Message passing algorithms

A#: a distributed A*

Distributed Cost-Optimal Planning

Loïg Jezequel

ENS Cachan Bretagne

November 13, 2012

Message passing algorithms

A#: a distributed A 00000000 Conclusion 00

Planning problems: overview



Goal

Find a *plan*: a sequence of actions (with minimal cost) moving the system from its initial state to one of its goal states

Message passing algorithms

A#: a distributed A'

Conclusion 00

Planning problems: search in graphs



A#: a distributed A* 00000000

Planning problems: resolution

Heuristic search

- A*-like algorithms: Hart et al. 1968
- Various heuristics: Bonet and Geffner 2001, Helmert et al. 2007, Karpas and Domshlak 2009, ...

Parallelism of actions (concurrency)

- GRAPHPLAN: Blum and Furst 1995
- Petri net unfolding: Hickmott et al. 2007, Bonet et al. 2008

Split problems into subproblems (Factored planning)

Amir and Engelhardt *2003*, Brafman and Domshlak 2006, Brafman and Domshlak 2008

Message passing algorithms

A#: a distributed A 00000000 Conclusion 00

Factored planning: principles



Each component is a planning problem with its own resources and actions

Goal

Find a set of *compatible* local plans: they can be *interleaved* into a global plan

Message passing algorithms

A#: a distributed A 00000000 Conclusion 00

Factored planning: principles



Each component is a planning problem with its own resources and actions

The components interact by resources and/or actions

Goal

Find a set of *compatible* local plans: they can be *interleaved* into a global plan

A#: a distributed A'

Factored planning: our contribution

Prior to this thesis, reasoning on the number of synchronizations:

- Absence of solution *can not be detected*
- Cost-optimality of plans *can not be achieved*

Our contribution

Two new approaches to *factored planning*, allowing to find *cost-optimal* plans with *distributed* algorithms

Top-down approach

Successive restrictions of the sets of local plans

Bottom-up approach

Progressive construction of a local plan per component

Message passing algorithms

A#: a distributed A'

Conclusion

Top-down approach

Factored cost-optimal planning using message passing algorithms

Message passing algorithms

A#: a distributed A'

Conclusion

Centralized planning problem = weighted automaton



Set of actions $\boldsymbol{\Sigma}$

The words are the plans

The words with minimal cost are the cost-optimal plans

Goal

Find a minimal cost word in a weighted automaton

Message passing algorithms

A#: a distributed A 00000000 Conclusion

Factored planning problem

Components are weighted automata

They interact by their shared actions: formalization using the notion of *synchronous product*

Goal

In $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, find a tuple (w_1, \ldots, w_n) of words which are all *compatible* and *minimize* the sum of their cost, *without computing* \mathcal{A}





Message passing algorithms

A#: a distributed A*

Conclusion

Factored planning problem: example



Centralized plans: $\beta d\gamma$ and $d\beta\gamma$ Factored/distributed/concurrent plan: $(\beta, \beta\gamma, d\gamma)$

Message passing algorithms

A#: a distributed A'

Conclusion

Projection: from global plans to local plans

Projection reduces a *global plan* to the actions of *a particular component*

 $\begin{array}{c} A \\ s_1 \\ b, 1 \\ s_3 \\ s_3 \\ \hline a, 0 \\ \end{array} \begin{array}{c} a, 0 \\ s_4 \\ s$

 $\Pi_{\Sigma'}$ corresponds to:

- $\textcircled{\ } \textbf{ Replace each action not in } \boldsymbol{\Sigma}' \textbf{ by } \boldsymbol{\varepsilon}$
- **2** Perform ε -reduction (to the left)
- (Minimize)



MPA: computing $\Pi_{\Sigma_i}(\mathcal{A})$ without computing \mathcal{A}

Central properties of the projection of $\mathcal{A} = \mathcal{A}_1 imes \cdots imes \mathcal{A}_n$

- any cost-optimal word w of A can be projected into a cost-optimal word w_i of Π_{Σi}(A), moreover c(w) = c_i(w_i)
- any cost-optimal word w_i of Π_{Σ_i}(A) is the projection of a cost-optimal word w of A, moreover c_i(w_i) = c(w)

Consequence

Taking the minimal cost word in each $\Pi_{\Sigma_i}(\mathcal{A})$ gives a cost-optimal global plan (hypothesis: it is unique)

Building the $\Pi_{\Sigma_i}(\mathcal{A})$ by local computations

Successive refinements of the A_i from the constraints imposed by their neighbours

Message passing algorithms

A#: a distributed A'

How to get the $\Pi_{\Sigma_i}(\mathcal{A})$: the message passing algorithms

Fundamental property (conditional independence)

$$\mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_1 \times \mathcal{A}_2) \equiv_{\mathcal{L}} \mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_1) \times \mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_2)$$

