

The MADP Toolbox 0.3

Frans A. Oliehoek
f.a.oliehoek@uva.nl

Matthijs T. J. Spaan
m.t.j.spaan@tudelft.nl

Philipp Robbel
robbel@mit.edu

João V. Messias
jmessias@isr.ist.utl.pt

May 19, 2014

Abstract

This is the manual documentation accompanying the version 0.3 release of the multiagent decision process (MADP) Toolbox. It is meant as a first introduction to the organization of the toolbox, and tries to clarify the approach taken to certain implementation details. In addition, it covers a few typical use cases. While suitable for a first introduction, this manual is by no means a substitution for the automatically generated (doxygen) reference.

Contents

1	Introduction	2
I	MADPs and Overview of the Toolbox	2
2	MADPs and Basic Notation	2
2.1	Discrete Time MASs	2
2.2	Basic MADP Components	3
2.3	Histories	3
3	Overview of the MADP Toolbox	4
3.1	MADP Libraries	4
3.1.1	The Base Library (<code>libMADPBase</code>)	4
3.1.2	The Parser Library (<code>libMADPParser</code>)	4
3.1.3	The Support Library (<code>libMADPSupport</code>)	4
3.1.4	The Planning Library (<code>libMADPPlanning</code>)	5
3.2	MADP Directory Structure	6
4	Using the MADP toolbox: An Example	6
5	Typical Uses	7
5.1	One-Shot Decision Making	7
5.2	Sequential Planning Algorithms	7
5.2.1	MultiAgentDecisionProcessInterface and PlanningUnits	7
5.2.2	Multiagent Planning	9
5.2.3	Planning for a Single Agent	10
5.3	Simulation and Reinforcement Learning	10
5.3.1	Simulations	10
5.3.2	The Agents Hierarchy	10
5.3.3	Reinforcement Learning	11
5.4	Specifying problems: file formats, etc.	11
5.4.1	Using the OpenMarkov Graphical Editor	11

5.4.2	Specifying & Parsing <code>.pomdp</code> & <code>.dpomdp</code> files	11
5.4.3	Specifying problems as a sub-class	12
6	Indices for discrete models	12
6.1	Enumeration of Joint Actions and Observations	14
6.2	Enumeration of (Joint) Histories	14
6.2.1	Observation Histories	14
6.2.2	Action Histories	15
6.2.3	Action-Observation Histories	15
6.2.4	Joint Histories	17
7	Joint Beliefs and History Probabilities	17
7.1	Theory	17
7.2	Implementation	18
8	Policies	18
9	The ProbModelXML format	19
9.1	Using OpenMarkov to design factored problems	20
9.2	Designing Event-Driven Models	21

1 Introduction

This text describes the Multiagent decision process (MADP) Toolbox, which is a software toolbox for scientific research in decision-theoretic planning and learning in multiagent systems (MASs). We use the term MADP to refer to a collection of mathematical models for multiagent planning: multiagent Markov decision processes (MMDPs) [6], decentralized MDPs (Dec-MDPs), decentralized partially observable MDPs (Dec-POMDPs) [2], partially observable stochastic games (POSGs) [10], etc.

The toolbox is designed to be rather general, potentially providing support for all these models, although so far most effort has been put in planning algorithms for discrete Dec-POMDPs. It provides classes modeling the basic data types of MADPs (e.g., action, observations, etc.) as well as derived types for planning (observation histories, policies, etc.). It also provides base classes for planning algorithms and includes several example applications using the provided functionality. For instance, applications that use JESP or brute-force search to solve problems (specified as `.dpomdp` files) for a particular planning horizon. In this way, Dec-POMDPs can be solved directly from the command line. Furthermore, several utility applications are provided, for instance one which empirically determines a joint policy's control quality by simulation.

This document is split in two parts. The first presents a mathematical model of the family of MADPs, which also introduces notation, gives a high-level overview of the toolbox and an example of how to use it. In the second part, some more specific design choices and mechanisms are explained.

Part I

MADPs and Overview of the Toolbox

2 MADPs and Basic Notation

As mentioned, MADPs encompass a number of different models. Here we briefly introduce the components of these mathematical models and some basic notation. For a more extensive introduction to these models, see, e.g., [25, 16].

2.1 Discrete Time MASs

An MADP is often considered for a particular finite number of discrete time steps t . When searching policies (*planning*) that specify h actions, this number is referred to as the (planning-) *horizon*. So typically

we look at time steps:

$$t = 0, 1, 2, \dots, h-2, h-1.$$

At each time step:

- The world is in a specific state $s \in \mathcal{S}$.
- Each agent receives an individual observation: a (noisy) observation of the environment's state.
- The agents take an action.

The individually selected actions form a joint action. After such a joint action, the system jumps to the next time step. In this jump the system's state may change stochastically, and this transition is influenced by the taken joint action. In MADPs (such as the Dec-POMDP), there are transition and observation functions describing the probability of state transitions and observations.

2.2 Basic MADP Components

More formally, a multiagent decision process (MADP) consists of a subset of the following components:

- A finite set of m agents.
- \mathcal{S} is a finite set of world states.
- The set $\mathcal{A} = \times_i \mathcal{A}_i$ is the set of *joint actions*, where \mathcal{A}_i is the set of actions available to agent i . Every time step, one joint action $\mathbf{a} = \langle a_1, \dots, a_m \rangle$ is taken. Agents do not observe each other's actions.
- $\mathcal{O} = \times_i \mathcal{O}_i$ is the set of joint observations, where \mathcal{O}_i is a finite set of observations available to agent i . Every time step one joint observation $\mathbf{o} = \langle o_1, \dots, o_m \rangle$ is received, from which each agent i observes its own component o_i .
- $b^0 \in \Delta(\mathcal{S})$, is the initial state distribution at time $t = 0$.¹
- A transition function that specifies the probabilities $P(s'|s, \mathbf{a})$.
- An observation function that specifies the probabilities $P(\mathbf{o}|\mathbf{a}, s')$.
- A set of reward functions $\{R_i\}$ that specify the payoffs of the agents.

The partially observable stochastic game (POSG) is the most general model in the MADP family. Dec-POMDPs are similar, but all the agents receive the same reward, so only 1 reward function is needed.

Unless stated otherwise, we use superscript for time indices. I.e., a_i^t denotes the agent i 's action at time $t = 2$.

2.3 Histories

Let us more formally consider what the history of the process is. An MADP history of horizon h specifies h time steps $t = 0, \dots, h-1$. At each of these time steps, there is a state s^t , joint observation \mathbf{o}^t and joint action \mathbf{a}^t . Therefore, when the agents will have to select their k -th actions (at $t = k-1$), the history of the process is a sequence of states, joint observations and joint actions, which has the following form:

$$(s^0, \mathbf{o}^0, \mathbf{a}^0, s^1, \mathbf{o}^1, \mathbf{a}^1, \dots, s^{k-1}, \mathbf{o}^{k-1}).$$

Here s^0 is the initial state, drawn according to the initial state distribution b^0 . The initial joint observation \mathbf{o}^0 is usually assumed to be the empty joint observation: $\mathbf{o}^0 = \mathbf{o}_\emptyset = \langle o_{1,\emptyset}, \dots, o_{m,\emptyset} \rangle$. Consequently in the MADP toolbox there is no initial observation.

An agent can only observe his own actions and observations. Therefore we introduce notions of histories from the perspective of an agent. We start with the *action-observation history* of agent i at time step t :

$$\vec{\theta}_i^t = (a_i^0, o_i^1, a_i^2, \dots, o_i^{t-1}, a_i^{t-1}, o_i^t)$$

¹We use $\Delta(X)$ to denote the infinite set of probability distributions over the finite set X .

note that the choice points for the agents are right before the action:

$$\vec{\theta}_i^k = \left(\uparrow_{t=0} a_i^0, o_i^1, \uparrow_{t=1} a_i^2, \dots, o_i^{k-1}, \uparrow_{t=k-1} a_i^{k-1}, o_i^k \right)$$

Therefore, when we write $\vec{o}_i^t = (o_i^1, \dots, o_i^{t-1}, o_i^t)$ for agent i 's *observation history* at time step t and $\vec{a}_i^t = (a_i^0, a_i^1, \dots, a_i^{t-1})$ for the *action history* of agent i at time step t . We can thus redefine the action-observation history as: $\vec{\theta}_i^t \equiv \langle \vec{o}_i^t, \vec{a}_i^t \rangle$. For time step $t = 0$, we have that $\vec{a}_i^0 = (()) = \vec{a}_\emptyset$ and $\vec{o}_i^0 = (()) = \vec{o}_\emptyset$ are empty sequences.

3 Overview of the MADP Toolbox

The MADP framework consists of several parts, grouped in different libraries. A brief overview of these ‘MADP libraries’ is given in Section 3.1. Also, there are a number of other libraries and software included to realize compilation with a minimum of effort. Therefore, Section 3.2 gives an overview of the entire directory structure.

3.1 MADP Libraries

The main part of MADP is the set of core libraries. These are briefly discussed here.

3.1.1 The Base Library (libMADPBase)

The base library is the core of the MADP toolbox. It contains:

- A representation of the basic elements in a decision process such as states, (joint) actions and observations.
- A representation of the transition, observation and reward models in a multiagent decision process. These models can also be stored in a sparse fashion.
- A uniform representation for MADP problems, which provides an interface to a problem’s model parameters.
- Auxiliary functionality regarding manipulating indices, exception handling and printing: `E`, `IndexTools`, `PrintTools`, `StringTools`, `TimeTools`, `VectorTools`. Some project-wide definitions are stored in the `Globals` namespace.

3.1.2 The Parser Library (libMADPParser)

The parser library depends on the base library, and contains a parser for Tony’s `.pomdp` file format as well as for `.dpomdp` files, which is a file format for problem specifications of discrete Dec-POMDPs, as well as a parser for models specified in the `ProbModelXML` format. A set of benchmark problem files in both formats can be found in the `problems/` directory.

