

Distributed Constraint Optimization

Michal Jakob

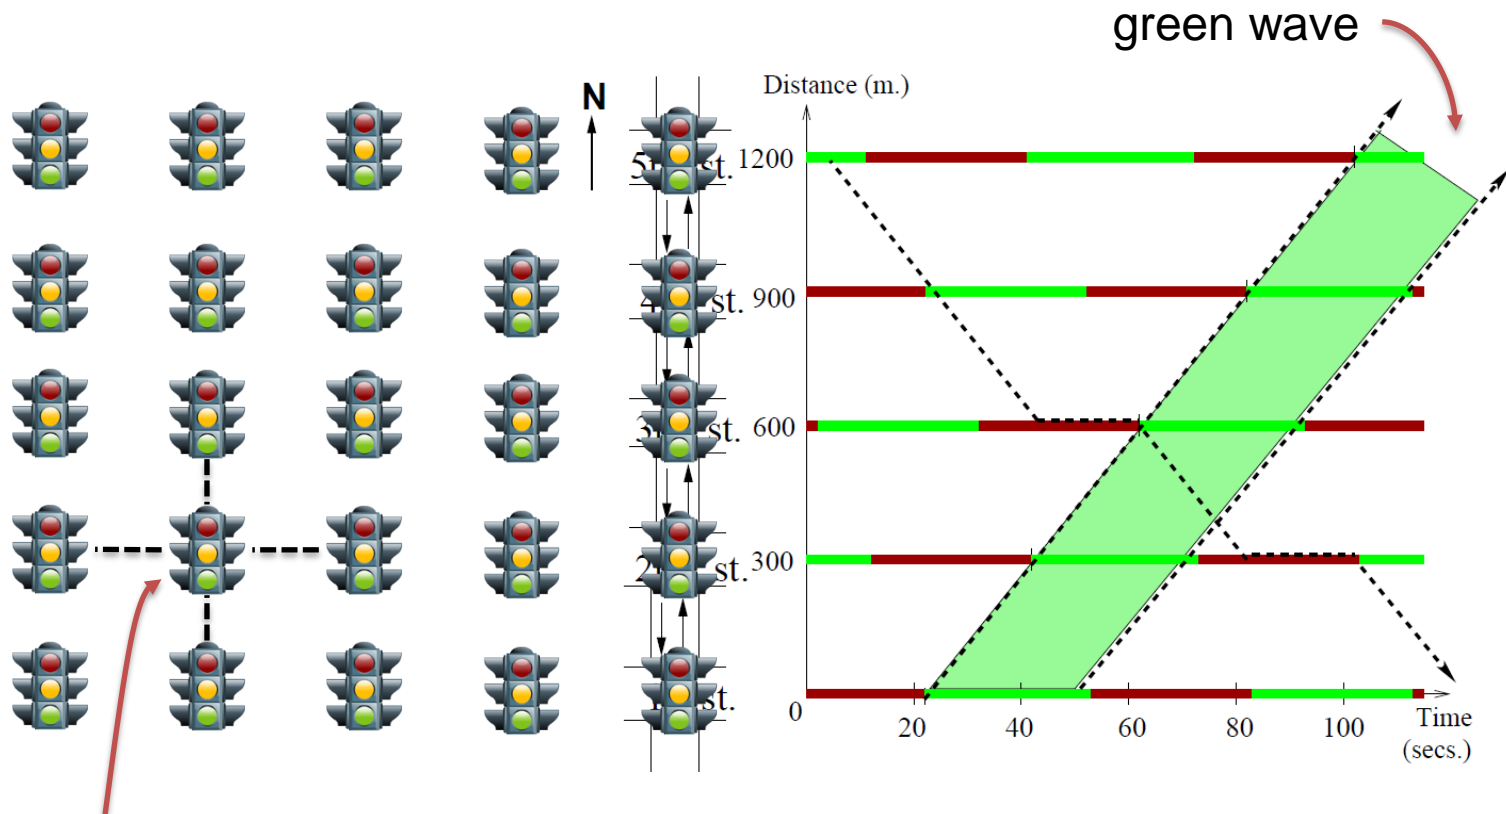
Agent Technology Center, Dept. of Cybernetics, FEE Czech Technical University

A4M33MAS Autumn 2010 - Lect. 12

(based partially on slides by Jay Modi)

Dynamic Traffic Light Scheduling

- Minimize traffic delay in a road grid by synchronizing lights



traffic light agent

- can communicate locally
- can adopt four synchronization strategies (NS, SN, WE, EW)






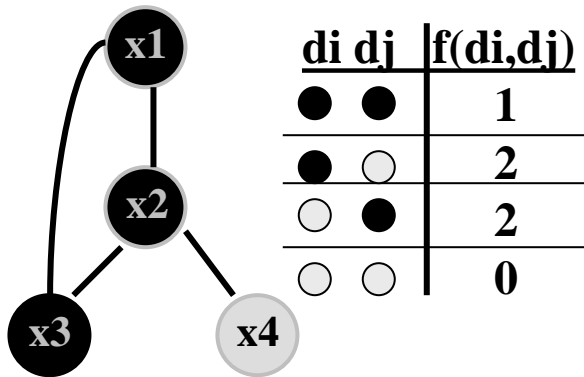
Constraint Optimization Problem (COP)

