## Outline

- 
- 
- 
- AND/OR search
- AND/OR search spaces
- Depth-first AND/OR branch and bound
- Best-first AND/OR search
- Advanced searches and tasks
- Exploiting parallelism
- Software


## Outline

- 
- 
- Bounds and heuristics
- AND/OR search
- AND/OR search spaces
- Depth-first AND/OR branch and bound
- Best-first AND/OR search
- Advanced searches and tasks


## Solution Techniques

## AND/OR search

Time: exp(treew
Space: linear
space: exp(treewidth)
Time: exp(treewidth)

Time: exp(treewidth) $\longrightarrow$
Space:exp(treewidth)


Inference: Elimination

## Classic OR Search Space



| A | B | $\mathrm{f}_{1}$ | A | C | $\mathrm{f}_{2}$ | A | E | $\mathrm{f}_{3}$ | A | F | $\mathrm{f}_{4}$ | B | C | $\mathrm{f}_{5}$ | B | D | $\mathrm{f}_{6}$ | B | E | $\mathrm{f}_{7}$ | C | D | $\mathrm{f}_{8}$ | E | F | $\mathrm{f}_{9}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{2}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{3}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{2}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{4}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{3}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ |
| $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{3}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{2}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{2}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{4}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ |
| $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{2}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{2}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ |
| $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{4}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{2}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{4}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{2}$ |

$$
\text { Objective function: } F^{*}=\min _{X} \sum_{i} f_{i}(X)
$$

## A



## The AND/OR Search Tree



Pseudo tree
[Freuder and Quinn, 1985]


IJCAI 2016
[Dechter and Mateescu, 2007]

## The AND/OR Search Tree



Pseudo tree


## Weighted AND/OR Search Tree



| A | B | $\mathrm{f}_{1}$ | A | C | $\mathrm{f}_{2}$ | A | E | $\mathrm{f}_{3}$ | A | F | $\mathrm{f}_{4}$ | B | C | $f_{5}$ |  | B | D |  | 6 | B | E | $\mathrm{f}^{\text {d }}$ |  |  | D | $\mathrm{f}_{8}$ | E | F | $\mathrm{f}_{9}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 2 | 0 | 0 | 3 | 0 | 0 | 0 | 0 | 0 | 2 | 0 | 0 | 0 | ) | 0 | 0 | 04 | 4 | 0 | 0 | 3 |  | 0 | 0 | 1 | 0 | 0 | 1 |
| 0 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 3 | 0 | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 12 | 2 | 0 | 1 | 2 |  | 0 | 1 | 4 | 0 | 1 | 0 |
| 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 2 | 1 | 0 | 0 | 1 | 0 | 2 | 2 | 1 | 0 | 01 | 1 | 1 | 0 | 1 |  | 1 | 0 | 0 | 1 | 0 | 0 |
| 1 | 1 | 4 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 2 | 1 | 1 | 4 | 4 | 1 | 1 | 10 | 0 | 1 | 1 |  |  | 1 | 1 | 0 | 1 | 1 | 2 |

Objective function: $F^{*}=\min _{X} \sum_{i} f_{i}(X)$


## Node Value (bottom-up evaluation)

OR - minimization
AND - summation

## AND/OR versus OR Spaces



## Complexity of AND/OR Tree Search

## AND/OR tree

## OR tree

## Space

$O(n)$
$O(n)$

## Time

$$
\begin{aligned}
& O\left(n d^{t}\right) \\
O\left(n d^{\left(w^{*} \log n\right)}\right) & O\left(d^{n}\right)
\end{aligned}
$$

[Freuder \& Quinn85], [Collin, Dechter \& Katz91], [Bayardo \& Miranker95], [Darwiche01]

$$
\begin{array}{ll}
d=\text { domain size } & n=\text { number of variables } \\
t=\text { depth of pseudo tree } & w^{*}=\text { induced width }
\end{array}
$$

## From Search Trees to Search Graphs

- Any two nodes that root identical sub-trees or sub-graphs can be merged


IJCAI 2016

## From Search Trees to Search Graphs

- Any two nodes that root identical subtrees or subgraphs can be merged


IJCAI 2016

## Merging Based on Contexts

- One way of recognizing nodes that can be merged (based on the graph structure)
- context $(X)=$ ancestors of $X$ in the pseudo tree that are connected to $X$ or to descendants of $X$



