# Exploiting Locality of Interaction in Factored Dec-POMDPs 

Frans A. Oliehoek Informatics Institute, University of Amsterdam<br>Kruislaan 403, 1098 SJ<br>Amsterdam, The Netherlands<br>faolieho@science.uva.nl

Matthijs T.J. Spaan<br>Institute for Systems Robotics Instituto Superior Técnico<br>Av. Rovisco Pais, 1, 1049-001<br>Lisbon, Portugal<br>mtjspaan@isr.ist.utl.pt<br>Nikos Vlassis<br>Department of Production<br>Engineering and Management<br>Technical University of Crete<br>Chania, Greece<br>vlassis@dpem.tuc.gr

Shimon Whiteson Informatics Institute, University of Amsterdam<br>Kruislaan 403, 1098 SJ<br>Amsterdam, The Netherlands<br>whiteson@science.uva.nl


#### Abstract

Decentralized partially observable Markov decision processes (Dec-POMDPs) constitute an expressive framework for multiagent planning under uncertainty, but solving them is provably intractable. We demonstrate how their scalability can be improved by exploiting locality of interaction between agents in a factored representation. Factored Dec-POMDP representations have been proposed before, but only for DecPOMDPs whose transition and observation models are fully independent. Such strong assumptions simplify the planning problem, but result in models with limited applicability. By contrast, we consider general factored Dec-POMDPs for which we analyze the model dependencies over space (locality of interaction) and time (horizon of the problem). We also present a formulation of decomposable value functions. Together, our results allow us to exploit the problem structure as well as heuristics in a single framework that is based on collaborative graphical Bayesian games (CGBGs). A preliminary experiment shows a speedup of two orders of magnitude.


## Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence - Multiagent systems

## General Terms

Algorithms, Performance, Theory, Experimentation

## Keywords

Planning under uncertainty, cooperative multiagent systems, decentralized POMDP.

## 1. INTRODUCTION

Planning under uncertainty is a central topic within artificial intelligence. Especially during the last two decades mod-
Cite as: Exploiting Locality of Interaction in Factored Dec-POMDPs, Oliehoek, Spaan, Whiteson and Vlassis, Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008), Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.
Copyright (C) 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.
els such as Markov decision processes (MDPs), for singleagent planning in a stochastic environment, and partially observable MDPs (POMDPs), for single-agent planning in an environment that is both stochastic and partially observable, have become popular. The decentralized POMDP (Dec-POMDP) is a natural extension to multiple agents that each have their individual observations, but have to optimize the same reward function.

Unfortunately, optimally solving Dec-POMDPs is NEXPcomplete [2], and the same holds for finding an $\epsilon$-approximate solution [14]. As a result, research has focused on special cases to overcome this complexity barrier. For instance models with more assumptions on observability and/or communication have been considered $[6,9]$. These approaches try to exploit independence between agents, but they either require the agents to observe the state of the system (as in a multiagent MDP), or to observe a subset of state variables (as in a factored Dec-MDP) and communicate at every stage. A recent approach [15] tries to minimize the amount of communication.

For the particular case of transition and observation independent Dec-POMDPs [1], exploitation of independence between agents has been suggested by Nair et al. [11] and extended by Varakantham et al. [19]. However, the assumption of transition and observation independence (TOI) severely limits the applicability of these models. In particular, TOI implies that the agents have disjoint sets of individual states and observations, and that one agent cannot influence the state transition nor the observation of another agent. In practice this means that many interesting tasks, such as two robots carrying a chair, cannot be modeled.

In this paper, we present a more general analysis of locality of interaction in factored Dec-POMDPs. In particular we show that the last stage in the process contains the highest degree of independence but, when moving back in time (towards the initial stage $t=0$ ), the scope of dependence grows. Nevertheless, we show that most of the independence is located-and can be exploited-where it is most needed, since the overwhelming majority of the computational effort of solving a Dec-POMDP is spent in the last stage [18].

We use this result to extend our previous work, in which we solved Dec-POMDPs by representing them as a series of Bayesian games (BGs), one for each stage [13, 12]. In
this setting, the increase in computational cost at the later stages results from the size of the BGs and grows doubly exponentially with $t$. Using locality of interaction, we introduce collaborative graphical BGs (CGBGs), which are less computationally expensive to solve in later stages where independence is high. We propose an extension of GMAA*, a generalized policy search algorithm, which uses CGBGs to exploit locality of interaction in the last stage of a DecPOMDP. A preliminary experiment shows a speedup of two orders of magnitude.

## 2. BACKGROUND

Here we introduce the regular (non-factored) Dec-POMDP, then we describe the extensions for the factored version.

### 2.1 Dec-POMDPs

Formally, a Dec-POMDP with $n$ agents is defined as a tuple $\left\langle\mathcal{A} g, \mathcal{S}, \mathcal{A}, T, R, \mathcal{O}, O, b^{0}, h\right\rangle$, where

- $\mathcal{A} g=\{1, \ldots, n\}$ is the set of agents.
- $\mathcal{S}$ is a finite set of states.
- $\mathcal{A}=\times_{i} \mathcal{A}_{i}$ is the set of joint actions, where $\mathcal{A}_{i}$ is the set of actions available to agent $i$. Every time step, the agents take one joint action $\mathbf{a}=\left\langle a_{1}, \ldots, a_{n}\right\rangle$, but agents do not observe each other's actions. We also write $\mathbf{a}_{\mathcal{I}}$ for the joint action of a subset of agents $\mathcal{I} \subseteq \mathcal{A} g$.
- The transition probabilities $P\left(s^{\prime} \mid s\right.$, a) are specified by the transition function $T$.
- $\mathcal{O}=\times_{i} \mathcal{O}_{i}$ is the set of joint observations. Every stage one joint observation $\mathbf{o}=\left\langle o_{1}, \ldots, o_{n}\right\rangle$ is received. Each agent $i$ only observes its own component $o_{i}$.
- $O$ is the observation function, which specifies the probability of joint observations: $P\left(\mathbf{o} \mid \mathbf{a}, s^{\prime}\right)$.
- $R$ is the immediate reward function, which maps states and joint actions to reals: $R(s, \mathbf{a})$.
- $b^{0} \in \mathcal{P}(\mathcal{S})$ is the initial state distribution at time $t=$ 0 , where $\mathcal{P}(S)$ denotes the infinite set of probability distributions over the finite set $S$.
- $h$ is the horizon of the problem.

At each stage $t=0 \ldots h-1$ the agents take an action and receive an observation. Their goal is to maximize the expected cumulative reward or return. The planning task entails finding a joint policy $\pi=\left\langle\pi_{1} \ldots \pi_{n}\right\rangle$, which specifies an individual policy $\pi_{i}$ for each agent $i$. Such an individual policy in general specifies an individual action for each action-observation history $\vec{\theta}_{i}^{t}=\left(a_{i}^{0}, o_{i}^{1}, \ldots, a_{i}^{t-1}, o_{i}^{t}\right)$, e.g., $\pi_{i}\left(\vec{\theta}_{i}^{t}\right)=a_{i}^{t}$. The set of $\vec{\theta}_{i}^{t}$ is denoted $\vec{\Theta}_{i}^{t}$. However, when only allowing deterministic or pure policies, a $\pi_{i}$ maps each observation history $\left(o_{i}^{1} \ldots o_{i}^{t}\right)=\vec{o}_{i}^{t} \in \overrightarrow{\mathcal{O}}_{i}^{t}$ to an action, e.g. $\pi_{i}\left(\vec{o}_{i}^{t}\right)=a_{i}^{t}$. For a more detailed introduction to DecPOMDPs see [12].

