A Hybrid Constraint Model for the Routing and Wavelength Assignment Problem

Helmut Simonis

Cork Constraint Computation Centre
Computer Science Department
University College Cork
Ireland

CP 2009

science foundation ireland
fondúireacht eolaíochta éireann

Cork Constraint Computation Centre

Main Points

- Well studied problem in network design
  - Survey gives over 60 references on this problem
- Hybridization of MIP and CP
- Uses explanation to find good relaxation
- More than 1000 times faster than best MIP
- Optimal for 99.8% of cases tested

Outline

1. Problem
2. Model
3. Results

Problem Definition

Routing and Wavelength Assignment (Demand Acceptance)

In an optical network, traffic demands between nodes are assigned to a route through the network and a specific wavelength. The route (called *lightpath*) must be a simple path from source to destination. Demands which are routed over the same link must be allocated to different wavelengths, but wavelengths may be reused for demands which do not meet. The objective is to find a combined routing and wavelength assignment which maximizes the number of accepted demands.

Problem Definition

Routing and Wavelength Assignment (Demand Acceptance)

In an optical network, traffic demands between nodes are assigned to a route through the network and a specific wavelength. The route (called *lightpath*) must be a simple path from source to destination. Demands which are routed over the same link must be allocated to different wavelengths, but wavelengths may be reused for demands which do not meet. The objective is to find a combined routing and wavelength assignment which maximizes the number of accepted demands.

This only one of the problem variants, but probably the most difficult

Example Network (NSF, 5 wavelengths)

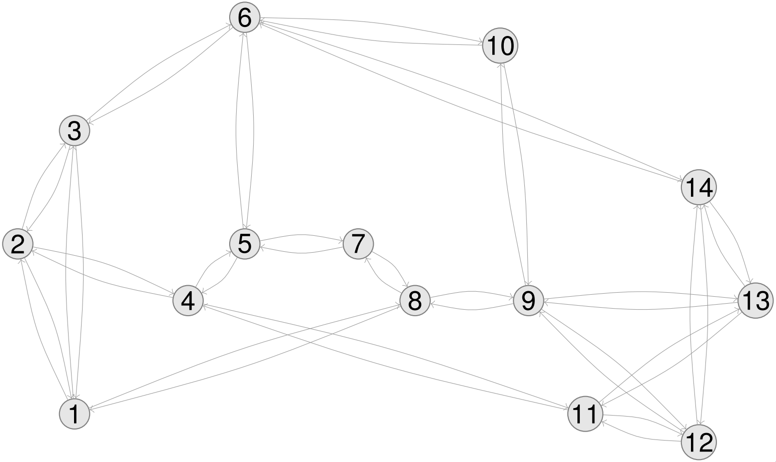

Lightpath from node 5 to node 13 ($5 \Rightarrow 13$)