## AND/OR Search Graph



| A |  | $\mathrm{f}_{\mathrm{ab}}$ | A | C |  | A | E |  | A | F | $\mathrm{faf}_{\text {a }}$ | B |  |  | B |  | D |  | B |  |  | C | D |  | E |  | ef |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 02 | 0 | 0 | 3 | 0 | 0 | 0 | 0 | 0 | 2 | 0 | 0 | 0 |  |  | 0 | 4 | 0 | 0 | 3 | 0 | 0 | 1 | 0 | 0 | 1 |
| 0 | 1 | 10 | 0 | 1 | 0 | 0 | 1 | 3 | 0 | 1 | 0 | 0 | 1 | 1 |  |  | 1 | 2 | 0 | 1 | 2 | 0 | 1 | 4 | 0 | 1 | 0 |
| 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 2 | 1 | 0 | 0 | 1 | 0 | 2 |  |  | 0 | 1 | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 0 |
| 1 | 1 | 14 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 2 | 1 | 1 | 4 |  |  | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 2 |

Objective function: $F^{*}=\min _{x} \sum_{\alpha} f_{\alpha}\left(x_{\alpha}\right)$

## OR

(A)


Cache table for $\mathbf{D}$

## How Big Is The Context?

- Theorem: The maximum context size for a pseudo tree is equal to the treewidth of the graph along the pseudo tree.

(C K H ABEJLNODPMFG)


## Complexity of AND/OR Graph Search

## AND/OR graph <br> OR graph

## Space

$O\left(n d^{w^{*}}\right)$
$O\left(n d^{p w^{*}}\right)$

## Time

$$
O\left(n d^{w^{*}}\right) \quad O\left(n d^{p w^{*}}\right)
$$

d = domain size
$\mathrm{w}^{*}=$ induced width
$\mathrm{n}=$ number of variables
pw* $=$ pathwidth

$$
w^{*} \leq p w^{*} \leq w^{*} \log n
$$ <br> \title{

All Four Search Spaces
} <br> \title{
All Four Search Spaces
}



Full OR search tree
126 nodes


Full ANDIOR search tree
54 AND nodes


Context minimal OR search graph
28 nodes


## Context minimal AND/OR search graph

18 AND nodes <br> \title{
All Four Search Spaces
} <br> \title{
All Four Search Spaces
}



Full OR search tree
126 nodes


Full ANDIOR search tree
54 AND nodes


Context minimal OR search graph
28 nodes


## Context minimal AND/OR search graph

18 AND nodes

# All Four Search Spaces 




Full OR search tree
126 nodes


Full ANDIOR search tree
54 AND nodes


Context minimal OR search graph
28 nodes


Context minimal AND/OR search graph
18 AND nodes
Any query is best computed over the c-minimal AO space

All Four Search Spaces


Full ANDIOR search tree
54 AND nodes


Context minimal AND／OR search graph
18 AND nodes
Any query is best computed over the c－minimal AO space

All Four Search Spaces


Any query is best computed over the c-minimal AO space

## The Impact of the Pseudo Tree



## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


How to find a good one?

(CDKBAOMLNPJHEFG)

## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


How to find a good one?


## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


How to find a good one?


## The Impact of the Pseudo Tree


(CKHABEJLNODPMFG)


How to find a good one?


## Outline

- 
- 
- Bounds and heuristics
- ANDIOR search
- Depth-first AND/OR branch and bound
- 
- 
- 
- 


## Classic Depth-First Branch and Bound

Each node is a COP sub-problem
 (defined by current conditioning)
$\mathrm{g}(\mathrm{n})$ : cost of the path from root to n

$$
\begin{array}{r}
\tilde{f}(n)=g(n)+\tilde{h}(n) \\
\quad \text { (lower bound) }
\end{array}
$$

$$
\text { Prune if } \tilde{f}(n) \geq U B
$$

$\tilde{h}(n)$ : under-estimates optimal cost below n
(UB) Upper Bound = best solution so far

