

Hierarchical Decision Making in Multi-Agent Systems using Answer Set Programming

Davy Van Nieuwenborgh^{1*}, Marina De Vos², Stijn Heymans³, and Dirk Vermeir¹

¹ Dept. of Computer Science
Vrije Universiteit Brussel, VUB
Pleinlaan 2, B1050 Brussels, Belgium
{dvnieuwe, dvermeir}@vub.ac.be

² Dept. of Computer Science
University of Bath
Bath, BA2 7AY, UK
mdv@cs.bath.ac.uk

³ Digital Enterprise Research Institute (DERI)
University of Innsbruck, Austria
stijn.heyman@deri.org

Abstract. We present a multi-agent formalism based on extended answer set programming. The system consists of independent agents connected via a communication channel, where knowledge and beliefs of each agent are represented by a logic program. When presented with an input set of literals from its predecessor, an agent computes its output as an extended answer set of its program enriched with the input, carefully eliminating contradictions that might occur.

It turns out that while individual agents are rather simple, the interaction strategy makes the system quite expressive: essentially a hierarchy of a fixed number of agents n captures the complexity class Σ_n^P , i.e. the n -th level of the polynomial hierarchy. Furthermore, unbounded hierarchies capture the polynomial hierarchy \mathcal{PH} . This makes the formalism suitable for modelling complex applications of MAS, for example cooperative diagnosis. Furthermore, such systems can be realized by implementing an appropriate control strategy on top of existing solvers such as DLV and SMOELS.

1 Introduction

In *answer set programming* ([30]) a logic program is used to intuitively describe the requirements that must be fulfilled by the solutions of a certain problem. The answer sets of the program, usually defined through (a variant/extension of) the stable model semantics [25], then correspond to the solutions of the problem. This technique has been successfully applied in problem areas such as planning [18, 30], configuration and verification [36], superoptimisation [3], diagnosis [17, 43], game theory [15] and multi-agent systems [2, 7, 16, 6, 10] where [6, 10] use answer set programming to reason about the behaviour of a group of agents, [2, 7, 16] use the formalism to model the reasoning capabilities, knowledge and beliefs of a single agent within a multi-agent

* Supported by the Flemish Fund for Scientific Research (FWO-Vlaanderen).

system. While [2] and [7] use the basic answer set semantics to represent an agent’s domain knowledge, [16] applies an extension of the semantics incorporating preferences among choices in a program.

The traditional answer set semantics, even in the absence of constraints, is not universal, i.e. some programs may not have any answer set at all. While natural, this poses a problem in cases where there are no exact solutions, but one would appreciate to obtain approximate ones, even if they violate some rules. For example, it is not acceptable that an airplane’s auto-pilot agent fails to work just because it has some contradictory readings regarding the outside temperature. To achieve this, the extended answer set semantics ([42]) allows problematic rules to be *defeated*: the rules $a \leftarrow$, $b \leftarrow$ and $\neg a \leftarrow b$ are clearly inconsistent and have no classical answer set, while both $\{a, b\}$ and $\{\neg a, b\}$ will be recognized as extended answer sets. In $\{a, b\}$, $\neg a \leftarrow b$ is defeated by $a \leftarrow$, while in $\{\neg a, b\}$, $a \leftarrow$ is defeated by $\neg a \leftarrow b$.

In this paper we use the extended answer set semantics to model the knowledge and beliefs of a single agent. Each agent reasons over two languages, one public and one private. This allows agents to dynamically decide which information they wish to share with others, with only public information being made available. Agents may then cooperate to select among the various possible solutions (extended answer sets) that are presented to them. In the case that an agent, using the extended answer set semantics, has a number of (approximate) solutions to a certain problem, it can rely upon other agents to sort out which solutions are the better ones. In the absence of any extended answer sets, the agent relies completely on the information received from the others. E.g., when a company has to make up an emergency evacuation plan for a building, one of the employees will make up all strategies that could be implemented for that building. However, as she is probably not aware of all current regulations about such strategies, her solutions are forwarded to the emergency services, who will only select those plans that are conforming to all legal requirements. These legal candidate plans are then presented to the firm’s management to select an optimal one (e.g. the cheapest) for implementation.

To deal with problems like the one described above, we propose a multi-agent framework that is capable of modelling hierarchical decision problems. To this end, we consider a sequence of agents $A_1 \dots A_n$, each having their private knowledge described by a logic program. Intuitively, an agent A_i communicates a solution she finds acceptable to the next agent A_{i+1} in the hierarchy. For such an A_i -acceptable solution, A_{i+1} computes a solution S that adds her knowledge to the given information. Provided that this new knowledge does not conflict with the information she received from her predecessor A_i , she passes this solution to the following agent in line, i.e. A_{i+2} . In case agent A_{i+1} is unable to provide any solutions of her own, she will simply pass on information she obtained from the previous agents higher up in the hierarchy. When her solution S conflicts with the solution offered by her predecessors, she sends S for verification to her predecessor A_i . If A_i is able to find another possible solution T that is consistent with S , the communication from A_i to A_{i+1} starts over again with T as a new input. In the case that none of the solutions of A_{i+1} survive the verification step, A_{i+1} has no other option than accepting the input from A_i and send it to A_{i+2} .

It turns out that, although the agents are relatively simple in complexity terms, such sequences of agents are rather expressive. More specifically, we show that such agent systems can capture the polynomial hierarchy, which make them suitable for encoding complex applications.

Computing the extended answer set semantics is located at the first level of the polynomial hierarchy. Problems located at this first level can be directly solved using the DLV [24] and SMODELS [34] answer set solvers. On the second level, only DLV remains to perform the job directly. However, by using a “guess and check” fixpoint procedure, SMODELS can indirectly be used to solve problems at the second level [4, 22, 45]. Beyond the second level, there are still interesting problems, such as the most expressive forms of diagnostic reasoning, i.e. subset-minimal diagnosis on disjunctive system descriptions [17] or preference-based diagnosis on ordered theories [43]. These are located at the third level of the polynomial hierarchy, together with sequences of weak constraints⁴ on disjunctive programs. For these problems, and problems located even higher in the polynomial hierarchy, no direct computational vehicle is available. The framework presented in this paper provides a means to effectively compute solutions for such problems with each agent using SMODELS or DLV to compute better solutions combined with an appropriate control strategy for the communication.

The remainder of the paper is organized as follows. In Section 2, we review the extended answer set semantics. Section 3 presents the definitions for hierarchical agents and agent systems. Section 4 discusses the complexity of the proposed semantics, while Section 5 compares it with related approaches from the literature. Finally, we conclude with directions for further research in Section 6.

2 Extended Answer Sets

In this section we provide a short overview of extended answer set semantics for simple logic programs [41]. A *term* is a constant or a variable, where the former will be written lower-case and the latter upper-case. An *atom* is of the form $p(t_1, \dots, t_n)$, $0 \leq n < \infty$, where p is an n -ary⁵ predicate name and t_i , $1 \leq i \leq n$, are terms. A *literal* is an atom a or a negated atom $\neg a$.

A *simple logic program* (SLP) is a finite set of *simple rules* of the form $\alpha \leftarrow \beta$ with $\alpha \cup \beta$ a set of literals and $|\alpha| \leq 1$. If $\alpha = \emptyset$, we call the rule a *constraint*. The set α is the *head* of the rule while β is called the *body*.

A *ground* atom, literal, rule, or SLP does not contain variables. Substituting every variable in a SLP P with every possible constant in P yields the ground SLP $gr(P)$. In what follows, we always assume ground SLPs and ground literals; to obtain the definitions for ungrounded SLPs, replace every occurrence of a SLP P by $gr(P)$, e.g., an extended answer set of an ungrounded SLP P is an extended answer set of $gr(P)$.

For a set of literals X , we use $\neg X$ to denote the set $\{\neg p \mid p \in X\}$ where $\neg\neg a \equiv a$. Further, X is said to be *consistent* if $X \cap \neg X = \emptyset$, i.e. X does not contain contradictory literals a and $\neg a$.

⁴ A weak constraint is a constraint that is “desirable” but may be violated if there are no other options to obtain an answer set.

⁵ We thus allow for 0-ary predicates, i.e., *propositions*.

The *Herbrand base* \mathcal{B}_P of a SLP P is the set of all ground atoms that can be formed using the language of P . The set of all literals that can be formed with P , i.e. $\mathcal{B}_P \cup \neg\mathcal{B}_P$, is denoted by \mathcal{L}_P . An *interpretation* I of P is any consistent subset of \mathcal{L}_P .