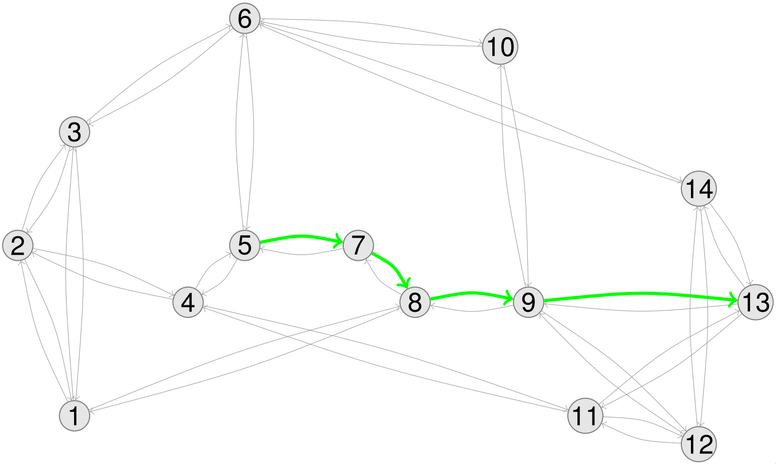

Conflict with demand 1 $\Rightarrow$ 12: Use different frequencies

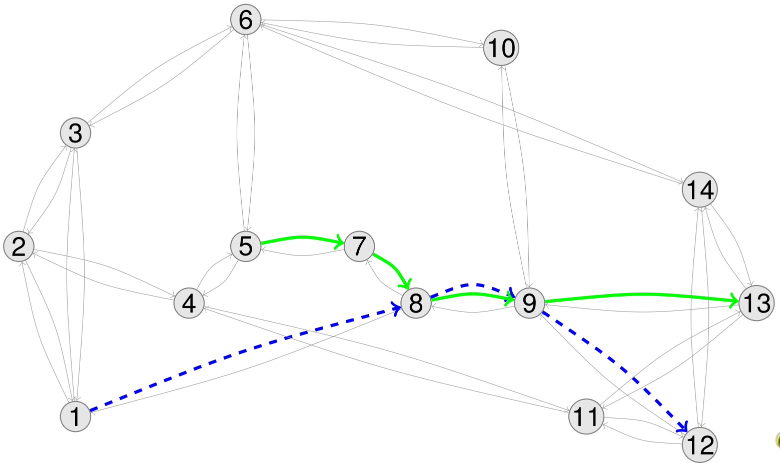

Conflict with demand 1 ⇒ 12: Use different path

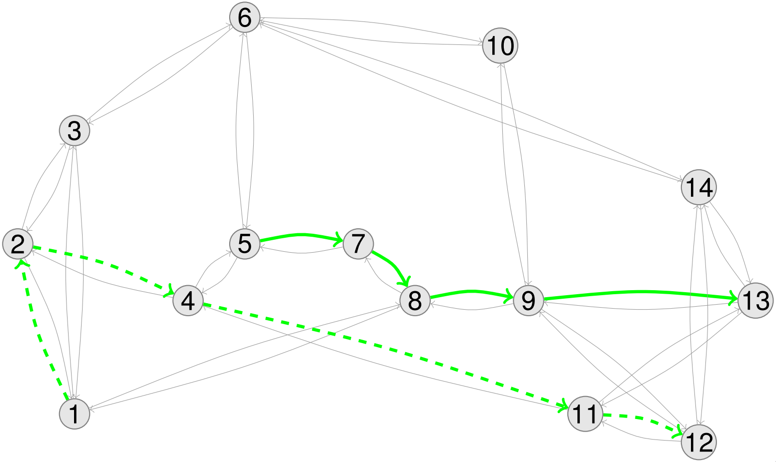

Conflict with demand 1 ⇒ 12: Reject demand

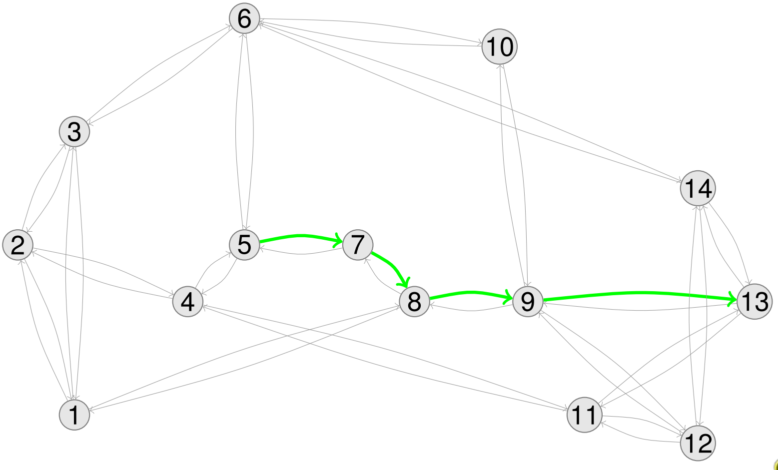

Demand Matrix (100 Demands)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1		1	1		1	1	1	1	1					
2	1		1						1			2		
3					2	2		1	1		1	1		1
4			1		3		1	1	1					
5		1				1				1		1		
6				2	2		1			1	1	1		
7				1		1		1		1	1	2		
8											1		1	
9		1	1		1						1	1		
10			2			1	1	1	1		1		1	2
11	1		1	1	1		2	1		2			1	
12		2				2	1	2			1		2	2
13			1		2						1	1		
14	1		1		1	1			3		1		1	

