

## (Lec 13) ASIC Layout: Routing by Maze Search

### ■ What you know

- ▼ Elementary placement by annealing
- ▼ Given the gates & wires: where do we put them to minimize estimated wire length

### ■ What you don't know

- ▼ How to actually **wire** the gates together
- ▼ Flavors of routing: global versus detailed, area versus region
- ▼ Our technical focus: area routing by maze routing

© R. Rutenbar, CMU 18-760, Spring99 1

## Copyright Notice

**© Rob A. Rutenbar 1999**  
**All rights reserved.**

**You may not make copies of this material in any form without my express permission.**

© R. Rutenbar, CMU 18-760, Spring99 2

## Where Are We?

### ■ Physical design--how to wire the placed gates...?

	M	T	W	Th	F	
Jan	11	12	13	14	15	1
	18	19	20	21	22	2
	25	26	27	28	29	3
Feb	1	2	3	4	5	4
	8	9	10	11	12	5
	15	16	17	18	19	6
	22	23	24	25	26	7
Mar	1	2	3	4	5	8
	8	9	10	11	12	9
	15	16	17	18	19	10
Break	22	23	24	25	26	11
	29	30	31	1	2	12
Apr	5	6	7	8	9	13
	12	13	14	15	16	14
	19	20	21	22	23	15
	26	27	28	29	30	16

Introduction  
 Advanced Boolean algebra  
 JAVA Review  
 BDDs  
 Formal verification  
 2-Level logic synthesis  
 Multi-level logic synthesis  
 Technology mapping  
 Placement  
**Routing**  
 Partitioning  
 Static timing analysis  
 Geometric data structures  
 Test and ATPG

© R. Rutenbar, CMU 18-760, Spring99 3


## Readings

### ■ Sarrafzadeh Book


- ▼ Section 3.1 is about maze routing and some other classical detailed area routing techniques.

© R. Rutenbar, CMU 18-760, Spring99 4

## Routing: The Problem



Dozens of circuits,  
1000s of blocks,  
100,000s of gates



Many meters of wire.


© R. Rutenbar, CMU 18-760, Spring99 5

## Routing: Global

■ Begin with the floorplan of the chip.

■ Problem:

- ▼ Route each net,
- ▼ ...within the given wire area
- ▼ ...without using too much wire, or causing unreasonable signal delays




© R. Rutenbar, CMU 18-760, Spring99 6

## Routing: Global Routing

### ■ Usually start with *global* or *coarse* routing

- ▼ Chop up chip into big regions
- ▼ Decide thru which regions the wires will go, but not exactly where each segment of each individual wire will go
- ▼ Idea is to **plan** global paths for the wires, so we know early we can fit them all in each region when we finally embed detailed rectangles

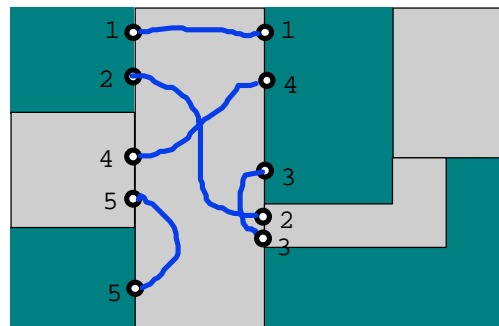


© R. Rutenbar, CMU 18-760, Spring99 7

## Routing: Global Routing

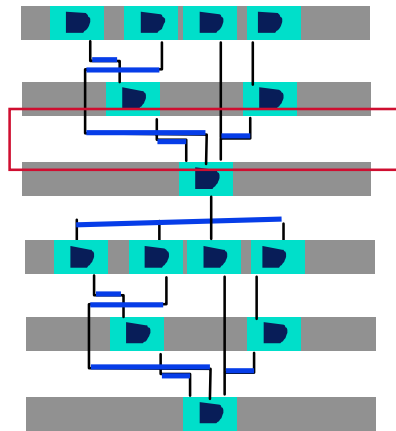
### ■ Result of global routing

- ▼ In each region of the chip, we know exactly which wires go thru that region, and we know where the pin IOs are to enter and exit
- ▼ Typical decomposition for ASICs is into rectangular channels, as below
- ▼ Signals only enter on the 2 opposite sides



© R. Rutenbar, CMU 18-760, Spring99 8

## For Row-Based Placements




### ■ Alternates logic & wiring

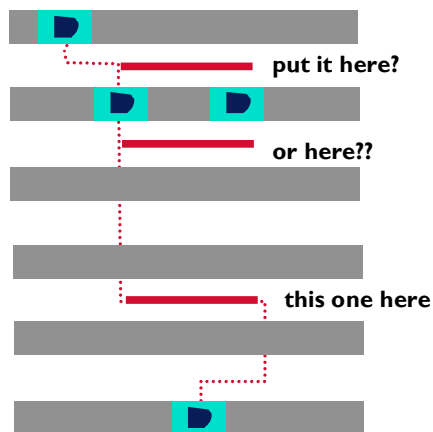
- ▼ Regions for wiring called **channels**, pins on top & bottom
- ▼ Used when you have only 2 or 3 layers of metal wiring
- ▼ Global routing determines where row-spanning signals cross the rows, and where the horizontal extent of signals are placed

© R. Rutenbar, CMU 18-760, Spring99 9

## Global Routing for Row-Based ASICs



Determines where each row gets crossed by row-spanning signal



Determines which channels horizontal parts of signal will use, when there is choice

© R. Rutenbar, CMU 18-760, Spring99 10

## Aside: Placement + Global Routing

### ■ Smart row-based placers do some global routing

- ▼ Helps decide if placement is good, by looking at where global routing wants to use space
- ▼ Routing can make rows wider if you need to add space to let signals cross the rows (depends on metal layer, use of pins in cells, etc)
- ▼ Routing can make layout taller if you need lots of tracks for wiring in each channel. If you make smarter decisions about where to put horizontal parts of the wiring in global routing, can get smaller layouts.