- Problem specification:
 - $X = \{x_1, \dots, x_m\}$ set of **variables**
 - $D = \{D_1, \dots, D_m\}$ set of **domains** for the variables, i.e. $x_i \in D_i$; if D_i is finite, let $D_i = \{v_{i,1}, \dots, v_{i,d(i)}\}$
 - $C = \{c_1, \dots, c_k\}$ set of **constraints** over X ; the constraint c_i is represented by a real-valued function $P_i(y_1, \dots, y_j) \rightarrow \mathbf{R}$, $\{y_1, \dots, y_j\} \subseteq X$ that determines the **degree of constraint violation (constraint value)** for a given assignment
- Solution:
 - find an **assignment** of variables $\{x_1, \dots, x_m\}$ such that the **sum of all constraint values is minimized**



Example

Constraint Graph



Distributed Constraint Optimization Problem (DCOP)

- $A = \{A_1, \dots, A_n\}$ set of agents
- Each agent A_i is responsible for one variable x_i
 - extension to multiple variables per agent possible
- Agent can communicate by sending messages
- Example



(D)COP with Binary Non-Negative Constraints

- Constraints involve **two variables** at maximum: $P(x_i, x_j)$
- Constraint values are **not negative**: $P(x_i, x_j) \geq 0$
- \Rightarrow addition of constraints **cannot decrease** the overall sum



Solution Algorithms

- Requirements on a good algorithm:
 - **terminates** in a finite number of steps
 - is **complete**: finds an optimum solution (always exists)
 - (is **sound**: the solution returned is indeed optimum)
- Existing algorithms
 - **Asynchronous Distributed Optimization (ADOPT)**
 - Optimal Asynchronous Partial Overlay (OptAPO)
 - Dynamic Parameter Optimization Problem (DPOP)
- Trade-offs between running time (number of cycles), number of messages and size of messages
- Recently, incomplete local search algorithms also proposed



Centralized Branch-and-Bound Algorithm

BRANCH-AND-BOUND-COP()

- 1 $c^* \leftarrow \infty$ ▷ Minimum cost found. Global variable.
- 2 $g^* \leftarrow \emptyset$ ▷ Best solution found. Global variable.
- 3 BRANCH-AND-BOUND-COP-HELPER(1, \emptyset)
- 4 **return** g^*

BRANCH-AND-BOUND-COP-HELPER(i, g)

- 1 **if** $i = n$
- 2 **then if** $P(g) < c^*$
- 3 **then** $g^* \leftarrow g$
- 4 $c^* \leftarrow P(g)$
- 5 **return**
- 6 **for** $v \in D_i$
- 7 **do** $g' \leftarrow g + \{x_i \leftarrow v\}$
- 8 **if** $P(g) < c^*$
- 9 **then** BRANCH-AND-BOUND-COP-HELPER($i + 1, g'$)



Desiderata for DCOP

Why is distributed important?


- Autonomy
- Communication cost
- Robustness (central point of failure)
- Privacy

Why is asynchrony important?

- Parallelism
- Robust to communication delays
- No global clock

Why are theoretical guarantees important?

- Optimal solutions feasible for special classes
- Bound on worst-case performance



loosely connected
communities

State of the Art in DCOP (~2004)

Why have previous distributed methods failed to provide asynchrony + optimality?

- Branch and Bound
 - **Backtrack condition** - when cost exceeds upper bound
 - **Problem** – sequential, synchronous
- Asynchronous Backtracking
 - **Backtrack condition** - when constraint is unsatisfiable
 - **Problem** - only hard constraints allowed
- **Observation** Previous approaches backtrack *only* when sub-optimality is proven

Adopt: Asynchronous Distributed Optimization

First key idea -- Weak backtracking

- Adopt's **backtrack condition** – when lower bound gets too high

Why lower bounds?


- allows asynchrony
- allows soft constraints
- allows quality guarantees

Any downside?

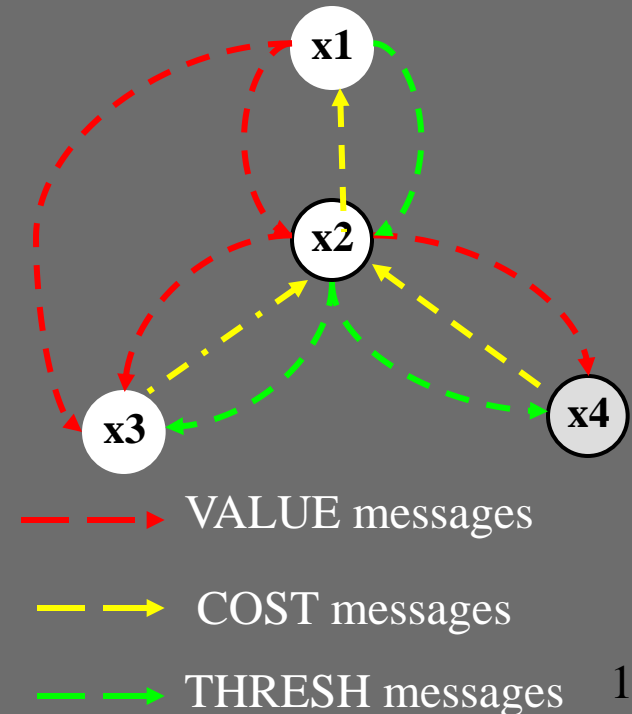
- backtrack **before** sub-optimality is proven
- solutions need revisiting
 - **Second key idea** -- Efficient reconstruction of abandoned solutions