The `.dpomdp` syntax is documented in `problems/example.dpomdp`. The format is based on Tony’s `.pomdp` file format, and the formal specification is found in `src/parser/dpomdp.spirit`. The parser uses the Boost Spirit library. Also, parsers for several transition-observation independent models are provided, which are derived from the `.dpomdp` parser.

The `ProbModelXML` format is an XML format and parsed using `libXML2`. The format is covered in detail by Arias et al. [1], and a detailed introduction is given in Section 9.

3.1.3 The Support Library (libMADPSupport)

The support library contains basic data types and support useful for planning, such as:

- A representation for (joint) histories, for storing and manipulating observation, action and action-observation histories.
- A representation for (joint) beliefs, both stored as a full vector as well as a sparse one.

- Functionality for representing (joint) policies, as mappings from histories to actions.
- Shared functionality for discrete MADP planning algorithms, collect in `PlanningUnitMADPDiscrete` and `PlanningUnitDecPOMDPDiscrete`. These classes compute, e.g., (joint) history trees, joint beliefs, and value functions.
- Implementation for various problems:
 - An implementation of the DecTiger problem [15] which does not use `dectiger.dpomdp`, see `ProblemDecTiger`.
 - Also an implementation of the Fire Fighting problem (`ProblemFireFighting`), as well as a factored version (`ProblemFireFightingFactored`) [19].
 - Implementation of factored problem domains `ProblemFireFightingGraph` and `ProblemAloha` [20, 24].
 - Fully observable versions of the firefighting problem: `ProblemFOBSFireFightingFactored` and `ProblemFOBSFireFightingGraph`.
- Functionality for handling command-line arguments is provided by `ArgumentHandlers`.

3.1.4 The Planning Library (`libMADPPanning`)

The planning library depends on the other libraries and contains functionality for planning algorithms, as well as some solution methods. In particular it contains

- MDP solution techniques: value iteration [28].
- POMDP solution techniques: Monahan’s algorithm [13] with incremental pruning [7], Perseus [26].
- Dec-POMDP solution algorithms:
 - Brute Force Search.
 - JESP (exhaustive and dynamic programming variations) [15].
 - Direct Cross-Entropy (DICE) Policy Search [18].
 - GMAA* type algorithms, in particular:
 - * MAA* [29],
 - * k -GMAA* (as well as forward sweep policy computation) [19],
 - * GMAA*-ELSI [20].
 - * GMAA*-Cluster [21] (also called GMAA*-IC [27])
 - * GMAA*-ICE [27].
 - DP-LPC [5]. (the implementation of this method is mostly inside `solvers/DP-LPC.cpp`).
- Functionality for building and solving collaborative Bayesian Games:
 - Random, Brute force search, Alternating Maximization, Cross-entropy optimization, BaGaBaB [22], and Max-Plus for regular CBGs.
 - Random, Non-serial dynamic programming (a.k.a. variable elimination) [20] and Max-Plus [23] for collaborative *graphical* BGs (CGBGs).
- Heuristic Q-functions: Q_{MDP} , Q_{POMDP} , and Q_{BG} [19]. Including ‘hybrid’ representations [27].
- A simulator class to empirically test the control quality of a solution, or perform evaluation of particular types of agents (e.g., reinforcement learning agents).

```

1 #include "ProblemDecTiger.h"
2 #include "JESPExhaustivePlanner.h"
3 int main()
4 {
5     ProblemDecTiger dectiger;
6     JESPExhaustivePlanner jesp(3,&dectiger);
7     jesp.Plan();
8     cout << jesp.GetExpectedReward() << endl;
9     cout << jesp.GetJointPolicy()->SoftPrint() << endl;
10    return(0);
11 }

```

Figure 1: A small example program that runs JESP on the DecTiger problem.

3.2 MADP Directory Structure

path	description
/	Package root.
/config	
/doc	Documentation.
/doc/html	The html documentation generated by doxygen. (perform ‘make htmldoc’ in root.)
/m4	M4 macros used by configure.
/problems	A number of problems in .dpomdp file format.
/src	C++ code
/src/base	The MADP base lib.
/src/boost	Included parts of the boost library (v. 1.42)
/src/include	This contains configuration .h files.
/src/libpomdp-solve	Tony Cassandra’s ‘pomdp-solve’ library. Used for pruning POMDP value vectors.
/src/libDAI	Library for Discrete Approximate Inference by Joris Mooij [14]. Used for max-plus implementation for Collaborative (graphical) Bayesian Games.
/src/parser	The MADP parser lib.
/src/planning	The MADP planning lib.
/src/solvers	Code for a number of executables that use the MADP libs to implement (Dec-)POMDP solvers.
/src/support	The MADP support lib.
/src/utils	Code for a number of executables that perform auxiliary tasks (e.g., printing problem statistics or evaluating a joint policy through simulation).

4 Using the MADP toolbox: An Example

Here we give an example of how to use the MADP toolbox. Figure 1 provides the full source code listing of a simple program. It uses exhaustive JESP to plan for 3 time steps for the DecTiger problem, and prints out the computed value as well as the policy. Line 5 constructs an instance of the DecTiger problem directly, without the need to parse `dectiger.dpomdp`. Line 6 instantiates the planner, with as arguments the planning horizon and a reference to the problem it should consider. Line 7 invokes the actual planning and lines 8 and 9 print out the results.

This is a simple but complete program, and in the distribution (in `src/examples`) more elaborate examples are provided which, for instance, demonstrate the command-line parsing functionality and the use of the `.dpomdp` parser. Furthermore, for each of the solution methods provided there is a program to use it directly.

5 Typical Uses

In this section we elaborate on and provide pointers to examples of typical ways in which the MADP toolbox can be used.

5.1 One-Shot Decision Making

MADP implements one-shot team decision making via Bayesian games with identical payoffs, also referred to as *collaborative Bayesian games*. These are implemented as different sub-classes of `BayesianGameIdenticalPayoffInterface`.

The current version of MADP includes a number of solvers for such collaborative BGs. For instance:

- `BGIP_SolverRandom` provides a random solution.
- `BGIP_SolverBruteForceSearch` a naive enumeration of all joint policies.
- `BGIP_SolverAlternatingMaximization` implements alternating maximization (a best-response hill-climbing).
- `BGIP_SolverCE` implements a cross-entropy optimization procedure [4] for optimizing the joint BG policy (see also [18]).
- `BGIP_SolverBranchAndBound`, the BaGaBaB method from [22] (the name is a bit confusing, really it performs A*).
- `BGIP_SolverMaxPlus` the max-plus solver from [23].

The usage of these solvers is illustrates in `examples/example_RandomBGs.cpp`. Additionally, there also are solvers for collaborative graphical Bayesian games:

- `BGCG_SolverRandom` provides a random solution.
- `BGCG_SolverNonserialDynamicProgramming` provides the exact solution via non-serial dynamic programming [3] (a.k.a. value iteration [9, 11] and bucket elimination [8]).
- `BGCG_SolverMaxPlus` provides a approximate solution via max-plus message passing [23]. This is based on (a older version of) the LibDAI library [14] which is included with MADP.

MADP also includes a (non-identical payoff) `BayesianGame` class. But so far, there is no solver for this class.

5.2 Sequential Planning Algorithms

Even though one-shot decision making is interesting on its own, the focus of the MADP toolbox lies on sequential decision making. Here we give a concise overview of the main components for planning for sequential decision settings.

5.2.1 MultiAgentDecisionProcessInterface and PlanningUnits

Two important sets of classes are those that represent actual multiagent decision process problems and those that represent planners.

The former classes inherit from `MultiAgentDecisionProcessInterface`, as illustrated in Figure 2. The figure indicates the relations between different models such as Dec-POMDPs, and POSGs, and shows that the toolbox separates interface classes from implementation. The figure also illustrates that the code offers opportunities to develop MADPs with continuous states, actions and observations, even though so far development has focused on problems with discrete sets (e.g., as represented by the `DecPOMDPDiscrete` class). Note that a number of included problems of type `FactoredDecPOMDPDiscrete` are shown in Figure 3.

The second important collection of classes pertain to planning. These classes all derive from the `PlanningUnit` base class. Part of this hierarchy is shown in Figure 4, centered around the `PlanningUnitMADPDiscrete` class. This class provides auxiliary functionality (e.g., generation of histories and conversion of history indices) for planners for discrete problems. The figure also shows that

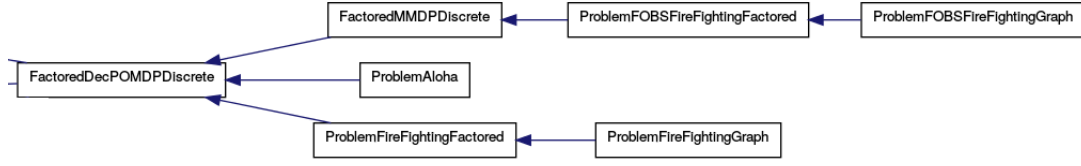


Figure 3: Included problems of type **FactoredDecPOMDPDiscrete**, the top branch are fully observable (a factored MMDP is the fully-observable special case of a factored Dec-POMDP), while the bottom two branches are actual Dec-POMDPs.