### 2.2 Factored Dec-POMDPs

A factored Dec-POMDP has a state space $\mathcal{S}=\mathcal{X}_{1} \times \ldots \times$ $\mathcal{X}_{|\mathcal{X}|}$ that is spanned by $\mathcal{X}=\left\{\mathcal{X}_{1}, \ldots, \mathcal{X}_{|\mathcal{X}|}\right\}$ a set of state variables, or factors. A state corresponds to an assignment
of values for all factors $s=\left\langle x_{1}, \ldots, x_{|\mathcal{X}|}\right\rangle$. In a factored DecPOMDP, the transition and observation model can be compactly represented by exploiting conditional independence between variables. In particular, the transition and observation model can be represented by a dynamic Bayesian network (DBN) [3]. In such a DBN arrows represent causal influence, and each node with incoming edges has a conditional probability table (CPT) associated with it. Although the size of these CPTs is exponential in the number of parents, the parents are typically a subset of all state factors and actions, leading to exponentially smaller tables than a flat model.

Similarly, the reward function often can be compactly represented by exploiting additive separability, which means that it is possible to decompose the reward function into the sum of 'smaller' local reward functions: $R=R^{1}+\cdots+R^{\rho}$. Here 'smaller' means that the functions are defined over a smaller number of state and action variables, i.e., the scope of these functions is smaller. We can represent the smaller functions in the DBN [3]. The reward nodes have conditional reward tables (CRTs) associated with them that represent the smaller reward functions. The local reward functions should not be confused with individual reward functions in a system with self-interested agents, such as partially observable stochastic games [7] and (graphical) BGs [16, 17]. In such models agents compete to maximize their individual reward functions, while we consider agents that collaborate to maximize the sum of local payoff functions.

Decision trees [3, 4] or algebraic decision diagrams [8] can be used to further reduce the size of representation (of both CPTs and CRTs) by exploiting context specific independence. E.g., the value of factor $x_{2}$ is of no influence when factor $x_{1}$ has a particular value. We will not consider these further enhancements in this paper.

### 2.3 Factored Firefighting problem

To illustrate the model we will introduce the factored firefighting (FFF) problem as a running example. It models a team of $n$ firefighters that have to extinguish fires in a row of $n_{H}=n+1$ houses. We will assume $n=3, n_{H}=4$ as illustrated in Figure 1 (Left). Each house $H$ is characterized by a fire level $f l_{H}$, an integer parameter in $\left[0, n_{f}\right)$. A state in FFF is an assignment of fire levels $s=\left\langle f l_{1}, f l_{2}, f l_{3}, f l_{4}\right\rangle$. At every time step, each agent $i$ can choose to fight fires at house $i$ or $i+1$. Figure 1 (Middle) shows the DBN for the problem. If a house $H$ is burning $\left(f l_{H}>0\right)$ and no firefighting agent is present, its fire level will increase by one point with probability 0.8 if any of its neighboring houses are burning, and with probability 0.4 if none of its neighbors are on fire. A house that is not burning can only catch fire with probability 0.8 if one of its neighbors is on fire. When two agents are in the same house, they will extinguish any present fire completely, setting the house's fire level to 0 . A single agent present at a house will lower the fire level by one point with probability 1 if no neighbors are burning, and with probability 0.6 otherwise. Each agent can only observe whether there are flames, $o_{i}=F$, or not, $o_{i}=N$, at its location. Flames are observed with probability 0.2 if $f l_{H}=0$, with probability 0.5 if $f l_{H}=1$, and with probability 0.8 otherwise. Initially, the fire level $f l_{H}$ of each house is drawn from a uniform distribution. The agents receive a reward of $-f l_{H}$ for each house $H$. In particular, the rewards are specified by the fire levels at the next time step


Figure 1: Left: An illustration of the FFF problem. Middle: A DBN modeling the transition, observation and reward function for the FFF problem. Right: The dynamics of the FFF problem over multiple stages. The scope of $Q^{1}$, illustrated by shading, increases when going back in time.
$\forall_{1 \leq H \leq 4} r^{H}\left(f l_{H}^{\prime}\right)=-f l_{H}^{\prime}$.
We can reduce these rewards to ones of the type $R^{H}(s, \mathbf{a})$ by considering the expectation over $f l_{H}^{\prime}$. For instance, for house 1

$$
R^{1}\left(f l_{\{1,2\}}, a_{1}\right)=\sum_{f l_{1}^{\prime}} P\left(f l_{1}^{\prime} \mid f l_{\{1,2\}}, a_{1}\right) r^{1}\left(f l_{1}^{\prime}\right),
$$

where $f l_{\{1,2\}}$ denotes $\left\langle f l_{1}, f l_{2}\right\rangle$. This formulation is possible because, as Figure 1 shows, $f l_{1}, f l_{2}$ and $a_{1}$ are the only variables that influence the probability of $f l_{1}^{\prime}$. Similarly, the other local reward functions are given by $R^{2}\left(f l_{\{1,2,3\}}, \mathbf{a}_{\{1,2\}}\right)$, $R^{3}\left(f l_{\{2,3,4\}}, \mathbf{a}_{\{2,3\}}\right)$ and $R^{4}\left(f l_{\{3,4\}}, a_{3}\right)$.

## 3. FACTORED Q-VALUE FUNCTIONS

We propose to exploit independence between agents. In this section, we first analyze this independence by considering Q-value functions that express the expected return. In particular, rather than using a single Q -value function that is defined over joint actions of all agents, we show that the value can be decomposed into $\rho$ local Q -value functions (corresponding to the $\rho$ local reward functions) involving only a subset of agents (although this subset may grow with each stage further from the last stage $t=h-1$ ). Such a decomposed Q -value function is called factored. In section 4 we will exploit such independence using a variation of BGs.

### 3.1 Scope and scope backup

As an example of a factored Q -value function, we consider the last stage of FFF. Clearly, $Q^{h-1}$ is defined to be equal to the immediate reward function, and thus can be decomposed in 4 local Q-value functions $Q=Q^{1}+\cdots+Q^{4}$. Where each $Q^{k}$ is defined over the same subset of variables as $R^{k}$. This subset of variables is called the scope of $Q^{k}$. In order to exploit independence between agents, we discriminate between the variables that pertain to state factors and those that pertain to agents (i.e., actions and observations). Exploitation of independence between state factors will be mentioned, but is left for future work.

Definition 3.1 (Scope) Let $\mathcal{H} \subseteq\{1, \ldots,|\mathcal{X}|\}$ be a subset of state factor indices, and let $\mathcal{I} \subseteq\{1, \ldots, n\}$ be a subset of agent indices, and let us write $\mathcal{X}_{\mathcal{H}}=x_{h \in \mathcal{H}} \mathcal{X}_{h}, \mathcal{A}_{\mathcal{I}}=$ $\times_{i \in \mathcal{I}} \mathcal{A}_{i}, \overrightarrow{\boldsymbol{\Theta}}_{\mathcal{I}}^{t}=x_{i \in \mathcal{I}} \overrightarrow{\boldsymbol{\Theta}}_{i}^{t}$. Then a function $q: \mathcal{X}_{\mathcal{H}} \times \overrightarrow{\boldsymbol{\Theta}}_{\mathcal{I}}^{t} \times$ $\mathcal{A}_{\mathcal{I}} \rightarrow \mathbb{R}$ has state factor scope $\mathbb{X}(q)=\mathcal{H}$ and agent scope $\mathbb{A}(q)=\mathcal{I}$.


Figure 2: Interaction graphs for $t=h-2$ (left) and $t=h-1$ (right) of the FFF problem.

This definition has the most general form of value function we will use, involving state-factors, action-observation histories and actions. It also applies when not all of these elements are used. For example, the function $R^{1}\left(f l_{\{1,2\}}, a_{1}\right)$ has scope $\mathbb{X}\left(R^{1}\right)=\{1,2\}$ and $\mathbb{A}\left(R^{1}\right)=\{1\}$.