Adopt Algorithm

- Agents are ordered in a tree
 - constraints between ancestors/descendents
 - no constraints between siblings




- Basic Algorithm:
 - choose value with min cost
 - Loop until **termination-condition true**:
 - When receive message:
 - choose value with min cost
 - send **VALUE** message to descendents
 - send **COST** message to parent
 - send **THRESHOLD** message to child




Weak Backtracking

- Suppose parent has two values, "white" and "black"

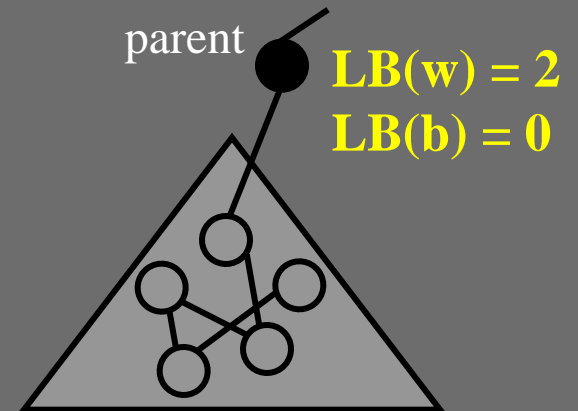
Explore "white" first



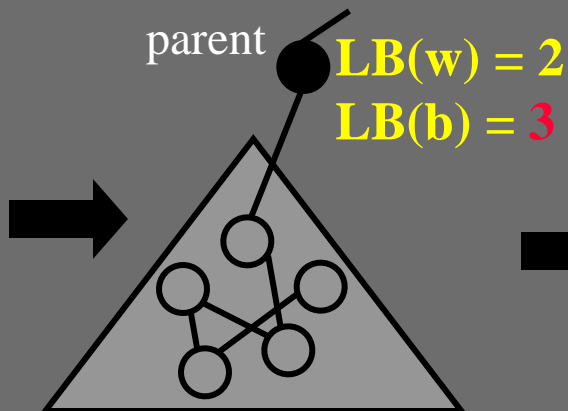
Receive cost msg



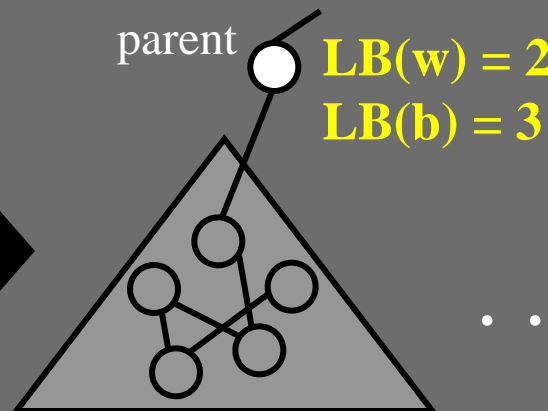
Now explore "black"



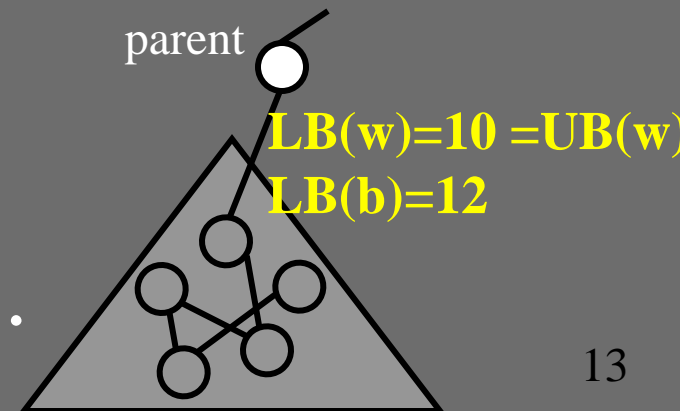
Receive cost msg




Go back to "white"



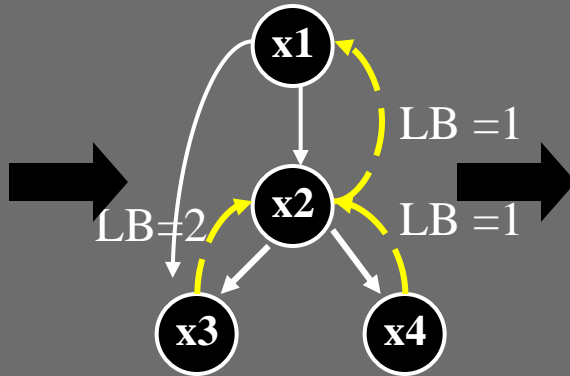
Termination Condition True




Example



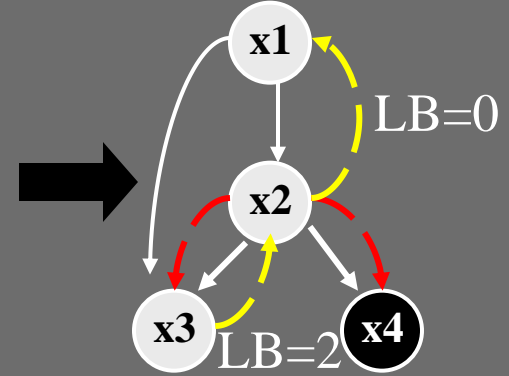
concurrently choose, send to descendents