A rule $r = a \leftarrow \beta \in P$ is *satisfied* by an interpretation I , denoted $I \models r$, if $a \in I$ whenever $\beta \subseteq I$, i.e. if r is *applicable* ($\beta \subseteq I$), then it must be *applied* ($\beta \cup \{a\} \subseteq I$). On the other hand, a constraint $\leftarrow \beta$ is satisfied if $\beta \not\subseteq I$, i.e. the constraint is not applicable. The rule r is said to be *defeated* w.r.t. I iff there exists an applied *competing rule* $\neg a \leftarrow \beta' \in P$. We use $P_I \subseteq P$ to denote the *reduct* of P w.r.t. I , i.e. $P_I = \{r \in P \mid I \models r\}$, the set of rules satisfied by I .

If an interpretation I satisfies all rules in P , i.e. $P_I = P$, I is called a *model* of P . A model I is a *minimal model* or *answer set* of P iff no other model J of P exists such that $J \subset I$. An *extended answer set* of P is any interpretation I such that I is an answer set of P_I and each unsatisfied rule in $P \setminus P_I$ is defeated. The set of all extended answer sets of a program P is denoted by $\mathcal{AS}(P)$.

Example 1. Consider the following SLP P about diabetes.

$$\begin{array}{lll} \text{hypoglycemia} \leftarrow & \text{sugar} \leftarrow \text{hypoglycemia} & \text{coke} \leftarrow \text{sugar} \\ \text{diabetes} \leftarrow & \neg\text{sugar} \leftarrow \text{diabetes} & \text{diet_coke} \leftarrow \neg\text{sugar} \end{array}$$

Clearly, while this program has no traditional answer sets, it does have two extended answer sets $I = \{\text{diabetes}, \text{hypoglycemia}, \text{sugar}, \text{coke}\}$ and $J = \{\text{diabetes}, \text{diet_coke}, \text{hypoglycemia}, \neg\text{sugar}\}$.

Note that the extended answer set semantics is universal for simple programs containing no constraints [41]. This is not the case for general, constraint allowing, simple programs, due to the fact that constraints cannot be defeated.

3 Hierarchical Agents

If humans want to share information or to have discussions in an effective manner they typically use the same language; without it, it would be nearly impossible to establish any communication. So it is only natural to assume that all agents in our framework “speak the same language” which we denote as \mathcal{AL} . Modeling an agent’s knowledge and beliefs, it might not always be a good idea to pass on the entire answer set, e.g., a manager is certainly not going to tell her employee that she cannot have a meeting on Monday because she wants to have an extended weekend in Paris. Instead she will simply say that Monday is out of the question. To allow this we need to perform some filtering on the answer set before it is passed to the next agent. Thus we consider agents that use two languages: a public language \mathcal{AL} used for communication and a private language \mathcal{AL}' for private reasoning purposes. The latter allows the manager in our example to tell her employee she cannot have the meeting on Monday, without giving her the underlying reason that she is in Paris for a trip. On the other hand, if it is a business trip, she could choose to communicate the reason. Information received from other agents will be assumed private by default. If it needs to be passed on, one simply adds a rule $l \leftarrow l'$ for each literal l' that could be received from the other agent. Summarizing,

an agent will receive input in the form of literals from $\mathcal{L}_{\mathcal{AL}}$, reason with a program over $\mathcal{L}_{\mathcal{AL}} \cup \mathcal{L}_{\mathcal{AL}'}$ and only communicate the part over $\mathcal{L}_{\mathcal{AL}}$ to the other agents. We will use $l' \in \mathcal{L}_{\mathcal{AL}'}$ to denote the private version of the literal $l \in \mathcal{L}_{\mathcal{AL}}$ and we have for $l' \in \mathcal{L}_{\mathcal{AL}'}$ that $l'' = l$ with $l \in \mathcal{L}_{\mathcal{AL}}$. We extend the notation as usual to a set $X \subseteq \mathcal{L}_{\mathcal{AL}} \cup \mathcal{L}_{\mathcal{AL}'}$, i.e. $X' = \{l' \mid l \in X\}$.

Definition 1. For an agent language \mathcal{AL} , a **hierarchical agent** A is a SLP such that $\mathcal{B}_A \subseteq \mathcal{AL} \cup \mathcal{AL}'$. For such an agent A and a set of literals $I \subseteq \mathcal{L}_{\mathcal{AL}}$, the **agent input**, we use $A(I)$ to denote the SLP $A \cup \{l' \leftarrow \mid l \in I\}$.

An interpretation $S \subseteq \mathcal{L}_{\mathcal{AL}}$ is an **agent answer set** w.r.t. the agent input I if

- $S = M \cap \mathcal{L}_{\mathcal{AL}}$ with $M \in \mathcal{AS}(A(I))$, or
- $S = I$ when $\mathcal{AS}(A(I)) = \emptyset$.

We use $\mathcal{AS}(A, I)$ to denote the set of all agent answer sets of A w.r.t. input I .

The first condition of the agent answer set definition ensures that the agent only communicates public information. The second condition makes the agent answer set semantics universal. In case our agent cannot produce an answer set, because of constraints being violated, she will assume the input as the solution. This makes sense in the context of hierarchical agents. E.g., as an employee, one should reschedule previously arranged meetings if one's boss cannot make it on the agreed time. The person with the least power to make changes should concede.

Example 2. Take $\mathcal{AL} = \{\text{hypoglycemia}, \text{diabetes}, \text{sugar}, \text{coke}, \text{diet_coke}\}$ and consider the following diabetes agent A .

$$\begin{array}{ll} \text{sugar}' \leftarrow \text{hypoglycemia}' & \neg \text{sugar}' \leftarrow \text{diabetes}' \\ \text{diet_coke} \leftarrow \neg \text{sugar}' & \text{coke} \leftarrow \text{sugar}' \end{array}$$

Intuitively, the above agent is set up to use information from a doctor agent concerning hypoglycemia and diabetes to decide if a patient needs to have diet coke or normal coke. In order to do so, she determines if the patient needs sugar or should not have sugar. The patient only needs to be told that she can have either a diet coke or a normal coke, hence diet coke and coke are the only literals in the public language.

Let $I_1 = \emptyset$, $I_2 = \{\text{diabetes}\}$ and $I_3 = \{\text{hypoglycemia}\}$ be three agent inputs. One can check that A has only one agent answer set w.r.t. I_1 which is $S_1 = \emptyset$. Similar, feeding both I_2 and I_3 as input to A results in a single agent answer set, i.e. $S_2 = \{\text{diet_coke}\}$ and $S_3 = \{\text{coke}\}$ respectively.

As mentioned before, information can be easily made public by adding a rule $l \leftarrow l'$ for each literal one wants to make public. Depending on the agent, a large number of such rules may be needed. To shorten the programs, we introduce the shorthand $\text{pass}(S)$ for $\{l \leftarrow l' \mid l \in S\}$, with $S \subseteq \mathcal{AL}$ the set of literals that need to be made public if derived in the private part.

Using a combination of public and private information, it is possible to easily encode that for example certain input information should be considered more important than the agent's own knowledge or vice versa.

Example 3. Consider the following employee agent A_1 :

$$\begin{array}{l} \mathbf{pass}(\{pay_rise, overworked\}) \\ overworked \leftarrow \hspace{15em} dislike_boss' \leftarrow overworked \\ happy \leftarrow \neg dislike_boss', pay_rise' \quad \neg dislike_boss' \leftarrow pay_rise' \end{array}$$

Obviously the agent will never publicly admit disliking her boss. Given $\{pay_rise\}$ as input, the agent produces two answer sets: $\{overworked, pay_rise\}$ and $\{overworked, pay_rise, happy\}$.

Now that we have defined a single hierarchical agent and the semantics that comes with it, we can start connecting them. As mentioned previously, we are interested in multi-agent systems where some agents have more authority than others, yet require information from others in order to make correct decisions. In the introduction, we discussed the situation of a company that needs to implement an emergency evacuation plan. Although a manager needs to approve the emergency plan, she does not need to verify legal issues or draw up the plans herself. She will stipulate the requirements that need to be fulfilled for her approval. So, in this case we have the employee being on top of the hierarchy generating all possible plans, followed by the legal office rejecting those plans which are not safe. Finally, these plans will be matched against the criteria set out by the manager. Since a plan is needed, she will be unable to reject them all.

We have a different situation when a head of department needs to arrange a meeting with her staff. Obviously she will allow her staff to have a say in the organisation, but at the end of the day her diary will take precedence over that of her staff. Here the head of department will be on top of the hierarchy to generate all possible dates she can have the meeting, which can then be verified by her staff.

The above two examples demonstrate that there can be a difference between the agent hierarchy and the hierarchy of the people/things modeled by the agents. The agent with the greatest power is the one generating all the candidate models. The effect of the other agents is inversely proportional to their distance to this initial (generator) agent.

In this paper, we restrict to linearly connected agents, since such systems are already capable of representing the most common forms of hierarchy.

