

Localizing & Bolstering Constraint Propagation in a Tree Decomposition

Shant Karakashian, Robert Woodward & Berthe Y. Choueiry



1. Paper & poster at AAAI 2013
2. PhD thesis of Shant Karakashian, ConSystLab, May 2013

Acknowledgments:

- Experiments conducted at UNL's Holland Computing Center
- NSF Award RI-111795

Outline

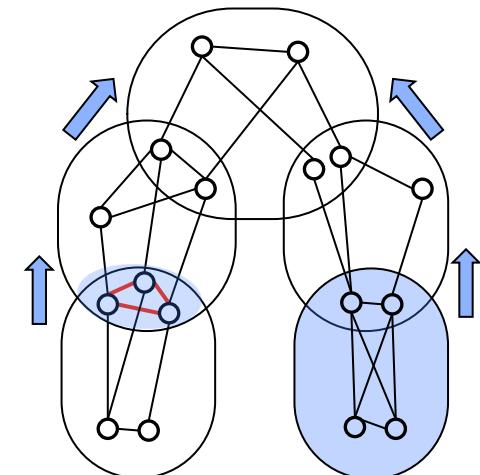
- Introduction & background information
- Key ideas
 - Localize consistency to clusters of a tree decomposition
 - Bolster propagation at separators
- Evaluation
 - Theoretical: Comparing resulting consistency properties
 - Empirical: Solving CSPs in a backtrack-free manner
- Conclusions & future work
 - ✧ Robert's RNIC on binary CSPs
 - ✧ Summary of contributions of Shant's thesis

Introduction

- One tractability condition links [Freuder 82]
 - Consistency level to
 - Width of the constraint network, a structural parameter
- Catch 22
 - Enforcing higher-levels of consistency may require adding constraints, which increases the width
 - Computing treewidth (\equiv induced width) is in NP-hard
- Goal
 - Exploit above condition to achieve practical tractability

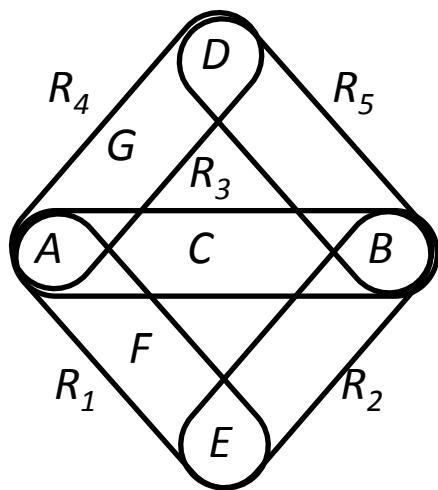
Our Approach

- Robert discussed consistency properties that
 - Preserve the structure of a problem instance
 - Enforce a (parameterized) consistency level
- Exploit a tree decomposition
 - **Localize** application of the consistency algorithm
 - **Steer** constraint propagation along the branches of the tree
 - **Add** redundant constraints at separators to enhance propagation



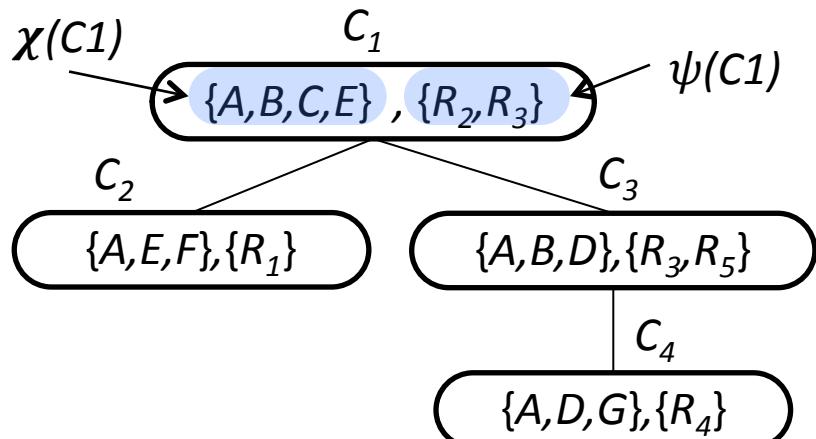
Tree Decomposition

- A tree decomposition:
 $\langle T, \chi, \psi \rangle$
 - T : a tree of clusters
 - χ : maps variables to clusters
 - ψ : maps constraints to clusters



Hypergraph

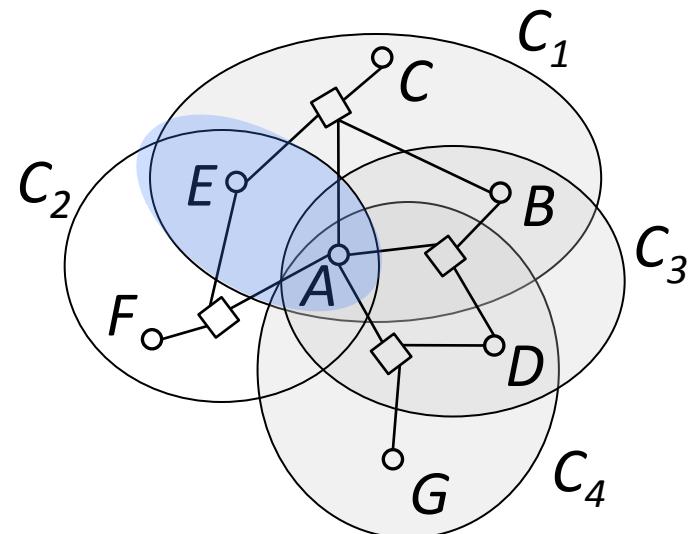
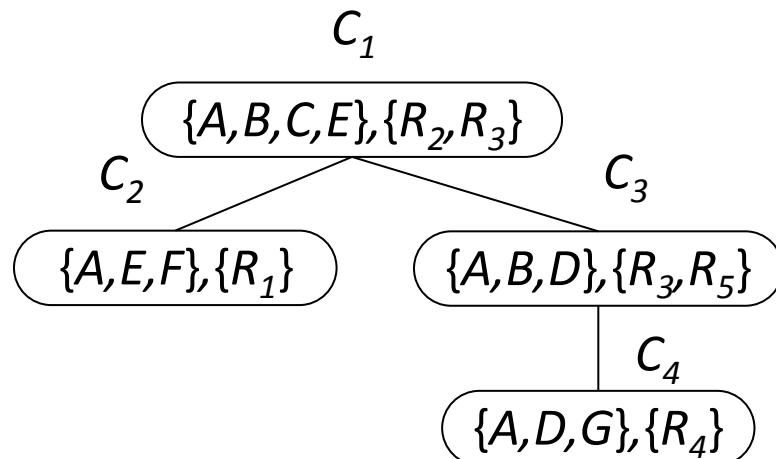
- Conditions
 - Each constraint appears in at least one cluster with all the variables in its scope
 - For every variable, the clusters where the variable appears induce a connected subtree