### ■ How?


- ▼ Can do some decent global routing inside an annealing-based placer
- ▼ Start global routing near the end, when you have OK evolving placement
- ▼ Look at row crossings, predicted congestion in channels, etc
- ▼ Try to evolve placement and global routing at same time.

© R. Rutenbar, CMU 18-760, Spring99 11

## Routing: Detailed Routing

### ■ Detailed routing follows global routing

- ▼ Detailed here means “actually put down the exact final rectangles that make each individual wire”
- ▼ In this case, you would use a Channel router, which wires up a channel-shaped rectangle with pins on the 2 opposite sides




© R. Rutenbar, CMU 18-760, Spring99 12


## Routing: Detailed Routing

- Different detailed routers for different regions:

Switchbox router  
for when rectangle has  
pins on all 4 sides




Channel router



© R. Rutenbar, CMU 18-760, Spring99 13

## Routing: Global + Detailed

- Repeat for each region until the whole chip is routed.
- Does it always work...?
  - ▼ Nope
  - ▼ Often get some unrouted nets which require some rework by hand.



© R. Rutenbar, CMU 18-760, Spring99 14

## Routing: Philosophy

### ■ CAD Conventional Wisdom:

- ▼ Divide & conquer a good idea on big, tightly coupled problems.

### ■ Use hierarchy where you can:

- ▼ Solve a **coarse** version of the problem.
- ▼ Refine to solve the **detailed** version.

### ■ Global router + detailed router

- ▼ Divide routable areas into coarse regions.
- ▼ For each wire, determine which regions it will pass through.
- ▼ Later, let a detailed router finish the wiring in each region.

© R. Rutenbar, CMU 18-760, Spring99 15

## Routing: Detailed Flavors

### ■ Two flavors for detailed router

#### ■ Area router: One-net-at-a-time, over a general area

- ▼ Completely routes all pins of a net, over as much of the wire area of the carrier as required.
- ▼ Routes one wire at a time; sensitive to wire ordering
- ▼ Why? Each wired net becomes an obstacle for future nets.

#### ■ Region router: Region-at-a-time

- ▼ Completely routes all nets constrained to pass through some region on the chip. Routes all nets simultaneously.

#### ■ Historically,

- ▼ First came the one-net-at-a-time routers. Example: maze routers
- ▼ Then came region-routers for the divide & conquer. Ex: channel route
- ▼ Now, people are moving back to maze routers again...


© R. Rutenbar, CMU 18-760, Spring99 16




## Why Area Routers Again?

### ■ Technology marches on

- ▼ Now have lots of layers of wiring.
- ▼ Don't have to only put wires between the blocks of the chip
- ▼ Now you can put wires over blocks of the chip
- ▼ Maze routers good at dealing with obstacles.



Chop up chip, cells and all, into regions for global routing




© R. Rutenbar, CMU 18-760, Spring99 17

## Area Routers Revisited




**Chunk of IC to be detail routed**  
**Pins appear on boundary or anywhere inside the region. Typically 20-100 wires to be embedded, over obstacles of circuits in blocks below. 3-6 wiring layers. Probably also have timing constraints.**



© R. Rutenbar, CMU 18-760, Spring99 18

## Typical Problem

IBM PowerPC 601



11.2 x 11.2 mm

### ■ Big processor chip

- ▼ Lots of layers of wiring
- ▼ Lots of big functional units
- ▼ Lots of signals, both locally inside the functional units, and globally unit-to-unit
- ▼ Could chop this up into regions and then try global + detailed routing

### ■ Still a big, very hard problem

© R. Rutenbar, CMU 18-760, Spring99 19

## So, What Will We Look At...?

### ■ Classical maze-style area-routers

- ▼ Can handle obstacles well...
- ▼ ..but they are sensitive to order in which wires routed.
- ▼ Can handle multiple wiring layers well...
- ▼ ...but get slower with more wiring layers.
- ▼ Can handle some timing constraints well...
- ▼ ...but we don't have time to do this. (Take 18-763...)

© R. Rutenbar, CMU 18-760, Spring99 20

## Routing: Maze Routers

### ■ Our Topics:

- ▼ History
- ▼ Basic Mechanics
  - ▼ Two-point nets in one layer - unit cost
  - ▼ Multipoint nets
  - ▼ Multiple layers
  - ▼ Weighted cost
- ▼ Design Variants
  - ▼ Vanilla scheme
  - ▼ Depth-first search (Rubin's scheme)

© R. Rutenbar, CMU 18-760, Spring99 21

## Maze Routing: History

### ■ 1961

- ▼ Lee, C. Y., "An algorithm for path connections and its applications", *IRE Trans. on Electronic Computers*, pp. 346-365, Sept. 1961.
- ▼ Chester Lee of Bell Labs invents the algorithm; gets famous for "Lee routers"

### ■ 1974

- ▼ Rubin, F., "The Lee path connection algorithm", *IEEE Trans. on Computers*, vol. c-23, no. 9, pp. 907-914, Sept. 1974.
- ▼ Frank Rubin comes up with a way to make it go much faster.