Application:

$$\begin{array}{rcl} \mathcal{A}_1 & & & \Sigma_1 \cap \Sigma_2 & & \mathcal{A}_2 & & \Sigma_2 \cap \Sigma_3 \\ \\ \Pi_{\Sigma_1}(\mathcal{A}) & = & & \Pi_{\Sigma_1}(\mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{A}_3) \\ & & \equiv_{\mathcal{L}} & & \Pi_{\Sigma_1}(\mathcal{A}_1) \times \Pi_{\Sigma_1}(\mathcal{A}_2 \times \mathcal{A}_3) \\ & & \equiv_{\mathcal{L}} & & \mathcal{A}_1 \times \Pi_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_2 \times \mathcal{A}_3) \\ & & \equiv_{\mathcal{L}} & & \mathcal{A}_1 \times \Pi_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_2 \times \Pi_{\Sigma_2 \cap \Sigma_3}(\mathcal{A}_3)) \end{array}$$

Message passing algorithms

A#: a distributed A*

How to get the $\Pi_{\Sigma_i}(\mathcal{A})$: the message passing algorithms

Fundamental property (conditional independence)

$$\mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_1 \times \mathcal{A}_2) \equiv_{\mathcal{L}} \mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_1) \times \mathsf{\Pi}_{\Sigma_1 \cap \Sigma_2}(\mathcal{A}_2)$$

Application:



Message passing algorithms

A#: a distributed A 00000000 Conclusion 00





Message passing algorithms

A#: a distributed A 00000000 Conclusion 00





Message passing algorithms

A#: a distributed A 00000000 Conclusion 00

Example



Message passing algorithms

A#: a distributed A'

Conclusion

Message passing algorithms: main results generalized¹



Theorem

If $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ has a *tree shaped interaction* graph, the message passing algorithm *converges* and returns $\mathcal{A}'_i \equiv_{\mathcal{L}} \prod_{\Sigma_i} (\mathcal{A})$ for each \mathcal{A}_i

¹Eric Fabre and Loïg Jezequel, *Distributed Optimal Planning: An Approach by Weighted Automata Calculus*, CDC 2009

Message passing algorithms

A#: a distributed A'

Conclusion

Distoplan: presentation

Distoplan

C++ implementation of the message passing on weighted automata, on top of $openFST^2$ and the HSP^* 's parser³

A benchmark: philosophers from IPC4



²http://www.openfst.org/
³Patrik Haslum, 4th IPC Booklet,2004

Message passing algorithms

A#: a distributed A'

Conclusion

Distoplan: presentation

Distoplan

C++ implementation of the message passing on weighted automata, on top of $openFST^2$ and the HSP^* 's parser³

A benchmark: philosophers from IPC4



²http://www.openfst.org/
³Patrik Haslum, 4th IPC Booklet,2004

Message passing algorithms

A#: a distributed A'

Conclusion

Distoplan⁴



Conclusions

- Practical solving of planning problems
- Can be more efficient than centralized search
- Difficulty: find decompositions

⁴Eric Fabre, Loïg Jezequel, Patrik Haslum, and Sylvie Thiébaux, Cost-Optimal Factored Planning: Promises and Pitfalls, ICAPS 2010

Message passing algorithms

A#: a distributed A 00000000 Conclusion

Extension 1: read arcs in networks of automata⁵



⁵Loïg Jezequel and Eric Fabre, *Networks of Automata with Read Arcs: A Tool for Distributed Planning*, IFAC World Congress 2011

Message passing algorithms

A#: a distributed A'

Conclusion

Extension 1: read arcs in networks of automata⁵



Read arcs

Automata \rightarrow automata with readings/writings on transitions

Theorem

The message passing algorithms extend to this setting with minor modifications

⁵Loïg Jezequel and Eric Fabre, *Networks of Automata with Read Arcs: A Tool for Distributed Planning*, IFAC World Congress 2011

Message passing algorithms

A#: a distributed A'

Conclusion 00

Extension 2: turbo planning⁶

Starting point

When interaction graphs contain cycles:



Existing solution

Tree decomposition of graphs:

- Not all parameters taken into account
- Tree-width can be huge

⁶Loïg Jezequel and Eric Fabre, *Turbo Planning*, WODES 2012



Planning problems	Message passing algorithms	A#: a distributed A* 00000000	C o
Extension 2	: turbo planning ⁶		

nclusion

Starting point

When interaction graphs contain cycles:



Turbo planning

Ignore cycles and perform *approximate planning*

Result: \mathcal{A}'_i such that $\mathcal{L}(\prod_{\Sigma_i}(\mathcal{A})) \subseteq \mathcal{L}(\mathcal{A}'_i) \subseteq \mathcal{L}(\mathcal{A}_i)$

⁶Loïg Jezequel and Eric Fabre, *Turbo Planning*, WODES 2012

Extension 2: convergence issues of turbo planning

As a constraint solving problem

- Convergence in (possibly) infinite time
- Convergence in finite time for words of small length

As an optimization problem

- Costs diverge in general
- Normalization:
 - Costs of optimal paths stabilize
 - Costs of other paths still diverge

Message passing algorithms

A#: a distributed A*

Conclusion

Extension 2: experiments on turbo planning



Results

- Fast convergence (always less than 5 iterations)
- Promising quality of the solutions found (70% of solutions cost less than 10% more than the optimal)