Tree decomposition

Tree Decomposition: Separators

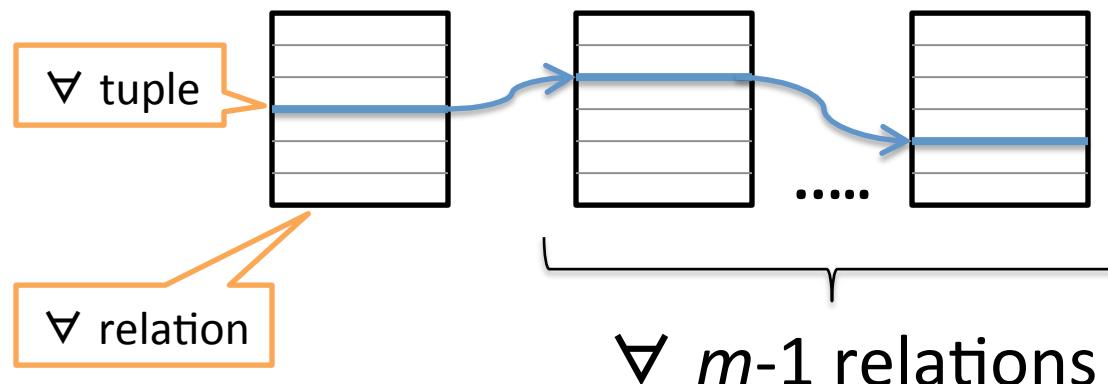
- A **separator** of two adjacent clusters is the set of variables associated to both clusters



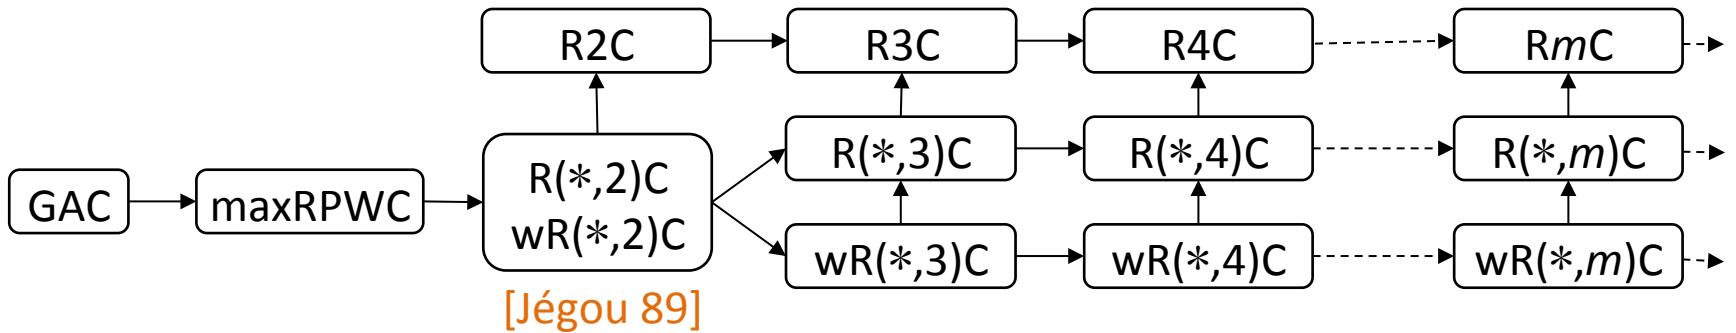
- **Width** of a decomposition/network
 - Treewidth = maximum number of variables in clusters - 1

Relational Consistency $R(*,m)C$

- A CSP is $R(*,m)C$ iff
 - Every tuple in a relation can be extended
 - to the variables in the scope of any $(m-1)$ other relations
 - in an assignment satisfying all m relations simultaneously
- $R(*,m)C \equiv$ Every set of m relations is minimal