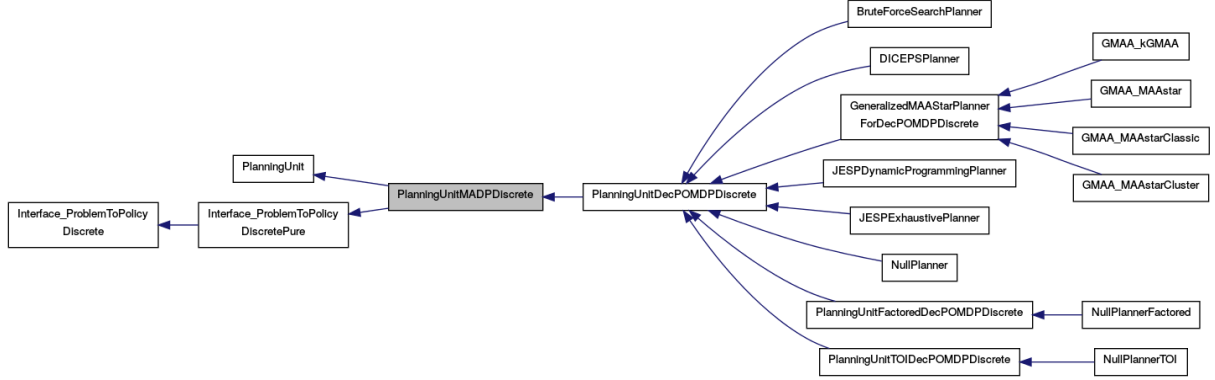


Figure 4: A part of the **PlanningUnit** class hierarchy.

the class implements the so called ‘**Interface_ProblemToPolicyDiscretePure**’. This is the mechanism by which the planner gets to know certain basic information about the problem it will be planning for, for more details, see Section 8.

5.2.2 Multiagent Planning

A typical program that performs multiagent planning has three main components: first, there is an ‘experiment’ (or ‘solver’) file which contains `main()`. This experiment instantiates both the MADP (e.g., a Dec-POMDP), and the planner (e.g., GMAA*), and subsequently performs the actual planning by calling `Plan()`.

An Example: `example_decTigerJESP.cpp`. For instance, let’s look at the `example_decTigerJESP.cpp` file, of which the important content was already shown in Figure 1. The file itself is the ‘experiment’ file that contains `main`. It instantiates an MADP—a **ProblemDecTiger**, see also Figure 2—on line 5. Next, it instantiates a planning unit—a **JESPExhaustivePlanner**, see Figure 4—and calls the `Plan()` method.

Solving fully-observable MADPs. MADP currently implements value iteration in `MDPValueIteration` (for flat models of limited scope) and supports finite and infinite horizon problems. Usage is best demonstrated by an example, such as in the file `./src/examples/MMDP_SolveAndSimulate.cpp`. This file also shows how an offline policy can then be simulated on a specific problem.

For larger, generally factored settings, *every* existing problem, be it a **FactoredMMDPDiscrete** or the fully-observable subset of any class derived from **FactoredDecPOMDPDiscrete**, can be exported into SPUDD format with a call to `FactoredDecPOMDPDiscrete::ExportSpuddFile("filename")`.²

²There currently exists no functionality to load and simulate policies from SPUDD in MADP, however (see the SPUDD package for this functionality).

5.2.3 Planning for a Single Agent

Of course, it is also possible to perform single-agent planning. From a modeling point of view, a single-agent model is just a multiagent one in which there happens to be just one agent. From a solution method perspective, however, this is different: most multiagent planning algorithms are not particularly suited for single agent planning. Methods suitable for single agents currently include `MDPValueIteration` (for MDPs) and `PerseusPOMDPPlanner`, `MonahanPOMDPPlanner` (for POMDPs).

5.3 Simulation and Reinforcement Learning

The toolbox also provides functionality to perform simulations for (teams of) agents interacting in a environment, as well as doing reinforcement learning. For instance, the following command performs 10000 simulations of the MMDP solution for the horizon 5 `GridSmall` problem:³

```
1 src/example$ ./example_MMDP_SolveAndSimulate ../../problems/GridSmall.dpomdp -h5 \  
2 --runs=10000  
3 Instantiating the problem...  
4 ...done.  
5 Avg rewards: < 3.00999 >
```

5.3.1 Simulations

Simulations are performed using the following classes:

- `Simulation` — base class for all simulations.
- `SimulationDecPOMDPDiscrete` — currently the only class that actually implements simulations.
- `SimulationResult` — class that stores the results of simulations.
- `SimulationAgent` — base class for agents that can interact in a simulation.

Currently, there is just one class, `SimulationDecPOMDPDiscrete`, that actually implements simulations. There are two modes in which it works:

1. By giving it a joint policy. In this mode, a Dec-POMDP policy will be simulated (e.g., to empirically test its quality).
2. By giving it a vector of `SimulationAgent` objects. In this mode, the `SimulationDecPOMDPDiscrete` will give all relevant information to each agent, and the agents return back an action.

The second mode can be used to simulate also environments that are not Dec-POMDPs. The trick is that `SimulationDecPOMDPDiscrete` simply provides all relevant the information (e.g., state, taken joint action, and/or received joint observation) to each `SimulationAgent`. For instance, it knows (via function overloading) that if it is dealing with agents of the type `AgentFullyObservable` it should provide them with the entire current state, while it will only give the individual observations to agents of type `AgentLocalObservations`. (e.g., see the different `GetAction` functions in `SimulationDecPOMDPDiscrete.cpp`). Since the simulator does not dictate anything about the inner working of the agents, this framework directly supports (reinforcement) learning agents.

5.3.2 The Agents Hierarchy

As may be clear by now, MADP provides a hierarchy of some different types of `SimulationAgent`. The current hierarchy is shown in Figure 5. It shows the class `AgentDecPOMDPDiscrete`, which is a superclass for all agents that make use of a (special case of a) `PlanningUnitDecPOMDPDiscrete`. It also shows that there are a number of different subclasses of agent: one for teams of agents with shared observations (i.e., for a POMDP or a ‘multiagent POMDP’), one for agents with just local observations (i.e., the ‘real Dec-POMDP setting’), one for agents with full observability, and one for teams of agents with delayed shared observations (i.e., the one-step delayed communication setting [17]).

³Note that this demonstrates the flexibility of the toolbox. Even though `GridSmall` is a Dec-POMDP benchmark, it can be treated as a fully observable (MMDP) problem.

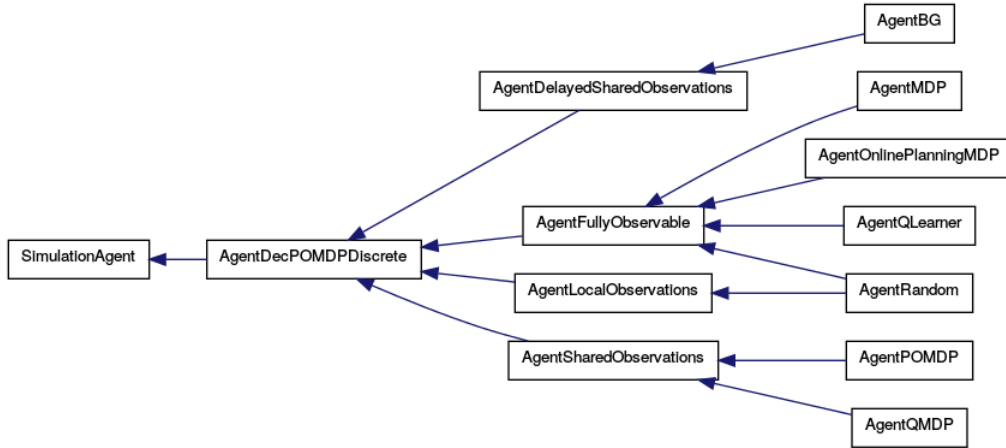


Figure 5: The `SimulationAgent` Hierarchy.

5.3.3 Reinforcement Learning

For *fully-observable* problems, this release of MADP includes a Q-learning agent which learns a joint policy in the joint state & action space. Both ϵ -greedy and Boltzmann exploration methods are currently implemented.

If multiple agents are specified in a problem, team learning can be performed without having to replicate the learning in each individual agent: a single agent can be designated as the only learning agent with the `SetFirstAgent()` call. In this case, only the specified agent learns a (sparse) Q-table while the other agents look up their respective action in the joint table.

Note that we currently do not handle (reinforcement learning style) episode ends. MADP simulations are always run as many times as specified in the `SimulationDecPOMDPDiscrete` horizon parameter. However, episode ends can be modeled with special sink states in the problem formulation, i.e., absorbing states that generate no more rewards until simulation end.

The program `./src/examples/example_MMDP.OnlineSolve.cpp` illustrates the use of the simple (joint) Q-learning agent. Only agent 0 is designated as the learning agent with the `SetFirstAgent()` call so that all other agents refer to its learned Q-table for action selection. The program first computes the off-line policy using value iteration and then performs Q-learning for a number of iterations. Finally, a single row from both resulting Q-tables is compared and displayed. An example run (e.g. on the fully-observable Tiger problem for a discount factor of 0.99) is as follows: `./example_MMDP.OnlineSolve -g 0.99 DT`

5.4 Specifying problems: file formats, etc.

There are three main ways to specify problems in MADP: in the `ProbModelXML` format, as a `.dpomdp` file, or as a sub-class of a suitable `MultiAgentDecisionProcess`.

5.4.1 Using the OpenMarkov Graphical Editor

OpenMarkov (<http://www.openmarkov.org/users.html>) is a GUI editor for creating Factored Models in the `ProbModelXML` format. Using this tool is treated in detail in Section 9.

5.4.2 Specifying & Parsing .pomdp & .dpomdp files

For single-agent POMDPs, it is possible to use Tony Cassandra's `.pomdp` file format, which is specified at <http://www.pomdp.org/code/pomdp-file-spec.shtml>, to specify problems. For non-factored (i.e., typically smaller) multiagent problems, an easy way to specify them is to create a `.dpomdp` file. It is easiest to get an understanding of this format by example: the code for the (in)famous decentralized tiger benchmark is illustrated in Figure 6. The figure clearly shows that there are parts to specify the number of agents, the states, the possible actions and observations, and finally the transition-, observation-

and reward model. A version of this file including comment is available in the `problems/` directory. This also contains `example.dpomdp` which provides even further information.

5.4.3 Specifying problems as a sub-class

While less portable, and arguably more complex, specifying your own problem as a sub-class is the most (run-time and space) efficient and gives you the most flexibility. In MADP, one typically implements a problem by deriving from the appropriate base class. We give a few examples here.