## Partial Solution Tree



Pseudo tree

( $A=0, B=0, C=0, D=0)$


Extension( $\left.T^{\prime}\right)$ - solution trees that extend $T^{\prime}$

## Exact Evaluation Function



$$
f^{*}\left(T^{\prime}\right)=w(A, 0)+w(B, 1)+w(C, 0)+w(D, 0)+v(D, 0)+v(F)
$$

## Heuristic Evaluation Function



| $B$ | $D$ | $E$ | $f_{3}(B D E)$ |
| :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 6 |
| 0 | 0 | 1 | 4 |
| 0 | 1 | 0 | 8 |
| 0 | 1 | 1 | 5 |
| 1 | 0 | 0 | 9 |
| 1 | 0 | 1 | 3 |
| 1 | 1 | 0 | 7 |
| 1 | 1 | 1 | 4 |



$$
f\left(T^{\prime}\right)=w(A, 0)+w(B, 1)+w(C, 0)+w(D, 0)+h(D, 0)+h(F)=12 \leq f^{*}\left(T^{\prime}\right)
$$

## AND/OR Branch and Bound Search

UB (best solution found so far)


## AND/OR Branch and Bound (AOBB)

- Each node n : heuristic lower bound $h(n)$ on $v(n)$
- EXPAND (top-down)
- Evaluate $f\left(T^{\prime}\right)$ and prune search if $f\left(T^{\prime}\right) \geq$ UB
- If not in cache, generate successors of the tip node n
- UPDATE (bottom-up)
- Update value of the parent $p$ of $n$
- OR nodes: minimization
- AND nodes: summation
- Cache value of $n$ based on context


## Breadth-Rotating AOBB

- AND/OR decomposition vs. depth-first search:
- Compromises anytime property of AOBB
- Breadth-Rotating AOBB:
- Combined breadth/depth-first schedule
- Maintains depth-first complexity
- Superior experimental results




## Mini-Bucket Heuristics for AND/OR Search

- The depth-first and best-first AND/OR search algorithms use $h(n)$ that can be computed:
- Static Mini-Bucket Heuristics
- Pre-compiled
- Reduced computational overhead
- Less accurate
- Static variable ordering
- Dynamic Mini-Bucket Heuristics
- Computed dynamically, during search
- Higher computational overhead
- High accuracy
- Dynamic variable ordering


## Outline

- 
- 
- Bounds and heuristics
- ANDIOR search
- 
- 
- Best-first AND/OR search
- Advanced searches and tasks


## Basic Heuristic Search Schemes

Heuristic function $\tilde{f}\left(\hat{x}_{p}\right)$ computes a lower bound on the best extension of partial configuration $\hat{x}_{p}$ and can be used to guide heuristic search.
We focus on:

1. Branch-and-Bound Use heuristic function $\tilde{f}\left(\hat{x}_{p}\right)$ to prune the depth-first search tree Linear space

- improve U from "above"



## 2. Best-First Search

Always expand the node with the lowest heuristic value $\tilde{f}\left(\hat{x}_{p}\right)$
Needs lots of memory

- improve L from "below"

$f(\hat{x})=U$


## AOBF: Best-First AND/OR Search



- Each node maintains a q-value $q(n)$ initially $q(n)=h(n)$
- Node q-values revised bottom-up after each node expansion
- Update current best partial solution subtree (a tip node expanded next)
- All expanded nodes are stored in memory
- Search terminates with optimal solution


## AOBF: Best-First AND/OR Search

- AO*-like traversal of the context minimal AND/OR graph
- All nodes expanded are stored in memory
- Each node maintains a q-value: $q(n)$
- best lower bound on optimal cost below $n$
- Node q-values are revised bottom-up after each expansion
- OR: minimization: $\quad q(n)=\min _{n^{\prime} \in \operatorname{succ}(n)}\left(w\left(n, n^{\prime}\right)+q\left(n^{\prime}\right)\right)$
- AND: summation: $\quad q(n)=\Sigma_{n^{\prime} \in \operatorname{succ}(n)} q\left(n^{\prime}\right)$
- (initially, $q(n)=h(n)-$ heuristic lower bound on cost below $n$ )
- Current best partial solution tree updated using efficient arcmarking mechanism
- OR nodes mark best AND successor, following cost revision
- Any of its tip nodes will be expanded at the next iteration