Now we can formalize the $e$-th component of the Q -value function for joint policy $\pi$ of the last stage $h-1$.

$$
\begin{equation*}
Q_{\pi}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{h-1}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{R}}^{h-1}, \mathbf{a}_{\mathbb{A}_{e}^{R}}^{h-1}\right) \equiv R^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{h-1}, \mathbf{a}_{\mathbb{A}_{e}^{R}}^{h-1}\right), e=1 \ldots \rho \tag{1}
\end{equation*}
$$

where $\mathbb{X}_{e}^{R}=\mathbb{X}\left(R^{e}\right)$ and $\mathbb{A}_{e}^{R}=\mathbb{A}\left(R^{e}\right)$. E.g., in the FFF problem $Q_{\pi}^{1}\left(f l_{\{1,2\}}, \vec{\theta}_{1}^{h-1}, a_{1}\right) \equiv R^{1}\left(f l_{\{1,2\}}, a_{1}\right)$. The scope of $Q_{\pi}^{1}$ is shown in the right diagram of Figure 1 at $h-1$.

This Q-value function can be used to define an interaction hyper-graph $G=\langle\mathcal{A} g, \mathcal{E}\rangle$. In this graph, the nodes correspond with agents and the $\rho$ hyper-edges $e \in \mathcal{E}$ (edges that can involve more than two nodes) correspond to the $\rho$ local Q-value functions $Q^{e, h-1}$ for stage $h-1 .^{1}$ In particular, an edge $e$ connects $\mathbb{A}\left(Q^{e, h-1}\right)$, the agents involved in $Q^{e, h-1}$. For example, the interaction hyper-graph for the last two stages of the FFF problem is shown in Figure 2. It is the independence represented by sparsity of such interaction graphs that we will exploit in section 4 to reduce computation costs.

Given $Q^{e, h-1}$, we consider the Q -value function at $h-$ 2, which will clearly depend on $P\left(\mathbf{x}_{\mathrm{X}_{e}^{R}}^{h-1}, \mathbf{o}_{\mathbb{A}_{e}^{R}}^{h-1} \mid \cdot\right)$ where $(\cdot)$ represents any factors of influence at stage $h-2$. These factors of influence are defined by the scope backup.

Definition 3.2 (Scope backup) Given a set $\mathcal{K}$ of particular state factors and observations of particular agents $\mathcal{K} \subseteq\left\{\mathcal{X}_{1}^{\prime}, \ldots, \mathcal{X}_{|\mathcal{X}|}^{\prime}, \mathcal{O}_{1}, \ldots, \mathcal{O}_{n}\right\}$, the scope backup operator $\Gamma(\mathcal{K})$ returns the subset of variables of the previous stage $\left\{\mathcal{X}_{1}, \ldots, \mathcal{X}_{|\mathcal{X}|}, \mathcal{A}_{1} \ldots, \mathcal{A}_{n}\right\}$ (i.e., from the left side of

[^0]the DBN ) that are ancestors of variables in $\mathcal{K}$. We discriminate the state-factor component and the agent component of the scope backup:
\[

$$
\begin{aligned}
& \Gamma^{\mathbb{X}}(\mathcal{K})=\left\{j \mid \exists_{Y \in \mathcal{K}} \text { s.t. } \mathcal{X}_{j}=\operatorname{ancestor}(Y)\right\}, \\
& \Gamma^{\mathbb{A}}(\mathcal{K})=\left\{j \mid \exists_{Y \in \mathcal{K}} \text { s.t. } \mathcal{A}_{j}=\operatorname{ancestor}(Y)\right\} .
\end{aligned}
$$
\]

For example, in FFF the probability of $f l_{\{1,2\}}^{h-1}, o_{1}^{h-1}$ depends only on $f l_{\{1,2,3\}}^{h-2}$ and actions $\mathbf{a}_{\{1,2\}}^{h-2}$ as shown at $h-2$ in Figure 1 (right). When the policy is fixed, the probability of actions $\mathbf{a}_{\{1,2\}}^{h-2}$ is determined by $\mathbf{o}_{\{1,2\}}^{h-2}$. In turn, to determine $P\left(f l_{\{1,2,3\}}^{h-2}, \mathbf{o}_{\{1,2\}}^{h-2} \mid \cdot\right)$ we need to consider all variables of influence at stage $h-3$. As indicated in Figure 1, this set encompasses all state factors and agents.

### 3.2 Decomposition of Value Functions

In order to more formally define factored value functions, we need to specify the probability of states and joint actionobservation histories. In order to do so, we will consider joint policies that are partially specified with respect to time. In particular, the past joint policy specifies a policy for each agent up to stage $t$, i.e., $\varphi^{t}=\left(\delta^{0} \ldots \delta^{t-1}\right)$, where $\delta^{t^{\prime}}=\left\langle\delta_{1}^{t^{\prime}}, \ldots, \delta_{n}^{t^{\prime}}\right\rangle$ is a joint decision rule for stage $t^{\prime}$. An individual decision rule for stage $t^{\prime}$ maps individual length- $t^{\prime}$ observation histories to actions $\delta_{i}^{t^{\prime}}: \overrightarrow{\mathcal{O}}_{i}^{t^{\prime}} \rightarrow \mathcal{A}_{i}$. Now we can recursively define the following joint distribution

$$
\begin{align*}
P\left(s^{t}, \overrightarrow{\boldsymbol{\theta}}^{t} \mid b^{0}, \varphi^{t}\right)= & P\left(\mathbf{o}^{t} \mid \mathbf{a}^{t-1}, s^{t}\right) \sum_{s^{t-1} \in \mathcal{S}} P\left(s^{t} \mid s^{t-1}, \mathbf{a}^{t-1}\right) \\
& P_{\varphi^{t}}\left(\mathbf{a}^{t-1} \mid \overrightarrow{\boldsymbol{\theta}}^{t-1}\right) P\left(s^{t-1}, \overrightarrow{\boldsymbol{\theta}}^{t-1} \mid b^{0}, \varphi^{t}\right) \tag{2}
\end{align*}
$$

When we are interested in a subset of state factors and agents we can marginalize as follows: ${ }^{2}$

$$
\begin{equation*}
P\left(\mathbf{x}_{f}^{t}, \overrightarrow{\boldsymbol{\theta}}_{g}^{t} \mid b^{0}, \varphi^{t}\right)=\sum_{\mathbf{x}_{\vec{f}}^{t}} \sum_{\overrightarrow{\boldsymbol{\theta}}_{\bar{g}}^{t}} P\left(\mathbf{x}_{f}^{t}, \mathbf{x}_{f}^{t}, \overrightarrow{\boldsymbol{\theta}}_{g}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\bar{g}}^{t} \mid b^{0}, \varphi^{t}\right) \tag{3}
\end{equation*}
$$

We will write $V^{t}(\pi)$ for the expected return of $\pi$ over stages $t \ldots(h-1)$. The value function $V^{h-1}(\pi)$ for the last stage $t=h-1$ can now be decomposed as

$$
\begin{align*}
& V^{h-1}(\pi)=\sum_{e \in \mathcal{E}} V^{e, h-1}(\pi)=\sum_{e \in \mathcal{E}} \sum_{\mathbf{x}_{\mathbb{X}_{e}^{R}}^{h-1}} \sum_{\overrightarrow{\mathbb{A}}_{e}^{h}} \sum_{i-1} \\
& P\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{h-1}, \overrightarrow{\boldsymbol{A}}_{\mathbb{A}_{e}^{R}}^{h-1} \mid b^{0}, \pi\right) Q_{\pi}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{h-1}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{R}}^{h-1}, \pi_{\mathbb{A}_{e}^{R}}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{R}}^{h-1}\right)\right) . \tag{4}
\end{align*}
$$

Such a decomposition is possible for every stage $t$.
Proposition 3.1 Given a decomposed immediate reward function, $V^{t}(\pi)$ of a finite-horizon Dec-POMDP is decomposable for any $t$.
Proof The proof is analogous to the one given in [11] and is omitted here.