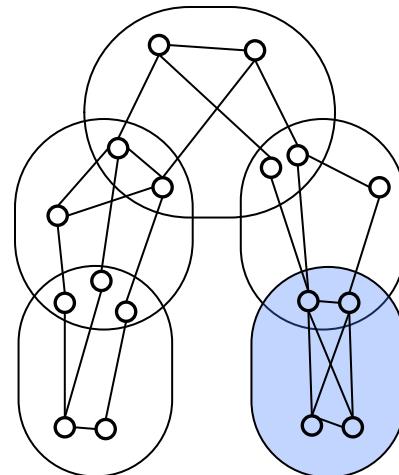
Characterizing $R(*,m)C$



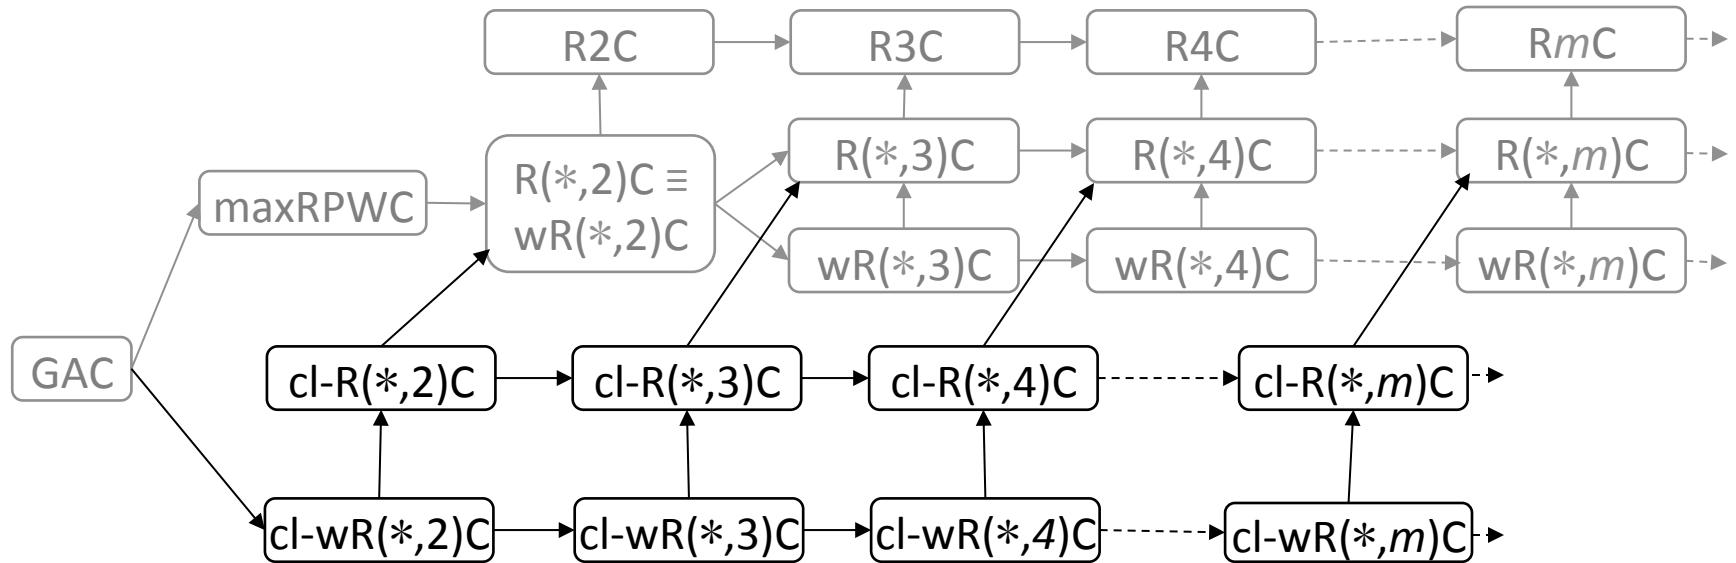
- GAC [Waltz 75]
- maxRPWC [Bessiere+ 08]
- RmC: Relational m Consistency [Dechter+ 97]

Localize Consistency

- Restricting $R(*,m)C$ to clusters: $\text{cl}\text{-}R(*,m)C$
- Localization $\text{cl}\text{-}R(*,m)C$
 - Fewer combinations of m relations
 - Reduces the enforced consistency level



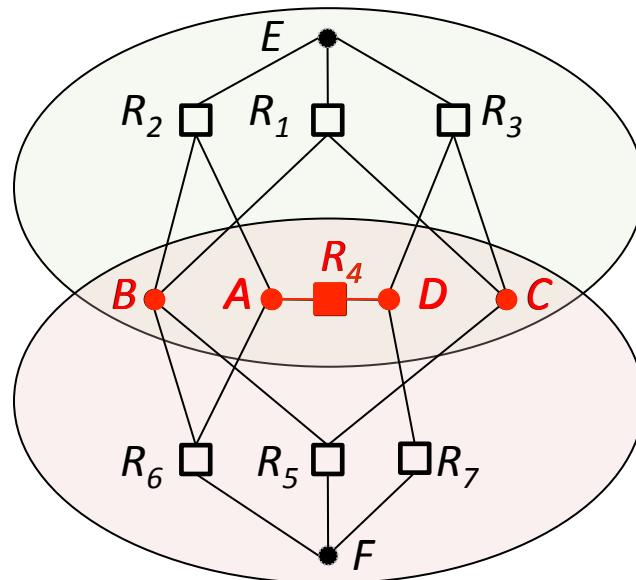
Characterizing cl-R(*,m)C



- GAC [Waltz 75]
- maxRPWC [Bessiere+ 08]
- RmC: Relational m Consistency [Dechter+ 97]

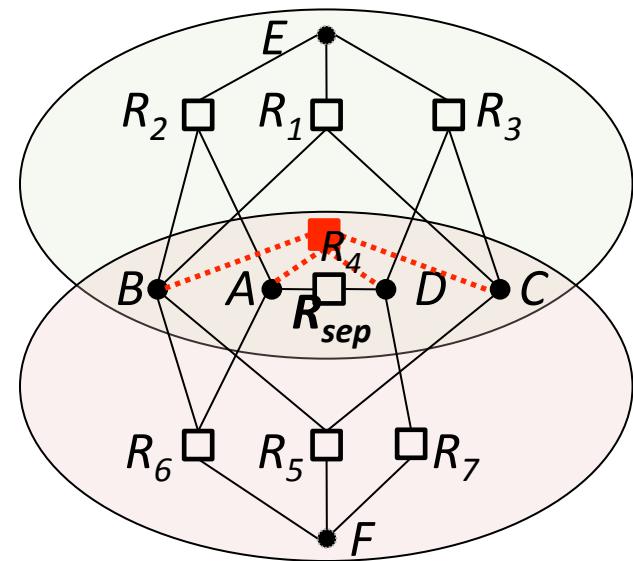
Communication Between Clusters

- Two clusters communicate via their separator
 - Constraints common to the two clusters
 - Domains of variables common to the two clusters