## Recursive Best-First Search

1. Branch-and-Bound (AOBB) Use heuristic function $\tilde{f}\left(\hat{x}_{p}\right)$ to prune the depth-first search tree Linear space

$f(\hat{x})=U$
2. Best-First Search (AOBF) Always expand the node with the lowest heuristic value $\tilde{f}\left(\hat{x}_{p}\right)$ Needs lots of memory

3. Recursive Best-First Search

Expand nodes best-first with the lowest heuristic value $\tilde{f}\left(\hat{x}_{p}\right)$
Linear or Bounded memory

## RBFAOO: Recursive Best-First AND/OR Search



- Threshold $\theta(n)$ of node $n$ is the second best cost to $q(n)$
- Search sub-problem A (blow n ) until revised $q(n)$ exceeds threshold
- Backtrack to node $m$ and discard (or cache) nodes just expanded
- Backup revised $q(n)$ as new threshold for $m$ and search sub-problem B


## RBFAOO: Recursive Best-First AND/OR Search

- AO*-like best-first search is transformed into depth-first search with a threshold
- Backtrack whenever: $q(n) \geq \theta(n)$
- Node q-values are updated in the usual manner
- OR: minimization: $\quad q(n)=\min _{n^{\prime} \in \operatorname{succ}(n)}\left(w\left(n, n^{\prime}\right)+q\left(n^{\prime}\right)\right)$
- AND: summation: $\quad q(n)=\Sigma_{n^{\prime} \in \operatorname{succ}(n) q\left(n^{\prime}\right)}$
- (initially, $q(n)=h(n)$ - heuristic lower bound on cost below $n$ )
- Context-based caching is used for efficiency
- Overestimation is used to avoid frequent node reexpansions


## Empirical Evaluation




Exact MAP inference. Grid and Pedigree benchmarks. Time limit 1 hour.

## Marginal MAP



## AND/OR Search for Marginal MAP

- New advances
- AND/OR Branch and Bound
- Best-First and recursive best-first AND/OR Search
- Anytime depth-first and best-first search
[Marinescu, Dechter, Ihler 2014,2015]; [Lee, Marinescu, Dechter, Ihler, 2016]
- Best-performing exact and anytime Marginal MAP solvers
- Heuristics based on Weighted Mini-Buckets (WMB)
- WMB-MM: single pass with cost-shifting by moment matching
- WMB-JG: iterative updates by message passing along the join-graph


## Searching for M Best Solutions

- New inference and search based algorithms for the task of finding the $m$ best solutions
- Search: m-A*, m-BB
- Inference: elim-m-opt, BE+m-BF
- Extended m-A* and m-BB to AND/OR search spaces for graphical models
- Algorithms: m-AOBB and m-AOBF
- Competitive and often superior to alternative (approximate) approaches based on LP relaxations
[Fromer and Globerson, 2009], [Batra, 2012]


## Searching for M Best Solutions



## Hybrid of Variable Elimination and Search

- Tradeoff space and time


## Search Basic Step: Conditioning

Variable Branching by Conditioning


## Search Basic Step: Conditioning

Variable Branching by Conditioning

Select a variable


## Search Basic Step: Conditioning

 Variable Branching by Conditioning

IJCAI 2015

## Search Basic Step: Conditioning

 Variable Branching by Conditioning

## The Cycle-Cutset Scheme

## Condition until Treeness

- Cycle-cutset
- i-cutset
- C(i)-size of i-cutset

<Tree part
(F) (E)
(C)
(D)


Space: $\exp (\mathrm{i})$, Time: $\mathbf{O}(\exp (\mathbf{i}+\mathrm{c}(\mathrm{i}))$
IJCAI 2015

## Eliminate First



IJCAI 2015

## Eliminate First



IJCAI 2015

## Eliminate First



IJCAI 2015

## Hybrid Variants

- Condition, condition, condition, ... and then only eliminate (w-cutset, cycle-cutset)
- Eliminate, eliminate, eliminate, ... and then only search
- Interleave conditioning and elimination steps (elim-cond(i), VE+C)