report lower bounds


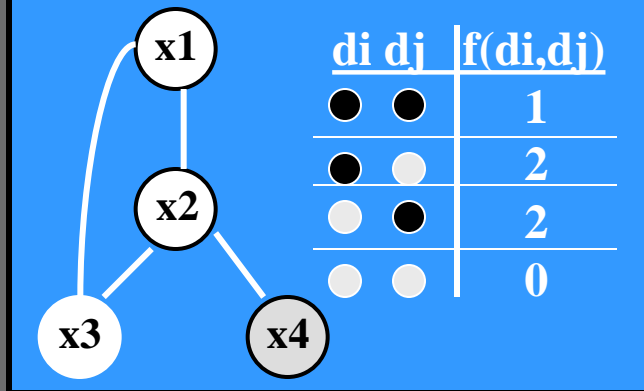


x1 switches value




x2, x3 switch value, report new lower bounds
Note: x3's cost message to x2 is obsolete since x1 has changed value, msg will be disregarded

Constraint Graph



x2, x3 report new lower bounds



optimal solution

Revisiting Abandoned Solutions

Problem


- reconstructing from scratch is **inefficient**
- remembering solutions is **expensive**

Solution


- *backtrack thresholds* – **polynomial space**
- control backtracking to **efficiently** re-search

Parent informs child of lower bound:

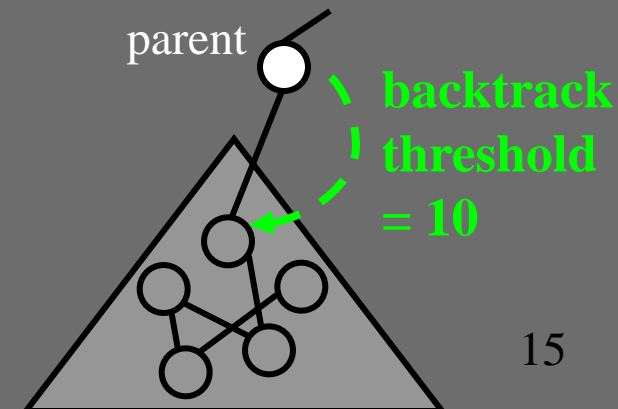
Explore “white” first



Now explore “black”



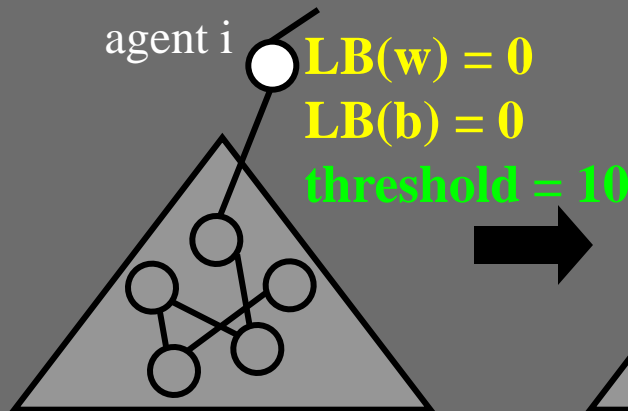
Return to “white”




Backtrack Thresholds

- Suppose agent i received threshold = 10 from its parent


Explore “white” first



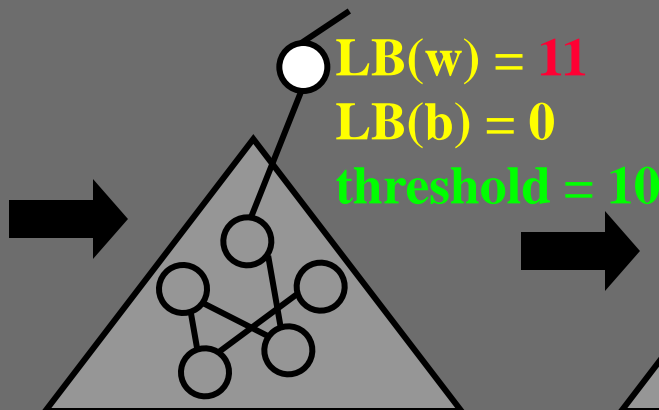
Receive cost msg




Stick with “white”



Receive more cost msgs




Now try black



Key Point: Don't change value until $LB(\text{current value}) > \text{threshold}$.