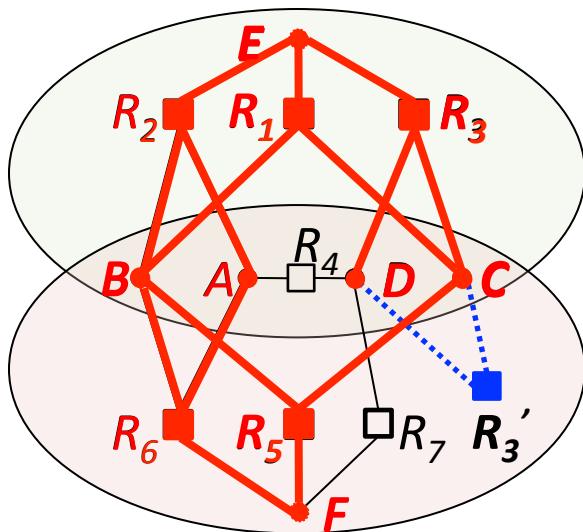


Bolstering Propagation at Separators

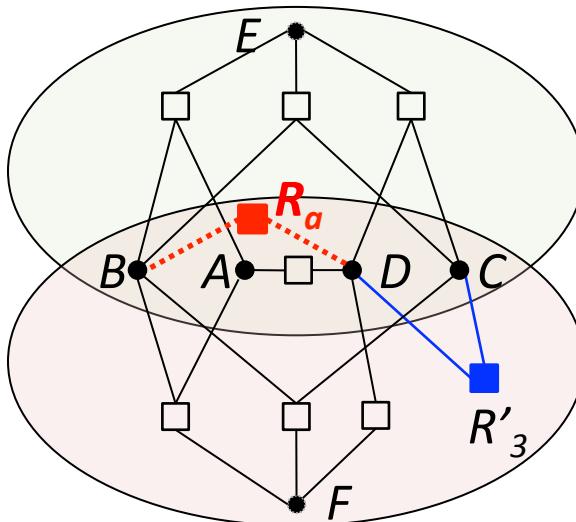
- For perfect communication between clusters
 - Ideally, add unique constraint
 - Space overhead, major bottleneck
- Enhance propagation by **bolstering**
 - **Projection** of existing constraints
 - Adding **binary** constraints
 - Adding **clique** constraints



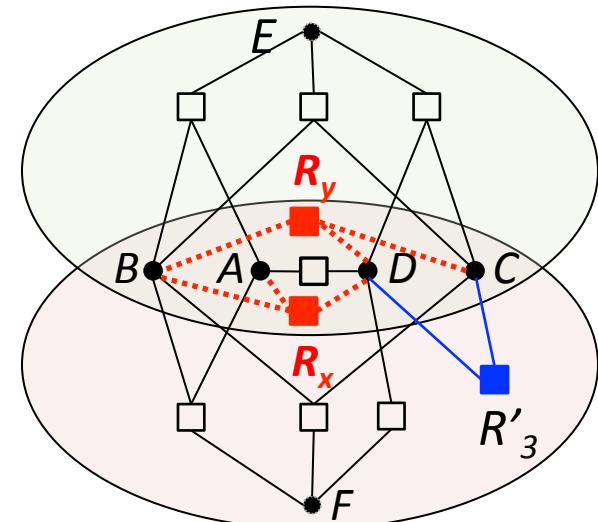
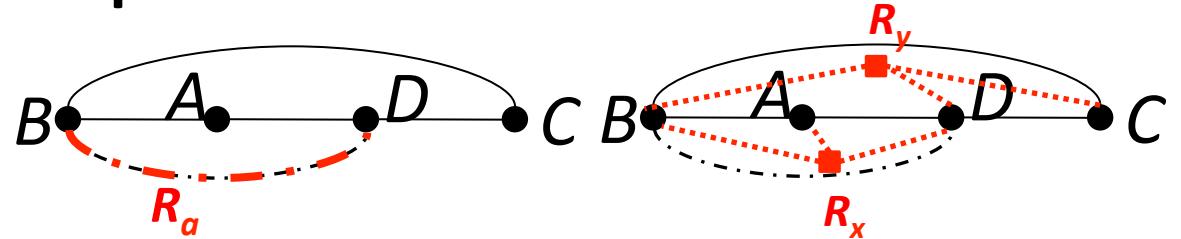
Bolstering Schemas: Approximate Unique Separator Constraint



