

Enumerating Tree Decompositions

Nofar Carmeli

Batya Kenig

Benny Kimelfeld

Technion – Israel Institute of Technology



Motivation

- Q1: Is there a manager with a relative in the company?

Works:

Employee	Project
Alice	A
Anna	A
Bob	B
Barak	B


Manages:

Employee	Project
Ester	A
Fady	B
Gil	C
Hava	D

Relative:

Emp1	Emp2
Barak	Hava
Anna	Ester
Carl	Clement
David	Dan

$\text{works}(\text{Emp1}, \text{Proj1}) \wedge \text{manages}(\text{Emp2}, \text{Proj2}) \wedge \text{relative}(\text{Emp1}, \text{Emp2})$



Motivation

- Q2: Is there an employee managed by a relative?

Works:

Employee	Project
Alice	A
Anna	A
Bob	B
Barak	B


Manages:

Employee	Project
Ester	A
Fady	B
Gil	C
Hava	D



Relative:

Employee	Employee
Barak	Hava
Anna	Ester
Carl	Clement
David	Dan

$\text{works}(\text{Emp1}, \text{Proj}) \wedge \text{manages}(\text{Emp2}, \text{Proj}) \wedge \text{relative}(\text{Emp1}, \text{Emp2})$

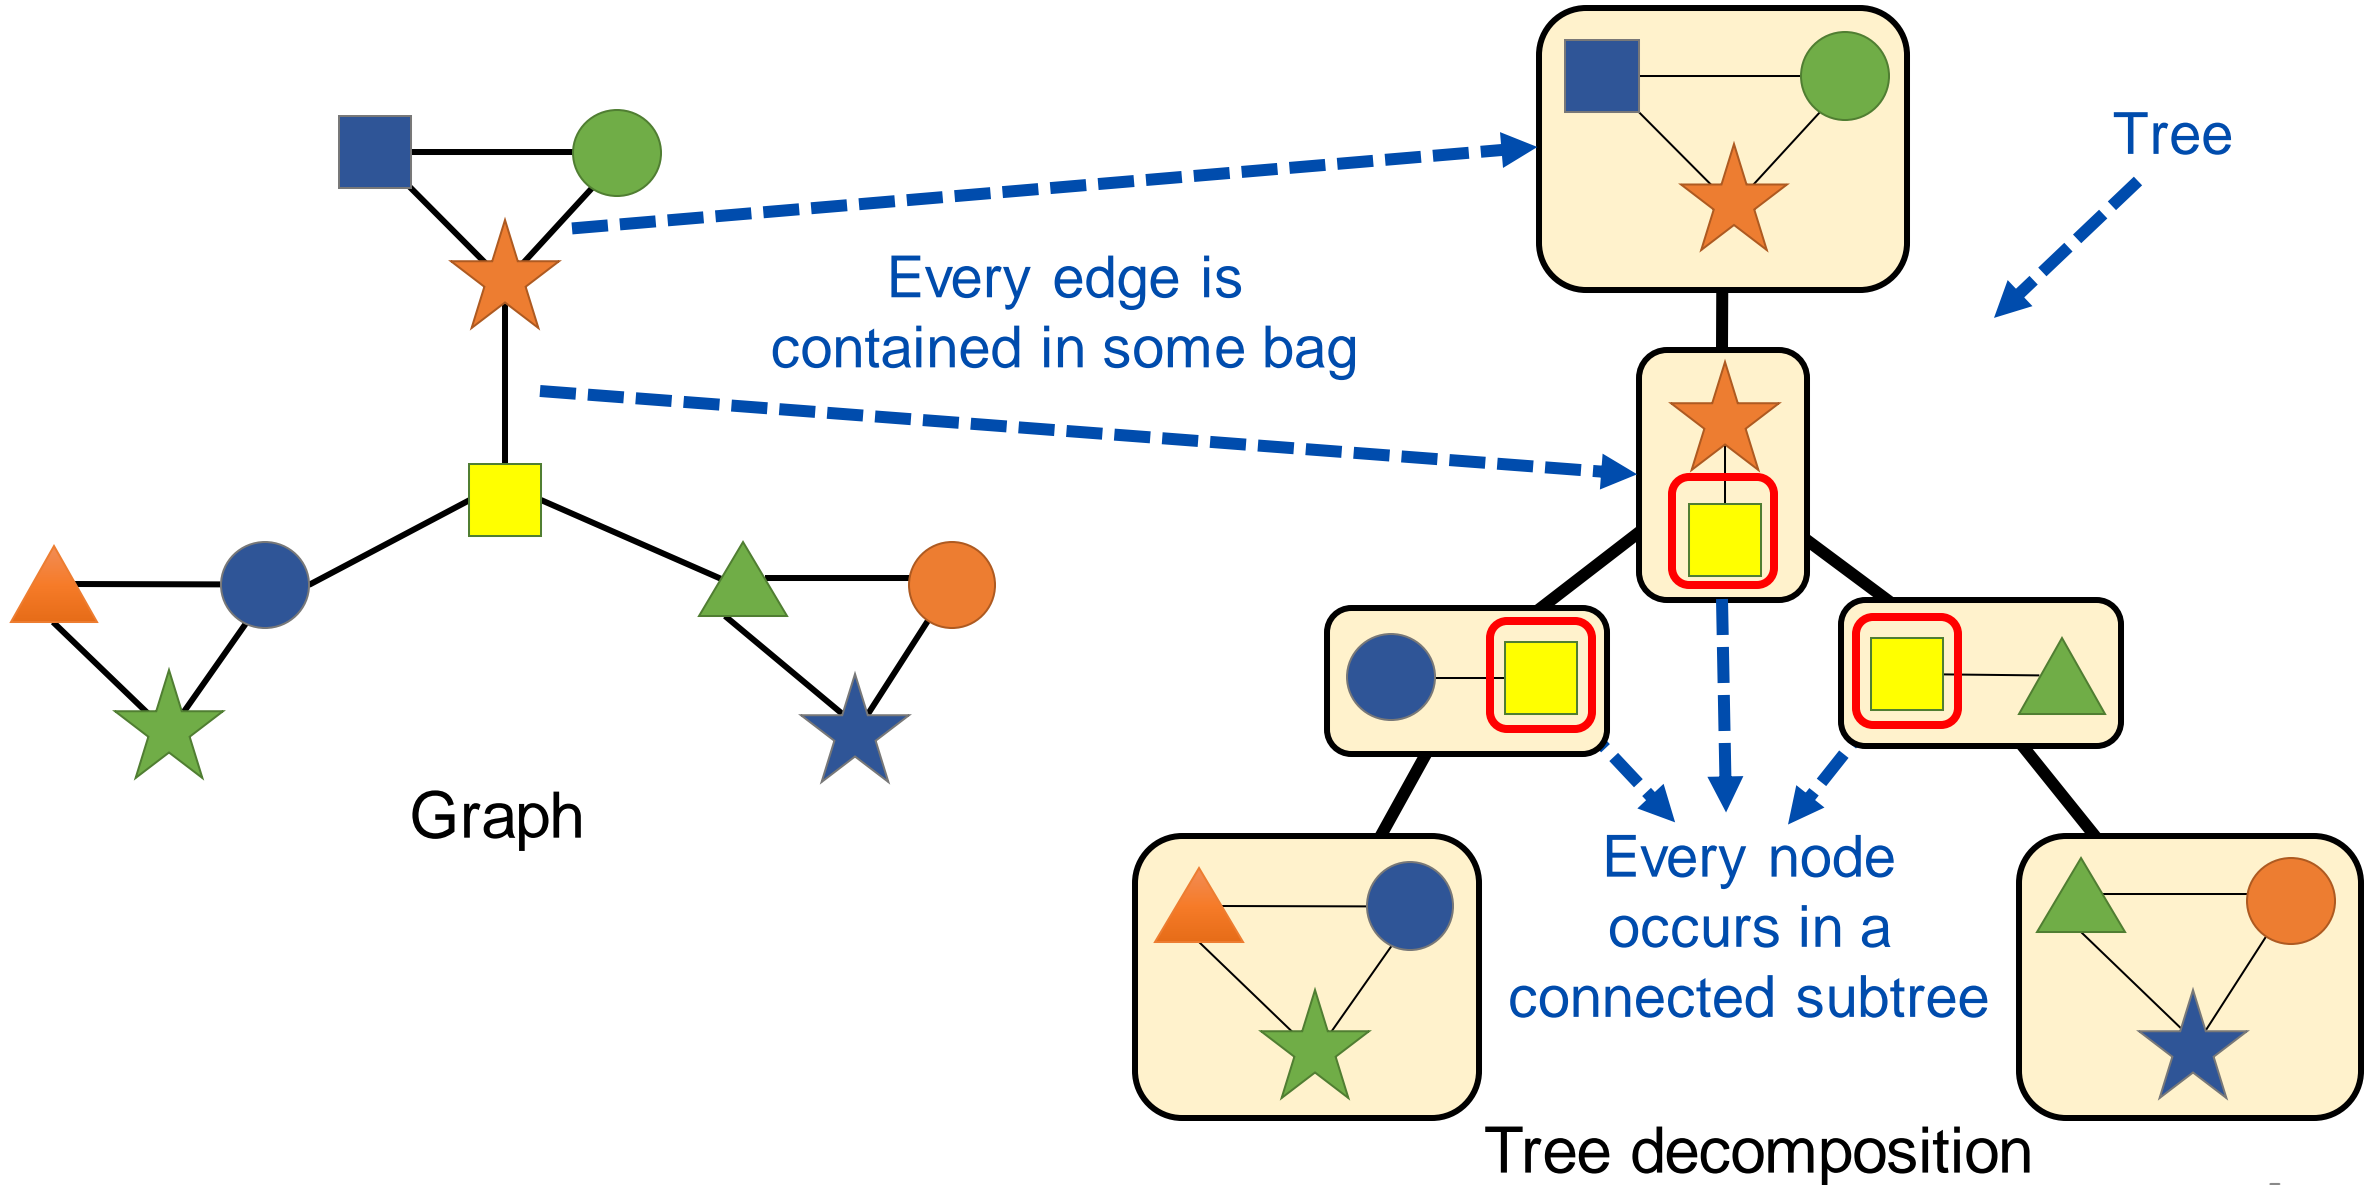


Motivation



- Evaluating a general conjunctive query is NP-complete [\[Chandra&Merlin77\]](#)
- Efficient algorithm for acyclic conjunctive queries [\[Yannakakis81\]](#)
- A **tree decomposition** allows applying Yannakakis's to general conjunctive queries [\[Chekuri&Rajaraman97\]](#)

Tree Decompositions





Tree Decompositions

- Many applications beyond join optimization:
 - Games
 - Nash equilibria computation [\[Gottlob+05\]](#)
 - Bioinformatics
 - prediction of RNA secondary structure [\[Zhao+06\]](#)
 - Probabilistic graphical models
 - statistical inference [\[Lauritzen&Spiegelhalter88\]](#)
 - Constraint-satisfaction problems [\[Kolaitis&Vardi00\]](#)
 - Weighted model counting [\[Li+08\]](#)
 - ...

Which TD to use?