Dec-POMDPs. For instance, to specify a Dec-POMDP with discrete states, actions and observations, one would inherit from `DecPOMDPDiscrete`. This class specifies the state, action and observation spaces, as well as the transition, observation and reward model. All that the derived class needs to do is actually construct these. For an example on how this works, see the `ProblemDecTiger` class.

Factored Dec-POMDPs. Similarly, factored Dec-POMDPs derive from `FactoredDecPOMDPDiscrete`, shown in Figure 3. The class `ProblemAloha` is a good example of a Factored Dec-POMDP model implemented as a class, and could serve as a template for new classes.

Fully-observable problems. Factored fully-observable problems can directly derive from the `FactoredMMDPDiscrete` class (which in turn derives from the partially-observable `FactoredDecPOMDPDiscrete` class). The benefit is that `FactoredMMDPDiscrete` includes convenience functions that shield the user from having to define an observation model. That is, construction of a factored MMDP is then no different than constructing a factored Dec-POMDP in MADP, except that observations do not have to be considered.⁴

The file `./src/tests/test_mmdp.cpp` gives an example of how a fully-observable version of the `ProblemFireFightingFactored` would be modeled in MADP. The code is equivalent to the partially-observable version, except that `ComputeObservationProb` and `SetOScopes` are omitted from the class declaration and that no calls to either `ConstructObservations` or `ConstructJointObservations` are performed. See the classes `./src/support/ProblemFOBS*` for further examples of fully-observable, factored problem domains already implemented. For demonstration, `test_mmdp.cpp` prints the entire observation model and exports the problem specification to the SPUDD file format as well.

Currently, there is no convenience class provided to specify *flat, non-factored* (but fully-observable) problems in MADP. Note, however, that it is easy to use the fully-observable subsets of already existing flat Dec-POMDP problems by simply ignoring the observation model. Alternatively, a factored problem with one (joint) agent could be set up according to the description above.

Background and Design of Particular Functionality

6 Indices for discrete models

Although the design allows for extensions, the MADP toolbox currently only provides implementation for discrete models. I.e., models where the sets of states, actions and observations are discrete. For such discrete models, implementation typically manipulates indices, rather than the basic elements themselves. The MADP toolbox provides such index manipulation functions. In particular, here we describe how individual indices are converted to and from joint indices.

⁴Internally, an observation model with one certain observation per (joint) state is implicitly maintained: Recall that in the general case $\mathcal{O} = \times_i \mathcal{O}_i$ is the set of joint observations. In the fully-observable MMDP, $\mathcal{O}_i = \mathcal{S}$ and $P(\mathbf{o}|\mathbf{a}, s')$ maps o_i to s' deterministically: all agents know the true state of the world with certainty.

```

agents: 2
discount: 1
values: reward
states: tiger-left tiger-right
start:
uniform

actions:
listen open-left open-right
listen open-left open-right

observations:
hear-left hear-right
hear-left hear-right

T: * :
uniform
T: listen listen :
identity

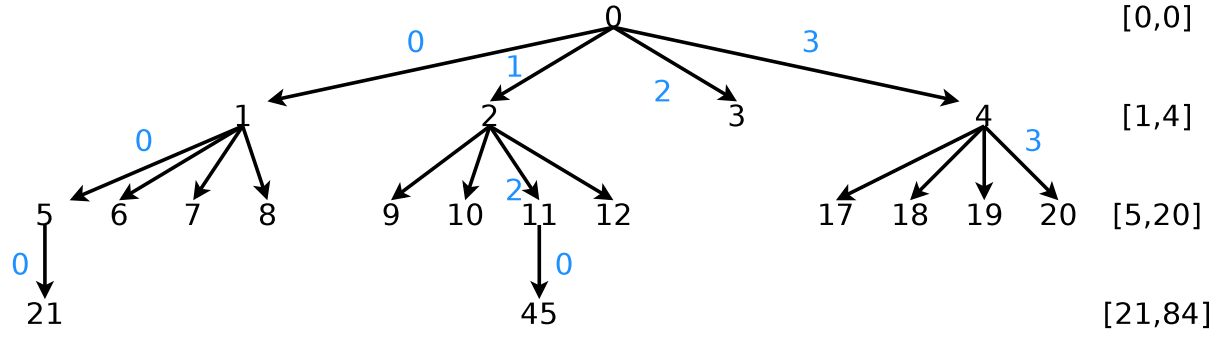
O: * :
uniform

O: listen listen : tiger-left : hear-left hear-left : 0.7225
O: listen listen : tiger-left : hear-left hear-right : 0.1275
O: listen listen : tiger-left : hear-right hear-left : 0.1275
O: listen listen : tiger-left : hear-right hear-right : 0.0225
O: listen listen : tiger-right : hear-right hear-right : 0.7225
O: listen listen : tiger-right : hear-left hear-right : 0.1275
O: listen listen : tiger-right : hear-right hear-left : 0.1275
O: listen listen : tiger-right : hear-left hear-left : 0.0225

R: listen listen: * : * : * : -2
R: open-left open-left : tiger-left : * : * : -50
R: open-right open-right : tiger-right : * : * : -50
R: open-left open-left : tiger-right : * : * : +20
R: open-right open-right : tiger-left : * : * : 20
R: open-left open-right: tiger-left : * : * : -100
R: open-left open-right: tiger-right : * : * : -100
R: open-right open-left: tiger-left : * : * : -100
R: open-right open-left: tiger-right : * : * : -100
R: open-left listen: tiger-left : * : * : -101
R: listen open-right: tiger-right : * : * : -101
R: listen open-left: tiger-left : * : * : -101
R: open-right listen: tiger-right : * : * : -101
R: listen open-right: tiger-left : * : * : 9
R: listen open-left: tiger-right : * : * : 9
R: open-right listen: tiger-left : * : * : 9
R: open-left listen: tiger-right : * : * : 9

```

Figure 6: The dectiger.dpomdp file.



1,2,... (joint) observation indices

1,2,... (joint) observation history indices

Figure 7: Illustration of the enumeration of (joint) observation histories. This illustration is based on a MADP with 4 (joint) observations.

6.1 Enumeration of Joint Actions and Observations

As a convention, joint actions $\mathbf{a} = \langle a_1, \dots, a_m \rangle$ are enumerated as follows

$$\begin{aligned}
 \langle 0, \dots, 0, 0 \rangle & \text{ --- } 0 \\
 \langle 0, \dots, 0, 1 \rangle & \text{ --- } 1 \\
 & \vdots \\
 \langle 0, \dots, 0, |\mathcal{A}_m| \rangle & \text{ --- } |\mathcal{A}_m| - 1 \\
 \langle 0, \dots, 1, 0 \rangle & \text{ --- } |\mathcal{A}_m| \\
 & \vdots \\
 \langle |\mathcal{A}_1|, \dots, |\mathcal{A}_m| - 1, |\mathcal{A}_m| \rangle & \text{ --- } |\mathcal{A}_1| \cdot \dots \cdot |\mathcal{A}_m| - 1.
 \end{aligned}$$

This enumeration is enforced by `ConstructJointActions` in `MADPComponentDiscreteActions`. The joint action index can be determined using the `IndividualToJointIndices` functions from `IndexTools.h`. This file also lists functions for the reverse operation.

Joint observation enumeration is analogous to joint action enumeration (and therefore the same functions can be used).

6.2 Enumeration of (Joint) Histories

Most planning procedures work with indices of histories. For example, `PolicyPureVector` implements a mapping not from observation histories to actions, but from indices (of typically observation-) histories to indices (of actions).

It is important to be able convert between indices of joint/individual action/observation histories and therefore that the method by which the enumeration is performed is clear. This is what is described in this section.

The number of such histories is dependent on the number of observations for each agent, as well as the planning history h . As a result the auxiliary functions for histories have been included in `PlanningUnitMADPDiscrete`. This class also provides the option to generate and cache joint (action-) observation histories, so that the computations described here do not have to be performed every time.

6.2.1 Observation Histories

Figure 7 illustrates how observation histories are enumerated. This enumeration is

- based on the indices of the observations of which they consist.
- breadth-first, such that smaller histories have lower indices and histories for a particular time step t occupy a closed range of indices (also indicated in figure 7).

We will now describe the conversion between observation history indices and observation indices in more detail.

Observation indices to observation history index. Let I_{o_i} denote the index of observation o_i . In order to convert a sequence of observation indices up to time step k for agent i $\left(I_{o_i^1}, I_{o_i^2}, \dots, I_{o_i^k}\right)$ ⁵ to an observation history index, the following formula can be used:

$$I_{\vec{o}_i^k} = \text{offset}_k + \left(I_{o_i^1} \cdot |\mathcal{O}_i|^{k-1} + I_{o_i^2} \cdot |\mathcal{O}_i|^{k-2} + \dots + I_{o_i^{k-1}} \cdot |\mathcal{O}_i|^1 + I_{o_i^k} \cdot |\mathcal{O}_i|^0\right),$$

I.e., $\left(I_{o_i^1}, I_{o_i^2}, \dots, I_{o_i^k}\right)$ is interpreted as a base- $|\mathcal{O}_i|$ number and offset by

$$\text{offset}_k = \sum_{j=0}^{k-1} |\mathcal{O}_i|^j - 1 = \frac{|\mathcal{O}_i|^k - 1}{|\mathcal{O}_i| - 1} - 1.$$

(One is subtracted, because the indices start numbering from 0.)

As an example, the sequence leading to index 45 in figure 7 is $(1, 2, 0)$ with $k = 3$ and $|\mathcal{O}_i| = 4$. We therefore get:

$$\begin{aligned} & \left[\frac{4^3 - 1}{4 - 1} - 1 \right] + [1 \cdot 4^2 + 2 \cdot 4^1 + 0 \cdot 4^0] = \\ & \frac{64 - 1}{3} + 16 + 8 + 0 = \\ & 21 + 24 = 45. \end{aligned}$$

This conversion is performed by `GetObservationHistoryIndex`.