Projection
 $\text{cl+proj-R}(*,m)\text{C}$

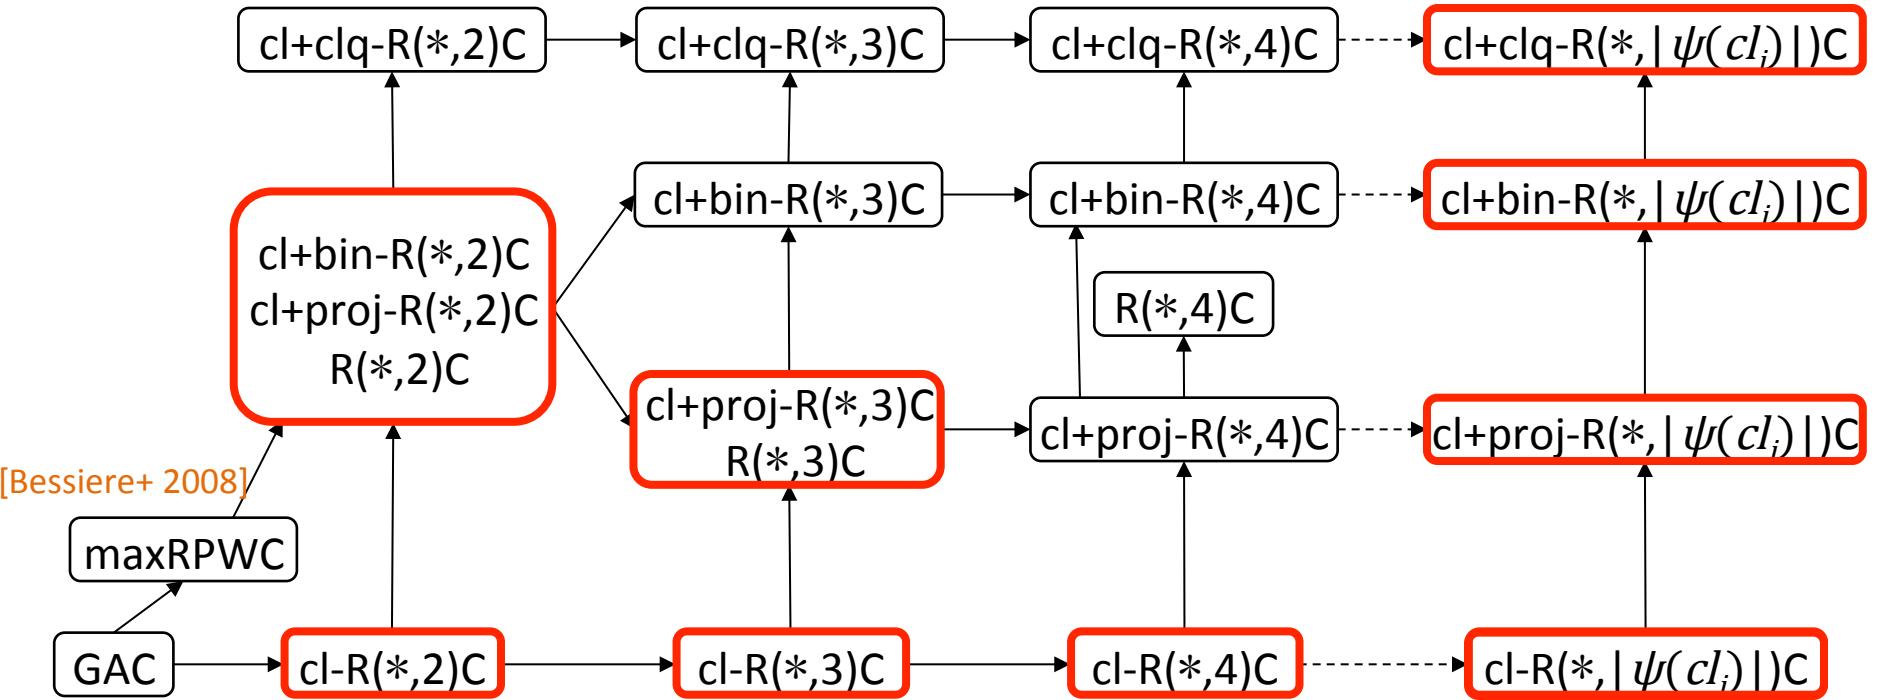


Binary constraints
 $\text{cl+bin-R}(*,m)\text{C}$



Clique constraints
 $\text{cl+clq-R}(*,m)\text{C}$

Resulting Consistency Properties



Experiments

- Domain-based filtering
 - GAC
 - maxRPWC
- Relation-based filtering
 - Global: $wR(*,m)C$ for $m=2,3,4$
 - Local: $cl\text{-}wR(*,m)C$ $m=2,3,4$ and $cl\text{-}R(*, |\psi(cl_i)|)C$
 - Bolstering
 - projection
 - binary
 - clique
- Backtrack search, full-lookahead, first solution
- Reported
 - Number of instances Completed, BT-free, Min(#NV), Fastest
 - Cumulative count of BT-Free as a function of treewidth