Formally, a *hierarchical agent system* (*HAS*) is a linear sequence of hierarchical agents $A = (A_1, \dots, A_n)$, where A_1 is the source agent, i.e. the agent that starts all communication. For a HAS A , we refer to the i -th agent as A_i , while we use $A_{<i}$ to denote the HAS consisting of the predecessors of A_i , i.e. $A_{<i} = (A_1, \dots, A_{i-1})$.

We assume for our theoretical model that agents are fully aware of the agents that they can communicate with (as the communication structure is fixed) and that they can communicate by passing sets of literals over the communication channels. When put to practice in an open multi-agent environment, an agent would first engage in establishing a community and the appropriate hierarchy before collaborating on establishing a consensus on the answer sets. Furthermore, one would expect the set of literals encapsulated in a communication protocol.

Each agent in a HAS is a separate entity with its own reasoning skills, knowledge and beliefs. Each agent has the right to remove or add information to the input as she sees fit. To reflect this, we introduce an *interpretation* for a HAS $A = (A_1, \dots, A_n)$ as a sequence of interpretations $I = (I_1, \dots, I_n)$, one for each agent, denoting the public

knowledge of each individual agent. For interpretations, we introduce the same notation I_i and $I_{<i}$ as we did for hierarchical agent systems. An interpretation I is consistent iff $\bigcup_{1 \leq i \leq n} I_i$ is consistent. Given a sequence I and a set S , we will write (I, S) to denote the new sequence obtained from concatenating I and S .

Example 4. Consider the HAS $A = (A_1, A_2)$ with $A_1 = \{\text{meeting} \leftarrow\}$ and $A_2 = \{\text{out_of_office} \leftarrow \text{meeting}'\}$. Then, $(\{\text{meeting}\}, \{\text{out_of_office}\})$ is an interpretation.

The solutions of such a hierarchical agent system, called hierarchical answer sets, are defined inductively. For a HAS A , we will use $\mathcal{AG}(A)$ to denote the set of all hierarchical answer sets of A .

Definition 2. Let \mathcal{AL} be an agent language.

- A **hierarchical answer set** of a HAS $A = (A_1)$ is a consistent interpretation $S = (S_1)$ such that $S_1 \in \mathcal{AS}(A_1, \emptyset)$.
- A **hierarchical answer set** of a HAS $A = (A_1, \dots, A_n)$ is a consistent interpretation $S = (S_1, \dots, S_n)$ such that $S_{<n}$ is a hierarchical answer set of $A_{<n}$, i.e. $S_{<n} \in \mathcal{AG}(A_{<n})$, and
 1. $S_n \in \mathcal{AS}(A_n, S_{n-1})$; or
 2. $S_n = S_{n-1}$ iff $\forall S' \in \mathcal{AG}(A_{<n}) \cdot \forall T \in \mathcal{AS}(A_n, S_{n-1}) \cdot (S', T)$ inconsistent.

The case of a single agent HAS is simple: hierarchical answer sets equal the agent's agent answer sets with empty input. The two conditions of the general case are the encoding of the principle that an agent either has to be able to augment the input in a consistent manner (condition 1) or convince itself that all the alternatives it can propose are inconsistent with solutions that are acceptable by its predecessors. In that case, the input will be accepted (condition 2). If not, the candidate will be rejected.

Example 5. Consider the following simple HAS $A = (A_1, A_2, A_3)$ with:

- the general director A_1 of a company containing the following rules⁶:

$$\text{monday} \oplus \text{tuesday} \oplus \text{friday} \leftarrow \quad \neg \text{wednesday} \leftarrow \quad \neg \text{thursday} \leftarrow$$

- the head of research A_2 containing the rules:

$$\text{monday} \oplus \text{thursday} \leftarrow \quad \neg \text{tuesday} \leftarrow \quad \neg \text{wednesday} \leftarrow \quad \neg \text{friday} \leftarrow$$

- the project manager A_3 containing the rules:

$$\text{friday} \oplus \text{wednesday} \leftarrow \quad \neg \text{monday} \leftarrow \quad \neg \text{tuesday} \leftarrow \quad \neg \text{thursday} \leftarrow$$

who attempt to arrange a meeting. The director agent produces three possible hierarchical answer sets for the HAS (A_1) , i.e.

- $(M_1) = (\{\text{monday}, \neg \text{tuesday}, \neg \text{wednesday}, \neg \text{thursday}, \neg \text{friday}\})$
- $(M_2) = (\{\neg \text{monday}, \text{tuesday}, \neg \text{wednesday}, \neg \text{thursday}, \neg \text{friday}\})$

⁶ In the following we will use rules of the form $a \oplus b \oplus c \leftarrow$ to denote the set of rules $\{a \leftarrow ; b \leftarrow ; c \leftarrow ; \neg a \leftarrow ; \neg b \leftarrow ; \neg c \leftarrow ; \leftarrow a, b ; \leftarrow a, c ; \leftarrow b, c ; \leftarrow \neg a, \neg b, \neg c\}$, i.e. an exclusive choice between a , b and c .

– $(M_3) = (\{\neg\text{monday}, \neg\text{tuesday}, \neg\text{wednesday}, \neg\text{thursday}, \text{friday}\})$

Let us now consider $A_{<3} = (A_1, A_2)$. When we feed A_2 with M_1 , we notice that M_1 is accepted. This means that (M_1, M_1) is a hierarchical answer set for $A_{<3}$. Any other answer set from A_2 with input M_1 leads to contradiction. When we use M_2 as input we have that $M_1 \in \mathcal{AS}(A_2, M_2)$ is clearly inconsistent with M_2 , but which is consistent with an acceptable solution of the predecessors, i.e. (M_1) . This implies that there is no hierarchical answer with M_2 as input for A_2 . The same is true when M_3 is used as input. As a result, we have $\mathcal{AG}(A_{<3}) = \{(M_1, M_1)\}$.

Now that we have the hierarchical answer sets for $A_{<3}$, we can define those of A . When we compute the answer sets of A_3 with M_1 as input, we obtain two answer sets: one assuming friday to be true and the other wednesday to be true. Both are inconsistent with (M_1, M_1) , so our project manager has no other option than to conform to M_1 herself, resulting in $\mathcal{AG}(A) = \{(M_1, M_1, M_1)\}$.

Now consider the rearranged HAS $B = (A_1, A_3, A_2)$, e.g. because A_3 has prior arrangements with customers who do not appreciate changes to their schedule. This change would result in a different hierarchical answer set, namely (M_3, M_3, M_3) .

Although we request that hierarchical answer sets are consistent, this does not mean that internal inconsistencies cannot appear. Further, the system also allows for cheating and/or lying.

Example 6. Consider the following HAS $A = (A_1, A_2)$ with $A_1 = \{a \leftarrow\}$ and $A_2 = \{b \leftarrow; \neg a' \leftarrow b; c \leftarrow \neg a'; a \leftarrow a'\}$. This HAS produces two hierarchical answer sets: $(\{a\}, \{a, b\})$ and $(\{a\}, \{b, c\})$. In the latter case, the agent A_2 knows that there would be a contradiction if she would admit $\neg a$, so she decides to ignore what she knows about $\neg a$ and only states the implication of $\neg a'$, i.e. the conclusion c .

Example 7. Consider the job selection procedure of a company. The first agent A_1 corresponds with the possible profiles of the applicants. Thus, each agent answer set of the agent below corresponds with a possible applicant's profile.

$male \oplus female \leftarrow \quad old \oplus young \leftarrow \quad experienced \oplus inexperienced \leftarrow$

The decision which applicant gets the job goes through a chain of decision makers. First, the agent A_2 of the human resources department implements company policy which stipulates that experienced persons should be preferred over inexperienced ones. Therefore, the agent passes through all of its input, except when it encounters a profile containing *inexperienced*, which it changes to *experienced*, intuitively implementing that an applicant with the same profile but *experienced* instead of *inexperienced* would be preferable. Further, the department is convinced that younger employees are ambitious.

pass($\{male, female, old, young, experienced\}$)
pass($\{\neg male, \neg female, \neg old, \neg young, \neg inexperienced\}$)
 $experienced \leftarrow inexperienced'$
 $\neg inexperienced \leftarrow inexperienced'$
 $ambitious \leftarrow young'$

On the next level of the decision chain, the financial department reviews the remaining candidates. As young and inexperienced persons tend to cost less, it has a strong desire to hire such candidates, which is implemented in the following agent A_3 .

$$\begin{array}{l}
\mathbf{pass}(\{male, female, young, inexperienced\}) \\
\mathbf{pass}(\{\neg male, \neg female, \neg old, \neg experienced\}) \\
inexperienced \leftarrow young', experienced' \quad \neg experienced \leftarrow young', experienced' \\
young \leftarrow young', experienced' \quad \neg old \leftarrow young', experienced' \\
young \leftarrow old', inexperienced' \quad \neg old \leftarrow old', inexperienced' \\
inexperienced \leftarrow old', inexperienced' \quad \neg experienced \leftarrow old', inexperienced' \\
inexperienced \leftarrow old', experienced' \quad \neg inexperienced \leftarrow old', experienced' \\
experienced' \leftarrow old', experienced' \quad \neg experienced' \leftarrow old', experienced' \\
young \leftarrow old', experienced' \quad \neg young \leftarrow old', experienced' \\
old \leftarrow old', experienced' \quad \neg old \leftarrow old', experienced' \\
\leftarrow old, experienced \quad cheaper \leftarrow inexperienced \\
\quad \quad \quad cheaper \leftarrow young
\end{array}$$

Intuitively, this agent handles the four possible cases: when the input profile is from a young and inexperienced person, nothing will be changed, indicating that the input cannot be improved. On the other hand, if only one of the properties is not as desired, e.g. *young* and *experienced*, then the only improvement would be a profile containing both *young* and *inexperienced*. Finally, a profile containing *old* and *experienced* has three possible improvements: the contradictory rules together with the constraint ensure that the agent answer sets proposed by A_3 will contain *young* or *inexperienced*, or both.