[^1]Let us write $\mathbb{X}_{e}, \mathbb{A}_{e}$ for the state factor and agent scope of the Q -value function under concern. That is, $\mathbb{X}_{e} \equiv \mathbb{X}\left(Q^{e, t}\right)$ and $\mathbb{A}_{e} \equiv \mathbb{A}\left(Q^{e, t}\right)$. We use similar shorthand $\mathbb{X}_{e}^{\prime}, \mathbb{A}_{e}^{\prime}$ for the next-stage Q -values. Now, the value of a joint policy $\pi$ for a particular stage is given by

$$
\begin{align*}
V^{t}(\pi)=\sum_{e \in \mathcal{E}} & \sum_{\mathbf{x}_{\mathbb{X}_{e}}^{t}} \sum_{\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}} P\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t} \mid b^{0}, \pi\right) \\
& Q_{\pi}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \pi_{\mathbb{A}_{e}}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}\right)\right)=\sum_{e \in \mathcal{E}} V^{e, t}(\pi) . \tag{5}
\end{align*}
$$

where, using shorthand notation $\Gamma^{\mathbb{X}}=\Gamma^{\mathbb{X}}\left(\mathbf{x}_{\mathbb{X}_{e}^{\prime}}^{t+1} \cup \mathbf{o}_{\mathbb{A}_{e}^{\prime}}^{t+1}\right)$ and $\Gamma^{\mathbb{A}}=\Gamma^{\mathbb{A}}\left(\mathbf{x}_{\mathbb{X}_{e}}^{t+1} \cup \mathbf{o}_{\mathbb{A}_{e}}^{t+1}\right)$,

$$
\begin{align*}
& Q_{\pi}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \mathbf{a}_{\mathbb{A}_{e}}\right)=R^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{t}, \mathbf{a}_{\mathbb{A}_{e}^{R}}\right)+\sum_{\mathbf{x}_{\mathbb{X}_{e}^{\prime}}^{t+1}} \sum_{\mathbf{o}_{\mathbb{A}_{e}^{\prime}}^{t+1}} \\
& \quad P\left(\mathbf{x}_{\mathbb{X}_{e}^{\prime}}^{t+1}, \mathbf{o}_{\mathbb{A}_{e}^{\prime}}^{t+1} \mid \mathbf{x}_{\Gamma^{\mathbb{X}}}^{t}, \mathbf{a}_{\Gamma^{\mathbb{A}}}\right) Q_{\pi}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{\prime}}^{t+1}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{\prime}}^{t+1}, \pi_{\mathbb{A}_{e}^{\prime}}^{t}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{\prime}}^{t+1}\right)\right) \tag{6}
\end{align*}
$$

Note that the scope of $Q_{\pi}^{e, t}$ is given by

$$
\begin{align*}
\mathbb{X}_{e} & \equiv \mathbb{X}_{e}^{R} \cup \Gamma^{\mathbb{X}}  \tag{7}\\
\mathbb{A}_{e} & \equiv \mathbb{A}_{e}^{R} \cup \Gamma^{\mathbb{A}} \cup \mathbb{A}_{e}^{\prime} \tag{8}
\end{align*}
$$

It is necessary that $\mathbb{A}_{e}^{\prime} \subseteq \mathbb{A}_{e}$, otherwise $\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{\prime}}^{t+1}$ can be undefined. For instance, for stage $h-2$, the scope of the value function corresponding to reward function 1 is given by

$$
\begin{aligned}
\mathbb{X}\left(Q_{\pi}^{1, h-2}\right) & =\mathbb{X}\left(R^{1}\right) \cup \Gamma^{\mathbb{X}}\left(\left\{f l_{1}, f l_{2}, o_{1}\right\}\right) \\
& =\left\{f l_{1}, f l_{2}\right\} \cup\left\{f l_{1}, f l_{2}, f l_{3}\right\}=\left\{f l_{1}, f l_{2}, f l_{3}\right\} \\
\mathbb{A}\left(Q_{\pi}^{1, h-2}\right) & =\mathbb{A}\left(R^{1}\right) \cup \Gamma^{\mathbb{A}}\left(\left\{f l_{1}, f l_{2}, o_{1}\right\}\right) \cup \mathbb{A}\left(Q_{\pi}^{1, h-1}\right) \\
& =\{1\} \cup\{1,2\} \cup\{1\}=\{1,2\}
\end{aligned}
$$

We can identify the scopes for the other $3(e=2 \ldots 4)$ value functions $Q_{\pi}^{e, h-2}$ in a similar way. The agent scopes define a new interaction graph for stage $h-2$ of the FFF problem, as illustrated in Figure 2. The figure clearly shows that $Q^{2}$ and $Q^{3}$ both involve all agents, which means that at this stage there is no reduction possible. However, our representation will allow us to exploit independence for the last stage, which is the bottleneck of computation. Also, to avoid duplicate work for stage $h-2$ it is possible to merge the Q-functions by adding them.

### 3.3 Locality of Interaction

When assuming transition and observation independence, the probability in (6) reduces easily: each variable $x_{i}^{t+1}$ (representing a local state for agent $i$ ) is only dependent on $x_{i}^{t}$ (its value at the previous state), and each $o_{i}$ is only dependent on $x_{i}^{t+1} .{ }^{3}$ This means that the scope does not expand and, as a result, Nair et al. [11] have a stationary interaction graph. Let $\mathcal{N}_{i}$ denote the neighbors of agent $i$, including $i$ itself, in this graph. They define the local neighborhood util$i t y$ of an agent $i$ as the expected return for all the edges that contain agent $i$ :

$$
\begin{equation*}
V\left(\pi_{\mathcal{N}_{i}}\right)=\sum_{s} b^{0}(s) \sum_{e \text { s.t. } i \in e} Q_{\pi}^{e, 0}\left(\mathbf{x}_{\mathbb{X}_{e}}^{0}, \overrightarrow{\boldsymbol{\theta}}_{\emptyset}, \pi_{\mathbb{A}_{e}}\left(\overrightarrow{\boldsymbol{\theta}}_{\emptyset}\right)\right) . \tag{9}
\end{equation*}
$$

[^2]Consequently they show that when an agent $j \notin \mathcal{N}_{i}$ changes its policy, $V^{t}\left(\pi_{\mathcal{N}_{i}}\right)$ is not affected, a property they refer to as locality of interaction.

In our more general case, such a notion of locality of interaction over full policies is not properly defined, simply because the interaction graph and hence agent $i$ 's neighborhood can be different at every stage. However, at a particular stage $t$, given a past joint policy $\varphi^{t}$ that has been followed, a future policy $\psi^{t}$ that will be followed and $\pi=\left(\varphi^{t}, \delta^{t}, \psi^{t}\right)$, we can define

$$
\begin{align*}
& V^{t}\left(\delta_{\mathcal{N}_{i}}^{t}\right)=\sum_{e \text { s.t. }} \sum_{i \in \mathbb{A}_{(Q}^{e}, t} \sum_{\mathbf{x}_{\mathbb{x}_{e}}^{t}} \sum_{\overrightarrow{\boldsymbol{A}}_{\mathbb{A}_{e}}^{t}} \\
&  \tag{10}\\
& P\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t} \mid b^{0}, \varphi^{t}\right) Q_{\pi_{\mathbb{A}_{e}}}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \delta_{\mathbb{A}_{e}}^{t}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}\right)\right) .
\end{align*}
$$

Consequently, locality of interaction holds within each stage. The definition of $\mathbb{A}_{e}(8)$ implies that the agent scope is nondecreasing when going back in time ${ }^{4}$, which means that the scope of dependency is typically larger in earlier stages (as also described in [10]) and the interaction graphs for those earlier stages are denser, as shown in Figure 2.