Characteristics of Benchmark Data

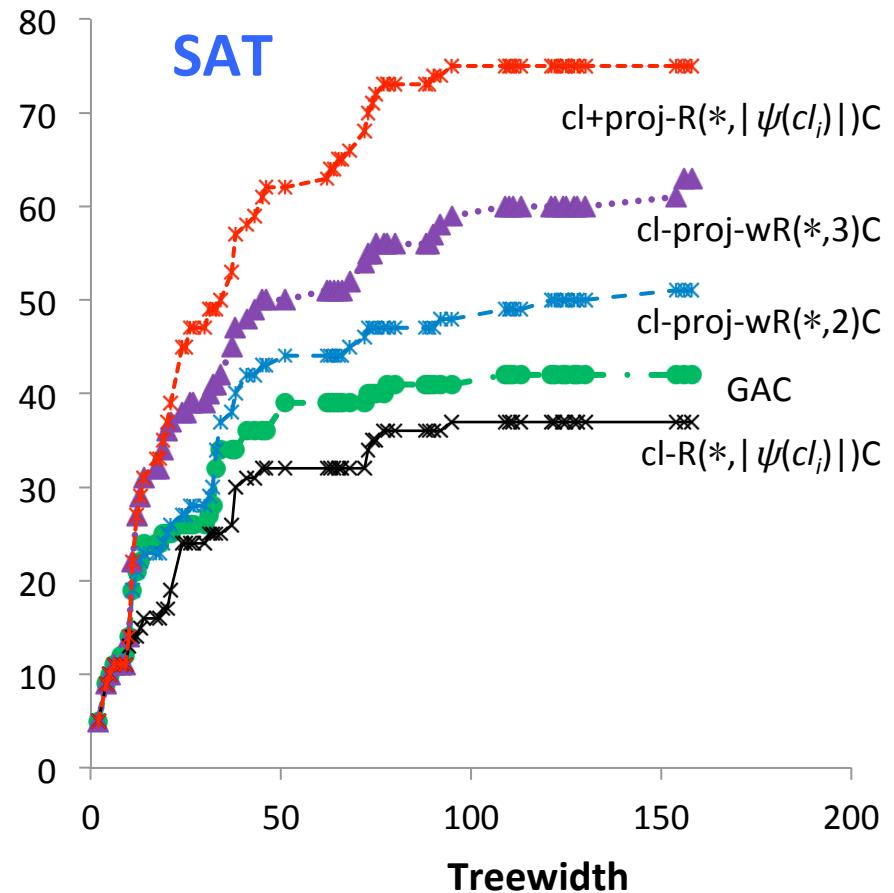
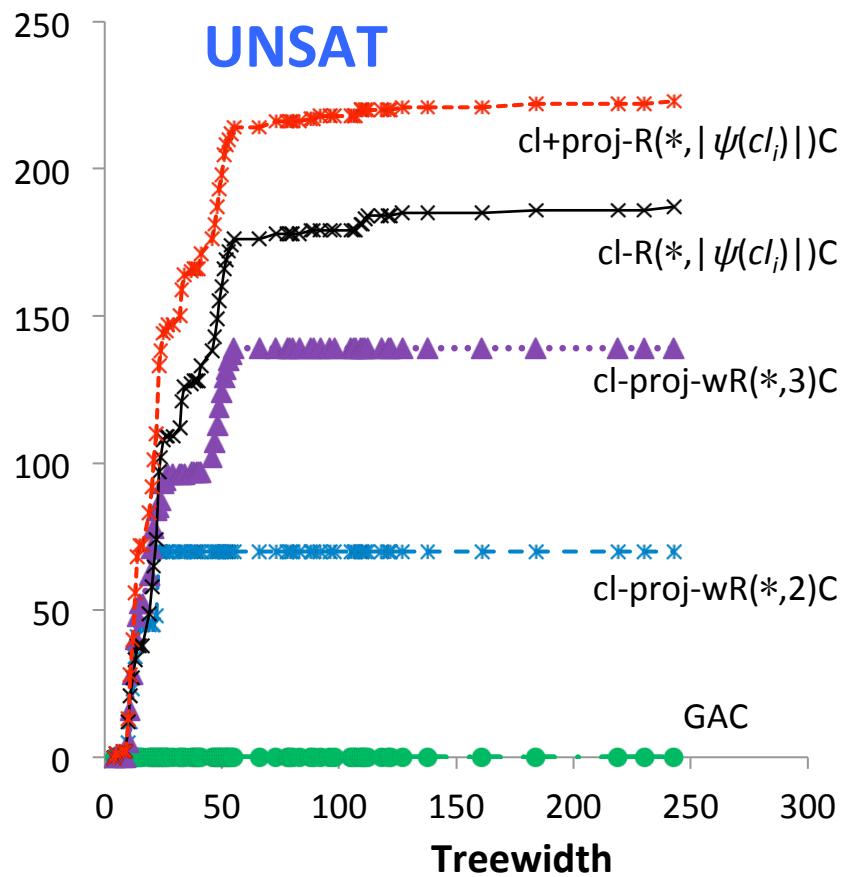
- UNSAT: 479 instances
- SAT: 200 instances

	max		median		mean	
	UNSAT	SAT	UNSAT	SAT	UNSAT	SAT
treewidth	243	158	33	18	43.45	34.44
largest sep.	214	157	28	16.5	39.02	31.33
max($ \psi(cl_i) $)						
local	1,243	211	16	8	109.54	18.35
projection	1,243	211	18	11	114.70	37.35
binary	1,243	653	24	12	199.50	80.65
clique	1,243	148	18	10	113.35	25.87
clique arity	48	26	7	4	7.40	5.97

Empirical Evaluations

		+ maxRPWC, $m=3,4$									
		wR(*,2)C						R(*, $ \psi(\text{cl}_i) $)C			
#inst.		GAC	global	local	Proj.	binary	clique	local	Proj.	binary	clique
Completed	UNSAT	167	170	167	172	169	162	285	286	282	271
	479	34.9%	35.5%	34.9%	35.9%	35.3%	33.8%	59.5%	59.7%	58.9%	56.6%
BT-Free	SAT	174	179	178	176	169	104	152	138	124	113
	200	87.0%	89.5%	89.0%	88.0%	84.5%	52.0%	76.0%	69.0%	62.0%	56.5%
Min(#NV)	UNSAT	0	70	39	70	70	74	187	223	223	213
	479	0.0%	14.6%	8.1%	14.6%	14.6%	15.4%	39.0%	46.6%	46.6%	44.5%
Fastest	SAT	44	55	37	53	52	38	39	77	71	58
	200	22.0%	27.5%	18.5%	26.5%	26.0%	19.0%	19.5%	38.5%	35.5%	29.0%
	UNSAT	17	73	43	72	72	77	220	249	248	239
	479	3.5%	15.2%	9.0%	15.0%	15.0%	16.1%	45.9%	52.0%	51.8%	49.9%
	SAT	47	64	37	62	61	39	83	111	100	79
	200	23.5%	32.0%	18.5%	31.0%	30.5%	19.5%	41.5%	55.5%	50.0%	39.5%