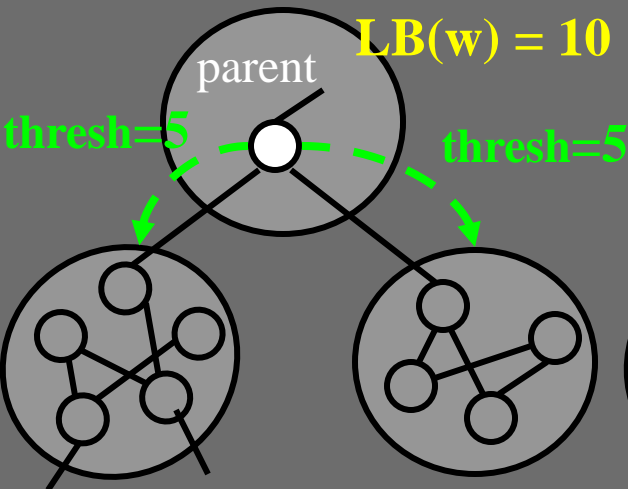
Backtrack thresholds with multiple children




How to correctly subdivide threshold?

Third key idea: Dynamically rebalance threshold

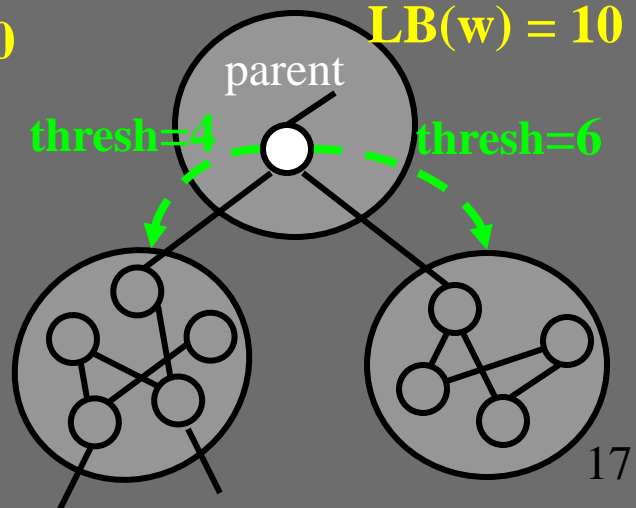
Time T_1



Time T_2



Time T_3



ADOPT (1)

RESET-VARIABLES(d, c)

- 1 $lower-bound[d, c] \leftarrow 0$
- 2 $t[d, c] \leftarrow 0$
- 3 $upper-bound[d, c] \leftarrow \infty$
- 4 $context[d, c] \leftarrow \{\}$

INITIALIZE()

- 1 $threshold \leftarrow 0$
- 2 $received-terminate \leftarrow \text{FALSE}$
- 3 $current-context \leftarrow \{\}$
- 4 $\forall d \in D_i, c \in children \text{ RESET-VARIABLES}(d, c)$
- 5 $x_i \leftarrow d \in D_i$ which minimizes: my cost plus $\sum_{c \in children} lower-bound[d, c]$
- 6 BACKTRACK()



ADOPT(2)

HANDLE-VALUE(j, x_j)

```
1  if  $\neg$  received-terminate
2      then  $current\text{-}context[j] \leftarrow x_j$ 
3           for  $d \in D_i, c \in children$  such that  $context[d, c]$  is
                                     incompatible with  $current\text{-}context$ 
4               do RESET-VARIABLES( $d, c$ )
5           MAINTAIN-THRESHOLD-INVARIANT()
6           BACKTRACK()
```

HANDLE-COST($k, context, lb, ub$)

```
1   $d \leftarrow context[i]$ 
2  delete  $context[i]$ 
3  if  $\neg$  received-terminate
4      then for  $(j, x_j) \in context$  and  $j$  is not my neighbor
5           do  $current\text{-}context[j] \leftarrow x_j$ 
6           for  $d' \in D_i, c \in children$  such that  $context[d, c]$  is
                                     incompatible with  $current\text{-}context$ 
7               do RESET-VARIABLES( $d', c$ )
8           if  $context$  compatible with  $current\text{-}context$ 
9               then  $lower\text{-}bound[d, k] \leftarrow lb$ 
10                   $upper\text{-}bound[d, k] \leftarrow ub$ 
11                   $context[d, k] \leftarrow context$ 
12                  MAINTAIN-CHILD-THRESHOLD-INVARIANT()
13                  MAINTAIN-THRESHOLD-INVARIANT()
14          BACKTRACK()
```



ADOPT (3)

BACKTRACK()

```
1  if threshold =  $\min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$ 
2     then  $x_i \leftarrow \arg \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$ 
3  elseif threshold <  $\text{cost}(x_i) + \sum_{c \in \text{children}} \text{lower-bound}[x_i, c]$ 
4     then  $x_i \leftarrow \arg \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{lower-bound}[d, c]$ 
5   $\forall k \in \text{neighbors} \wedge k$  has lower priority  $k.\text{HANDLE-VALUE}(i, x_i)$ 
6  MAINTAIN-ALLOCATION-INVARIANT()
7  if threshold =  $\min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$  and
   (received-terminate or I am root)
8     then current-context[i]  $\leftarrow x_i$ 
9          $\forall c \in \text{children} c.\text{HANDLE-TERMINATE}(\text{current-context})$ 
10        exit
11  parent.HANDLE-COST(current-context,
    $\min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{lower-bound}[d, c]$ ,
    $\min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$ )
```



ADOPT (4)

MAINTAIN-THRESHOLD-INVARIANT()

- 1 if $threshold < \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{lower-bound}[d, c]$
- 2 then $threshold \leftarrow \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{lower-bound}[d, c]$
- 3 if $threshold > \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$
- 4 then $threshold \leftarrow \min_{d \in D_i} \text{cost}(d) + \sum_{c \in \text{children}} \text{upper-bound}[d, c]$