However, even when there is no independence in transitions and observations (i.e., the scope includes all factors and agents with just one backup), the decomposed value function for the last stage still has limited scope, allowing for a significant speedup of the last stage, and thus for the entire computation.

## 4. FACTORED DEC-POMDPS VIA GRAPHICAL BAYESIAN GAMES

A non-factored Dec-POMDP can be represented by a series of Bayesian games, one for each time step [13]. Consequently, a policy can be found by solving the BGs for stage $0,1, \ldots, h-1$ consecutively, a procedure we refer to as forward-sweep policy computation. However, the cost of solving these BGs grows doubly-exponentially with the horizon. The number of joint policies for a BG for the last stage, and thus the cost of optimally solving such a BG, is

$$
\begin{equation*}
O\left(\left|\mathcal{A}_{*}\right|^{n\left(\left|\mathcal{O}_{*}\right|^{h-1}\right)}\right) \tag{11}
\end{equation*}
$$

where $\mathcal{A}_{*}$ and $\mathcal{O}_{*}$ denote the largest individual action and observation sets. In this section we provide tools to exploit the independence between agents as represented by the interaction graphs, reducing the agent component $n$ in (11). First, we show that a factored Dec-POMDP with additively separable rewards can be modeled using a series of collaborative graphical BGs. Then we show how such CGBGs can be solved efficiently.

### 4.1 Collaborative Graphical Bayesian Games

Here we define the collaborative graphical BG as a graphical $\mathrm{BG}[16,17]$, in which the agents try to optimize the sum of local rewards, rather than an individual payoff function.

Definition 4.1 A collaborative graphical $B G(C G B G)$, is a tuple $\left\langle\mathcal{A} g,\left\langle\mathcal{A}_{1}, \ldots, \mathcal{A}_{n}\right\rangle, \boldsymbol{\Theta}, P(\boldsymbol{\Theta}),\left\langle u_{1}, \ldots, u_{\rho}\right\rangle\right\rangle$, with:

[^3]| $Q^{e=1, t=1}\left(\vec{\theta}_{1}^{1}, a_{1}^{1}\right)$ |  |  |
| :---: | :--- | :--- |
| $\vec{\theta}_{1}^{1}$ |  |  |
| $\left(H_{1}, F\right)$ | $H_{1}$ | -0.25 |
|  | $H_{2}$ | -1.10 |
| $\left(H_{1}, N\right)$ | $H_{1}$ | -0.14 |
|  | $H_{2}$ | -0.79 |


| $Q^{e=2, t=1}\left(\overrightarrow{\boldsymbol{\theta}}_{\{1,2\}}^{1}, \mathbf{a}_{\{1,2\}}^{1}\right)$ |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\vec{\theta}_{2}^{1}$ | $\left(H_{2}, F\right)$ |  | $\left(H_{2}, N\right)$ |  |
| $\vec{\theta}_{1}^{1}$ |  | $H_{2}$ | $H_{3}$ | $H_{2}$ | $H_{3}$ |
| $\left(H_{1}, F\right)$ | $H_{1}$ | -0.55 | -1.60 | -0.50 | -1.50 |
|  | $H_{2}$ | 0 | -0.55 | 0 | -0.50 |
| $\left(H_{1}, N\right)$ | $H_{1}$ | -0.16 | -1.10 | -0.14 | -1.00 |
|  | $H_{2}$ | 0 | -0.16 | 0 | -0.14 |



| $\vec{\theta}_{3}{ }^{1}$ | $Q^{e=3, t=1}\left(\overrightarrow{\boldsymbol{\theta}}_{\{2,3\}}^{1}, \mathbf{a}_{\{2,3\}}^{1}\right)$ |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\vec{\theta}_{2}^{1}$ | $\left(H_{2}, F\right)$ |  | $\left(H_{2}, N\right)$ |  |
|  |  | $\mathrm{H}_{2}$ | $\mathrm{H}_{3}$ | $\mathrm{H}_{2}$ | $\mathrm{H}_{3}$ |
| $\left(H_{4}, F\right)$ | $\mathrm{H}_{3}$ | -1.10 | 0 | -0.71 | 0 |
|  | $\mathrm{H}_{4}$ | -1.90 | -1.10 | -1.70 | -0.71 |
| $\left(H_{4}, N\right)$ | $\mathrm{H}_{3}$ | -1.00 | 0 | -0.58 | 0 |
|  | $\mathrm{H}_{4}$ | -1.90 | -1.00 | $-1.60$ | -0.58 |

Figure 3: The CGBG for stage $t=1$ of the horizon $h=2$ FFF problem for past joint policy $\varphi^{1}=$ $\left\langle H_{1}, H_{2}, H_{4}\right\rangle$. The shading indicates an arbitrary policy for agent 2.

- $\mathcal{A} g=\{1, \ldots, n\}$ is the set of agents.
- $\mathcal{A}_{i}$ is the set of actions of agent $i . \mathcal{A}=\times_{i=1}^{n} \mathcal{A}_{i}$ is the set of joint actions.
- $\Theta=x_{i} \Theta_{i}$ is the set of joint types over which a probability function $P(\boldsymbol{\Theta})$ is given (and is common knowledge).
- $u_{k}$ are the local payoff functions.

A CGBG is collaborative, which means that all agents try to maximize the same payoff. It is also graphical, which means that the payoff can be decomposed into a sum of local payoff functions $u_{k}$, each of which can depend on only a subset of agents, just like in factored Dec-POMDPs. Formally, a local payoff function is specified as a mapping from the types and actions of a subset of agents to reals: $u_{k}: \boldsymbol{\theta}_{\mathcal{I}} \times \mathbf{a}_{\mathcal{I}} \rightarrow \mathbb{R}$, where $\mathcal{I}=\mathbb{A}\left(u_{k}\right)$ is the subset of agent indices that participate in payoff function $k$.

As for factored Dec-POMDPs, it is possible to construct a hyper-graph $G=\langle\mathcal{A} g, \mathcal{E}\rangle$ from the set of payoff functions such that the agent scope of each local payoff function corresponds to a hyper-edge $e \in \mathcal{E}$. We will write $u_{e}$ as the local payoff function corresponding to edge $e$ and $\mathbb{A}_{e}^{u}=\mathbb{A}\left(u_{e}\right)$.

Finally, the goal is to maximize the expected sum of rewards:

$$
\begin{equation*}
\beta^{*}=\underset{\beta}{\arg \max } \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} P(\boldsymbol{\theta}) \sum_{e \in \mathcal{E}} u_{e}\left(\boldsymbol{\theta}_{\mathbb{A}_{e}^{u}}, \beta_{\mathbb{A}_{e}^{u}}\left(\boldsymbol{\theta}_{\mathbb{A}_{e}^{u}}\right)\right), \tag{12}
\end{equation*}
$$

where $\left.\beta_{\mathbb{A}_{e}^{u}}\left(\boldsymbol{\theta}_{\mathbb{A}_{e}^{u}}\right)\right)=\left\langle\beta_{i}\left(\theta_{i}\right)\right\rangle_{i \in \mathbb{A}_{e}^{u}}$ is the joint action of that subset resulting from application of the individual BG-policies $\beta_{i}$. It can be shown that $\beta^{*}$ constitutes a Pareto-optimal Bayes-Nash equilibrium [12, 20].