Cumulative Count of Instances Solved w/o Backtracking



Acknowledgment: Charts suggested by Rina Dechter

Conclusions & Future Work

- Adapted $R(*,m)C$ to a tree decomposition of the CSP
 - Localizing $R(*,m)C$ to the clusters
 - Bolstering separators to strengthen the enforced consistency
- Directions for future work
 - $R(*,m)C$ on non-table constraints via domain filtering
 - Automating the selection of a consistency property and of a bolstering scheme
 - On clusters/seperators
 - During search
 - Modify the structure of a tree decomposition to improve performance (e.g., merging clusters [Fattah & Dechter 1996])
 - Improve our algorithms for $R(*,m)C$

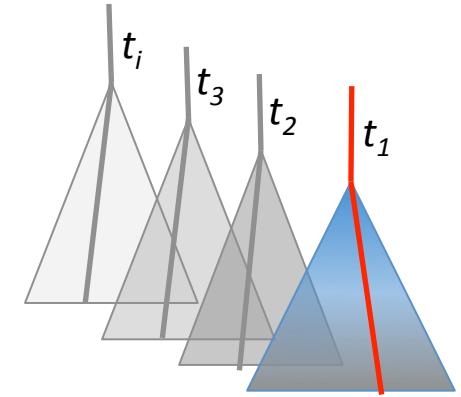
Thank You for Your Attention

Shant's Thesis

- Question
 - Practical tractability of CSPs exploiting the condition linking
 - the level of consistency
 - to the width of the constraint graph
- Solution
 - Introduced a parameterized consistency property $R(*,m)C$
 - Designed algorithms for implementing it
 - `PERTUPLE` and `ALLSOL`
 - `Hybrid` algorithms
 - Adapted $R(*,m)C$ to a tree decomposition of the CSP
 - `Localizing` $R(*,m)C$ to the clusters
 - Strategies for `guiding propagation` along the structure
 - `Bolstering` separators to strengthen the enforced consistency
 - Improved the BTD algorithm for solution counting, `WITNESSBTD`
- + Two incidental results in appendices

Algorithms for Enforcing R(*,m)C

- PERTUPLE
 - For each tuple find a solution for the variables in the $m-1$ relations
 - Many satisfiability searches
 - Effective when there are many solutions
 - Each search is quick & easy
- ALLSOL
 - Find all solutions of problem induced by m relations, & keep their tuples
 - A single exhaustive search
 - Effective when there are few or no solutions
- Hybrid Solvers (portfolio based)



[+Scott]

Hybrid Solver

- Choose between PERTUPLE & ALLSOL
- Parameters to characterize the problem
 - κ predicts if instance is at the phase transition
 - `relLinkage` approximates the likelihood of a tuple at the overlap two relations to appear in a solution
- Classifier built using Machine Learning
 - C4.5
 - Random Forests

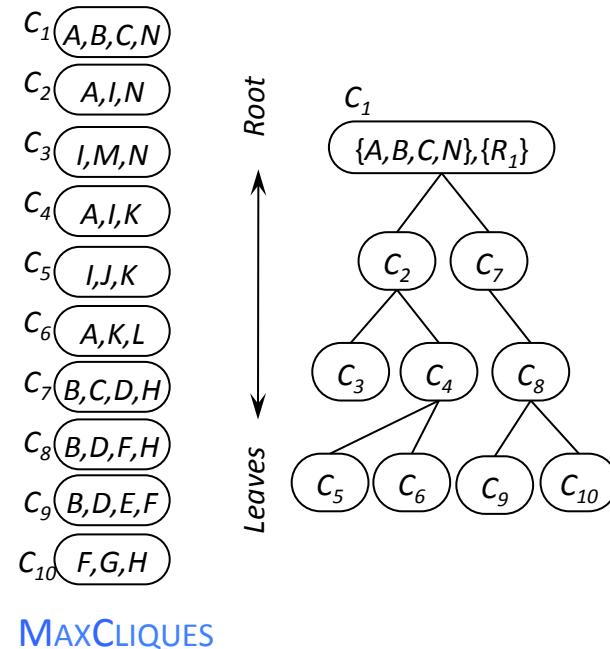
Empirical Evaluations: Hybrid

Task: compute the minimal CSP

Partition	#Instances solved by...					Average CPU sec				
	ALLSOL	PERTUPLE	SOLVER _{C4.5}	SOLVER _{RF}	#Instances	ALLSOL	PERTUPLE	SOLVER _{C4.5}	SOLVER _{RF}	
\mathcal{A}	5,777	5,776	5,777	5,777	5,776	1.27	4.97	2.14	2.27	
\mathcal{P}	10,095	15,457	15,439	14,012	10,095	109.61	5.21	7.72	31.53	
$\mathcal{A} \cup \mathcal{P}$	15,872	21,333	21,216	19,789	15,871	70.18	5.12	5.69	20.88	

Structure-Guided Propagation

- Orderings
 - RANDOM: FIFO/arbitrary
 - STATIC, PRIORITY, DYNAMIC: Leaves \leftrightarrow root
- Structure-based propagation
 - STATIC
 - Order of MAXCLIQUEs
 - PRIORITY:
 - Process a cluster once in each direction
 - Select most significantly filtered cluster
 - DYNAMIC
 - Similar to PRIORITY
 - But may process ‘active’ clusters more than once