- A graph can have many TDs

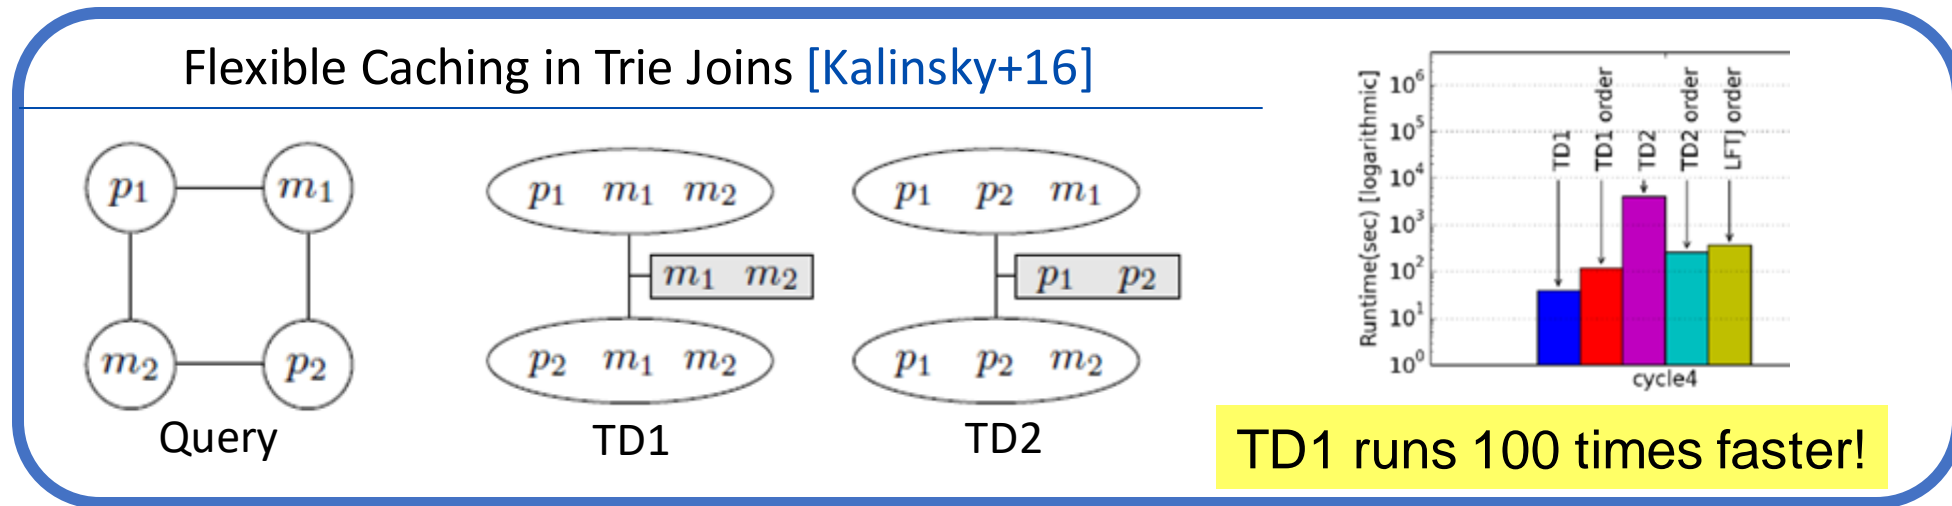


Tree decompositions

- We want the 'best' decomposition
- Common – minimize the cardinality of the largest bag (smallest width)

Which TD to use?

- Smallest width is NP-hard [Arnborg+87]
- Common: Use heuristics
- Width isn't enough



- Different applications – different requirements

TD enumeration is needed

- Related work:
 - Query plans using generalized hypertree decompositions [\[Tu&Ré15\]](#)
 - Generate all, choose one
 - No complexity guarantees
 - Works for small graphs
 - Improving the efficiency of dynamic programming on tree decompositions using machine learning [\[Abseher+15\]](#)
 - Heuristically generate a pool, choose using machine learning
 - Limited pool, may not contain the best
 - Can we enumerate the TDs with efficiency guarantees?

Goal




Problem: Enumerating **all** TDs of a graph


1. Complexity guarantees
2. Effective practical solution

There can be exponentially many TDs!