Finally, the management has the final call in the selection procedure. As the current team of employees is largely male, the management prefers the new worker to be a woman, as described by the next agent A_4 , which is similar to A_2 .

$$\begin{array}{l}
\mathbf{pass}(\mathcal{AL} \setminus \{male, \neg female, ambitious, cheaper\}) \\
female \leftarrow male' \quad \neg male \leftarrow \neg female' \quad \leftarrow female'
\end{array}$$

One can check that the system (A_1) has eight hierarchical answer sets, among them are

- $$\begin{array}{l}
(M_1) = (\{experienced, \neg inexperienced, male, \neg female, young, \neg old\}) , \\
(M_2) = (\{experienced, \neg inexperienced, male, \neg female, old, \neg young\}) , \\
(M_3) = (\{experienced, \neg inexperienced, female, \neg male, young, \neg old\}) , \\
(M_4) = (\{experienced, \neg inexperienced, female, \neg male, old, \neg young\}) , \\
(M_5) = (\{inexperienced, \neg experienced, female, \neg male, young, \neg old\}) .
\end{array}$$

However, only four of these will survive agent A_2 , i.e. $\mathcal{AG}((A_1, A_2)) = \{(M_1, M_1 \cup \{ambitious\}), (M_2, M_2), (M_3, M_3 \cup \{ambitious\}), (M_4, M_4)\}$, which fits the human resource policy to drop inexperienced people. Feeding M_5 as input to A_2 yields one agent answer set $M_3 \cup \{ambitious\}$, which is consistent with $(M_3) \in \mathcal{AG}((A_1))$, making (M_5, M_5) unacceptable as a solution for the system. Similarly, when agent A_3 is taken into account, only $(M_1, M_1 \cup \{ambitious\}, M_1 \cup \{cheaper\})$ and $(M_3, M_3 \cup$

$\{\textit{ambitious}\}, M_3 \cup \{\textit{cheaper}\})$ are contained in $\mathcal{AG}((A_1, A_2, A_3))$. Considering the last agent A_4 , the HAS (A_1, A_2, A_3, A_4) yields a single hierarchical answer set,

$$(M_3, M_3 \cup \{\textit{ambitious}\}, M_3 \cup \{\textit{cheaper}\}, M_3) ,$$

which fits our intuition that, if possible, a woman should get the job.

4 Complexity

We briefly recall some relevant notions of complexity theory (see [31] for an introduction). The class \mathcal{P} (\mathcal{NP}) represents the problems that are deterministically (nondeterministically) decidable in polynomial time, while \textit{coNP} contains the problems whose complements are in \mathcal{NP} . The polynomial hierarchy, denoted \mathcal{PH} , is made up of three classes of problems, i.e. Δ_k^P , Σ_k^P and Π_k^P , $k \geq 0$, which are defined as follows:

1. $\Delta_0^P = \Sigma_0^P = \Pi_0^P = \mathcal{P}$; and
2. $\Delta_{k+1}^P = \mathcal{P}^{\Sigma_k^P}$, $\Sigma_{k+1}^P = \mathcal{NP}^{\Sigma_k^P}$, $\Pi_{k+1}^P = \textit{co}\Sigma_{k+1}^P$.

The class $\mathcal{P}^{\Sigma_k^P}$ ($\mathcal{NP}^{\Sigma_k^P}$) represents the problems decidable in deterministic (non-deterministic) polynomial time using an oracle for problems in Σ_k^P , where an oracle is a subroutine capable of solving Σ_k^P problems in unit time. Note that $\Delta_1^P = \mathcal{P}$, $\Sigma_1^P = \mathcal{NP}$ and $\Pi_1^P = \textit{coNP}$. The class \mathcal{PH} is defined by $\mathcal{PH} = \bigcup_{k=0}^{\infty} \Sigma_k^P$. Finally, the class \textit{PSPACE} contains the problems that can be solved deterministically by using a polynomial amount of memory and unlimited time.

A decision problem D is called *hard* in a complexity class C of the polynomial hierarchy if any other problem from this class can be reduced to it by a polynomial time reduction⁷. A decision problem D is called *complete* in a complexity class C if both D is in C and D is hard in C .

A quantified boolean formula (QBF) is an expression of the form $Q_1 X_1 Q_2 X_2 \dots Q_k X_k \cdot G$, where $k \geq 1$, G is a Boolean expression over the atoms of the pairwise nonempty disjoint sets of variables X_1, \dots, X_k and the Q_i 's, for $i = 1, \dots, k$ are alternating quantifiers from $\{\exists, \forall\}$. When $Q_1 = \exists$, the QBF is k -existential, when $Q_1 = \forall$ we say it is k -universal. We use $\textit{QBF}_{k,\exists}$ ($\textit{QBF}_{k,\forall}$) to denote the set of all valid k -existential (k -universal) QBFs.

Deciding, for a given k -existential (k -universal) QBF ϕ , whether $\phi \in \textit{QBF}_{k,\exists}$ ($\phi \in \textit{QBF}_{k,\forall}$) is a Σ_k^P -complete (Π_k^P -complete) problem. When we drop the bound k on the number of quantifiers, i.e. considering $\textit{QBF}_{\exists} = \bigcup_{i \in \mathbb{N}} \textit{QBF}_{i,\exists}$, we have a hard problem for \textit{PSPACE} .

The following results shed some light on the data complexity of the hierarchical answer set semantics for hierarchical agent systems, i.e. we measure the complexity with respect to the size of the facts, while the rules in the different hierarchical agents are fixed. Note that in the case of data complexity, for an ungrounded hierarchical agent P , the size of $\textit{gr}(P)$ is polynomial in terms of the size of the facts.

First, we consider the case where the number of agents in the hierarchy is fixed by some number n .

⁷ The only exception is \mathcal{P} -hardness, where logarithmic space reductions have to be used.

Theorem 1. *The problem of deciding, given a HAS $(A_i)_{i=1,\dots,n}$, with n fixed, and a literal $l \in \mathcal{L}_{\mathcal{AL}}$, whether there exists a hierarchical answer set I containing l is Σ_n^P -complete. On the other hand, deciding whether every hierarchical answer set contains l is Π_n^P -complete.*

Proof Sketch. *Membership Σ_n^P :* It is shown, by induction, in [38] that checking if an interpretation I for A is not a hierarchical answer set of A , i.e. $I \notin \mathcal{AG}(A)$, is in Σ_{n-1}^P . The main result follows by

- guessing an interpretation I containing l ; and
- checking that it is not the case that $I \notin \mathcal{AG}(A)$.

As the latter is in Σ_{n-1}^P , the problem itself can be done by an $\mathcal{NP}^{\Sigma_{n-1}^P}$ algorithm, i.e. the problem is in Σ_n^P .

Hardness in Σ_n^P : To prove this, we provide a reduction of deciding validity of QBFs by means of a HAS. Let $\phi = \exists X_1 \forall X_2 \dots Q X_n \cdot G \in \text{QBF}_{n,\exists}$, where $Q = \forall$ if n is even and $Q = \exists$ otherwise. We assume, without loss of generality [37], that, if n is even, G is in disjunctive normal form, i.e. $G = \bigvee_{c \in C} C$ where C is a set of sets of literals over $X_1 \cup \dots \cup X_n$ and each $c \in C$ has to be read as a conjunction; or that, if n is odd, G is in conjunctive normal form, i.e. $G = \bigwedge_{c \in C} C$ where C is a set of sets of literals over $X_1 \cup \dots \cup X_n$ and each $c \in C$ has to be read as a disjunction.