### ■ 1983

- ▼ Hightower, D., "The Lee router revisited", *ICCAD*, pp. 136-139, 1993.
- ▼ Dave Hightower, who originally got famous for coming up with an alternative to the Lee-router (a Hightower line-probe router) that was faster and used a lot less memory, undergoes a spiritual conversion and implements a killer maze router. Trick is: now machines have enough (real and virtual) memory to do big maze routing tasks.

© R. Rutenbar, CMU 18-760, Spring99 22

## Maze Routing: Strategy

### ■ Strategy

- ▼ One net at a time - completely wire one net.
- ▼ Optimize path - find the **best** wiring path.

### ■ Problems


- ▼ Early nets wired may block path of later nets.
- ▼ Optimal choice for one net may block later nets.

### ■ Solutions

- ▼ Careful net ordering.
- ▼ Careful optimization to include impact on later wiring.
- ▼ No Guarantees.
- ▼ (How do people really do it today? Let router remove blocking nets or shove them aside; called **ripup/reroute** or **shove-aside** routing.)

© R. Rutenbar, CMU 18-760, Spring99 23

## Maze Router: Basic Idea




### ■ Given:

- ▼ Grid - each square represents where one wire can cross.
- ▼ A source and target.

### ■ Problem:

- ▼ Find shortest path connecting source and target.

wires can:




© R. Rutenbar, CMU 18-760, Spring99 24


## Maze Routing: Expansion

S

Start at the source.



Find all new cells that are reachable at distance 1, ie, all paths that are just 1 unit in total length - mark all with distance.




Using the distance 1 cells, find all new cells which are reachable at distance 2.

Repeat until the target is found.

© R. Rutenbar, CMU 18-760, Spring99 25

## Maze Router: Expansion




### Strategy

- ▼ Expand one cell at a time until all of the shortest paths from S to T are found.
- ▼ Expansion creates a wavefront of paths that search broadly out from source cell until target is hit

© R. Rutenbar, CMU 18-760, Spring99 26

## Maze Router: Backtrace




### ■ Now what? Backtrace

- ▼ Select a shortest-path (any shortest-path) back to the source and mark its cells so they can't be used again.
- ▼ Since there are many paths back, optimization information can be used to select the best one.
- ▼ Here, just follow the path costs in the cells in descending order...

© R. Rutenbar, CMU 18-760, Spring99 27

## Maze Router: Clean-Up




### ■ Now what? Clean-up

- ▼ Clean up the grid for the next net, leaving the S to T route as an obstacle.
- ▼ Now, ready to route the next net with the obstacles from the previously routed net in place in the grid.

© R. Rutenbar, CMU 18-760, Spring99 28

## Maze Router: Blockage



### ■ Blockages

- ▼ All future nets must route around this blockage.

© R. Rutenbar, CMU 18-760, Spring99 29

## Classical Maze Router

### ■ Three main steps:

#### ■ Expansion

- ▼ Breadth-first-search to find all paths from source to target.

#### ■ Backtrace

- ▼ Walk the shortest path back to the source and mark path cells as used.

#### ■ Clean-Up

- ▼ Erase all distance marks from other grid cells before the next net is routed.

© R. Rutenbar, CMU 18-760, Spring99 30

## Maze Router: Concerns

### ■ Storage

- ▼ Do we need a really big grid to represent a big routing problem?
- ▼ What info required in each cell of this grid?

### ■ Complexity

- ▼ Do we really have to search the whole grid each time we add a wire?

### ■ Technology

- ▼ Just 1 wiring layer? How do we do 2 layers? 3? 4? 6??
- ▼ Complex wire widths or spacings?

### ■ 2 issues here

- ▼ Applications of basic algorithm
- ▼ Implementation issues for the basic algorithm

© R. Rutenbar, CMU 18-760, Spring99 31

## Applications: Multipoint Nets

### ■ Multipoint Nets

- ▼ One source -> **Many** targets.
- ▼ You get this with any net that represents fanout


### ■ Strategy

- ▼ Use maze route algorithm to find path from source to nearest target.
- ▼ Relabel all cells on the path as sources and rerun maze router using all sources simultaneously.
- ▼ Repeat for each segment.

© R. Rutenbar, CMU 18-760, Spring99 32



## Multipoint Nets



■ **Given:**

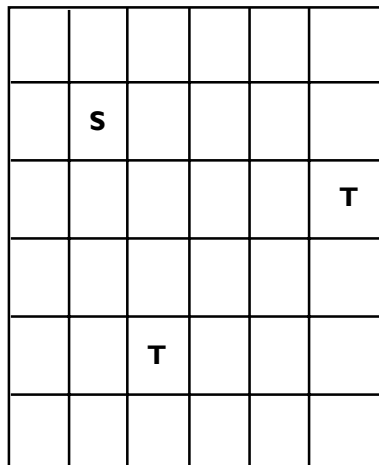
- ▼ A source and many targets.

■ **Problem:**

- ▼ Find shortest path connecting source and targets.