MAINTAIN-ALLOCATION-INVARIANT()

- 1 while $threshold > \text{cost}(x_i) + \sum_{c \in \text{children}} t[x_i, c]$
- 2 do $chosen \leftarrow c' \in \text{children}$ such that $\text{upper-bound}[x_i, c'] > t[x_i, c']$
- 3 $t[x_i, chosen] \leftarrow t[x_i, chosen] + 1$
- 4 while $threshold < \text{cost}(x_i) + \sum_{c \in \text{children}} t[x_i, c]$
- 5 do $chosen \leftarrow c' \in \text{children}$ such that $\text{lower-bound}[x_i, c'] < t[x_i, c']$
- 6 $t[x_i, chosen] \leftarrow t[x_i, chosen] - 1$
- 7 $\forall_{c \in \text{children}} c.\text{HANDLE-THRESHOLD}(t[x_i, chosen], \text{current-context})$

MAINTAIN-CHILD-THRESHOLD-INVARIANT()

- 1 for $d \in D_i, c \in \text{children}$
- 2 do if $\text{lower-bound}[d, c] > t[d, c]$
- 3 then $t[d, c] \leftarrow \text{lower-bound}[d, c]$
- 4 for $d \in D_i, c \in \text{children}$
- 5 do if $\text{upper-bound}[d, c] < t[d, c]$
- 6 then $t[d, c] \leftarrow \text{upper-bound}[d, c]$



ADOPT (5)

HANDLE-THRESHOLD(t , $context$)

```
1  if  $context$  is compatible with  $current-context$ 
2    then  $threshold \leftarrow t$ 
3         MAINTAIN-THRESHOLD-INVARIANT()
4         BACKTRACK()
```


HANDLE-TERMINATE($context$)

```
1   $received-terminate \leftarrow \text{TRUE}$ 
2   $current-context \leftarrow context$ 
3  BACKTRACK()
```




Evaluation of Speedups

GraphColor, Link Density 2



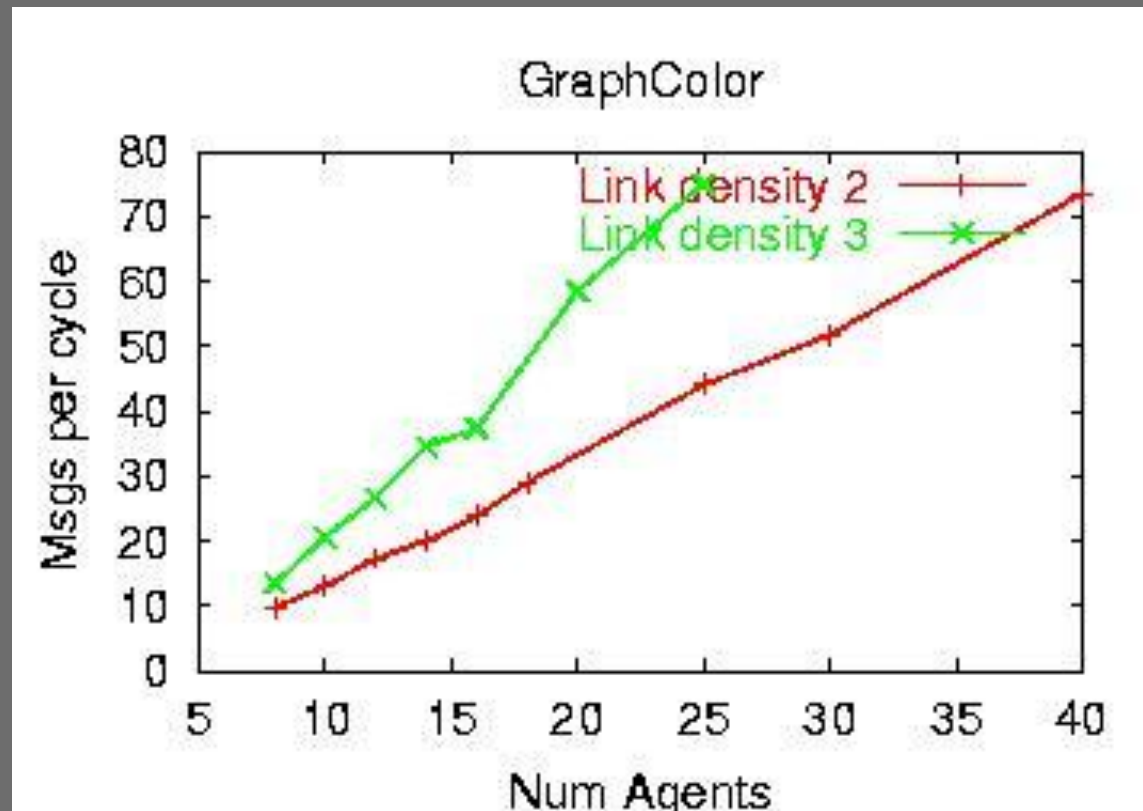
GraphColor, Link Density 3



Conclusions

- Adopt's lower bound search method and parallelism yields significant efficiency gains
- Sparse graphs (density 2) solved **optimally, efficiently** by Adopt.