In what follows, we will use the following shorthand notations for certain sets of rules. To check satisfiability of G in case n is even, we use $P_{\text{even}}^{\text{sat}}$ to denote the set of rules $\{\text{sat}' \leftarrow c' \mid c \in C\}$. On the other hand, satisfiability of G in case n is odd is checked by using $P_{\text{odd}}^{\text{sat}}$, which contains the rules⁸

- $\{\text{notsat}' \leftarrow c'_- \mid c \in C\}$,
- $\{\leftarrow c'_-, \neg \text{notsat}' \mid c \in C\}$,
- $\neg \text{notsat}' \leftarrow$,
- $\text{sat}' \leftarrow \neg \text{notsat}'$,

where c_- denotes the version of c where the literals are negated and read conjunctively, i.e. $c = a \vee b \vee \neg d$ results in $c_- = \neg a \wedge \neg b \wedge d$. One can check that in both $P_{\text{even}}^{\text{sat}}$ and $P_{\text{odd}}^{\text{sat}}$ only sat' is derived iff G is valid w.r.t. the chosen truth values for $X_1 \cup \dots \cup X_n$.

Further, we use P_{\forall}^i to denote the program containing the rules

- **pass**($\{x, \neg x \mid x \in X_j \wedge 1 \leq j < i\}$),
- $\{x' \leftarrow ; \neg x' \leftarrow \mid x \in X_j \wedge i \leq j \leq n\}$,
- $P_{\text{even}}^{\text{sat}}$ or $P_{\text{odd}}^{\text{sat}}$,
- $\{\leftarrow \text{sat}' ; \neg \text{sat} \leftarrow ; \leftarrow \neg \text{sat}\}$.

Similarly, we use P_{\exists}^i to denote the program

- **pass**($\{x, \neg x \mid x \in X_j \wedge 1 \leq j < i\}$),
- $\{x' \leftarrow ; \neg x' \leftarrow \mid x \in X_j \wedge i \leq j \leq n\}$,
- $P_{\text{even}}^{\text{sat}}$ or $P_{\text{odd}}^{\text{sat}}$,
- $\{\neg \text{sat}' \leftarrow ; \leftarrow \neg \text{sat}' ; \leftarrow \text{sat}\}$.

⁸ The rules use an encoding in the extended answer set semantics for negation as failure [39].

The HAS $A_\phi = (A_1, \dots, A_n)$ corresponding to ϕ is defined by the following hierarchical agents:

- A_1 contains the rules $\{x' \leftarrow ; \neg x' \leftarrow \mid x \in X_j \wedge 1 \leq j \leq n\}$ and either P_{even}^{sat} or P_{odd}^{sat} ;
- if n is even, then $A_i = P_{\forall}^{n+2-i}$ when i even and $A_i = P_{\exists}^{n+2-i}$ when $i > 1$ odd;
- if n is odd, then $A_i = P_{\exists}^{n+2-i}$ when i even and $A_i = P_{\forall}^{n+2-i}$ when $i > 1$ odd.

Obviously, the above construction can be done in polynomial time. Intuitively, the hierarchical agent A_1 has agent answer sets for every possible combination of the X_i 's and if such a combination makes G valid, then the corresponding agent answer set also contains the atom sat . The intuition behind the hierarchical agent P_{\forall}^i is that it tries to disprove, for the received input, the validity of the corresponding \forall , i.e. for a given input combination over the X_j 's making G satisfied, the hierarchical agent P_{\forall}^i will try to find a combination, keeping the X_j 's with $j < i$ fixed, making G false. On the other hand, the hierarchical agent P_{\exists}^i will try to prove the validity of the corresponding \exists , i.e. for a given combination making G false it will try to compute a combination, keeping the X_j 's with $j < i$ fixed, making G satisfied.

Instead of giving the formal proof for the above construction, we illustrate, for clarity, the construction and the working of the HAS A_ϕ on an example and refer the reader to [38] for the actual proof.

Consider

$$\phi = \exists x \cdot \forall y \cdot \exists z \cdot \forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z) \vee w .$$

The hierarchical agent A_1 contains the following rules.

$$\begin{array}{cccc} x' \leftarrow & \neg x' \leftarrow & y' \leftarrow & \neg y' \leftarrow \\ z' \leftarrow & \neg z' \leftarrow & w' \leftarrow & \neg w' \leftarrow \\ sat' \leftarrow x', \neg y', z' & sat' \leftarrow y', \neg z' & sat' \leftarrow w' & \end{array}$$

We have 16 possible agent answer sets for $A_1(\emptyset)$, i.e. $I_1 = \{x, y, z, w, sat\}$, $I_2 = \{x, y, z, \neg w\}$, $I_3 = \{x, y, \neg z, w, sat\}$, $I_4 = \{x, y, \neg z, \neg w, sat\}$, $I_5 = \{x, \neg y, z, w, sat\}$, $I_6 = \{x, \neg y, z, \neg w, sat\}$, $I_7 = \{x, \neg y, \neg z, w, sat\}$, $I_8 = \{x, \neg y, \neg z, \neg w\}$, $I_9 = \{\neg x, y, z, w, sat\}$, $I_{10} = \{\neg x, y, z, \neg w\}$, $I_{11} = \{\neg x, y, \neg z, w, sat\}$, $I_{12} = \{\neg x, y, \neg z, \neg w, sat\}$, $I_{13} = \{\neg x, \neg y, z, w, sat\}$, $I_{14} = \{\neg x, \neg y, z, \neg w\}$, $I_{15} = \{\neg x, \neg y, \neg z, w, sat\}$ and $I_{16} = \{\neg x, \neg y, \neg z, \neg w\}$.

Clearly, for $1 \leq i \leq 16$, (I_i) is a hierarchical answer set of (A_1) .

The second hierarchical agent A_2 is defined by P_{\forall}^4 and thus contains the following rules.

$$\begin{array}{ccc} \mathbf{pass}(\{x, \neg x, y, \neg y, z, \neg z\}) \leftarrow & w' \leftarrow & \neg w' \leftarrow \\ sat' \leftarrow x', \neg y', z' & sat' \leftarrow y', \neg z' & sat' \leftarrow w' \\ \leftarrow sat' & \neg sat \leftarrow & \leftarrow \neg sat \end{array}$$

Now, feeding I_1 to A_2 yields I_2 as the single agent answer set, which is clearly inconsistent with I_1 , yielding that (I_1, I_2) cannot be a hierarchical answer set of the HAS (A_1, A_2) . Further, for I_2 we have that $I_2 \in \mathcal{AG}(A_1)$, yielding that (I_2, I_2) is clearly consistent, which implies that (I_1, I_1) is neither a hierarchical answer set of (A_1, A_2) .

On the other hand, $A_2(I_3)$ yields I_3 as the single agent answer set which is clearly consistent with itself, yielding that (I_3, I_3) is a hierarchical answer set of (A_1, A_2) , i.e. A_2 passes the input I_3 as a result as it cannot disprove $\forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z) \vee w$ for the chosen truth value of x, y, z in I_3 .

In case of the input I_2 , the agent program $A_2(I_2)$ has no extended answer sets and I_2 is returned as the single agent answer set. As I_2 is consistent with itself, (I_2, I_2) is a hierarchical answer set of (A_1, A_2) .

One can check in similar ways that $\mathcal{AG}((A_1, A_2))$ contains 11 interpretations, i.e. $\mathcal{AG}((A_1, A_2)) = \{(I_2, I_2), (I_3, I_3), (I_4, I_4), (I_5, I_5), (I_6, I_6), (I_8, I_8), (I_{10}, I_{10}), (I_{11}, I_{11}), (I_{12}, I_{12}), (I_{14}, I_{14}), (I_{16}, I_{16})\}$. It is not difficult to see that for each of these hierarchical answer sets it holds that $\forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z) \vee w$ when x, y and z are taken as in the interpretation iff the literal *sat* is contained in the hierarchical answer set.

The third hierarchical agent A_3 is given by P_{\exists}^3 and thus contains the following rules.

$$\begin{array}{llll} \text{pass}(\{x, \neg x, y, \neg y\}) \leftarrow & z' \leftarrow & \neg z' \leftarrow & \text{sat}' \leftarrow x', \neg y', z' \\ & w' \leftarrow & \neg w' \leftarrow & \text{sat}' \leftarrow y', \neg z' \\ \neg \text{sat}' \leftarrow & \leftarrow \neg \text{sat}' & \leftarrow \text{sat} & \text{sat}' \leftarrow w' \end{array}$$

When providing A_3 with the input I_2 , we have $\mathcal{AS}(A_3, I_2) = \{I_1, I_3, I_4\}$, none of them consistent with I_2 . However, for both I_3 and I_4 there is a $T \in \mathcal{AG}((A_1, A_2))$ such that I_3 and I_4 is consistent with T , i.e. $T = (I_3, I_3)$ and $T = (I_4, I_4)$ respectively. As a result $(I_2, I_2, I_2) \notin \mathcal{AG}((A_1, A_2, A_3))$.