Color	Distance
Blue	1
Green	2
Pink	3
Yellow	4
Red	≥ 5

Outline

Early References on Problem

- Rajiv Ramaswami and Kumar N. Sivarajan.
  Routing and wavelength assignment in all-optical networks.
  IEEE/ACM Trans. Netw., 3(5):489–500, 1995.
- Dhritiman Banerjee and Biswanath Mukherjee.
  A practical approach for routing and wavelength assignment in large wavelength-routed optical networks.
  IEEE Journal on Selected Areas in Communications, 14(5):903–908, June 1996.

Solution Approaches

- Greedy heuristics
- Optimization algorithm for complete problem

Solution Approaches

- Greedy heuristics
- Optimization algorithm for complete problem

Discussion of MIP Models

Brigitte Jaumard, Christophe Meyer, and Babacar Thiongane.
ILP formulations for the routing and wavelength assignment problem: Symmetric systems.
In M. Resende and P. Pardalos, editors, *Handbook of Optimization in Telecommunications*, pages 637–677. Springer, 2006.

Brigitte Jaumard, Christophe Meyer, and Babacar Thiongane.
Comparison of ILP formulations for the RWA problem.
Optical Switching and Networking, 4(3-4):157–172, 2007.

MIP Model Ideas

- Naive model
  - Decide whether/how to route each demand
  - 0/1 variables $Y_d^\lambda, X_{de}^\lambda$
  - Problem: Full symmetries
- Source aggregation
  - Decide how to route all demands from common source
  - Integer variables Y_{sd}, 0/1 variables X_{se}^λ
  - Removes many symmetries, but not all

Our Solution Approach

- Greedy heuristics
- Optimization algorithm for complete problem
- Decomposition into two problems
  - Route maximal number of demands
  - Assign wavelengths

Step 1: Route Maximal Number of Demands

- Ignore wavelengths
- Capacity constraints on all links
- Solve as MIP problem
- Source aggregation
- Find DAG to supply (all) demands with shared source
- Maximize number of accepted demands

Notation

- y_{sd}, integer number of accepted demands from s to d
- z_{se}, integer capacity used on edge e to satisfy demands sourced in s
- C, number of available wavelengths, edge capacity
- P_{sd}, requested number of demands from s to d
- T_s, total number of requested demands sourced from s
- D_s, nodes which have a requested demand sourced in s

Model (Step 1)

$$\max \sum_{s \in N} \sum_{d \in D_s} y_{sd}$$

s.t.

$$y_{sd} \in \{0, 1 ... P_{sd}\}, z_{se} \in \{0, 1 ... T_s\}$$

$$\forall e \in E: \quad \sum_{s \in N} z_{se} \leq C$$

$$\forall s \in N: \quad \sum_{e \in \mathrm{In}(s)} z_{se} = 0$$

$$\forall s \in N, \forall d \in D_s: \quad \sum_{e \in \mathrm{In}(d)} z_{se} = \sum_{e \in \mathrm{Out}(d)} z_{se} + y_{sd}$$

$$\forall s \in N, \forall n \neq s, n \notin D_s: \quad \sum_{e \in \mathrm{In}(n)} z_{se} = \sum_{e \in \mathrm{Out}(n)} z_{se}$$

Observation

- Optimal cost is upper bound for full problem
- LP Relaxation is also upper bound for full problem
- No 0/1 variables in model
- Source aggregation has massive impact on efficiency
  - Much better than treating each demand on its own
  - Reason 1: Reduced number of variables
  - Reason 2: Avoids symmetries due to multiple demands between nodes