Observation history index to observation indices The inverse is given by a standard division procedure:

$$\begin{aligned} 45 - 21 &= 24 \xrightarrow{\%4} 0 \\ &\xrightarrow{/4} 6 \xrightarrow{\%4} 2 \\ &\xrightarrow{/4} 1 \end{aligned}$$

6.2.2 Action Histories

Individual action histories are enumerated exactly the same way as observation histories: a sequence of actions indices up to time step k $\left(I_{a_i^0}, I_{a_i^1}, \dots, I_{a_i^{k-1}}\right)$ can be converted to an action history index by:

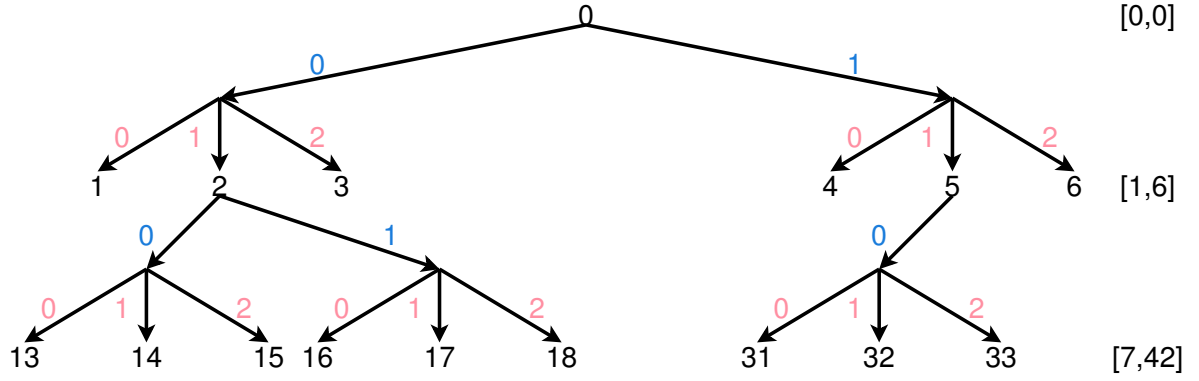
$$I_{\vec{a}_i^k} = \text{offset}_k + \left(I_{a_i^0} \cdot |\mathcal{A}_i|^{k-1} + \dots + I_{a_i^{k-1}} \cdot |\mathcal{A}_i|^0\right).$$

6.2.3 Action-Observation Histories

Enumeration of action-observation histories follows the same principle as for observation histories (and action histories), but have a complicating factor. ‘Action-observations’ are no data type and indices are not clearly defined.

Therefore, in order to implement action-observation histories an enumeration of action-observation is assumed. Let an action observation history $\vec{\theta}_i^k = (o_i^0, a_i^0, o_i^1, a_i^1, o_i^2, \dots, a_i^{k-1}, o_i^k)$ be characterized by

⁵Note, we assume the initial observation o_i^0 to be empty. I.e. the sequence of indices $\left(I_{o_i^1}, I_{o_i^2}, \dots, I_{o_i^k}\right)$ corresponds to the following sequence of observations: $(o_{i,\emptyset}, o_i^1, o_i^2, \dots, o_i^k)$.



- 1,2,... (joint) action indices
- 1,2,... (joint) observation indices
- 1,2,... (joint) action-observation history indices

Figure 8: Illustration of the enumeration of (joint) action-observation histories. This illustration is based on a MADP with 2 (joint) actions and 3 (joint) observations.

its indices $(I_{a_i^0}, I_{o_i^1}, I_{a_i^1}, I_{o_i^2}, \dots, I_{a_i^{k-1}}, I_{o_i^k})$ (again we assume no initial observation). We can group these indices as $(\langle I_{a_i^0}, I_{o_i^1} \rangle, \langle I_{a_i^1}, I_{o_i^2} \rangle, \dots, \langle I_{a_i^{k-1}}, I_{o_i^k} \rangle)$, such that each $\langle I_{a_i^{t-1}}, I_{o_i^t} \rangle$ corresponds to an action-observation. Clearly, there are $|\mathcal{A}_i| \cdot |\mathcal{O}_i|$ action-observations. Let's denote an action-observation with θ_i and its index with I_{θ_i} , corresponding to action a_i and observation o_i , we then have that:

$$I_{\theta_i} = I_{a_i} \cdot |\mathcal{O}_i| + I_{o_i}.$$

`ActionAndObservation_to_ActionObservationIndex` from `IndexTools.h` performs this computation. The inverse operation is performed by `ActionObservation_to_ActionIndex` and `ActionObservation_to_ObservationIndex`,

Now these indices are defined, the same procedure for observation histories can be used as illustrated in fig. 8. I.e.,

$$\begin{aligned} I_{\theta_i^k} = & \text{offset}_k + I_{\theta_i^1} \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^{k-1} + I_{\theta_i^2} \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^{k-2} + \dots \\ & + I_{\theta_i^{k-1}} \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^1 + I_{\theta_i^k} \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^0, \end{aligned}$$

Note that

$$\begin{aligned} I_{\theta_i^t} \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^{k-t} &= (I_{a_i^t} \cdot |\mathcal{O}_i| + I_{o_i^t}) \cdot (|\mathcal{A}_i| \cdot |\mathcal{O}_i|)^{k-t} \\ &= I_{a_i^t} \cdot |\mathcal{O}_i|^{k-t+1} \cdot |\mathcal{A}_i|^{k-t} + I_{o_i^t} \cdot |\mathcal{O}_i|^{k-t} \cdot |\mathcal{A}_i|^{k-t} \end{aligned}$$

As an example, index 32 corresponds to index sequence (1, 1, 0, 1) which, in action-observation indices, corresponds with (4, 1) and thus:

$$\begin{aligned} (6^0 + 6^1) + 4 \cdot 6^{2-1} + 1 \cdot 6^{2-2} &= \\ 7 + 24 + 1 &= 32 \end{aligned}$$

Alternatively we can use the sequence (1, 1, 0, 1) directly:

$$\begin{aligned}
& 7 + (1 \cdot 3^{2-1+1} \cdot 2^{2-1} + 1 \cdot 3^{2-1} \cdot 2^{2-1}) + (0 \cdot 3^{1-1+1} \cdot 2^{1-1} + 1 \cdot 3^{1-1} \cdot 2^{1-1}) = \\
& 7 + (1 \cdot 3^2 \cdot 2^1 + 1 \cdot 3^1 \cdot 2^1) + (0 \cdot 3^1 \cdot 2^0 + 1 \cdot 3^0 \cdot 2^0) = \\
& 7 + (1 \cdot 9 \cdot 2 + 1 \cdot 3 \cdot 2) + (0 \cdot 3 \cdot 1 + 1 \cdot 1 \cdot 1) = \\
& 7 + 18 + 6 + 1 = \\
& 7 + 25 = 32
\end{aligned}$$

6.2.4 Joint Histories

Joint observations are enumerated in the same way as individual observation histories, only now using the indices of joint observations rather than individual observations.

I.e., figure 7 also illustrates how joint observation histories are enumerated. And in order to convert a sequence of joint observations indices up to time step k ($I_{\mathbf{o}^1}, \dots, I_{\mathbf{o}^k}$) to an observation history index, the following formula can be used:

$$I_{\vec{\mathbf{o}}} = \text{offset}_k + \left(I_{\mathbf{o}^0} \cdot |\mathcal{O}|^{k-1} + I_{\mathbf{o}^1} \cdot |\mathcal{O}|^{k-2} + \dots + I_{\mathbf{o}^{k-1}} \cdot |\mathcal{O}|^1 + I_{\mathbf{o}^k} \cdot |\mathcal{O}|^0 \right),$$

I.e., $(I_{\mathbf{o}^1}, \dots, I_{\mathbf{o}^k})$ is interpreted as a base- $|\mathcal{O}|$ number and offset by

$$\text{offset}_{k+1} = \sum_{j=0}^{k-1} |\mathcal{O}|^j - 1 = \frac{|\mathcal{O}|^k - 1}{|\mathcal{O}| - 1} - 1.$$

Indices for joint action histories and joint action-observation histories are computed in the same way. The action-observation functions (`ActionObservation_to_ActionIndex`, etc.) can also be used for joint action-observations.

`PlanningUnitMADPDDiscrete` also provides functions to convert joint to individual history indices `JointToIndividualObservationHistoryIndices`, etc.

7 Joint Beliefs and History Probabilities

Planning algorithms for MADPs will typically need the probabilities of particular joint action observation histories, and the probability over states they induce (called joint beliefs). `PlanningUnitMADPDDiscrete` provides some functionality for performing such inference, which we discuss here.

7.1 Theory

Let $P_{\pi}(\mathbf{a}^t | \vec{\theta}^t)$ denote the probability of \mathbf{a} as specified by π , then $P(s^t, \vec{\theta}^t | \pi, b^0)$ is recursively defined as

$$P(s^t, \vec{\theta}^t | \pi, b^0) = \sum_{s^{t-1} \in \mathcal{S}} P(s^t, \vec{\theta}^t | s^{t-1}, \vec{\theta}^{t-1}, \pi) P(s^{t-1}, \vec{\theta}^{t-1} | \pi, b^0). \quad (1)$$

with

$$P(s^t, \vec{\theta}^t | s^{t-1}, \vec{\theta}^{t-1}, \pi) = P(\mathbf{o}^t | \mathbf{a}^{t-1}, s^t) P(s^t | s^{t-1}, \mathbf{a}^{t-1}) P_{\pi}(\mathbf{a}^{t-1} | \vec{\theta}^{t-1}).$$

For stage 0 we have that $\forall_{s^0} P(s^0, \vec{\theta}_0 | \pi, b^0) = b^0(s^0)$.