On the other hand, feeding I_3 as input to A_3 will yield the single agent answer set I_3 , implying that $(I_3, I_3, I_3) \in \mathcal{AG}((A_1, A_2, A_3))$.

In case of the input I_{14} , we have $\mathcal{AS}(A_3, I_{14}) = \{I_{13}, I_{15}\}$ and none of them consistent with I_{14} . Further, neither for I_{13} nor I_{15} there is a $T \in \mathcal{AG}((A_1, A_2))$ such that I_{13} or I_{15} is consistent with T , yielding that (I_{14}, I_{14}, I_{14}) is a hierarchical answer set of (A_1, A_2, A_3) in this case, i.e. A_3 passes the input I_{14} as a result because it proved that $\exists z \cdot \forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z)$ does not hold for the chosen truth value of x and y in I_{14} . In a similar way one can check that also $(I_3, I_3, I_3) \in \mathcal{AG}((A_1, A_2, A_3))$.

One can check in similar ways that $\mathcal{AG}((A_1, A_2, A_3))$ contains 8 interpretations, i.e. $\mathcal{AG}((A_1, A_2, A_3)) = \{(I_3, I_3), (I_4, I_4), (I_5, I_5), (I_6, I_6), (I_{11}, I_{11}), (I_{12}, I_{12}), (I_{14}, I_{14}), (I_{16}, I_{16})\}$. Again, it is not difficult to see that for each of these hierarchical answer sets it holds that $\exists z \cdot \forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z) \vee w$ when x and y are taken as in the interpretation iff the literal *sat* is contained in the hierarchical answer set.

The final hierarchical agent A_4 is given by P_{\forall}^2 , and by similar reasoning as we did for A_2 , one can check that $\mathcal{AG}((A_1, A_2, A_3, A_4))$ contains 6 interpretations, i.e. $\mathcal{AG}((A_1, A_2, A_3, A_4)) = \{(I_3, I_3), (I_4, I_4), (I_5, I_5), (I_6, I_6), (I_{14}, I_{14}), (I_{16}, I_{16})\}$. Again, for each hierarchical answer set in $\mathcal{AG}((A_1, A_2, A_3, A_4))$ it holds that $\forall y \cdot \exists z \cdot \forall w \cdot (x \wedge \neg y \wedge z) \vee (y \wedge \neg z)$ for x taken as in the interpretation iff the literal *sat* is contained in the hierarchical answer set. From this it follows that ϕ is valid iff there exists a hierarchical answer set $I \in \mathcal{AG}((A_1, A_2, A_3, A_4))$ such that the literal *sat* is contained in the interpretation. In our example, I_3 is such a hierarchical answer set and one can check that ϕ holds when assuming x is true.

Π_n^P -completeness: To show this result, we consider in [38] the complementary decision problem and show that it is Σ_n^P -complete, from which the result follows. \square

While the previous result handles the cases where the number of agents in the hierarchy is fixed, we can generalize the results to arbitrary hierarchies.

Theorem 2. *Given a HAS $(A_i)_{i=1,\dots,n}$ and a literal $l \in \mathcal{L}_{\mathcal{AL}}$, the problem of deciding whether there exists a hierarchical answer set I containing l is PSPACE-complete.*

Proof Sketch. Membership PSPACE: Intuitively, each agent in the hierarchy needs the space to represent a single HAS-interpretation (used for computing the agent answer sets), while the HAS itself needs the space to represent a hierarchical answer set. Now, the algorithm will place a possible solution in the latter allocated space, and will use the former allocated space to check if it is indeed a hierarchical answer set. Thus, an algorithm for a hierarchy of n programs, needs maximum $n + 1$ times the space to represent a HAS-interpretation, which is clearly polynomial in space, from which membership to PSPACE follows.

Hardness PSPACE: Clearly, the hardness proof of Theorem 1 can be generalized to validity checking of arbitrary quantified boolean formulas, from which hardness readily follows. \square

While the previous results describe the complexity of reasoning with the presented framework, they do not give a clear picture on the expressiveness [27] of the system, i.e. whether each problem that belongs to a certain complexity class can be expressed in the framework. This is because a formalism F being complete for a particular class only implies that each instance of a problem in that class can be reduced in polynomial time to an instance of F such that the yes/no answer is preserved.

However, completeness does not imply that the polynomial time reduction itself from an instance of the problem to an instance in F is expressible in F^9 .

In this context, one says that a formalism *captures* a certain complexity class iff the formalism is in the class and every problem in that class can be expressed in the formalism. The latter part is normally proved by taking an arbitrary expression in a normal (or general) form¹⁰ for the particular complexity class and by showing that it can be expressed in the formalism.

By using the results from [21, 20], the following normal forms for the complexity classes Σ_k^P , with $k \geq 2$ and k even, and Π_k^P , with $k \geq 2$ and k odd, can be obtained. First, we have to consider a signature $\sigma = (O, F, P)$, with O finite and $F = \emptyset$, i.e. we do not allow function symbols. A finite database over σ is any finite subset of the Herbrand Base over σ . Secondly, we have three predicates that do not occur in P , i.e. *succ*, *first* and *last*. Enumeration literals are literals over the signature $(O, \emptyset, \{succ, first, last\})$ that satisfy the conditions:

- *succ* describes an enumeration of the elements in O ; and
- *first* and *last* contain the first and last element in the enumeration respectively.

⁹ A good example of this fact is the query class *fixpoint*, which is PTIME-complete but cannot express the simple query *even(R)* to check if $|R|$ is even. See e.g. [12, 1] for a more detailed explanation on the difference between completeness and expressiveness (or capturing).

¹⁰ A normal (or general) form of a complexity class is a form to which every problem in the class can be reduced [27]. E.g. the polynomial hierarchy is the set of languages expressible by second-order logic, i.e. each problem of the polynomial hierarchy can be reduced to a second-order logical formula. Note that not every complexity class necessarily has a general form.

Intuitively, succ is a binary predicate such that $\text{succ}(x, y)$ means that y is the successor of x . Further, first and last are unary predicates.

A collection S of finite databases over the signature $\sigma = (O, \emptyset, P)$ is in Σ_k^P , with $k \geq 2$ and k even, iff there is a second order formula of the form

$$\phi = Q_1 U_{1, \dots, m_1}^1 Q_2 U_{1, \dots, m_2}^2 \dots Q_k U_{1, \dots, m_k}^k \exists \bar{x} \cdot \theta_1(\bar{x}) \vee \dots \vee \theta_l(\bar{x}) ,$$

where $Q_i = \exists$ if i is odd, $Q_i = \forall$ if i is even, U_{1, \dots, m_i}^i ($1 \leq i \leq k$) are finite sets of predicate symbols and $\theta_i(\bar{x})$ ($1 \leq i \leq l$) are conjunctions of enumeration literals or literals involving predicates in $P \cup \{U_{1, \dots, m_1}^1, U_{1, \dots, m_2}^2, \dots, U_{1, \dots, m_k}^k\}$ such that for any finite database w over σ , $w \in S$ iff w satisfies ϕ .

Similarly, a collection S of finite databases over the signature $\sigma = (O, \emptyset, P)$ is in Π_k^P , with $k \geq 2$ and k odd, iff there is a second order formula of the form

$$\phi = Q_1 U_{1, \dots, m_1}^1 Q_2 U_{1, \dots, m_2}^2 \dots Q_k U_{1, \dots, m_k}^k \exists \bar{x} \cdot \theta_1(\bar{x}) \vee \dots \vee \theta_l(\bar{x}) ,$$

where $Q_i = \forall$ if i is odd, $Q_i = \exists$ if i is even, U_{1, \dots, m_i}^i ($1 \leq i \leq k$) are finite sets of predicate symbols and $\theta_i(\bar{x})$ ($1 \leq i \leq l$) are conjunctions of enumeration literals or literals involving predicates in $P \cup \{U_{1, \dots, m_1}^1, U_{1, \dots, m_2}^2, \dots, U_{1, \dots, m_k}^k\}$ such that for any finite database w over σ , $w \in S$ iff w satisfies ϕ .

Again, we first consider the case in which the number of agents in the hierarchy is fixed by a number $n \in \mathbb{N}$.

Theorem 3. *The hierarchical answer set semantics for hierarchical agent systems with a fixed number n of agents captures Σ_n^P .*

Proof Sketch. Membership Σ_n^P : The result follows directly from the membership part of the proof of Theorem 1.

Capture Σ_n^P : This proof is a generalization of the technique used in the hardness proof of Theorem 1. Further, the construction of the agents, especially the first one, is based on the proofs in [21], where it is shown that disjunctive logic programming under the brave semantics captures Σ_2^P .