Finding Accepted Demands

- Solution to MIP does not tell how demands are routed
- Program required to convert source “tree” into sets of paths
- Conversion not deterministic, may allow different solutions
- Solution may contain loops, these need to be removed

Source Model Solution

Color	Type
	Source
	Sink
	Unreached
	Chosen Link

Source Node 10

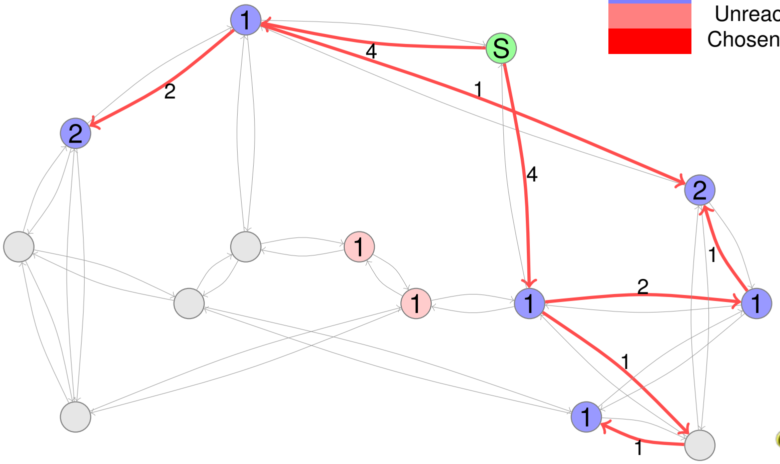

Comparison

Demand Matrix

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1		1	1		1	1	1	1	1					
2	1		1						1			2		
3					2	2		1	1		1	1		1
4			1		3		1	1	1					
5		1				1				1		1		
6				2	2		1			1	1	1		
7				1		1		1		1	1	2		
8											1		1	
9		1	1		1						1	1		
10			2			1	1	1	1		1		1	2
11	1		1	1	1		2	1		2			1	
12		2				2	1	2			1		2	2
13			1		2						1	1		
14	1		1		1	1			3		1		1	

Accepted Demands

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1		1	1		1	1	1	1	0/1					
2	1		1						0/1			2		
3					1/2	2		1	1		1	1		1
4			1		3		1	0/1	1					
5		1				1				1		1		
6				2	2		1			1	1	1		
7				1		1		1		1	1	2		
8											1		1	
9		1	1		0/1						1	1		
10			2			1	0/1	0/1	1		1		1	2
11	1		1	1	0/1		0/2	1		2			1	
12		1/2				2	1	2			1		2	2
13			1		0/2						1	1		
14	1		1		0/1	1			3		1		1	

Observations

- Accepted demands do not always use shortest path
- Tendency to reject demands with larger minimal distance
- These use more resources
- Not compensated in objective function
- Not fair

Resource Requirements

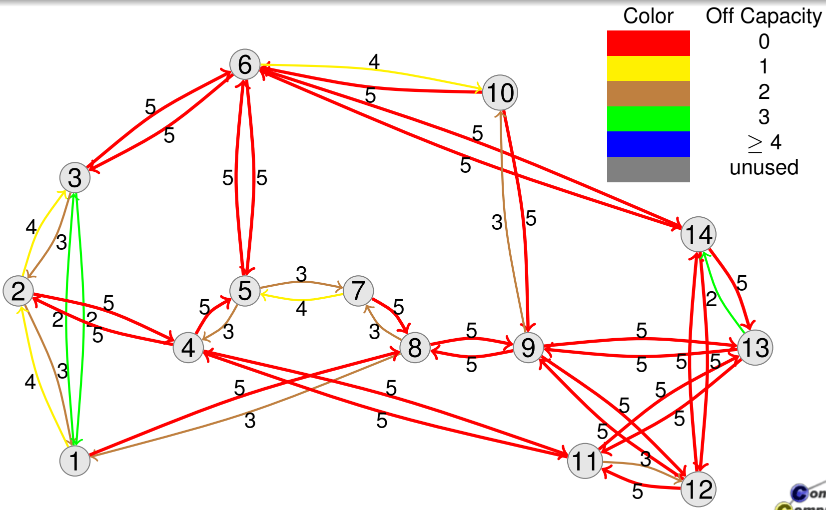

Step 2: Assign Wavelengths

- For each accepted demand, find frequency
- All demands routed over a link compete for frequencies
- Graph coloring problem
- Graph given as sets of cliques
- Solve with finite domains
- If solution found, then optimal for complete problem