Since we tend to think in joint beliefs $b^{\vec{\theta}^t}(s^t) \equiv P(s^t | \vec{\theta}^t, \pi, b^0)$, we can also represent the distribution (1) as:

$$P(s^t, \vec{\theta}^t | \pi, b^0) = P(s^t | \vec{\theta}^t, \pi, b^0) P(\vec{\theta}^t | \pi, b^0) \quad (2)$$

The joint belief $P(s^t | \vec{\theta}^t, \pi, b^0)$ The joint belief $P(s^t | \vec{\theta}^t, \pi, b^0)$ is given by:

$$\begin{aligned}
P(s^t | \vec{\theta}^t, \pi, b^0) &= \frac{P(\mathbf{o}^t | \mathbf{a}^{t-1}, s^t) \sum_{s^{t-1}} P(s^t | s^{t-1}, \mathbf{a}^{t-1}) P(s^{t-1} | \vec{\theta}^{t-1}, \pi, b^0)}{\sum_{s^t} P(\mathbf{o}^t | \mathbf{a}^{t-1}, s^t) \sum_{s^{t-1}} P(s^t | s^{t-1}, \mathbf{a}^{t-1}) P(s^{t-1} | \vec{\theta}^{t-1}, \pi, b^0)} \\
&= \frac{P(s^t, \mathbf{o}^t | \mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)}{P(\mathbf{o}^t | \mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)}
\end{aligned} \quad (3)$$

where $P(\mathbf{o}^t | \mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0) = P(\vec{\theta}^t | \mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)$.

Algorithm 1 $[b^{\vec{\theta}^t}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}})] = \text{GetJAOHProbs}(\vec{\theta}^t, \pi, b^{\vec{\theta}^{t'}}, \vec{\theta}^{t'})$

```

1: if  $\vec{\theta}^t = \vec{\theta}^{t'}$  then
2:   return  $[b^{\vec{\theta}^t} = b^{\vec{\theta}^{t'}}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}}) = 1]$ 
3: end if
4: if  $\vec{\theta}^t$  not an extension of  $\vec{\theta}^{t'}$  then
5:   return  $[b^{\vec{\theta}^t} = \vec{0}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}}) = 0]$ 
6: end if
7:  $\vec{\theta}^{t''} = (\vec{\theta}^{t'}, \mathbf{a}^{t'}, \mathbf{o}^{t'+1})$  {consist. with  $\vec{\theta}^t$ }
8:  $[b^{\vec{\theta}^{t''}}, P(\vec{\theta}^{t''}|\mathbf{a}^{t'}, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}})] = b^{\vec{\theta}^{t'}}$ .Update( $\mathbf{a}^{t'}, \mathbf{o}^{t'+1}$ ) {belief update, see (3)}
9:  $[b^{\vec{\theta}^t}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t''}, b^{\vec{\theta}^{t''}})] = \text{GetJAOHProbs}(\vec{\theta}^t, \pi, b^{\vec{\theta}^{t''}}, \vec{\theta}^{t''})$ 
10:  $P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}}) = P(\vec{\theta}^t|\pi, \vec{\theta}^{t''}, b^{\vec{\theta}^{t''}})P(\vec{\theta}^{t''}|\mathbf{a}^{t'}, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}})P_\pi(\mathbf{a}^{t'}|\vec{\theta}^{t'}, \pi)$ 
11: return  $[b^{\vec{\theta}^t}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}})]$ 

```

The probability of an history $P(\vec{\theta}^t|\pi, b^0)$ The second part of (2) is given by

$$\begin{aligned}
P(\vec{\theta}^t|\pi, b^0) &= P(\vec{\theta}^t|\vec{\theta}^{t-1}, \pi, b^0)P(\vec{\theta}^{t-1}|\pi, b^0) \\
&= P(\vec{\theta}^t|\mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)P_\pi(\mathbf{a}^{t-1}|\vec{\theta}^{t-1}, \pi, b^0)P(\vec{\theta}^{t-1}|\pi, b^0) \\
&= P(\mathbf{o}^t|\mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)P_\pi(\mathbf{a}^{t-1}|\vec{\theta}^{t-1}, \pi, b^0)P(\vec{\theta}^{t-1}|\pi, b^0)
\end{aligned} \tag{4}$$

where $P(\mathbf{o}^t|\mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)$ is the denominator of (3).

7.2 Implementation

Since computation of $P(\vec{\theta}^t|\pi, b^0)$ is interwoven with the computation of the joint belief through $P(\mathbf{o}^t|\mathbf{a}^{t-1}, \vec{\theta}^{t-1}, \pi, b^0)$, it is impractical to separately evaluate (3) and (4).

Rather we define a function

$$[b^{\vec{\theta}^t}, P(\vec{\theta}^t|\pi, b^0)] = \text{GetJAOHProbs}(\vec{\theta}^t, \pi, b^0)$$

Because in many situations an application might evaluate similar $\vec{\theta}^t$ (i.e., ones with an identical prefix), a lot of computation will be redundant. To give the user the possibility to avoid this, we also define

$$[b^{\vec{\theta}^t}, P(\vec{\theta}^t|\pi, \vec{\theta}^{t'}, b^{\vec{\theta}^{t'}})] = \text{GetJAOHProbs}(\vec{\theta}^t, \pi, b^{\vec{\theta}^{t'}}, \vec{\theta}^{t'})$$

which returns the probability and associated joint belief of $\vec{\theta}^t$, *given* that $\vec{\theta}^{t'}$ (and associated joint belief $b^{\vec{\theta}^{t'}}$) are realized (i.e., given that $P(\vec{\theta}^{t'}) = 1$).

8 Policies

Here we discuss some properties and the implementation of policies. Policies are plans for agents that specify how they should act in each possible situation. As a result a policy is a mapping from these ‘situations’ to actions. Depending on the assumption on the observability in a MADP, however, these ‘situations’ might be different. Also we would like to be able to reuse the implementations of policies for problems with a slightly different nature, for instance (Bayesian) games.

In the MADP toolbox, the most general form of a policy is a mapping from a domain (`PolicyDomain`) to (probability distributions over) actions. Currently we have only considered policies for discrete domains, which are mappings from indices (of these histories) to indices (of actions). It is typically still necessary to know what type of indices a policy maps from in order to be able to reuse our implementation of

policies. To this end a discrete policy maintains its `IndexDomainCategory`. So far there are four types of index-domain categories: `TYPE_INDEX`, `OHIST_INDEX`, `OAIST_INDEX` and `STATE_INDEX`.

As said `PolicyDiscrete` class represents the interface policies for discrete domains. `PolicyDiscretePure` is the interface for a pure (deterministic) policy. A class that actually implements a policy is `PolicyPureVector`. This class also implements a function to get and set the index of the policy (pure policies over a finite domain are enumerable). Joint policies are represented by similarly named classes `JointPolicyDiscrete`, `JointPolicyPureVector`, etc.

In order to instantiate a (joint) policy, it needs to know several things about the problem it is defined over. We already mentioned the index domain category, but there is other information needed as well (the number of agents, the sizes of their domains, etc.). To provide this information, each problem for which we want to construct a (joint) policy has to implement the `Interface.ProblemToPolicyDiscretePure`.

9 The ProbModelXML format

MADP can parse problem files written in the *ProbModelXML* format [1]. This format is useful for the definition of Factored MDPs / POMDPs / Dec-POMDPs, or any other class of problems that can be represented graphically as a two-time-slice Dynamic Bayesian Network (DBN). You can find the complete specification of the ProbModelXML format at: <http://www.cisiad.uned.es/ProbModelXML/>. You should save your ProbModelXML files with the “.pgmx” extension.

Since ProbModelXML supports many different probabilistic graphical models and concepts that are outside of the scope of MADP, there are some constraints on what can be interpreted by the MADP ProbModelXML parser:

Network Types: The MDP, POMDP, and DEC_POMDP formats are supported. The network type is actually inferred by the parser, so you can always safely specify DEC_POMDP as the network type, even if you are designing a less general model. This is also valid for centralized multiagent models (MMDPs, MPOMDPs);

Variables: Only discrete, finite-domain variables (state factors, actions, observations and reward factors) are currently supported. You can only specify variables for time “0” and “1”, which respectively represent time steps t and $t+1$ in the two-time-slice DBNs. In ProbModelXML terminology, model variables can be defined as follows:

- State Factors: *Chance* nodes, that should be defined both at time 0 and 1. The *Potential* of a state factor node at time 0 is its initial distribution, and at time 1 it is its Conditional Probability Distribution (CPD);
- Actions: *Decision* nodes, that should be defined both at time 0 and 1.
- Reward Factors: *Utility* nodes, that can be defined either at time 0 or at time 1 (but not both), depending on whether you want to represent $R(s^t, \mathbf{a})$ or $R(\mathbf{a}, s^{t+1})$.
- Observations: ProbModelXML does not explicitly recognize “observations” as a separate type of variable. Instead, observations are defined implicitly as time 1 “chance” nodes that link to the actions (at time 1) of their respective agents. Examples are shown in the following section. The *Potential* of an observation node defined in this way is its CPD.

CPDs: CPDs can only be defined as **Table**, **Tree/ADD**, or **Uniform**. Note that, internally, MADP only supports CPDs defined as tables, but you can still use decision trees or ADDs in the ProbModelXML representation - just keep in mind that they’ll be “flattened” into tables;

Inference options: Even though ProbModelXML provides some options for probabilistic inference, this is ignored in MADP, since belief propagation is handled internally.

The MADP problems folder contains some examples of problem files written in ProbModelXML. The “Dec-Tiger” problem file (“DTPGMX.pgm”) can be referred to as a starting point to understand the general layout of this file format, as it includes additional descriptive comments.