© R. Rutenbar, CMU 18-760, Spring99 33

## Multipoint Nets




■ **First...**

- ▼ Run maze route to find the closest target.
- ▼ Start at source, go till you find ANY target.

© R. Rutenbar, CMU 18-760, Spring99 34

## Multipoint Nets




### ■ Second...

- ▼ Backtrace and relabel the **whole route** as sources for the next pass.
- ▼ We will expand this entire set of source cells to find the net segment of the net
- ▼ Idea is we will look for paths of length 1 away from this whole set of sources, then length 2, 3, etc.
- ▼ Go till you hit another target

© R. Rutenbar, CMU 18-760, Spring99 35

## Multipoint Nets




### ■ Trick

- ▼ Expand simultaneously from all of these sources to find the shortest path from the existing route to the next target.

© R. Rutenbar, CMU 18-760, Spring99 36

## Multipoint Nets



### ■ Finally

- ▼ Do usual cleanup
- ▼ Mark all of the segment cells as used and clean-up the grid.
- ▼ Now, have embedded a multipoint net, and rendered it as an obstacle for future nets

© R. Rutenbar, CMU 18-760, Spring99 37

## Application: Multiple Layer Routing

### ■ How -- mechanically do we handle multiple layers?

- ▼ Parallel grids, vertically stacked, one for each layer.
- ▼ Use vias to access other layers.
- ▼ Label cell as to whether a via is permitted at this location.

### ■ New expand


- ▼ **Out** on each layer,
- ▼ **Up/down** at each via.

© R. Rutenbar, CMU 18-760, Spring99 38


## Multiple Layers

■ 2 parallel grids, vertically stacked

▼ Expansion can now go UP/DOWN; vias can go where “v” mark is




Layer 1




Layer 2 © R. Rutenbar, CMU 18-760, Spring99 39

## Multiple Layers

■ Backtrace & cleanup as usual



Layer 1



Layer 2

© R. Rutenbar, CMU 18-760, Spring99 40


## Aside: About Vias

### ■ Vias are how you move layer to layer

- ▼ Electrical connects between separate layers of physical wiring
- ▼ “Vertical” electrical connection

### ■ Geometric issue #1: Size

- ▼ On chips, vias are now usually a lot bigger than wire widths are
- ▼ So you have to be careful where you assume you can put them
- ▼ You can't put vias as close to each other as you can put wires




© R. Rutenbar, CMU 18-760, Spring99 41


## Aside: About Vias

### ■ Geometric Issues #2: Vertical stacking

- ▼ Relevant in multi-layer metal process, and in PC boards
- ▼ Can you put multiple vias connecting different sets of layers directly on top of each other, in a so-called **stack**?
- ▼ In some processes yes, in some no. Router has to deal with this.




5-layer metal cross section  
from IBM PowerPC



Metal 1 to metal 2 vias

© R. Rutenbar, CMU 18-760, Spring99 42

## Implementation Issues: Non-unit grids



### Old problem

- ▼ Each cell in grid cost the same to cross it with a wire
- ▼ Cost == 1, unit-cost
- ▼ Is this necessary?

### Now

- ▼ Given grid, Source and target
- ▼ Weights for each cell.


### Problem:

- ▼ Find **minimum cost** path connecting source and target.

shaded cells cost 3

© R. Rutenbar, CMU 18-760, Spring99 43

## Weighted Grids



### Many good applications


- ▼ Make the router avoid electrically sensitive areas of IC
- ▼ After global routing, weight cells with lots of potential wiring congestion higher, so router tries to avoid them
- ▼ Can make different layers have different expense to use
- ▼ Can make different vias have different expense to use
- ▼ Can make different directions of expansion have different expense, eg, you want metal 2 mostly vertical, so left-right expansions cost more...

### Expansion...?

- ▼ Always expand next **cheapest** partial path

© R. Rutenbar, CMU 18-760, Spring99 44


## Subtle Search Issues with NonUnit Costs



What cost does this get, and why?

Cost = 4 = cheapest path to this cell from the Source cell

We expand the cell to north to reach this cell, and we add this cell to the search wavefront at cost=4, reached from the north.



Expand this cell...

Reach this cell at cost = 1 + 3 = 4...

© R. Rutenbar, CMU 18-760, Spring99 45

## Maze Routing: Mid-Point Summary

### ■ What do we know?

- ▼ Grid-based expansion, one net at a time
- ▼ Can use costs in grid to get different effects
- ▼ Can deal with multiple wiring layers, multi-point nets

### ■ What don't you know?

- ▼ Real implementation issues
- ▼ Data structures for grid, for the search
- ▼ Depth-first expansion techniques for speed
- ▼ Subtle interactions between cost strategy and search strategy
- ▼ Expanding a cell vs reaching a cell vs multiple-reaching....

### ■ Next topics: *serious implementation issues.*

© R. Rutenbar, CMU 18-760, Spring99 46

## Implementation Concerns

### ■ Representation

- ▼ How do we store the routing grid?
- ▼ What do we need in each cell?
- ▼ How do we represent the state of the advancing path search process?

### ■ Algorithms

- ▼ We have a serial computer: we can process one cell at a time...
- ▼ ...so, which cell is next to “label” in the search process
- ▼ Does the order matter?
- ▼ How can we do this as fast as possible?