Graph Coloring Problem

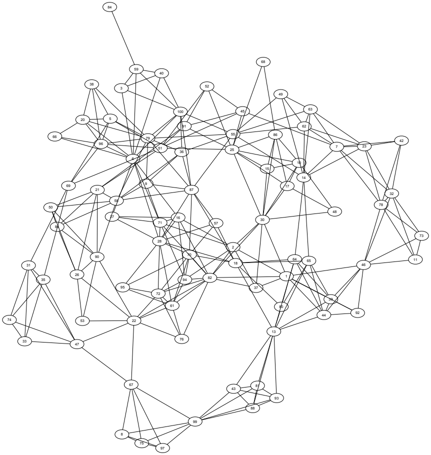

Model (Step 2)

- X_d finite domain variable $1..C$ for each accepted demand
- One `alldifferent` constraint for each edge
- Many `alldifferent` constraints are at capacity
- Possible to improve model

Graph Coloring Solution

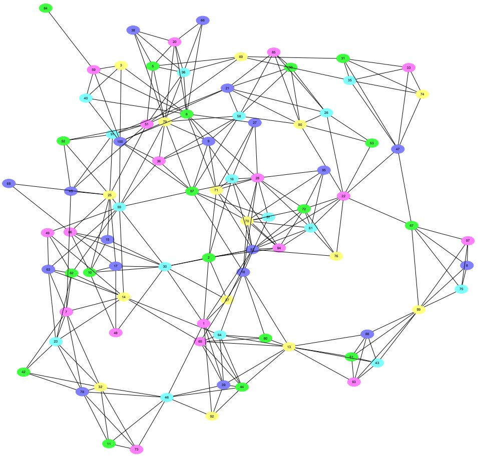

Color	Wavelength
	1
	2
	3
	4
	5

Observation

- All demands could be assigned to frequencies
- Optimal solution to complete problem

What Happens If No Solution Found

- Problem infeasible
  - Remove some demand and try again until solution found
  - Possibly sub-optimal solution of high quality
  - Different solution to MIP problem may lead to optimal solution
- No solution found within time limit
  - Try harder!
  - Improve reasoning and/or search technique
  - Special techniques to show infeasibility