However, we first have to consider the case where $n = 1$ separately, as the general form discussed above only holds for $n \geq 2$. In [39] we have shown that the extended answer set semantics coincides with the classical answer set semantics. As the latter is already proven in the literature (e.g. in [1]) to capture $\Sigma_1^P = \mathcal{NP}$, the same holds for the former, yielding that also the hierarchical answer set semantics for hierarchies of a single agent captures \mathcal{NP} .

To prove that any problem of Σ_n^P , with $n \geq 2$, can be expressed in a HAS of n agents under the hierarchical answer set semantics, we consider two cases, i.e. one where n is even and from which the result follows directly; and one where n is odd and the result follows from the fact that we solve the complementary problem (and $\Sigma_k^P = \text{co}\Pi_k^P$).

In case n is even, we show a construction of a HAS $\langle A_i \rangle_{i=1, \dots, n}$ such that a finite database w satisfies the formula

$$\phi = \exists U_{1, \dots, m_1}^1 \forall U_{1, \dots, m_2}^2 \dots \forall U_{1, \dots, m_n}^n \exists \bar{x} \cdot \theta_1(\bar{x}) \vee \dots \vee \theta_l(\bar{x}) ,$$

with everything defined as in the general form for Σ_n^P described before, iff $\langle A_i \rangle_{i=1, \dots, n}$ has a hierarchical answer set containing *sat*.

The first agent A_1 in the hierarchy contains, beside the facts that introduce the database w (as w'), the following rules:

- For the enumeration of the predicates $U_{1, \dots, m_1}^1, U_{1, \dots, m_2}^2, \dots, U_{1, \dots, m_n}^n$, we have, for $1 \leq i \leq n$ and $1 \leq k \leq m_i$, the rules:

$$U_k^{i'}(\overline{w_k^i}) \leftarrow \qquad \neg U_k^{i'}(\overline{w_k^i}) \leftarrow$$

- To introduce the linear ordering, we need a set of rules similar to the ones used in Section 2.1.13. of [1] (see the technical report [38] for a detailed description). This set of rules has the property that when a linear ordering is established, the literal *linear'* is derived.
- To check satisfiability, we use, for $1 \leq i \leq l$, the rules $sat' \leftarrow \theta_i'(\overline{x}), linear'$.

The other agents of the hierarchy are defined, similar to the hardness proof of Theorem 1 by using two skeletons P_{\forall}^i and P_{\exists}^i . First, both skeletons have the following set of rules P^i in common

- **pass**($\{U_k^j(\overline{w_k^j}), \neg U_k^j(\overline{w_k^j}) \mid (1 \leq j < i) \wedge (1 \leq k \leq m_i)\}$) ,
- **pass**($\{facts\ of\ the\ linear\ ordering\}$) ,
- **pass**(w) ,
- $\{U_k^{j'}(\overline{w_k^j}) \leftarrow ; \neg U_k^{j'}(\overline{w_k^j}) \leftarrow \mid (i \leq j \leq n) \wedge (1 \leq k \leq m_i)\}$, and
- $\{sat' \leftarrow \theta_i'(\overline{x}), linear' \mid 1 \leq i \leq l\}$.

Now, we define the programs P_{\forall}^i and P_{\exists}^i as $P_{\forall}^i = P^i \cup \{ \leftarrow sat' ; \neg sat \leftarrow ; \leftarrow \neg sat \}$ and $P_{\exists}^i = P^i \cup \{ \leftarrow \neg sat' ; \leftarrow \neg sat' ; \leftarrow sat \}$ respectively.

Besides A_1 , the remaining agents in the hierarchy are defined by: $A_i = P_{\forall}^{n+2-i}$ when i even and $A_i = P_{\exists}^{n+2-i}$ when $i > 1$ odd.

It is not difficult to see (similar to the hardness proof of Theorem 1 that the above constructed HAS will only generate, for a given input database w , hierarchical answer sets that contain *sat* iff ϕ is satisfied.

Finally, for the case where n is odd, we can reuse the agents defined above for the even case, i.e. the HAS $\langle A_i \rangle_{i=1, \dots, n}$ is defined, besides A_1 , as $A_i = P_{\exists}^{n+2-i}$ when i even and $A_i = P_{\forall}^{n+2-i}$ when $i > 1$ odd.

This time, it is not so hard to see that ϕ is satisfied iff all hierarchical answer sets of the above constructed HAS, for a given input database w , contain *sat*. As this proves capturing of Π_n^P for the complementary problem we want to show expressiveness for, the result follows. \square

When we drop the fixed number of agents in the hierarchy, the above result can be easily generalized to arbitrary HAS.

Corollary 1. *The hierarchical answer set semantics for hierarchical agent systems captures¹¹ \mathcal{PH} , i.e. the polynomial hierarchy.*

The above result yields that the presented framework is able to encode each problem in the polynomial hierarchy, making the framework useful for complex knowledge reasoning tasks by agents.

¹¹ Note that while the semantics captures \mathcal{PH} , it can never be complete for it as the hierarchy would then collapse.

5 Relationships to Other Approaches

In [5], answer set optimization (ASO) programs are presented. Such ASO programs consist of a generator program and a sequence of optimizing programs. To perform the optimization, the latter programs use rules of the form $c_1 < \dots < c_n \leftarrow \beta$ which intuitively read: when β is true, making c_1 true is the most preferred option and only when c_1 cannot be made true, the next best option is to make c_2 true, ... Solutions of the generator program that are optimal w.r.t. the first optimizing program and, among those, are optimal w.r.t. the second optimizing program, and so on, are called preferred solutions for the ASO program.

The framework of ASO programming can be simply adapted to the setting of agents, i.e. just consider the generator program as agent A_1 and the optimizing programs as agents A_2, \dots, A_n . The resulting semantics is very similar to our approach. However, ASO programs are far more limited w.r.t. their expressiveness. It turns out that the expressiveness of an ASO program does not depend on the length of the sequence of optimizing programs, but it is always Σ_2^P -complete. This yields that ASO programs can easily be captured by the presented agent systems in this paper using two single agents. How these two agents simulating ASO programs can be constructed is subject to further research.

Weak constraints were introduced in [8] as a relaxation of the concept of a constraint. Intuitively, a weak constraint is allowed to be violated, but only as a last resort, meaning that one tries to minimize the number of violated constraints. Additionally, weak constraints are allowed to be hierarchically layered by means of a sequence of sets of weak constraints. Intuitively, one first chooses the answer sets that minimize the number of violated constraints in the first set of weak constraints in the sequence, and then, among those, one chooses the answer sets that minimize the number of violated constraints in the second set, etc.

Again, this approach can be “agentized” in a straightforward manner and will look similar to our approach. This time the complexity of such a system, independent of the number of sets of weak constraints, is at most Δ_3^P -complete. Thus, using the presented agent system from Section 3 with three single agents will suffice to capture the most expressive form of that formalism.

In [26, 40], hierarchies of preferences on a single program are presented. The preferences are expressible on both the literals and the rules in that program. It is shown that for a sequence of n agents the complexity of the system is Σ_{n+1}^P -complete. The semantics proposed in Section 3 is a generalization of that approach: instead of using one global program with agents only using preferences on that program, we equip each agent with her own, in general different, program and let her implement whatever optimizing strategy she wants. To simulate a hierarchy of n preference relations, we need $n+1$ optimizing agents: the first one will correspond with the global program, while the rest will correspond to the n preference relations. The system described in Example 7 can be seen as a translation of such a preference hierarchy. Intuitively, agent A_2 describes the preference relation¹² $experienced < inexperienced$, while A_3 implements the relation $young < old ; inexperienced < experienced$. Finally, A_4 corresponds to

¹² The expression $a < b$ means a is preferred upon b .

the single preference $female < male$. Further, the current approach also allows each agent to express its program by using any syntactical and semantical extension of answer set programming (not only preferences, but e.g. cardinality constraints), as long as the extension can be transformed into the (extended) answer set semantics¹³. Finally, [26, 40] only contained complexity results for the given semantics, while the current work also shows important expressiveness results.

In [33], a theory for coordinating agents is presented. When two agents A and B , which are represented by extended disjunctive programs, coordinate their answer sets, they can either opt for generating the union of their answer sets or the intersection. Their coordination act results in a new program that has the desired answer sets as outcome. Our approach is very different, our aim is to compromise on answer sets without changing the internal knowledge of our agents, i.e. the programs. Our agents share answer sets and not programs. Furthermore, in HAS a hierarchy is established, giving more power to some agents while in [33] agents are considered equal.