© R. Rutenbar, CMU 18-760, Spring99 47

## Big Idea: Search Wavefront

### ■ One big goal

- ▼ Efficient storage: big layout needs a big grid
- ▼ Want to put as little in each cell as possible

### ■ Idea: Don't actually store path costs in grid cells

- ▼ Big costs -> many bits per cell.
- ▼ Only the cells most recently labeled during search will be used to expand the search for new paths
- ▼ These cells constitute the **search wavefront**

### ■ Wavefront is important:

- ▼ Store wavefront list with all needed info about each wavefront cell.
- ▼ Mark a few bits in the actual grid cells just to indicate how you found the path to this cell - i.e., remember the predecessor cell.

© R. Rutenbar, CMU 18-760, Spring99 48



## Example Wavefront for Simple Search

	4	3	4		
4	3	2	3	4	
3	2	S	2	3	4
4	3	2	3	4	
	4	3	4	T	
		4			

### ■ Wavefront is...

- ▼ The frontier of the active search for new paths
- ▼ The neighbors of the new cells worth looking at to try to extend the evolving path search
- ▼ The only cells we need to look at to decide how to continue the search process

### ■ Implication

- ▼ Don't store the path cost numbers *in* the grid
- ▼ Just store the wavefront cells themselves in a special data structure

© R. Rutenbar, CMU 18-760, Spring99 49

## More Complex Wavefront

After expanding 1st S cell

		2			
	2	S	4		
		2			
				T	

After expanding 3 neighbors of S

		3			
	3	2	3		
3	2	S	4		
	3	2	5		
		3		T	

© R. Rutenbar, CMU 18-760, Spring99 50

## More Complex Wavefront

	4	3	4		
4	3	2	3	4	
3	2	S	4		
4	3	2	5		
	4	3	4	T	
		4			

### ■ What wavefront is...

- ▼ Set of cells already reached in the expansion process...
- ▼ ...which have neighbors we have not already reached
- ▼ Indexed by cost of cells reached (== costs of paths that start at source and end at this cell)
- ▼ Expanded in cost order, cheapest cells before more expensive cells

© R. Rutenbar, CMU 18-760, Spring99 51

## Outline of Expansion Algorithm

### ■ Cheapest-cell-first search

- ▼ Variant of Dijkstra's algorithm
- ▼ Assume wavefront is a cost-indexed list of cells you have already visited during search process, and "labeled" with path cost

### ■ How does the wavefront grow?

- ▼ Pull out a cheapest cell C from the wavefront
- ▼ Look at the neighbors N1, N2, ... of cell C you have not visited yet
- ▼ Compute the cost of expanding this path to reach these new cells N1...
- ▼ Add these new cells N1, N2, ... to the wavefront data structure (indexed by their cost)
- ▼ Remove cell C from the wavefront
- ▼ Repeat with the next cheapest cell on wavefront...

© R. Rutenbar, CMU 18-760, Spring99 52

## Maze Router: Terminology

■ We need some terminology or we'll get confused

■ Wavefront


■ Reached

■ Expanded

■ Dijkstra's Approach

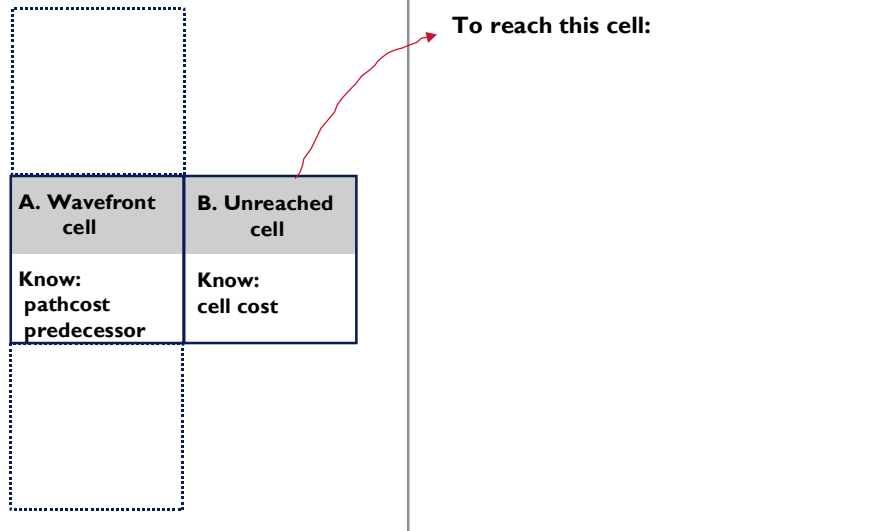
© R. Rutenbar, CMU 18-760, Spring99 53

## Illustrating the Terminology



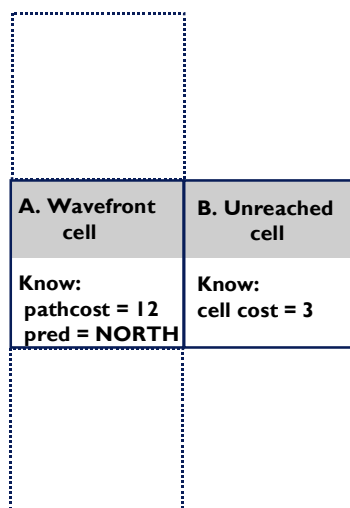
© R. Rutenbar, CMU 18-760, Spring99 54

## Reaching a New Cell During Search



© R. Rutenbar, CMU 18-760, Spring99 55

## Reaching A New Cell



### ■ To reach cell B

- ▼ Grab cell A, pathcost=12, from wavefront
- ▼ See unreached neighbor B
- ▼ Compute cost to reach B is  $12 + 3 = 15$
- ▼ Add this cell to wavefront