Explaining Infeasibility

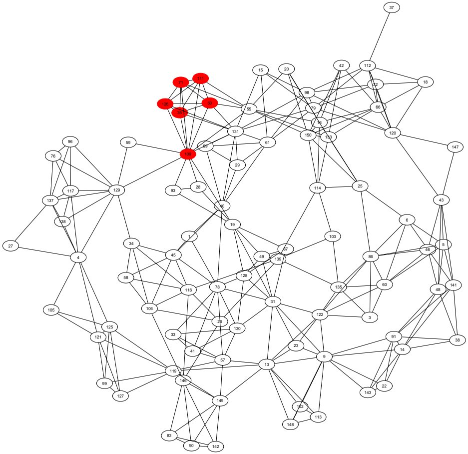

Solution Approach

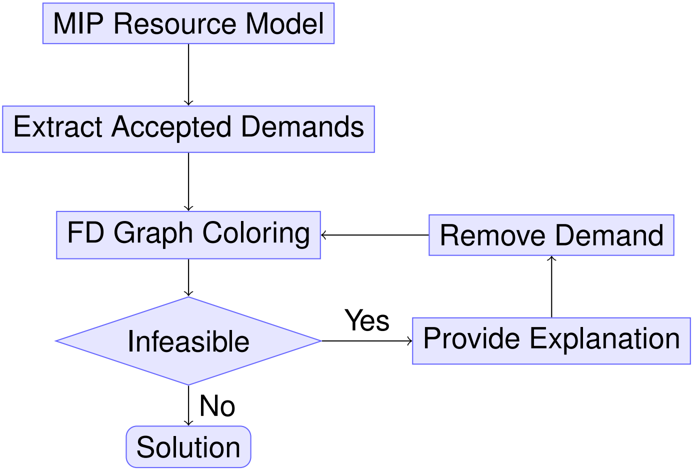

Explanations

- Ad-hoc: Find pattern which show infeasibility
  - Find large cliques
  - If clique is larger than number of colors, problem is infeasible
  - This is simple for graphs given
- General explanation techniques
  - Find subset of variables/constraints which is infeasible

Explanation Method Used: QuickXPlain (Junker 2001)

- Find minimal subset of constraints which is infeasible
- *Conflict set*
- Works when overall problem fails without search
- But our problem does not fail initially
- Requires some technical trick to use QuickXPlain

Outline

Tools Used

- ECLiPSe 6.0
- Coin-OR CLP/CBC MIP solver via eplex
- Developed as part of ECLiPSe ELearning Course
  - Gift grant from Cisco Systems / Silicon Valley Community Foundation
  - `http://4c.ucc.ie/~hsimonis/ELearning/index.htm`

Benchmarks

- Fixed network structure

nsf	14 nodes, 42 edges
eon	20 nodes, 78 edges
mci	19 nodes, 64 edges
brezil	27 nodes, 140 edges

- Random network structure
  - Sizes from 30 to 100 nodes
  - Edge density 0.25
  - 500 demands, 30 wavelengths

What about Search?

- Multiple phases
- Start with incomplete, credit based search
  - Increasing number of credit a,a^2,a^3
- Complete search with timeout
  - First try with BC `alldifferent`
  - Then use GAC `alldifferent`
- We don't rely on search to prove infeasibility

Overall Distribution of Solutions

Type	Technique	Count
Infeasible	clique	50
	preassign	38
Feasible	credit total	59962
	of that, credit a units	58861
	of that, credit a^2 units	940
	of that, credit a^3 units	161
	complete search, BC alldifferent	25
	complete search, GAC alldifferent	12

Compared to MIP Model for Complete Problem

Network	Dem.	λ	Opt.	Hybrid Model Avg FD	Avg Time	Max Time	Full MIP Avg Opt	Avg Time	Max Time
brezil	500	15	98	483.84	0.92	1.34	483.86	1218.40	14103.84
brezil	600	20	100	590.96	1.00	1.34	590.96	6076.81	87767.95
brezil	700	25	98	695.48	1.01	1.80	*695.48*	*13623.15*	*78463.89*
brezil	800	25	99	781.37	1.44	11.47	781.39	7567.68	15456.50
eon	500	20	100	471.56	0.65	0.77	471.56	352.21	585.45
eon	600	25	100	574.80	0.82	1.13	574.80	1411.67	2877.88
eon	700	30	100	677.35	1.05	1.81	677.35	1727.52	3568.13
eon	800	35	100	779.17	1.28	1.94	779.17	2485.64	4116.11
mci	500	25	100	486.38	0.80	2.28	486.38	1023.16	1664.31
mci	600	30	100	585.18	1.27	29.81	585.18	1621.30	2895.88
mci	700	35	100	684.00	1.30	3.53	684.00	1987.23	3428.41
mci	800	40	100	782.86	1.68	5.21	782.86	2316.88	4402.44
nsf	500	35	100	495.20	0.50	0.60	495.20	82.85	173.19
nsf	600	40	100	588.63	0.66	0.98	588.63	155.90	373.63
nsf	700	45	100	678.44	0.86	1.35	678.44	205.82	586.61
nsf	800	45	100	727.15	0.95	1.56	727.15	173.53	410.97