Number of Messages



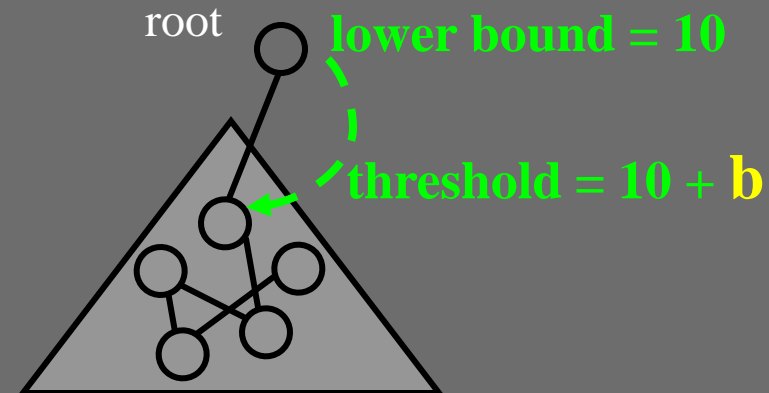
Conclusion

- Communication grows linearly
 - only local communication (no broadcast)

Bounded error approximation

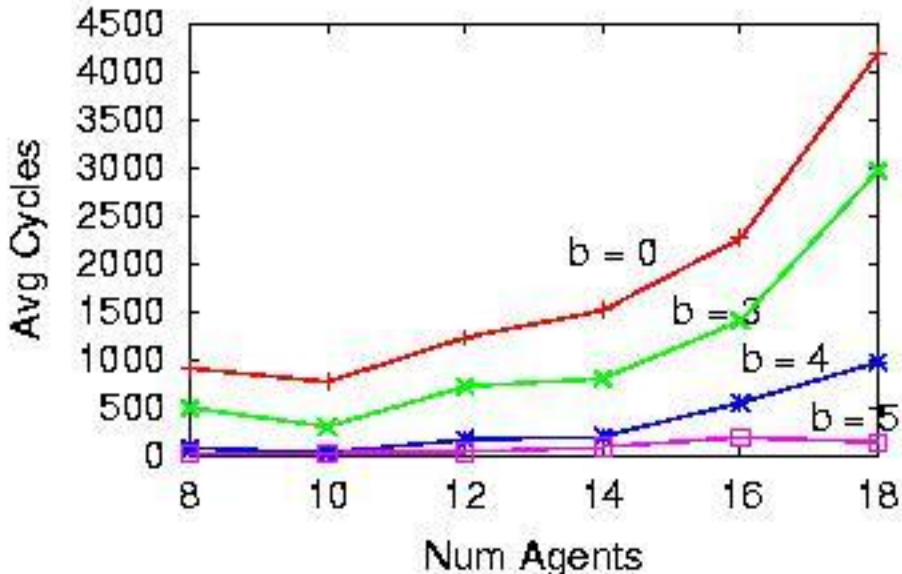
- *Motivation* Quality control for approximate solutions
 - *Problem* User provides error bound **b**
 - *Goal* Find any solution **S** where
$$\text{cost}(\mathbf{S}) \leq \text{cost}(\text{optimal soln}) + \mathbf{b}$$
-

- *Fourth key idea:* Adopt's **lower-bound based** search method naturally leads to bounded error approximation!




Evaluation of Bounded Error

GraphColor, Link Density 3



GraphColor, Link Density 3 (18 agents)



Conclusion

- Time-to-solution decreases as b is increased.
- Plus: Guaranteed worst-case performance!

Optimality of Adopt

- For finite DCOPs with binary non-negative constraints, Adopt is **guaranteed to terminate** with the **globally optimal solution**



Adopt summary – Key Ideas

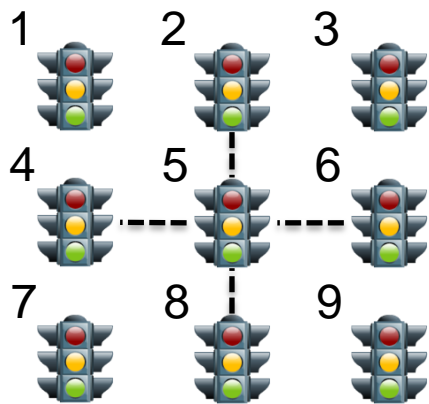
- **Optimal, asynchronous** algorithm for DCOP
 - polynomial space at each agent
- Weak Backtracking
 - **lower bound** based search method
 - Parallel search in independent subtree
- Efficient reconstruction of abandoned solutions
 - **backtrack** thresholds to control backtracking
- Bounded error approximation
 - sub-optimal solutions **faster**
 - **bound** on worst-case performance



Traffic Light Control DCOP

Formulation

- $\rho_{i,j}$ density of vehicle in the lane $i \rightarrow j$
 $\beta_{i,j} = \rho_{i,j} / \sum_k \rho_{k,j}$ fraction of traffic at intersection j coming from i
 $\gamma_{i,j}$ 1.5 (if i and j synchronized), 2 otherwise
 $\tau_{i,j} \in \{0, 1, 1.5, 2\}$ degree of alignment of synchronization with traffic