Which TDs to Generate?




Graph



Tree decomposition



Better tree
decomposition



Better tree
decomposition

Proper TDs

- We define “proper” TDs
- Intuitively, in a proper TD you cannot:
 - Split bags
 - Remove bags


Goal: Enumerating all **proper** TDs of a graph

Problem: exponentially many TDs, what is an “efficient” algorithm?


Efficiency of enumeration algorithms

[Johnson, Papadimitriou, Yannakakis 88]


polynomial total time



incremental polynomial time



polynomial delay



The main theoretical result

Main Theorem:



Given a graph, it is possible to enumerate in incremental polynomial time:

- The proper tree decompositions
- The *minimal triangulations*


Goal: Enumerate Proper TDs

- Chord: An edge between two non-adjacent nodes in a cycle
- Chordal graph: Every cycle of length > 3 has a chord

Chord:



Not Chordal




Chordal

Goal: Enumerate Proper TDs


- Chord: An edge between two non-adjacent nodes in a cycle
- Chordal graph: Every cycle of length > 3 has a chord
- Finding proper TDs of a chordal graph is easy
 - The bags are the maximal cliques
 - These TDs can be enumerated in polynomial delay

[Jordan02][Gavril74] [Yamada+10]




Goal: Enumerate Proper TDs


- Triangulation of a graph: Adding edges to make it chordal
- Minimal triangulation:
Adding a proper subset of the edges does not make it chordal



Graph




Triangulation



Minimal triangulation

Goal: Enumerate Proper TDs

- A bijection:
classes of bag equivalent proper TDs \leftrightarrow min triangulations



Goal: Enumerate Proper TDs


Goal: Enumerating all **proper TDs** of a graph




Goal: Enumerating all **min triangulations** of a graph

Goal: Enumerate Minimal Triangulations

- Minimal Separator:
Removing these nodes separates some u and v
No proper subset separates u and v
- Crossing separators:
One of them separates nodes of the other



- Minimal separators:



- Crossing separators:




- Parallel separators:



Goal: Enumerate Minimal Triangulations

- A bijection [Parra&Scheffler97]:
minimal triangulations \leftrightarrow maximal sets of non crossing
minimal separators



Goal: Enumerate Proper TDs

Goal: Enumerating all **min triangulations** of a graph



Goal: Enumerating all **max independent sets** of a graph

Goal: Enumerate Maximal Independent Sets




Enumerating max independent sets can be done in polynomial delay [Johnson+88]



Problem:
The graph may be of exponential size!

Challenge:
Solve without generating the graph



The Algorithm (Enumerating max independent sets)



```
1:  $J := \text{Extend}(x, \emptyset)$ 
2: print  $J$ 
3:  $\mathcal{Q} := \{J\}$ 
4:  $\mathcal{P} := \emptyset$ 
5:  $\mathcal{V} := \emptyset$ 
6: iterator :=  $A_V(x)$ 
7: while  $\mathcal{Q} \neq \emptyset$  do
8:    $J := \mathcal{Q}.\text{pop}()$ 
9:    $\mathcal{P}.\text{push}(J)$ 
10:  for all  $v \in \mathcal{V}$  do
11:     $J_v := \{v\} \cup \{u \in J \mid \neg A_E(x, v, u)\}$ 
12:     $K := \text{Extend}(x, J_v)$ 
13:    if  $K \notin \mathcal{Q} \cup \mathcal{P}$  then
14:      print  $K$ 
15:       $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 
16:  while  $\mathcal{Q} = \emptyset$  and iterator.hasNext() do
17:     $v := \text{iterator.next}()$ 
18:     $\mathcal{V} := \mathcal{V} \cup \{v\}$ 
19:    for all  $J' \in \mathcal{P}$  do
20:       $J'_v := \{v\} \cup \{u \in J' \mid \neg A_E(x, v, u)\}$ 
21:       $K := \text{Extend}(x, J'_v)$ 
22:      if  $K \notin \mathcal{Q} \cup \mathcal{P}$  then
23:        print  $K$ 
24:         $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 
```