Message passing algorithms ○○○○○○○○○

A#: a distributed A*

Conclusion

Extension 2: experiments on turbo planning





Results

- Fast convergence (always less than 5 iterations)
- Promising quality of the solutions found (70% of solutions cost less than 10% more than the optimal)

Open question

Theoretical explanation of this efficiency

Message passing algorithms

A#: a distributed A*

Conclusion 00

Bottom-up approach

Cost-optimal planning using a distributed version of A*

Message passing algorithms

A#: a distributed A* ●○○○○○○○ Conclusion

A*: a best-first search algorithm



Rank of a node

Most promising node: $s^* = argmin_s(g(s) + h(s))$ Always expand from s^* first

Message passing algorithms

A#: a distributed A* ●○○○○○○○ Conclusion

A*: a best-first search algorithm



Rank of a node

Most promising node: $s^* = argmin_s(g(s) + h(s))$ Always expand from s^* first

Message passing algorithms

A#: a distributed A* ●○○○○○○○ Conclusion 00

A*: a best-first search algorithm



Termination

When s_f is the most promising node

Message passing algorithms

A#: a distributed A* 0 = 0 = 0 = 0

Conclusion 00

A#: intuition



Goal

A pair of compatible paths

Message passing algorithms

A#: a distributed A* ○●○○○○○○ Conclusion 00

A#: intuition



Goal

A pair of compatible paths

Idea

Parallel constrained searches

Agent φ

A* with information from $\overline{\varphi}$

A#: a distributed A* ○○●○○○○○

Compatible final states

The problem

- Two automata (not sharing actions)
- A colouring function on final states

Goal: find a path in each automaton such that:

- They both reach final states of the same colour
- The sum of their costs is minimal among such paths



Message passing algorithms

A#: a distributed A*

Conclusion 00

Compatible final states: ranking



Locally

One heuristic h per color

Message passing algorithms

A#: a distributed A*

Conclusion 00

Compatible final states: ranking



Locally

One heuristic h per color

Externally

information \overline{H} from $\overline{\varphi}$

Rank(s)

$$g(s) + \min_{c}(h(s, c) + \overline{H}(c))$$

Message passing algorithms

A#: a distributed A* ○○○○●○○○ Conclusion 00

Compatible final states: termination



\overline{G}

Achieved best costs for colors

Termination

s a goal with color c_s such that: $g(s) + \overline{G}(c_s)$ lowest rank

Message passing algorithms

A#: a distributed A* ○○○○○●○○ Conclusion 00

Compatible final states: results⁷

Theorem

When executed by φ on any CFS problem:

- A# terminates,
- A# is sound,
- A# is complete,

assuming that φ has access to \overline{G} and \overline{H}

⁷Loïg Jezequel and Eric Fabre, A#: A Distributed Version of A^* for Factored Planning, CDC 2012

Message passing algorithms

A#: a distributed A* ○○○○○●○○ Conclusion 00

Compatible final states: results⁷

Theorem

When executed by φ on any CFS problem:

- A# terminates,
- A# is sound,
- A# is complete,

assuming that φ has access to \overline{G} and \overline{H}

Theorem

 \overline{G} and \overline{H} can be computed by $\overline{\varphi}$ along its own execution of A#

⁷Loïg Jezequel and Eric Fabre, $A \neq: A$ Distributed Version of A^* for Factored Planning, CDC 2012

A#: a distributed A* ○○○○○○●○

A#: from CFS to factored planning

CCP and factored planning with two components

- Colour = sequence of (shared) actions
- Number of colours cannot be bounded locally

Consequence: computation of h, H, and G more difficult



A#: a distributed A* ○○○○○○●○

A#: from CFS to factored planning

CCP and factored planning with two components

- Colour = sequence of (shared) actions
- Number of colours cannot be bounded locally

Consequence: computation of h, H, and G more difficult



Theorem: termination

As soon as the considered factored planning problem has a solution, A# terminates

Planning	problems

Message passing algorithms

A#: a distributed A* ○○○○○○○● Conclusion

A#: extension to larger interaction graphs

From the point of view of agent φ_i , any factored planning problem has two components:

 \mathcal{A}_i — $\prod_{k \neq i} \mathcal{A}_k$

A#: a distributed A* ○○○○○○○●

A#: extension to larger interaction graphs

From the point of view of agent φ_i , any factored planning problem has two components:

Theorem

If the interaction graph is a tree \overline{H} and \overline{G} can be constructed using only information (messages) from the neighbours of φ_i



Message passing algorithms

A#: a distributed A 00000000 Conclusion

Conclusion and perspectives

Planning	problems

Message passing algorithms

A#: a distributed A*

Conclusion

Main contribution

Two planning algorithms allowing: *distributed planning* and *cost-optimal planning*

First approach

- Message passing algorithms + weighted automata calculus
- Implementation in Distoplan
- Reading variables
- Turbo planning (approximate methods for factored planning)

Second approach

- Distributed version of A*
- Proof of validity and implementability

Message passing algorithms

A#: a distributed A 00000000 Conclusion ○●

Perspectives

One needs *benchmarks* for factored planning

How to automatically decompose planning problems?

Is it possible to benefit from *local concurrency*?

When/why turbo planning works?