## Interleaving Conditioning and Elimination



## Interleaving Conditioning and Elimination



IJCAI 2015

## Interleaving Conditioning and Elimination



IJCAI 2015

## Interleaving Conditioning and Elimination



IJCAI 2015

## Interleaving Conditioning and Elimination



IJCAI 2015

## Interleaving Conditioning and Elimination



IJCAI 2015

## Interleaving Conditioning and Elimination



IJCAI 2015

## Outline

- Exploiting parallelism
- Distributed and parallel search


## New Advances

- Parallel AOBB, first of its kind
- Runs on computational grid
- Extends parallel tree search paradigm
- Two variants with different parallelization logic [Otten and Dechter, 2012]
- Parallel shared-memory RBFAOO
- Parallelization of the sequential RBFAOO [Kishimoto, Marinescu, Botea, 2015]
- Parallel dovetailing for AOBB, RBFAOO, SPRBFAOO
- Towards large-scale MAP/MMAP inference [Kishimoto, Marinescu, Botea, 2016 IJCAI 2016


## Parallel AOBB Illustrated

- Master process applies partial conditioning to obtain parallel subproblems.



## Fixed-depth Parallel AOBB

- Algorithm receives cutoff depth $d$ as input:
- Expand nodes centrally until depth $d$.
- At depth d, submit to grid job queue.
- Explored sub-problem search spaces potentially very unbalanced.



## Variable-depth Parallel AOBB

- Given sub-problem count $p$ and estimator $N$ :
- Iteratively deepen frontier until size p reached:
- Pick sub-problem $n$ with largest estimate $N(n)$ and split.
- Submit sub-problems into job queue by descending complexity estimates.
- Hope to achieve better sub-problem balance.




## Parallel Scaling Summary

- Plot speedup against CPU count.
- Trade off load balancing vs. overhead:
- \#subproblems $\approx 10 \times$ \#CPUs



## SPRBFAOO: Parallel Shared-Memory RBFAOO

- All threads start from the root with the same search strategy and with one shared cache table
- The virtual $q$-value $v q(n)$ for node n is used to control parallel search
- Initially, $v q(n)$ is set to $q(n)$
- When a thread examines $n, v q(n)$ is incremented by a small value $\zeta$
- When all threads finish examining $n, v q(n)$ is set to $q(n)$
- An effective load balancing is obtained without any sophisticated schemes, while promising portions of the search space are examined
- The algorithm guarantees solution optimality


## Empirical Evaluation




(up to 7-fold speedup with 12 threads)




## Parallel Dovetailing

- Simple distributed scheme
- Launch in parallel $m$ instances of the inference algorithm (AOBB, AOBF, RBFAOO, SPRBFAOO) each one solving the same problem instance but with a different input parameter configuration
- e.g., parameter configuration = pseudo tree (ordering)


12 cores


12 nodes $\times 12$ cores $=144$ cores

## Outline

- 
- Inference
- Bounds and heuristics


- Software
- UAI probabilistic inference competitions


## Software

- aolib
- http://graphmod.ics.uci.edu/group/Software (standalone AOBB, AOBF solvers)
- daoopt
- https://github.com/lotten/daoopt
(distributed and standalone AOBB solver)
- merlin
- https://developer.ibm.com/open/merlin
(standalone WMB, AOBB, AOBF, RBFAOO solvers) open source, BSD license


## UAI Probabilistic Inference Competitions

- 2006
- 2008
- 2011
- 2014

(daoopt)

(merlin)

MPE/MAP

MMAP

## Summary

- Only a few principles
- Inference and search should be combined
- Time-space tradeoff
- AND/OR search should be used
- Caching in search should be used
- Parallel search should be used if distributed and/or shared-memory environments are available



## For publication see:

 http://www.ics.uci.edu/~dechter/publications.html

## Thank You

Kalev Kask Irina Rish
Bozhena Bidyuk
Robert Mateescu
Radu Marinescu
Vibhav Gogate
Emma Rollon
Lars Otten
Natalia Flerova