- Redesign of an algorithm for hereditary graph properties [\[Cohen+08\]](#)
- Assuming:
 - Efficiently enumerating nodes
 - Efficiently checking edges
 - Efficiently extending an independent set
 - Polynomial size of max independent sets
- Extends all nodes in the direction of all independent sets.
- Runs in incremental poly time

The Algorithm (Enumerating max independent sets)

```
1:  $J := \text{Extend}(x, \emptyset)$ 
2: print  $J$ 
3:  $\mathcal{Q} := \{J\}$ 
4:  $\mathcal{P} := \emptyset$ 
5:  $\mathcal{V} := \emptyset$ 
6: iterator :=  $A_V(x)$ 
7: while  $\mathcal{Q} \neq \emptyset$  do
8:    $J := \mathcal{Q}.\text{pop}()$ 
9:    $\mathcal{P}.\text{push}(J)$ 
10:  for all  $v \in \mathcal{V}$  do
11:     $J_v := \{v\} \cup \{u \in J \mid \neg A_E(x, v, u)\}$ 
12:     $K := \text{Extend}(x, J_v)$ 
13:    if  $K \notin \mathcal{Q} \cup \mathcal{P}$  then
14:      print  $K$ 
15:       $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 
16:  while  $\mathcal{Q} = \emptyset$  and iterator.hasNext() do
17:     $v := \text{iterator.next}()$ 
18:     $\mathcal{V} := \mathcal{V} \cup \{v\}$ 
19:    for all  $J' \in \mathcal{P}$  do
20:       $J'_v := \{v\} \cup \{u \in J' \mid \neg A_E(x, v, u)\}$ 
21:       $K := \text{Extend}(x, J'_v)$ 
22:      if  $K \notin \mathcal{Q} \cup \mathcal{P}$  then
23:        print  $K$ 
24:         $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 
```

- In our case, extending = triangulating
- We can use **any** triangulation or tree decomposition algorithm
- First result = algorithm's result



Goal: Enumerating max independent sets



Goal: Enumerating all **max independent sets** of a graph



Find a **single minimal triangulation**



Solution Summary



Experiments



- Goals: check efficiency and quality
- C++ implementation
- Triangulation algorithms:
 - MCS-M [\[Berry+02\]](#)
 - LB-Triang [\[Berry+06\]](#) with min fill heuristics
- Benchmarks:
 - DuncCap [\[Tu&Ré15\]](#)
 - Heuristics (First result)

Experiments

- Datasets:
 - Database queries
 - TPC-H (LogicBlox translation)
 - 2-19 nodes, 1-46 edges
 - Probabilistic graphical models
 - UAI inference challenge
 - 60-1039 nodes, 135-1696 edges
 - Random
 - 30-200 nodes, 131-13955 edges

Experiments


- A single run (UAI, 414 nodes, 801 edges, MCS-M, 30 minutes)



- Queries, completed within 5 seconds
 - 11 graphs: triangulated
 - 9 graphs: 2-5 triangulations
 - 1 graph: 588 triangulations
 - 1 graph: 700 triangulations

Experiments

- Random (30 minutes)



- Probabilistic graphical models (30 minutes)

alg.	measure	avg #results	avg #≤first	avg min	avg %improv	max %improv
MCS-M	width	33635.0	12733.4	20.2	2.6%	26.3%
MCS-M	fill	33635.0	12724.9	2043.8	14.4%	55.8%
LB-T(fill)	width	11998.3	4744.1	18.5	3.4%	20.7%
LB-T(fill)	fill	11998.3	1013.6	965.8	2.2%	27.6%

Future Work

- Practical
 - Parallelized implementation
 - Heuristics for ranked enumeration
- Theoretical
 - Polynomial delay
 - Restricted versions

Questions?