Plan run by agent i	Plan run by agent j	τ
\oplus	\oplus	0
\ominus	\oplus	1
\oplus	\ominus	1.5
\ominus	\ominus	2

- + synchronized in direction of max traffic
- synchronized in other directions


Constraint cost

$$f_{i,j} = \beta_{i,j} \tau_{i,j} \gamma_{i,j}$$

from: R. Junges and A. L. C. Bazzan. *Evaluating the performance of DCOP algorithms in a real world, dynamic problem*. In AAMAS 2008, pages 599–606, 2008



Results - Complexity



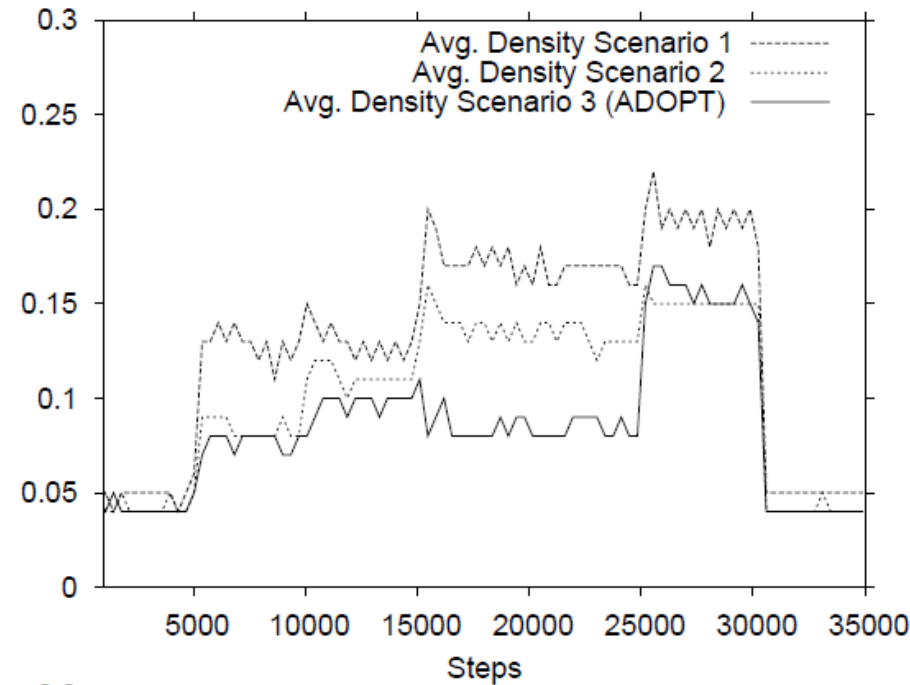
Alg.	Nb. of Msgs	Msgs size (bytes)	Total (MB)
ADOPT	358962.4 ± 1690.4	125.40 ± 20.79	43958
OptAPO	4015.66 ± 68.70	93.45 ± 8.15	366
DPOP	322	69286.78 ± 721.12	21787



Results - Quality

Net.	Experiment	Stopped Veh.	Density
3x3	Case 1	10.21 ± 2.28	0.29 ± 0.06
	Case 2	6.72 ± 1.37	0.23 ± 0.03
	Case 3 (ADOPT)	1.88 ± 0.53	0.11 ± 0.02
	Case 3 (OptApo)	1.97 ± 0.60	0.11 ± 0.01
	Case 3 (DPOP)	2.42 ± 0.46	0.13 ± 0.02
5x5	Case 1	9.22 ± 2.36	0.28 ± 0.04
	Case 2	6.95 ± 1.82	0.23 ± 0.03
	Case 3 (ADOPT)	1.78 ± 0.80	0.11 ± 0.01
	Case 3 (OptApo)	1.90 ± 0.59	0.12 ± 0.01
	Case 3 (DPOP)	3.08 ± 0.40	0.14 ± 0.02
7x7	Case 1	8.91 ± 2.29	0.27 ± 0.04
	Case 2	5.82 ± 1.31	0.21 ± 0.02
	Case 3 (ADOPT)	1.97 ± 0.50	0.12 ± 0.01
	Case 3 (OptApo)	1.91 ± 0.49	0.12 ± 0.01
	Case 3 (DPOP)	2.90 ± 0.62	0.14 ± 0.01
9x9	Case 1	8.05 ± 2.17	0.14 ± 0.01
	Case 2	5.35 ± 0.97	0.10 ± 0.01
	Case 3 (ADOPT)	2.04 ± 0.77	0.12 ± 0.01
	Case 3 (OptAPO)	1.98 ± 0.85	0.11 ± 0.01
	Case 3 (DPOP)	2.88 ± 0.72	0.13 ± 0.02

Avg. Density - 49 agents



Case 1 = worst case

Case 2 = static optimum

Case 3 = dynamic optimization



Conclusion

- Distributed constraint optimization is a general widely applicable model
- **Optimal (complete) asynchronous** algorithms exist for DCOPs with **binary, non-negative** constraints
 - ADOPT, OptAPO, DPOP
 - Different computation and communication complexity profiles
- ADOPT
 - optimal, asynchronous algorithm for DCOP
 - polynomial space at each agent
- Reading:
 - [Vidal] – Chapter 2
 - P. J. Modi et al. ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees. *Artificial Intelligence Journal*, 161(1-2):149–180, January 2005.