Increasing Number of Demands

Network	Dem.	λ	Opt.	Avg LP	Avg MIP	Avg FD	Max Gap	Avg MIP Time	Max MIP Time	Avg FD Time	Max FD Time
eon	800	30	100	741.78	741.78	741.78	0.00	0.15	0.17	0.83	1.61
eon	900	40	100	880.59	880.59	880.59	0.00	0.14	0.16	1.18	2.17
eon	1000	40	100	950.36	950.36	950.36	0.00	0.15	0.17	1.37	3.42
eon	1100	50	100	1082.61	1082.61	1082.61	0.00	0.14	0.16	1.71	2.83
eon	1200	50	100	1156.38	1156.38	1156.38	0.00	0.15	0.17	2.07	5.92
eon	1300	50	100	1219.82	1219.82	1219.82	0.00	0.16	0.17	2.22	5.24
eon	1400	60	100	1361.47	1361.47	1361.47	0.00	0.15	0.16	2.92	4.94
eon	1500	60	99	1428.78	1428.78	1428.77	1.00	0.15	0.17	4.22	106.97
eon	1600	70	100	1565.90	1565.90	1565.90	0.00	0.15	0.16	3.89	8.48
eon	1700	70	100	1637.47	1637.47	1637.47	0.00	0.16	0.17	4.58	13.59
eon	1800	80	100	1769.86	1769.86	1769.86	0.00	0.15	0.16	5.19	8.81
eon	1900	80	99	1844.46	1844.46	1844.45	1.00	0.15	0.17	7.23	163.41
eon	2000	90	100	1972.66	1972.66	1972.66	0.00	0.15	0.17	6.34	9.61

Random Networks (Edge Density 0.25, 100 Runs Each)

Network	Dem.	λ	Opt.	Avg LP	Avg MIP	Avg FD	Avg MIP Time	Max MIP Time	Avg FD Time	Max FD Time
r30	500	30	100	391.82	391.82	391.82	0.45	0.55	0.12	0.16
r40	500	30	100	424.58	424.58	424.58	1.07	1.23	0.14	0.17
r50	500	30	100	437.69	437.69	437.69	2.13	2.38	0.09	0.13
r60	500	30	100	447.21	447.21	447.21	3.92	4.34	0.08	0.16
r70	500	30	100	453.41	453.41	453.41	6.78	7.50	0.10	0.17
r80	500	30	100	457.65	457.65	457.65	10.75	11.95	0.10	0.17
r90	500	30	100	464.69	464.69	464.69	16.08	17.45	0.08	0.22
r100	500	30	100	466.67	466.67	466.67	22.74	25.22	0.09	0.25

Observations

- MIP and LP relaxation of phase 1 are very good bounds
- Solved to optimality in most cases
- Simple decomposition quite effective
- Good solution even if initial graph coloring infeasible
- Special structure of graph coloring helps FD model

Why a Hybrid Model?

- Not possible with MIP alone
  - Monolithic MIP does scale poorly
  - Decomposition trick does not help MIP
- Not possible with CP alone
  - CP has nothing to estimate number of demands that can be accepted

In Defense of Constraint *Languages*

- Success story for Constraint *Programming*
- Programming required, modelling languages are not enough
  - MIP solution extraction
  - Explanation generation
  - Variable selection strategy

A Simpler Problem Variant

Helmut Simonis.
Solving the static design routing and wavelength assignment problem.
CSCLP 2009, Barcelona, Spain, June 2009.

- Decomposition still works well
- But competition from specialized graph coloring codes

Conclusions

- Combination of MIP and FD solver in problem decomposition
- Each doing what they do best
  - MIP: optimal solution, select items to include
  - FD: find feasible solution, explain infeasibility
- Hybrid model produces very high quality results
- Proven optimality in over 99.85% of problems tested
- Near optimal solutions by relaxation
- Much ($> 10^3$) faster than monolithic MIP solution