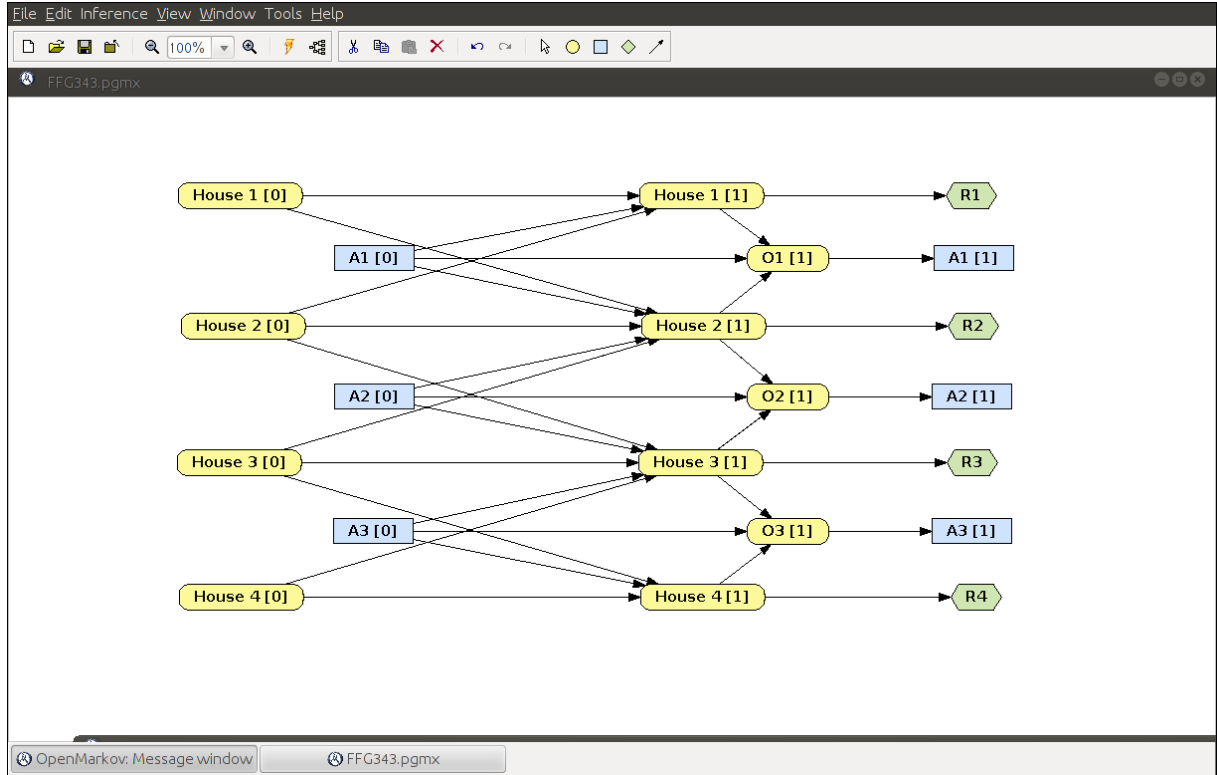


Figure 9: The Graphical Factored Firefighters (FFG) problem in OpenMarkov.

9.1 Using OpenMarkov to design factored problems

OpenMarkov (<http://www.openmarkov.org/users.html>) is a Java-based graphical editor for ProbModelXML files. Although using OpenMarkov is not strictly necessary to design ProbModelXML factored problems, it is the recommended option for this purpose, since it is far more intuitive and less time-consuming than coding the .pgmx file directly. As of the time of writing, OpenMarkov requires Java 7. The MADP ProbModelXML parser has been tested with the latest (0.1.4) version of OpenMarkov. Note that since OpenMarkov is an independent project, and it is not authored or maintained by the MADP community, later versions are not guaranteed to be immediately compatible.

After you download OpenMarkov, move it to your MADP folder and rename it to something simpler (e.g. “openmarkov.jar”). Then you can start it by typing:

```
~/madp$ java -jar openmarkov.jar
```

After OpenMarkov is loaded, try to open one of the .pgmx problem files in the MADP problems/ folder, for instance, FFG343.pgm. Your view should then be similar to Figure 9.

In this view, *chance* nodes, shown in yellow, correspond to the state and observation variables of the problem; *decision* nodes, shown in blue, correspond to the actions of each agent; and *utility* nodes, shown in green, correspond to local reward factors. The bracketed number next to each node ([0] or [1]) represents the time-slice that it belongs to (t or $t + 1$ respectively). Notice that observation variables (O1, O2, O3) are also chance nodes, but they are only defined at time 1 and they link only to the actions of the agents that they belong to (A1, A2, A3, respectively). This symbolizes that the actions of those agents at time $t + 1$ “depend” on the values of these variables, although the way through which that dependency is manifested is not explicitly represented in the model. For fully observable problems, there is no need to specify observations, since all chance nodes are assumed to be observable. In that case, you can either omit the time 1 action nodes or represent them as “orphaned” nodes.

Right-click a chance node and view its “Node properties”. There, you can view and edit its name, time slice, and domain values (Figure 10). Other options are not relevant for MADP.

Likewise, right-click a time 1 chance node and select “Edit probability”. Your view should be similar to Figure 11. In the Tree/ADD view, branches are represented by the white labeled boxes under each

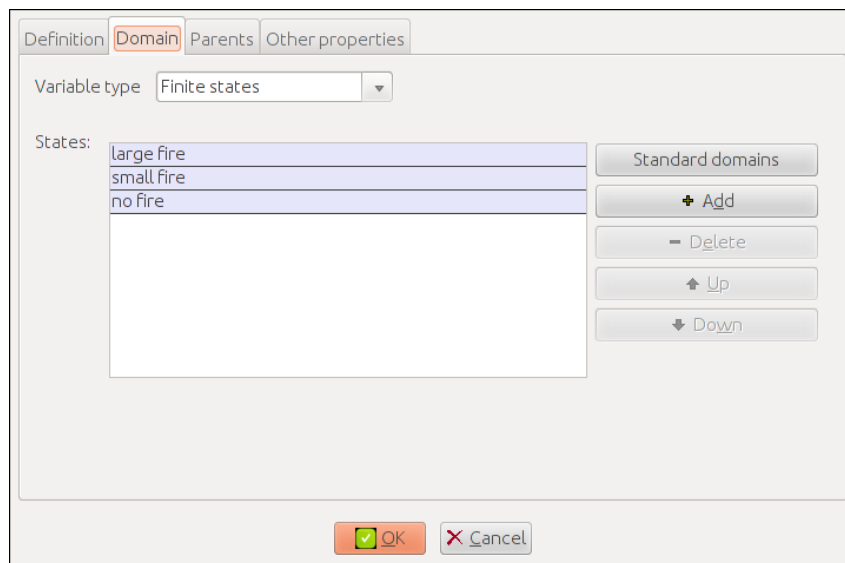


Figure 10: The domain values of a chance node in OpenMarkov. The variable type should always be “Finite states”. You can add, delete, or rearrange the domain values. The bottom-most value will have index 0.

of the colored “variable” elements, each of them assigned to a particular value or set of values of that variable. You can expand or contract each branch by clicking to the left of its box. You can also right-click each branch to add or remove values (*a.k.a.* “states” in OpenMarkov), and right-click variable nodes to change their assignment (only possible if there are valid alternatives). In leaf nodes, you can right-click to “edit potential”, i.e. define the CPD at that point. To define a CPD as an ADD, you can assign a label to a branch (“set label”), and then you can bind subsequent equivalent branches to that label, so that you don’t have to re-define them (“set reference”).

A CPD can also be specified as a table (Figure 12). In that case, the various combinations of the parent variables will be shown at the top, and each row of the table corresponds to one specific value of the dependent variable (shown on the left). This implies that all columns should sum to 1.

Again, in order to define the initial state distribution for your problem, you should define the *potentials* for all state variables in time slice 0. The potentials at time slice 1 encode the CPDs of all relevant variables, which are assumed to be stationary (do not depend on the absolute time index).

Rewards (*a.k.a.* *utilities*) can be defined in the same way as CPDs – the only difference is that they can have real-valued outcomes outside of the $[0, 1]$ interval.

To create a new problem file (“File→ New”), select “Dec-POMDP” as the network type, and simply follow the above guidelines. Here’s what you **can’t** do:

- You can’t have time slices with indexes greater than 1;
- You can’t have anti-causal influences, *i.e.* variables at time 1 that influence variables at time 0;
- You can’t have non-stationary (dynamic) CPDs;
- You can’t have ADDs with external references (that is, you cannot define a label in an ADD and use it in another ADD);
- In partially-observable problems, you can’t have a different number of observations and decision nodes – this should always be equal to the number of agents (an exception is discussed in the next subsection).

9.2 Designing Event-Driven Models

A special kind of model that is supported in MADP is the Event-Driven MPOMDP model [12]. In these models, state factor transitions are typically *asynchronous*, that is, only one state variable changes

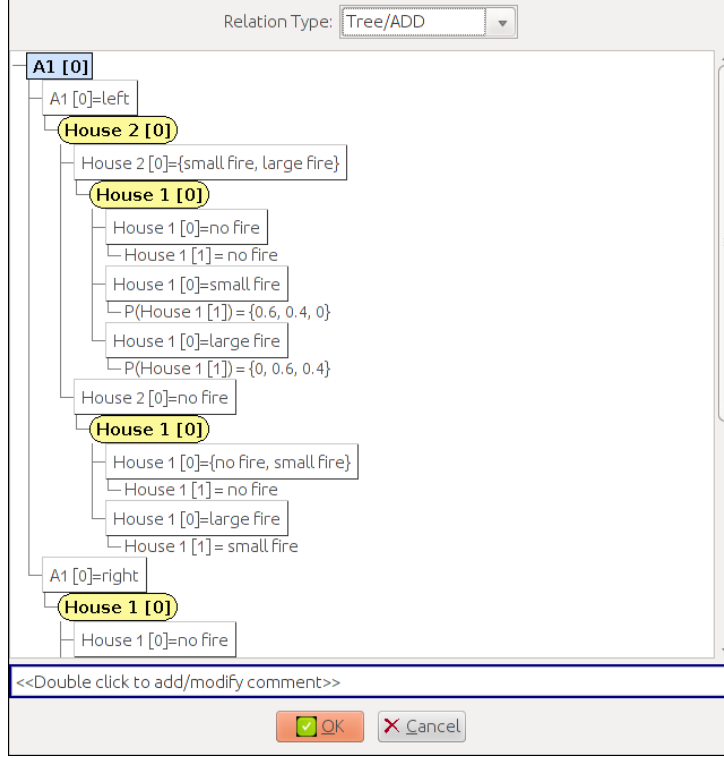


Figure 11: A $t + 1$ CPD specified as a tree.