- ▼ Mark grid cell B as “reached” (only takes a few bits) so we don’t try to put it on wavefront again (ie, reach it again)

© R. Rutenbar, CMU 18-760, Spring99 56

## Basic Maze Routing Algorithm

```
wavefront_structure = { source cell }
while (we have not hit target) {
  if ( wavefront == empty )
    quit -- no path to be found

  C = get lowest cost cell on wavefront_structure
  if ( C == target ) {
    backtrack path in grid
    cleanup
    return -- we found a path
  }
  foreach ( unreached neighbor N of cell C ) {
    mark N cell in grid as reached
    compute cost to reach it = pathcost of C + cellcost of N
    mark N cell in grid with predecessor direction back to cell C from N
    add this cell N to wavefront
  }
  delete cell C from wavefront
}
```

© R. Rutenbar, CMU 18-760, Spring99 57

## Data Structure Issues


### ■ 2 key structures

#### ■ Routing grid

- ▼ Hold cells of area to route, costs of each cell, blockages
- ▼ Mark these cells to know what cells you have already reached
- ▼ Mark predecessor in here too

#### ■ Wavefront


- ▼ Hold active cells to expand
- ▼ Cell info has pathcost, predecessor information
- ▼ Indexed on pathcost
- ▼ Always expand cheapest cell (ie, cheapest partial path) next



© R. Rutenbar, CMU 18-760, Spring99 58

## Data Structure Implementation

### Grid of cells:




Each grid cell stores:

```
int cellCost
char predecessor
↕
↔
↕
boolean reached
```

A 2-D array is just fine here

### Wavefront list:



Each wavefront cell stores:

```
grid *whichCell
int pathCost
```


Need something clever here-- want fast insert/delete in cost here. Dumb linked list isn't going to be fast enough...

© R. Rutenbar, CMU 18-760, Spring99 59

## Wavefront Option: Heap

### Store cells of wavefront in a heap

- ▼ (or a priority queue -- consult your favorite data structures book)
- ▼ Classical data structure designed for fast insertion and retrieval of lowest cost data item.
- ▼ All operations (add, delete, etc.) have  $O(\lg N)$  time complexity for  $N$  objects.
- ▼ Most routers do it like this.



min cost item is always at the top.

Insert "bubbles" the items in heap around to ensure the cheapest item always on top.

Ditto for delete.

© R. Rutenbar, CMU 18-760, Spring99 60


## Wavefront Option: Cost-Indexed Array

- Basically a big hash table

- ▼ Index is path cost
- ▼ In each bin of table we insert the cells in wavefront at that cost
- ▼ Get almost constant-time insert/delete

- Problem

- ▼ It's a BIG table -- 1 bucket for each possible path cost.



© R. Rutenbar, CMU 18-760, Spring99 61

## Remember How Expand Algorithm Works...

	4	3	4		
4	3	2	3	4	
3	2	S	4		
4	3	2	5		
	4	3	4	T	
		4			

- Always expand next cheapest cell

- ▼ Implication is you cannot have an arbitrarily wide "spread" in the values of "live" cells on wavefront
- ▼ If CMAX is the maximum cell cost value you can see in the grid
- ▼ ...then the biggest spread of cost values you can see in the wavefront is CMAX + 1

- In this example


- ▼ Biggest spread in costs you can see is 4, since biggest cell has cost=3,  
 $3 + 1 = 4$  different cost values possible in wavefront at any time

© R. Rutenbar, CMU 18-760, Spring99 62

## Remember How Expand Algorithm Works

■ Said another way...

Cost bins --->



The worst thing that can ever happen is you expand a “cheapest” cell down here to reach a cell of cost = CMAX...

Ex: expand S cost = 1

...and reinsert it up here at cost s+CMAX

Ex: reach west neighbor at pathcost=4

© R. Rutenbar, CMU 18-760, Spring99 63


## Wavefront Implementation

■ So, you can do this with only CMAX+1 bins in array

▼ It's a **circular** array: when you empty the lowest cost bin, it means you can reuse it for the next MAX pathcost value you need.

New Cmin is previous Cmin + 1 list.

Old Cmin goes to the end of the array.



© R. Rutenbar, CMU 18-760, Spring99 64



## Wavefront Data Structures

- **Historically, saw both kinds of structures in real routers**
  - ▼ Hashed cost array
  - ▼ Heap-based
  - ▼ Today, heaps seem to be dominating; just a lot more flexible on what they allow you to do with costs
- **But, there are still some subtle interactions with the search algorithm to discuss**
- **Question**
  - ▼ What constraints have to met for this simple expand-cheapest-cell-next strategy to get the best path?

© R. Rutenbar, CMU 18-760, Spring99 65

## Plain Maze Routing Revisited

- **Key assumptions**
  - ▼ Always expand cheapest cell next
  - ▼ Expand each cell just once
  - ▼ Reach each cell just once
  - ▼ Guaranteed to find the min cost path (we hope...)
- **Question**
  - ▼ What are the constraints on the cost function used for paths (ie, pathcost of a cell as it is reached) so that above stuff holds...?