### 4.2 A stage as a Collaborative Graphical BG

Here we explain how a CGBG can model a stage of a factored Dec-POMDP. For each stage $t$, we want to find a joint decision rule which is specified by $\delta^{t} \equiv \beta^{t, *}$ the solution of the CGBG for that stage. Also, the private information each agent $i$ holds is its action-observation history $\vec{\theta}_{i}^{t}$, which therefore naturally corresponds to agent $i$ 's type: $\theta_{i} \equiv \vec{\theta}_{i}^{t}$. The past joint policy $\varphi^{t}$ is available when planning, therefore
the probability of joint types is specified by $P\left(\overrightarrow{\boldsymbol{\theta}}^{t} \mid b^{0}, \varphi^{t}\right)=$ $\sum_{s^{t}} P\left(s^{t}, \overrightarrow{\boldsymbol{\theta}}^{t} \mid b^{0}, \varphi^{t}\right)$ which is given by (2).

What is left to be determined is the set of payoff functions. For now, we will assume that there is a factored Q -value function $Q^{t}\left(\overrightarrow{\boldsymbol{\theta}}^{t}, \mathbf{a}\right)$ for the Dec-POMDP that can be used as the payoff function. We will describe such factored Q -value functions in section 4.4.

In Figure 3, a CGBG for stage $t=1$ of the horizon $h=2$ FFF problem is shown. The past joint policy $\varphi^{1}$ is a joint decision rule for the first stage, which is equivalent to joint action $\left\langle H_{1}, H_{2}, H_{4}\right\rangle$.

### 4.3 Solving the CGBGs

Given a past joint policy $\varphi^{t}$, the solution of a CGBG is

$$
\begin{align*}
\beta^{t, *} & =\underset{\beta^{t}}{\arg \max } \sum_{e \in \mathcal{E}} \sum_{\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}} P\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t} \mid \varphi^{t}, b^{0}\right) Q^{e, t}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \beta_{\mathbb{A}_{e}}^{t}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}\right)\right) \\
& =\underset{\beta^{t}}{\arg \max } \sum_{e \in \mathcal{E}} V_{\varphi^{t}}^{e, t}\left(\beta_{\mathbb{A}_{e}}^{t}\right) \tag{13}
\end{align*}
$$

where $\mathbb{A}_{e}=\mathbb{A}\left(Q^{e, t}\right)$ and where the probability can be computed as a marginal of (2). In contrast to solving a regular BG, the maximization of (13) does not require enumeration of all joint BG policies. For example, in Figure 3, agent 1 and 3 can separately compute their best response to the shaded policy of agent 2 .

More formally, we can optimally solve the CGBG using non-serial dynamic programming (NDP), also known as variable elimination (VE) [6, 20]. We will illustrate this technique for the CGBG of Figure 3. Let us write $V^{e}$ for $V_{\varphi^{1}}^{e, t=1}$ with $\varphi^{1}=\left\langle H_{1}, H_{2}, H_{4}\right\rangle$, the maximization we are then performing is

$$
\begin{equation*}
\max _{\beta}\left[V^{1}\left(\beta_{1}\right)+V^{2}\left(\beta_{\{1,2\}}\right)+V^{3}\left(\beta_{\{2,3\}}\right)+V^{4}\left(\beta_{3}\right)\right], \tag{14}
\end{equation*}
$$

where $\beta_{\{1,2\}}=\left\langle\beta_{1}, \beta_{2}\right\rangle$. Now, by isolating the functions in which agent 1 participates

$$
\max _{\beta_{\{2,3\}}}\left[V^{3}\left(\beta_{\{2,3\}}\right)+V^{4}\left(\beta_{3}\right)+\max _{\beta_{1}}\left[V^{1}\left(\beta_{1}\right)+V^{2}\left(\beta_{\{1,2\}}\right)\right]\right],
$$

it is possible to split the problem in two smaller ones. NDP continues by isolating a next agent, etc. The time needed by NDP is exponential in the induced width of the interaction graph [6, 20]. Let us denote this width $w$, then we have that the complexity of solving the last stage CGBG is

$$
\begin{equation*}
O\left(n \cdot\left|\mathcal{A}_{*}\right|^{w\left(\left|\mathcal{O}_{*}\right|^{h-1}\right)}\right) \tag{15}
\end{equation*}
$$

yielding an exponential speedup over (11) as long as $n \gg w$.
Conceptually, the $V_{\varphi^{e}, t}^{e, t}$ functions from (14) define a graphical normal-form game, which is then solved by NDP. This normal-form game does not need to be constructed explicitly, however. Alternatively, it is possible to reduce to a standard GBG by considering the local neighborhood utility, defined in a similar fashion as (9), as the individual payoff function for agent $i$. Therefore all methods that will find a Bayes-Nash equilibrium for GBGs, e.g. [16, 17], can be used to find a local optimum. A different option is to solve the CGBGs using max-plus [9, 20], which only is optimal for tree-shaped graphs, but in practice works well on graphs with cycles. Also, it may be possible to employ upper bounds over BG-policies that are partially specified with respect to individual BG-policies as in SPIDER [19].

### 4.4 Factored Q-value Functions

Here we will consider the payoff functions for the CGBG for each stage, which constitute factored Q -value function for the factored Dec-POMDP. In particular, we will consider a normative description of the optimal Q -value function, which demonstrates that such a series of CGBGs exactly models a Dec-POMDP.

In Section 3.2, we derived an expression for the expected value of a given policy. This description also holds for an optimal joint policy. Let us assume we are given an optimal joint policy $\pi^{*}$. Then substitution in (5) and (6) yields a normative description of the optimal value function. Here we transform the resulting Q-value function $Q_{\pi^{*}}^{e}\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{A}}_{\mathbb{A}_{e}}^{t}, \mathbf{a}_{\mathbb{A}_{e}}\right)$ to the form $Q_{\pi *}^{e}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \mathbf{a}_{\mathbb{A}_{e}}\right)$.

It is possible to split $P\left(\mathbf{x}_{\mathbb{X}_{e}}^{t}, \overrightarrow{\boldsymbol{\theta}}_{\mathrm{A}_{e}}^{t} \mid b^{0}, \pi\right)$ in a marginal and a conditional. Also $\pi^{*}=\left(\varphi^{t, *}, \delta^{t, *}, \psi^{t}\right)$ with $\varphi^{t, *}, \psi^{t}$ the optimal past and future joint policy. This allows us to write:

$$
\begin{equation*}
V^{t}\left(\pi^{*}\right)=\sum_{e \in \mathcal{E}} \sum_{\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}} P\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t} \mid b^{0}, \varphi^{t, *}\right) Q_{\pi^{*}}^{e}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \delta_{\mathbb{A}_{e}}^{t, *}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}\right)\right) \tag{16}
\end{equation*}
$$

with

$$
\begin{align*}
& Q_{\pi^{*}}^{e}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, \mathbf{a}_{\mathbb{A}_{e}}\right)=\sum_{\mathbf{x}_{\mathbb{X}_{e}}^{t}} P\left(\mathbf{x}_{\mathbb{X}_{e}}^{t} \mid \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, b^{0}, \pi^{*}\right) R^{e}\left(\mathbf{x}_{\mathbb{X}_{e}^{R}}^{t}, \mathbf{a}_{\mathbb{A}_{e}^{R}}\right)+ \\
& \sum_{\mathbf{o}_{\mathbb{A}_{e}^{\prime}}^{t+1}} P\left(\mathbf{o}_{\mathbb{A}_{e}^{\prime}}^{t+1} \mid \overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}}^{t}, b^{0}, \pi, \mathbf{a}_{\mathbb{A}_{e}}\right) Q_{\pi^{*}}^{e}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{\prime}}^{t+1}, \pi_{\mathbb{A}_{e}^{\prime}}^{t+1}\left(\overrightarrow{\boldsymbol{\theta}}_{\mathbb{A}_{e}^{\prime}}^{t+1}\right)\right) . \tag{17}
\end{align*}
$$

the Q -value function of optimal policy $\pi^{*}$.
Theorem 4.1 When using (17) as the payoff functions for the series of CGBGs representing a Dec-POMDP, forwardsweep policy computation yields an optimal policy.
Proof Following the reasoning from [12], we only have to show that, given that the optimal past policy $\varphi^{t, *}$ is followed, using (17) as the payoff function of a CGBG yields $\delta^{t, *}$. When using this factored Q -value function, the solution of that CGBG-given by (13) - maximizes the quantity from (16). Now, because the optimal policy per definition maximizes the expected value, $\delta^{t, *}$ per definition maximizes (16) and as a result we have that $\beta^{t, *}$ found by (13) is optimal.