Relation Type:

Table

Reorder variables

S[1]	left	left	left	left	left	left	left	left	le
A2[0]	listen	listen	listen	open left	open left	open left	open right	open right	open
A1[0]	listen	open left	open right	listen	open left	open right	listen	open left	open
right	0.15	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.
left	0.85	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.

<<Double click to add/modify comment>>

OK

Cancel

Figure 12: A $t + 1$ CPD specified as a table (from “problems/DTPGMX.pgm”).

between t and $t + 1$. Event-Driven MPOMDPs are not as straightforward to represent as two-time-slice DBNs as standard (Dec-)POMDPs due to this property. However, they can still be modeled by considering a “virtual” state-factor that encodes the prior probability, at time t , that each factor will be affected by the transition at time $t + 1$. These factors are special in that they are influenced by time t variables, but they subsequently influence time $t + 1$ variables, that is, they establish intra-slice dependencies at time $t + 1$. An example of an Event-Driven POMDP model is shown in (Figure 13).

Another characteristic of event-driven models is that observations depend on *transitions*, as opposed to *states*. This means that the (only) observation node at time $t + 1$ contains both time $t + 1$ and time t parents.

Event-Driven POMDPs are the only type of model that can accept a different number of action nodes and observation nodes (since there is always only one observation node). If you include more than one decision node at time 0, you should include the following element in your .pmgx file (under the ProbNet element):

```
<AdditionalProperties>
  <Property name="EventDriven" value="1" />
</AdditionalProperties>
```

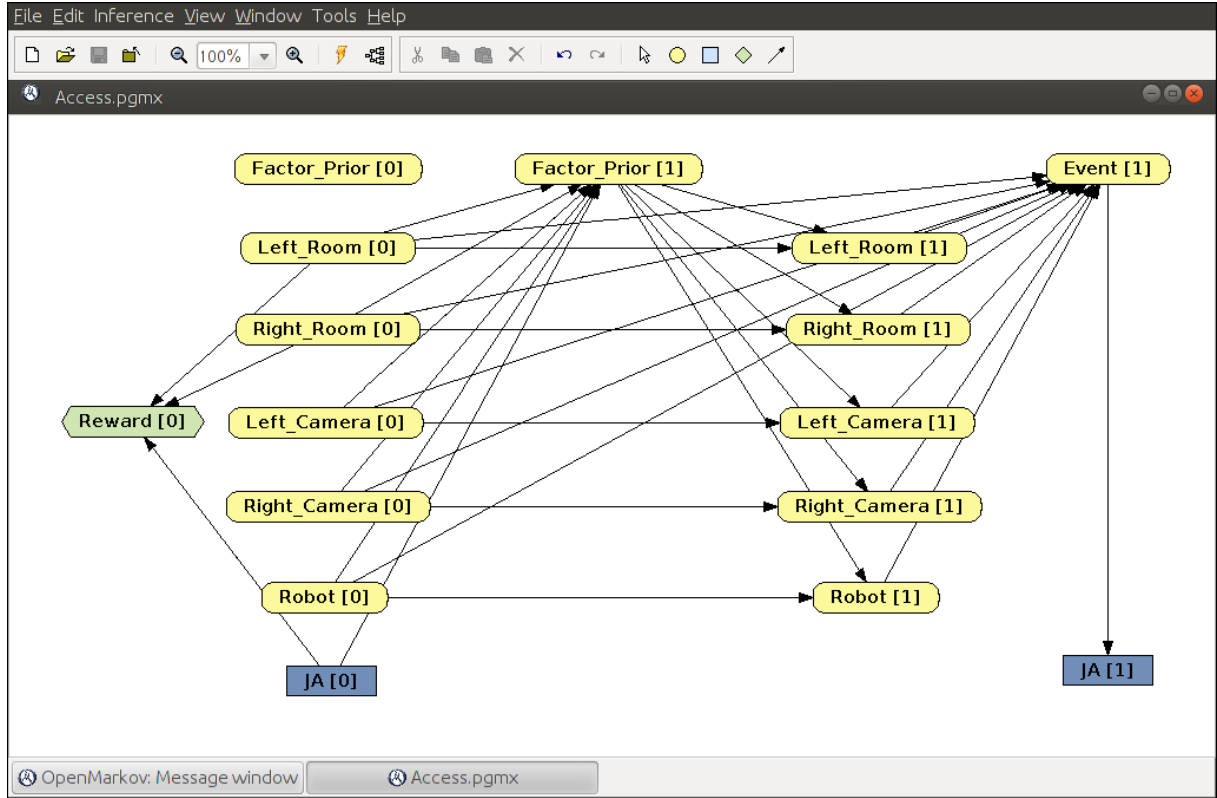


Figure 13: The “Access” problem, an Event-Driven MPOMDP, modeled in OpenMarkov.

References

- [1] M. Arias, F. J. Díez, and M. P. Palacios. ProbModelXML. A format for encoding probabilistic graphical models. Technical report cisiad-11-02, UNED, Madrid, Spain, 2011.
- [2] Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4): 819–840, 2002.
- [3] Umberto Bertele and Francesco Brioschi. *Nonserial Dynamic Programming*. Academic Press, Inc., 1972.
- [4] Pieter-Tjerk de Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinstein. A tutorial on the cross-entropy method. *Annals of Operations Research*, 134(1):19–67, 2005.
- [5] Abdeslam Boularias and Brahim Chaib-draa. Exact dynamic programming for decentralized POMDPs with lossless policy compression. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 2008.
- [6] Craig Boutilier. Planning, learning and coordination in multiagent decision processes. In *Proc. of the 6th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 195–210, 1996.
- [7] Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 54–61. Morgan Kaufmann, 1997.
- [8] Rina Dechter. Bucket elimination: a unifying framework for processing hard and soft constraints. *Constraints*, 2(1):51–55, 1997.
- [9] Carlos Guestrin, Daphne Koller, and Ronald Parr. Multiagent planning with factored MDPs. In *Advances in Neural Information Processing Systems 14*, pages 1523–1530, 2002.

- [10] Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. Dynamic programming for partially observable stochastic games. In *Proceedings of the National Conference on Artificial Intelligence*, pages 709–715, 2004.
- [11] Jelle R. Kok and Nikos Vlassis. Collaborative multiagent reinforcement learning by payoff propagation. *Journal of Machine Learning Research*, 7:1789–1828, 2006.
- [12] Joao V. Messias, Matthijs T. J. Spaan, and Pedro U. Lima. Asynchronous execution in multiagent POMDPs: Reasoning over partially-observable events. In *AAMAS’13 Workshop on Multi-agent Sequential Decision Making under Uncertainty (MSDM)*, pages 9–14, May 2013.
- [13] George E. Monahan. A survey of partially observable Markov decision processes: theory, models and algorithms. *Management Science*, 28(1), January 1982.
- [14] Joris M. Mooij. libDAI: A free and open source C++ library for discrete approximate inference in graphical models. *Journal of Machine Learning Research*, 11:2169–2173, August 2010.
- [15] Ranjit Nair, Milind Tambe, Makoto Yokoo, David V. Pynadath, and Stacy Marsella. Taming decentralized POMDPs: Towards efficient policy computation for multiagent settings. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 705–711, 2003.
- [16] Frans A. Oliehoek. Decentralized POMDPs. In Marco Wiering and Martijn van Otterlo, editors, *Reinforcement Learning: State of the Art*, volume 12 of *Adaptation, Learning, and Optimization*, pages 471–503. Springer Berlin Heidelberg, Berlin, Germany, 2012.
- [17] Frans A. Oliehoek, Matthijs T. J. Spaan, and Nikos Vlassis. Dec-POMDPs with delayed communication. In *Proceedings of the AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)*, May 2007.
- [18] Frans A. Oliehoek, Julian F.P. Kooi, and Nikos Vlassis. The cross-entropy method for policy search in decentralized POMDPs. *Informatica*, 32:341–357, 2008.
- [19] Frans A. Oliehoek, Matthijs T. J. Spaan, and Nikos Vlassis. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32:289–353, 2008.
- [20] Frans A. Oliehoek, Matthijs T. J. Spaan, Shimon Whiteson, and Nikos Vlassis. Exploiting locality of interaction in factored Dec-POMDPs. In *Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems*, pages 517–524, May 2008.
- [21] Frans A. Oliehoek, Shimon Whiteson, and Matthijs T. J. Spaan. Lossless clustering of histories in decentralized POMDPs. In *Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems*, pages 577–584, May 2009.
- [22] Frans A. Oliehoek, Matthijs T. J. Spaan, Jilles Dibangoye, and Christopher Amato. Heuristic search for identical payoff Bayesian games. In *Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems*, pages 1115–1122, May 2010.
- [23] Frans A. Oliehoek, Shimon Whiteson, and Matthijs T. J. Spaan. Exploiting structure in cooperative Bayesian games. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 654–664, August 2012.
- [24] Frans A. Oliehoek, Shimon Whiteson, and Matthijs T. J. Spaan. Approximate solutions for factored Dec-POMDPs with many agents. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems*, pages 563–570, 2013.
- [25] Sven Seuken and Shlomo Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *Autonomous Agents and Multi-Agent Systems*, 17(2):190–250, 2008.
- [26] Matthijs T. J. Spaan and Nikos Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.

- [27] Matthijs T. J. Spaan, Frans A. Oliehoek, and Chris Amato. Scaling up optimal heuristic search in Dec-POMDPs via incremental expansion. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 2027–2032, 2011.
- [28] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, March 1998.
- [29] Daniel Szer, François Charpillet, and Shlomo Zilberstein. MAA*: A heuristic search algorithm for solving decentralized POMDPs. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 576–583, 2005.