?




Penn ESE 535 Spring 2015 -- DeHon

48

Slack Budgeting

- Divide slack among the paths
 - Slack of 3
 - Example: give slack 1 to first link
2 to second




Penn ESE 535 Spring 2015 -- DeHon

[So / FPGA 2008]

49

Slack Budgeting

- Divide slack among the paths
- Each net now has delay target
- Reject any path exceeding delay target
- Reduce to congestion negotiation



Penn ESE 535 Spring 2015 -- DeHon

[So / FPGA 2008]

50

Slack Budgeting

- Can often find lower delay routes that VPR
- Takes 10x as long
 - Mostly in slack budgeting
- Solution depends on slack budget
 - Not exploiting full freedom to re-allocate slack among links

Penn ESE 535 Spring 2015 -- DeHon

[So / FPGA 2008]

51

Delay Target Routing

- Similar high-level idea
- Just set target for Pathfinder cost
 - Rather than allowing to float

Penn ESE 535 Spring 2015 -- DeHon

52

Delay Target

- Cost=

$$(1-W(\text{edge})) \cdot \text{delay} + W(\text{edge}) \cdot \text{congest}$$
- $W(\text{edge}) = \text{Slack}(\text{edge}) / D_{\text{target}}$
 - Previously: denominate was D_{max}
- Compute Slack based on D_{target}
 - can be negative
- $W(\text{edge}) = \text{Max}(\text{min}W, \text{Slack}(\text{edge}) / D_{\text{target}})$
 - $\text{min}W > 0$

Penn ESE 535 Spring 2015 -- DeHon



53

Delay Target Routing

- Does allow slack to be used on any of the gate→gate connections on path
 - ...but not being that deliberate/efficient about the allocation
- Doesn't require time for slack allocation



Penn ESE 535 Spring 2015 -- DeHon

54




- ### Run Time?
- Route $|E|$ edges
 - Each path search $O(|E_{\text{graph}}|)$ worst case
 - ...generally less
 - Iterations?
- Penn ESE 535 Spring 2015 -- DeHon 57

- ### Quality and Runtime Experiment
- For Synthetic netlists on HSRA
 - Expect to be worst-case problems
 - Congestion only
 - Quality = # channels
 - Number of individual route trials limited (measured) as multiple of nets in design
 - (not measuring work per route trial)
-
- Penn ESE 535 Spring 2015 -- DeHon 58



Conclusions?

- Iterations increases with N
- Quality degrade as we scale?



Penn ESE 535 Spring 2015 -- DeHon

61


Techniques to Accelerate

(already in use in data just shown)

Penn ESE 535 Spring 2015 -- DeHon

62

Inefficient?




What is inefficient about this search?

How might we do better?

Penn ESE 535 Spring 2015 -- DeHon

63

Inefficient?



What if we only searched for minimum length paths?


How would we do that?

Downside?

Penn ESE 535 Spring 2015 -- DeHon

64

Recall Search Example




Not possible to use minimum Manhattan distance route.

Penn ESE 535 Spring 2015 -- DeHon

65


Only Search Minimum Length



Penn ESE 535 Spring 2015 -- DeHon

66



Minimum Search



What is the minimum we need to search (if uncongested)?

What would that search look like?

BFS vs. DFS





Search Ordering

- Default: breadth first search for shortest
 - $O(\text{total-paths})$
- Alternately: use A*:
 - estimated costs/path length, prune candidates earlier
 - can be more depth first
 - (search promising paths as long as know can't be worse)



BFS → A*

- Start at source
- Put src node in priority queue with cost 0
 - Priority queue orders by cost
 - $\text{Cost} = \Sigma(\text{path so far}) + \text{min path to dest}$ Maybe show a parameter here for tuning ho
- While (not found sink)
 - Pop least cost node from queue
 - Get: current_node, current_cost
 - Is this sink? → found
 - For each outgoing edge
 - Push destination onto queue
 - with cost $\text{current_cost} + \text{edge_cost}$

BFS vs. A*





Single-side, Directed (A*)




Only expand search windows as prove necessary to have longer route.

Search: one-side vs. two-sides



Searching

- In general:
 - greedy/depth first searching
 - find a path faster
 - may be more expensive
 - (not least delay, congest cost)
 - tradeoff by weighting
 - estimated delay on remaining path vs. cost to this point
 - control greediness of router
 - More greedy is faster at cost of less optimal paths (wider channels)
 - 40% W → 10x time reduction [Tessier/thesis'98]



Searching

- Use A* like search
 - Always expanded (deepen) along shortest ... as long as can prove no other path will dominate
 - Uncongested: takes $O(\text{path-length})$ time
 - Worst-case reduces to breadth-first
 - $O(\text{total-paths})$

Summary

- Finding short path easy/well known
- **Complication:** need to route set of signals
 - who gets which path?
 - Arbitrary decisions earlier limit options later
- **Idea:** iterate/relax using congestion history
 - update path costs based on congestion
 - Cost adaptive to route
 - reroute with new costs
- Accommodate delay and congestion

Big Ideas

- Exploit freedom
- Technique:
 - Graph algorithms (BFS, DFS)
 - Search techniques: A*
 - Iterative improvement/relaxation
 - Adaptive cost refinement

Admin

- Assignment 4 due Wednesday
- Reading for Wednesday on web
- Spring Break next week
- Reading for Monday after break
 - On Blackboard