The implication of this theorem is that modeling a factored Dec-POMDP using CGBGs is exact. Unfortunately, as in the non-factored case, computation of the optimal Qvalue function is impractical and seems to require the optimal policy.

## 5. HEURISTIC POLICY SEARCH

Because it is impractical to compute (16), it is not practical to apply forward-sweep policy computation to compute an optimal policy. However, for the last-and most computationally demanding - stage we in fact do have the factored optimal Q-value function available: it is given by the immediate reward function. We propose a policy search method based on Generalized MAA* (GMAA*) [12] that employs a factored Q-value function for, and thus exploits locality of interaction in, the last stage.

GMAA* consists of three procedures. The first procedure iterates over a pool of partial joint policies, pruning this

```
Algorithm 1 GMAA*
    \(\mathrm{v}^{\star}:=-\infty\)
    \(\overline{\mathrm{P}}:=\left\{\varphi^{0}=()\right\}\)
    repeat
        \(\varphi^{t}:=\operatorname{Select}(\mathrm{P})\)
        \(\Phi_{\text {Next }}:=\operatorname{Next}\left(\varphi^{t}\right)\)
        if \(\Phi_{\text {Next }}\) contains full policies \(\Pi_{\text {Next }} \subseteq \Phi_{\text {Next }}\) then
            \(\pi^{\prime}:=\arg \max _{\pi \in \Pi_{\mathrm{Next}}} V(\pi)\)
            if \(V\left(\pi^{\prime}\right)>\underline{\mathrm{v}}^{\star}\) then
                \(\underline{\mathrm{v}}^{\star}:=V\left(\bar{\pi}^{\prime}\right)\)
                \(\pi^{\star}:=\pi^{\prime}\)
                \(\mathrm{P}:=\left\{\varphi \in \mathrm{P} \mid \widehat{V}(\varphi)>\underline{\mathrm{v}}^{\star}\right\}\{\) prune P\(\}\)
            end if
            \(\Phi_{\text {Next }}=\Phi_{\text {Next }} \backslash \Pi_{\text {Next }}\) \{remove full policies \(\}\)
        end if
        \(\mathrm{P}:=\left(\mathrm{P} \backslash \varphi^{t}\right) \cup\left\{\varphi \in \Phi_{\text {Next }} \mid \widehat{V}(\varphi)>\underline{\mathrm{v}}^{\star}\right\}\)
    until P is empty
```

pool whenever possible. This is the core of GMAA* and is fixed, while the other two procedures can be performed in many ways. The second procedure, Select, selects which policy to process next and thus determines the type of search (e.g., depth-first, breadth-first, A*-like). The third procedure, Next, constructs a new set of (partial) joint policies based on the selected policy.

The core procedure of GMAA* (see Algorithm 1) maintains a set of partial joint policies $\varphi$ together with their heuristic values $\widehat{V}(\varphi)$. This 'policy pool' P is initialized with a completely unspecified joint policy $\varphi^{0}=()$ and the maximum lower bound (found so far) $\underline{v}^{\star}$ is set to $-\infty$. $\pi^{\star}$ denotes the best joint policy found so far. First, Select selects a partial joint policy $\varphi$ from P. We assume that, as in MAA* [18], the partial policy with the highest heuristic value is selected. Next, the selected policy is processed by the policy search operator Next, which returns a set of (partial) joint policies $\Phi_{\text {Next }}$ and their heuristic values. When Next returns one or more full policies $\pi \in \Phi_{\text {Next }}$, the provided values $\widehat{V}(\pi)=V(\pi)$ are a lower bound for an optimal joint policy, which can be used to prune the search space. Any found partial joint policies $\varphi \in \Phi_{\text {Next }}$ with a heuristic value $\widehat{V}(\varphi)>\underline{\mathrm{v}}^{\star}$ are added to P . Then process is repeated until the policy pool is empty.

The original MAA* can be seen as an instance of the generalized case with a Next operator that, given a partial joint policy $\varphi^{t}=\left(\delta^{0}, \ldots, \delta^{t-1}\right)$, constructs and returns all possible $\varphi^{t+1}=\left(\delta^{0}, \ldots, \delta^{t-1}, \delta^{t}\right)$ that are consistent with $\varphi^{t}$ along with their heuristic values. (This corresponds with constructing a regular BG for stage $t$ and evaluating all BGpolicies $\delta^{t}$.) Because the MAA* Next operator returns all 'children' ( $\varphi^{t+1}$ consistent with $\varphi^{t}$ ), it is guaranteed to find an optimal joint policy, as long as the used heuristic is an upper bound to the optimal value function [18].

We propose GMAA*-ELSI, a method for exploiting the last-stage independence (ELSI), which is described by Algorithms 1 and 2 jointly. For stages $0, \ldots, h-2$ regular MAA* is used, while for the last stage $t=h-1$, a CGBG is constructed rather than a regular BG. This CGBG is then solved using NDP and the solution of this CGBG used to construct a full-length joint policy $\pi$.

## 6. EXPERIMENTS

In Table 1 we compare results for GMAA*-ELSI with regular GMAA* (that solves a regular BG even for the last

```
Algorithm \(2 \operatorname{Next}\left(\varphi^{t}\right)\)-ELSI
    if \(t<h-1\) then
        \(\{\) Perform regular MAA* for \(t=0, \ldots, h-2\}\)
        \(\Phi^{t+1}=\) construct all policies \(\varphi^{t+1}\) as in [18, 13]
        return \(\Phi^{t+1}\)
    else
        \{For \(t=h-1\) exploit independence \}
        \(G=\operatorname{ConstructCGBG}\left(\varphi^{h-1}, b^{0}\right)\)
        \(\delta^{h-1}=\operatorname{NDPSolve}(G)\)
        \(\pi=\left(\varphi^{h-1}, \delta^{h-1}\right)\)
        return \(\pi\)
    end if
```

|  | $V^{*}$ | $T_{\text {GMAA }}$ | $T_{\text {GMAA*-ELSI }}$ |
| ---: | ---: | ---: | ---: |
| $h=2$ | -5.213685 | $\leq 0.01 \mathrm{~s}$ | 0.03 s |
| $h=3$ | -6.654551 | 0.15 s | 0.46 s |
| $h=4$ | -7.462685 | 4834.07 s | 12.67 s |

Table 1: Results for the FFF for horizon 2, 3, and 4. Column $V^{*}$ indicates the value of an optimal policy, $T_{\text {GMAA }}$ the computation time for the plain GMAA* algorithm, and $T_{\text {GMAA*-ELSI }}$ the computation time for our proposed method.
stage) for FFF, introduced in Section 2.3. We used $\mathrm{Q}_{\mathrm{BG}}$ [13] as the heuristic Q -value function for all time steps of GMAA*, and for time steps $0-h-2$ of GMAA*-ELSI. As both methods are optimal, they compute identical optimal policies, whose values are shown in the first column. The remaining two columns show computation time of both methods. For the low horizons, we can see that GMAA* is faster, as the Bayesian games to be solved are small, and not worth the overhead caused by the construction and solution of CGBGs in GMAA*-ELSI's last time step. However, for horizon 4 we can see the dramatic speedup provided by exploiting the local interactions in this factored Dec-POMDP.