© R. Rutenbar, CMU 18-760, Spring99 66

## Pathcost Constraints

- **Basic constraint: *Consistency***

- ▼ The cost of adding a cell to a path (reaching a cell) is independent of the path itself
- ▼ It does **NOT** matter how you reached this new cell, it still adds the same cost to the path
- ▼ Guarantees we reach it once (from a cheapest path) and thus expand it just once

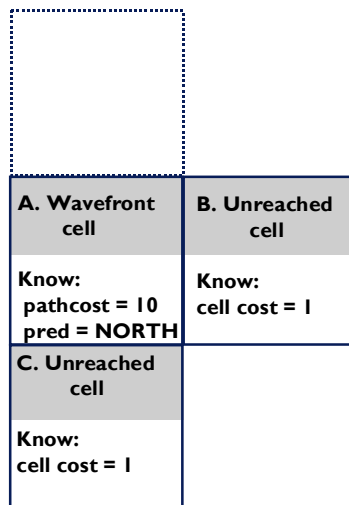
- **It's actually easy to create a cost function that is *inconsistent*, and violates all these nice properties**

© R. Rutenbar, CMU 18-760, Spring99 67

## Inconsistent Cost Function

- **Penalize paths with bends**

- ▼ Still store 1 cost inside each cell
- ▼ But, now add another cost when you reach a cell that requires a turn (a bend) from the direction that reached the expanding cell



Suppose bend penalty = 2

© R. Rutenbar, CMU 18-760, Spring99 68

## Inconsistent Cost Function


■ Try this example with bend penalty = 2

- ▼ Don't mark the "reached" bit in each grid cell when you reach the cell
- ▼ Allow search to revisit previously reached cells...

SOURCE		
cost=1	cost=1	cost=3
cost=2	cost=1	TARGET
		cost=1

© R. Rutenbar, CMU 18-760, Spring99 69

## Inconsistent Cost Example



© R. Rutenbar, CMU 18-760, Spring99 70

## Inconsistent Cost Function: Implications

### ■ Notice what happened

- ▼ Reached same cell, later, at a higher cost, but it was ultimately on the cheaper overall source-to-target path

### ■ Implications

- ▼ You will reach cells multiple times at different costs.
- ▼ You will have same cell in wavefront multiple times at different costs.
- ▼ Cannot guarantee you need only  $C_{MAX}+1$  hash bins in array
- ▼ Can still expand cheapest first, but cannot quit when you reach target

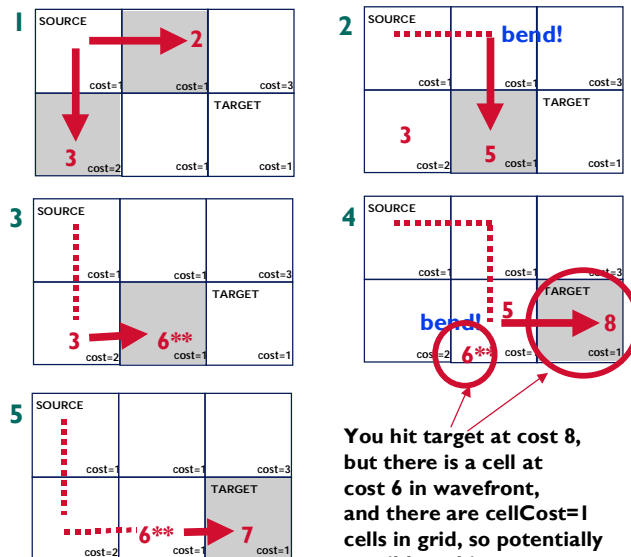
### ■ Termination of search?

- ▼ Cannot quit until each cell in wavefront has a cost so big that it is **NOT POSSIBLE** to reach target any cheaper than current cheapest path
- ▼ May reach, expand lot more cells with an inconsistent cost function...
- ▼ ..but you can do a lot of cool things with such functions

© R. Rutenbar, CMU 18-760, Spring99 71

## Termination of Search

- Cannot quit until no cell in the wavefront has a cost that could lead to a cheaper path to target



© R. Rutenbar, CMU 18-760, Spring99 72

## Expansion Process, Revisited

### ■ Problem:

- ▼ Expand lots of cells to find one path to the target.
- ▼ CPU time is proportional to # of cells you search.
- ▼ No attempt to search in *direction of target* first.

### ■ Questions:

- ▼ How do you search **toward** the target?
- ▼ Can we do this and still keep guarantees of reaching the target with the minimum cost path?

© R. Rutenbar, CMU 18-760, Spring99 73

## Motivation for Smart Search

	4	3	4		
4	3	2	3	4	
3	2	S	2	3	4
4	3	2	3	4	
	4	3	4	T	
		4			

Expanding away from the target seems to be a waste of time.

Searching toward the target in the shaded region.

© R. Rutenbar, CMU 18-760, Spring99 74

## Smarter Search: Rubin's Scheme

### ■ Two parts:

- ▼ Add predictor function to the cost.
- ▼ Direct the search toward the target

### ■ Plain maze router


- ▼ You add a cell to the wavefront with a cost that measures partial cost of the path, source-to-target

### ■ Rubin's Scheme

- ▼ You add a cell to the wavefront with a cost that estimates the entire source-to-target cost of the path
- ▼ Trick: estimate this as  $\text{pathcost}(\text{source to cell}) + \text{predictor}(\text{cell to target})$
- ▼ (We will see this exact same idea again, when we do Static Timing analysis; this predictor will be called the **ESPERANCE** of a path...)