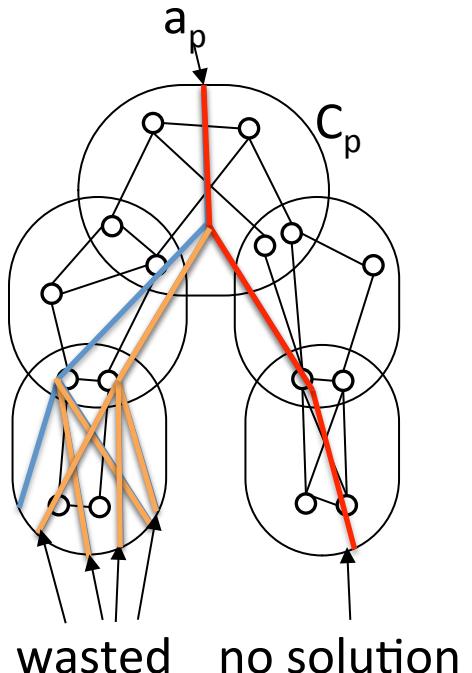


Solution Counting

- BTD is used to count the number of solutions
- Using COUNT algorithm on trees
- Witness-based solution counting
- Find a witness solution before counting

[Favier+ 09]

[Dechter+ 87]



Empirical Evaluations

- Task: Count solutions
 - BTD versus WITNESSBTD
 - GAC: Pre-processing & full look-ahead

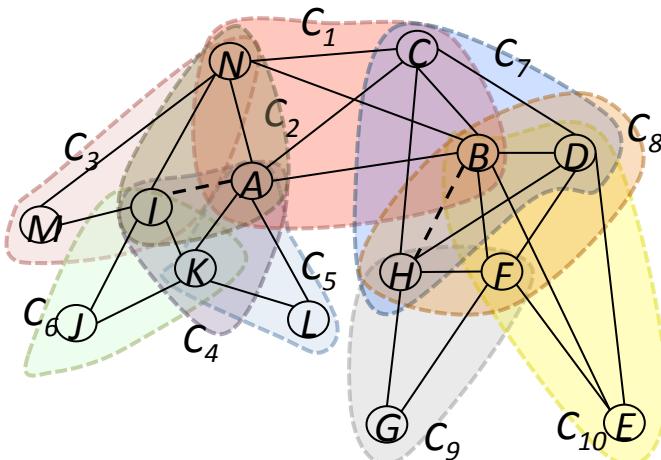
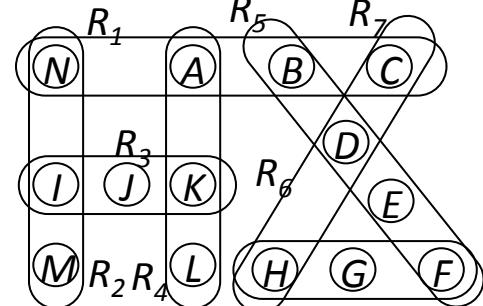
		BTD	WITNESSBTD
Avg. #NV	UNSAT (89)	1,437,909.79	734,983.25
	SAT (101)	4,785,737.57	4,735,136.28
Avg. Time (sec)	UNSAT (89)	145.55	123.86
	SAT (101)	1,151.35	1,148.46

Empirical Evaluations: WITNESSBTD

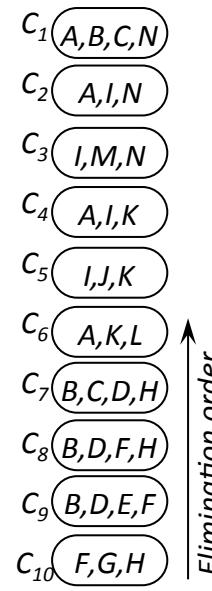
		wR(*,3)C						R(*, ψ(cl _i))C				
		#inst.	GAC	global	local	Proj.	binary	clique	local	Proj.	binary	clique
	+ maxRPWC, m=2,4											
Completed	UNSAT	200	192	248	237	236	220	302	290	286	277	
	479	41.8%	40.1%	51.8%	49.5%	49.3%	45.9%	63.0%	60.5%	59.7%	57.8%	
BT-Free	SAT	111	85	112	100	96	81	110	90	88	78	
	200	55.5%	42.5%	56.0%	50.0%	48.0%	40.5%	55.0%	45.0%	44.0%	39.0%	
Min(#NV)	UNSAT	0	97	104	139	139	131	186	221	222	212	
	479	0.0%	20.3%	21.7%	29.0%	29.0%	27.3%	38.8%	46.1%	46.3%	44.3%	
Fastest	SAT	15	42	17	47	47	45	25	61	61	52	
	200	7.5%	21.0%	8.5%	23.5%	23.5%	22.5%	12.5%	30.5%	30.5%	26.0%	
	UNSAT	2	101	111	145	145	136	235	263	264	244	
	479	0.4%	21.1%	23.2%	30.3%	30.3%	28.4%	49.1%	54.9%	55.1%	50.9%	
	SAT	19	42	23	52	49	49	57	74	73	65	
	200	9.5%	21.0%	11.5%	26.0%	24.5%	24.5%	28.5%	37.0%	36.5%	32.5%	
	UNSAT	100	26	120	71	23	23	189	126	58	52	
	479	20.9%	5.4%	25.1%	14.8%	4.8%	4.8%	39.5%	26.3%	12.1%	10.9%	
	SAT	73	20	20	18	9	9	27	15	9	9	
	200	36.5%	10.0%	10.0%	9.0%	4.5%	4.5%	13.5%	7.5%	4.5%	4.5%	

Generating a Tree Decomposition

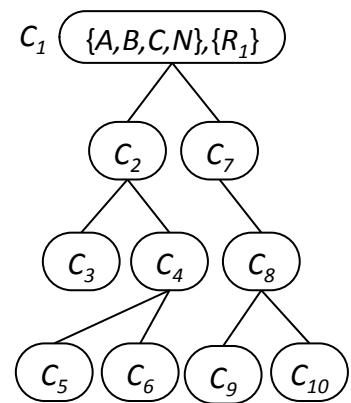
- Triangulate the primal graph using min-fill
- Identify the maximal cliques using MAX CLIQUES
- Connect the clusters using JOIN TREE
- Add constraints to clusters where their scopes appear



min-fill



MAX CLIQUES



JOIN TREE