## 7. DISCUSSION AND FUTURE WORK

In this paper we presented a framework for exploiting locality of interaction in factored Dec-POMDPs. By formalizing the dependence propagation through time, and analyzing the dependencies per stage, it is possible to construct an interaction graph for each stage. The independence represented by this interaction graph can then be exploited by efficiently solving the corresponding collaborative graphical Bayesian game (CGBG). Also, we presented a formulation of factored Q -value functions that can be used as the payoff function for these CGBGs, and have shown how GMAA*, an existing heuristic search method, can be extended to exploit the independence in the last stage by using CGBGs.

The analysis of this paper strengthens our belief that general Dec-POMDPs should be tackled per stage, rather than trying to find individual full-length policies in an iterative procedure. It also shows that solution of CGBGs scale exponentially in the induced width of the interaction graph $(w)$, but only polynomially in the number of agents $(n)$. In accordance, a preliminary experimental evaluation shows that the resulting algorithm GMAA*-ELSI obtains a speedup of two orders of magnitude. To our knowledge, these are the first experimental results for general (factored) Dec-POMDPs with more than 2 agents.

There are two main directions for future research. The first is the identification and evaluation of approximate fac-
tored Q-value functions. Analogous to our previous work [13, 12], one idea is to first use dynamic programming to compute a Q-function for the underlying (factored) MDP or POMDP and then 'project' this (typically non-decomposable) function onto an interaction graph. Such projection operators can possibly be integrated in the dynamic programming backups [9]. Using such an approach, it will be possible to exploit independence at all stages, by computing CGBGs for stages $0 \ldots h-1$.

The second main research direction is the unification of two previous heuristic search methods for Dec-POMDPs. On the one hand MAA* [18] searches over joint policies which are partially specified with respect to time (e.g., a policy only specified for the first stage). On the other hand SPIDER [19] searches over joint policies that are partially specified w.r.t. the individual agents (e.g., a joint policy only specifying an individual policy for agent 1). The algorithm proposed here takes a step towards this unification: intuitively GMAA* corresponds to the former, while solution of the CGBGs with NDP corresponds to the latter. Future algorithms may truly unify the two approaches.

## Acknowledgments

The research reported here is part of the Interactive Collaborative Information Systems (ICIS) project, supported by the Dutch Ministry of Economic Affairs, grant nr: BSIK03024. This work was partially supported by Fundação para a Ciência e a Tecnologia (ISR/IST pluriannual funding) through the POS_Conhecimento Program that includes FEDER funds and through grant PTDC/EEA-ACR/73266/2006.

## 8. REFERENCES

[1] R. Becker, S. Zilberstein, V. Lesser, and C. V. Goldman. Solving transition independent decentralized Markov decision processes. Journal of Artificial Intelligence Research (JAIR), 22:423-455, December 2004.
[2] D. S. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. Math. Oper. Res., 27(4):819-840, 2002.
[3] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. Journal of Artificial Intelligence Research, 11:1-94, 1999.
[4] C. Boutilier, R. Dearden, and M. Goldszmidt. Stochastic dynamic programming with factored representations. Artificial Intelligence, 121(1-2):49-107, 2000.
[5] X. Boyen and D. Koller. Tractable inference for complex stochastic processes. In Proc. of Uncertainty in Artificial Intelligence, pages 33-42. San Francisco: Morgan Kaufmann, 1998.
[6] C. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored MDPs. In Advances in Neural Information Processing Systems 14, pages 1523-1530, 2002.
[7] E. A. Hansen, D. S. Bernstein, and S. Zilberstein. Dynamic programming for partially observable stochastic games. In AAAI '04: Proceedings of the Nineteenth National Conference on Artificial Intelligence, pages 709-715, 2004.
[8] E. A. Hansen and Z. Feng. Dynamic programming for POMDPs using a factored state representation. In Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems, pages 130-139, Apr. 2000.
[9] J. R. Kok and N. Vlassis. Collaborative multiagent reinforcement learning by payoff propagation. Journal of Machine Learning Research, 7:1789-1828, 2006.
[10] D. Koller and R. Parr. Computing factored value functions for policies in structured MDPs. In Proc. Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), pages 1332-1339, 1999.
[11] R. Nair, P. Varakantham, M. Tambe, and M. Yokoo. Networked distributed POMDPs: A synthesis of distributed constraint optimization and POMDPs. In Proc. of the National Conference on Artificial Intelligence, pages 133-139, 2005.
[12] F. A. Oliehoek, M. T. J. Spaan, and N. Vlassis. Optimal and approximate Q-value functions for decentralized POMDPs. Journal of Artificial Intelligence Research, 2008. (to appear).
[13] F. A. Oliehoek and N. Vlassis. Q-value functions for decentralized POMDPs. In Proc. of Int. Joint Conference on Autonomous Agents and Multi Agent Systems, pages 833-840, May 2007.
[14] Z. Rabinovich, C. V. Goldman, and J. S. Rosenschein. The complexity of multiagent systems: the price of silence. In Proc. of Int. Joint Conference on Autonomous Agents and Multi Agent Systems, pages 1102-1103, 2003.
[15] M. Roth, R. Simmons, and M. Veloso. Exploiting factored representations for decentralized execution in multi-agent teams. In Proc. of Int. Joint Conference on Autonomous Agents and Multi Agent Systems, pages 467-463, May 2007.
[16] S. Singh, V. Soni, and M. Wellman. Computing approximate Bayes-Nash equilibria in tree-games of incomplete information. In EC '04: Proceedings of the 5th ACM conference on Electronic commerce, pages 81-90, 2004.
[17] V. Soni, S. Singh, and M. Wellman. Constraint satisfaction algorithms for graphical games. In Proc. of Int. Joint Conference on Autonomous Agents and Multi Agent Systems, pages 423-430, 2007.
[18] D. Szer, F. Charpillet, and S. Zilberstein. MAA*: A heuristic search algorithm for solving decentralized POMDPs. In Proc. of Uncertainty in Artificial Intelligence, 2005.
[19] P. Varakantham, J. Marecki, Y. Yabu, M. Tambe, and M. Yokoo. Letting loose a SPIDER on a network of POMDPs: Generating quality guaranteed policies. In Proc. of Int. Joint Conference on Autonomous Agents and Multi Agent Systems, 2007.
[20] N. Vlassis. A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan \& Claypool Publishers, 2007.


[^0]:    ${ }^{1}$ We omit the stage index to Q-functions if it is implicit in their arguments, as in (1).

[^1]:    ${ }^{2}$ When there is independence between the histories in the joint probability distribution this may be exploited. However, such independence quickly vanishes due to the influence of the transition and observation model. The result is a space of joint probability distributions that is exponential in the number of state factors, which typically renders the problem intractable, although approximations with bounded error are possible [5]. We focus on settings where exact evaluations of (2) are feasible.

[^2]:    ${ }^{3}$ In fact the probabilities can also depend on an 'external' uninfluenceable state, but this is of no consequence.

[^3]:    ${ }^{4}$ In contrast, the state scope is not necessarily nondecreasing. It is, however, when the probability of all factors $x_{i}^{\prime}$ is dependent on their $x_{i}$ value in the previous stage, a property that often holds in the real world.