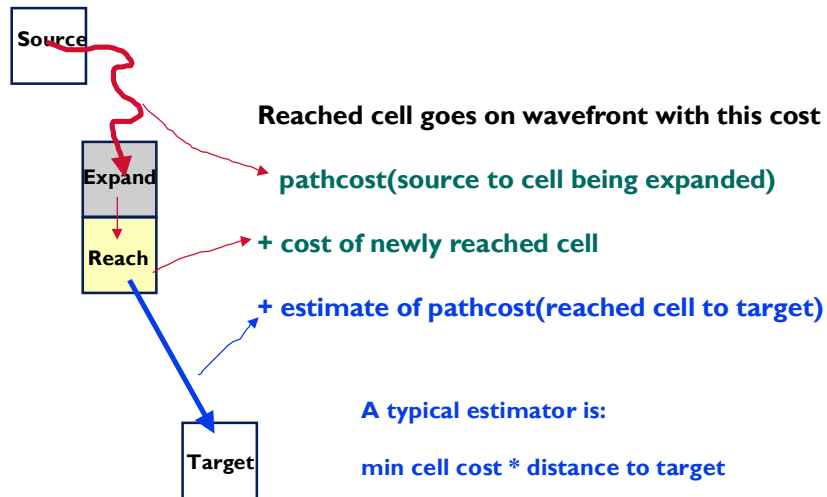
© R. Rutenbar, CMU 18-760, Spring99 75

## Plain Maze Routing



© R. Rutenbar, CMU 18-760, Spring99 76

## Add A Depth-First Predictor



© R. Rutenbar, CMU 18-760, Spring99 77

## Technical Results

### ■ Depth first predictor

- ▼ If the predictor is always a lower bound on how much pathcost you will really add to get to the target...
- ▼ ...you will still get the min cost path, guaranteed


### ■ What does it do?

- ▼ It alters the **order** in which we expand cells
- ▼ It prefers to expand cells that are closer to the target first
- ▼ It does this in a very geometrically stylized way

### ■ Look at an example...

© R. Rutenbar, CMU 18-760, Spring99 78

## Rubin Expansion Example



### Observe

▼ It prefers to stay inside the bounding box of the source-target rectangle before it expands other cells.


▼ How do we know which cell to expand in order inside the box?



▼ Several heuristics which all basically say: don't turn unless you have to, and prefer to expand the cells that are actually closest to target, first

© R. Rutenbar, CMU 18-760, Spring99 79

## Some Subtleties



### Will again reach a cell multiple times with different costs

▼ Suppose you really expand cheapest first, and among those of same cost, closest to target

▼ Can shoot directly toward target...

▼ ...but you reach cells early with suboptimal costs


▼ Can reach time again later

You will first reach this cell as NORTH neighbor of cell x, but cost will be big because your path goes thru cell x, which is bad path here

© R. Rutenbar, CMU 18-760, Spring99 80



## Rubin Expansion Example

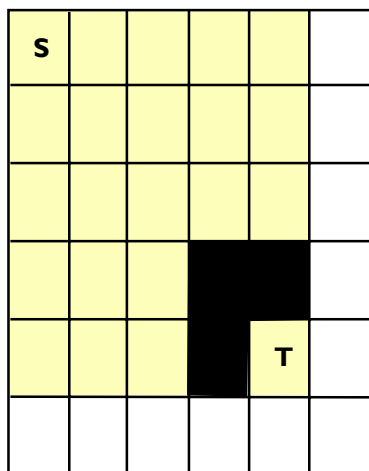


### ■ Works great...

- ▼ Until you get a case like this with the target blocked inside the source to target rectangle.
- ▼ Problem now is that it explores the whole rectangle before it tries anyplace else.
- ▼ Might it not be faster to search outside the rectangle if the rectangle is VERY big...?

© R. Rutenbar, CMU 18-760, Spring99 81

## Another Tweak



### ■ Can insist that cells closer to target *always* be cheaper

- ▼ Add cell to wavefront as
  - $\text{pathcost}(\text{src} \rightarrow \text{cell})$
  - $+\text{cost of cell}$
  - $+ K \cdot \text{estimated cost}(\text{cell} \rightarrow \text{target})$

### ■ K is just a fudge factor

- ▼ Forces cells closer to target to be cheaper.
- ▼ Typically small (like 1.1)
- ▼ Try  $K=2$  in this example for effect

### ■ Effects

- ▼ Faster search, smaller search, but lose guarantees of minimum soln

© R. Rutenbar, CMU 18-760, Spring99 82

## Area Routing By Maze Routing: Summary

### ■ Been around a LONG time

- ▼ Very flexible cost-based search
- ▼ Extremely flexible, can be recast to attack many problems
- ▼ Zillions of tweaks for speed, space, etc.
- ▼ Still widely used, but now often with rather more sophisticated representations of “space” than a 2D grid.

### ■ Remaining problems

- ▼ Still routes one net at a time. Early nets block later nets.
- ▼ Lots of iterative improvement strategies here (I didn't talk about)
- ▼ Great if there IS a path; if not, will spend a long time to prove to you that there is NO path

© R. Rutenbar, CMU 18-760, Spring99 83