[14, 16] also present a multi-agent framework, LAIMA, where the agents are represented as logic programs. LAIMA allows agents to be connected in any sort of way, including loops. The HAS system presented in this paper places the most influential agents at the start of the sequence of agents, providing a top-down approach, while in LAIMA more power is given to agents where communication ends, as they can completely change the answer set even if this causes a contradiction. A more important difference is the absence of failure feedback in LAIMA. It is exactly this feedback that provides the increase in complexity for HAS.

In the Minerva architecture [29], the authors build their agents out of subagents that work on a common knowledge base written as a MDLP (Multi-Dimensional Logic Programs) which is an extension of Dynamic Logic Programming. Our agents do not work with a common knowledge base; each agent decides what she wants to keep private or make available. Minerva does not restrict itself to modeling the beliefs of agents, but allows for full BDI-agents that can plan towards a certain goal. It would be interesting to see if this can also be incorporated in HAS. The complexity of the language used for representing the knowledge is similar. MDLP can be translated to extended logic programs. The procedure given [13] for DLP can easily be extended to MDLP. It is the failure feedback between agents that gives HAS its expressive power, not the expressiveness of the representation language. It would be very interesting to see MDLP used as the representation language of the agent HAS, giving users more flexibility in expressing preferences.

6 Conclusions and Directions for Further Research

We presented a framework suitable for solving hierarchical decision problems using simple logic programming agents that cooperate via a sequential communication channels. The resulting semantics turns out to be very expressive, as it essentially captures the polynomial hierarchy, thus enabling further complex applications. The framework

¹³ Due to the results in [39], normal answer set programming can be reduced to extended answer set programming.

could be used to develop implementations for diagnostic systems at the third level of the polynomial hierarchy [17, 19, 43].

Future work comprises the development of a dedicated implementation of the approach, using existing answer set solvers, e.g. DLV [24] or SMOBELS [34], possibly in a distributed environment. Such an implementation will use a control structure that communicates candidate solutions between consecutive agents. For the implementation of this control loop and the communication between the agents, we foresee the use of JADE [28] and Protégé [32] in much the same way as it has been done for the LAIMA system [14].

In the context of an implementation, it is also interesting to investigate which conditions an agent has to fulfil in order for it not to lift the complexity up one level in the polynomial hierarchy, yielding possible optimizations of the computation and communication processes.

Once the system is implemented we will have the opportunity to work on larger applications. One of our goals, is to try to incorporate the ALIAS [9] system, an agent architecture for legal reasoning based on abductive logic, into ours. The Carell multi-agent system [44] for allocation organs and tissue would be an interesting test case.

In terms of integration it would be nice to see how HAS could possibly work together with agents written for the Dali [11] or Socs [23, 35] platforms, two agent platforms using logic programming languages to model the agents.

At present, we only work with a linear sequence of communication channels. We plan to look into a broader class of communication structures, like for example trees or more generally, a (strict) partial ordering of agents.

Finally, we would like to experiment with the language(s) used for our agents. The definitions of hierarchical agent system and the corresponding hierarchical answer set do not rely on the formalisms used to represent the agents but on the generation of answer sets. Even the generation of these can be debated, as it suffices that agents communicate sets of literals.

References

- [1] C. Baral. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge Press, 2003.
- [2] C. Baral and M. Gelfond. Reasoning agents in dynamic domains. In *Logic-based artificial intelligence*, pages 257–279. Kluwer Academic Publishers, 2000.
- [3] M. Brain, T. Crick, M. De Vos, and J. Fitch. Toast: Applying answer set programming to superoptimisation. In *Proceedings of the 22nd International Conference on Logic Programming (ICLP2006)*, volume 4079 of *LNCS*, pages 270–284. Springer, 2006.
- [4] G. Brewka, I. Niemela, and T. Syrjanen. Implementing ordered disjunction using answer set solvers for normal programs. In *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA 2002)*, volume 2424 of *LNAI*, pages 444–455. Springer, 2002.
- [5] G. Brewka, I. Niemelä, and M. Truszczynski. Answer set optimization. In G. Gottlob and T. Walsh, editors, *IJCAI*, pages 867–872. Morgan Kaufmann, 2003.
- [6] F. Buccafurri and G. Caminiti. A social semantics for multi-agent systems. In C. Baral, G. Greco, N. Leone, and G. Terracina, editors, *LPNMR*, volume 3662 of *Lecture Notes in Computer Science*, pages 317–329. Springer, 2005.

- [7] F. Buccafurri and G. Gottlob. Multiagent compromises, joint fixpoints, and stable models. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part I*, volume 2407 of *LNCS*, pages 561–585. Springer, 2002.
- [8] F. Buccafurri, N. Leone, and P. Rullo. Strong and weak constraints in disjunctive datalog. In *Proceedings of the 4th International Conference on Logic Programming (LPNMR '97)*, pages 2–17, 1997.
- [9] A. Ciampolini and P. Torroni. Using abductive logic agents for modeling the judicial evaluation of criminal evidence. *Applied Artificial Intelligence*, 18:251–275, 2004.
- [10] O. Cliffe, M. De Vos, and J. Padget. Specifying and analysing agent-based social institutions using answer set programming. In *Selected revised papers from the workshops on Agent, Norms and Institutions for Regulated Multi-Agent Systems (ANIREM) and Organizations and Organization Oriented Programming (OOOP) at AAMAS'05*, volume 3913 of *LNCS*, pages 99–113. Springer, 2006.
- [11] S. Costantini and A. Tocchio. Context-based Commonsense Reasoning in the DALI Logic Programming Language. In *Proceedings of the 4th International and Interdisciplinary Conference, Context 2003*, volume 2680 of *LNAI*. Springer, 2003.
- [12] E. Dantsin, T. Eiter, G. Gottlob, and A. Voronkov. Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3):374–425, 2001.
- [13] M. De Vos. Implementing Ordered Choice Logic Programming using Answer Set Solvers. In *Third International Symposium on Foundations of Information and Knowledge Systems (FoIKS'04)*, volume 2942, pages 59–77, Vienna, Austria, Feb. 2004. Springer Verlag.
- [14] M. De Vos, T. Crick, J. Padget, M. Brain, O. Cliffe, and J. Needham. A Multi-agent Platform using Ordered Choice Logic Programming. In *Declarative Agent Languages and Technologies (DALT'05)*, 2005.
- [15] M. De Vos and D. Vermeir. Choice Logic Programs and Nash Equilibria in Strategic Games. In J. Flum and M. Rodríguez-Artalejo, editors, *Computer Science Logic (CSL'99)*, volume 1683 of *LNCS*, pages 266–276, Madrid, Spain, 1999. Springer Verslag.
- [16] M. De Vos and D. Vermeir. Extending Answer Sets for Logic Programming Agents. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):103–139, Sept. 2004. Special Issue on Computational Logic in Multi-Agent Systems.
- [17] T. Eiter, W. Faber, N. Leone, and G. Pfeifer. The diagnosis frontend of the dlv system. *AI Communications*, 12(1-2):99–111, 1999.
- [18] T. Eiter, W. Faber, N. Leone, G. Pfeifer, and A. Polleres. The DLV^k planning system. In *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA 2002)*, volume 2424 of *LNAI*, pages 541–544. Springer, 2002.
- [19] T. Eiter and G. Gottlob. The complexity of logic-based abduction. *Journal of the Association for Computing Machinery*, 42(1):3–42, 1995.
- [20] T. Eiter, G. Gottlob, and Y. Gurevich. Normal forms for second-order logic over finite structures, and classification of np optimization problems. *Annals of Pure and Applied Logic*, 78(1-3):111–125, 1996.
- [21] T. Eiter, G. Gottlob, and H. Mannila. Disjunctive datalog. *ACM Transactions on Database Systems*, 22(3):364–418, 1997.
- [22] T. Eiter and A. Polleres. Towards automated integration of guess and check programs in answer set programming. In *Proceedings of the 7th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2004)*, volume 2923 of *LNCS*, pages 100–113. Springer, 2004.
- [23] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Logic-based agent communication protocols. In *Advances in Agent Communication*, volume 2922 of *LNAI*, pages 91–107. Springer, 2004.
- [24] W. Faber and G. Pfeifer. dlv homepage. <http://www.dbai.tuwien.ac.at/proj/dlv/>.

- [25] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings of the Fifth International Conference and Symposium on Logic Programming*, pages 1070–1080. The MIT Press, 1988.
- [26] S. Heymans, D. Van Nieuwenborgh, and D. Vermeir. Hierarchical decision making by autonomous agents. In *Proceedings of 9th European Conference on Logics in Artificial Intelligence (JELIA2004)*, volume 3229 of *LNCS*, pages 44–56. Springer, 2004.
- [27] N. Immerman. *Descriptive Complexity*. Springer, 1999.
- [28] Jade: <http://jade.tilab.com/>.
- [29] J. A. Leite, J. J. Alferes, and L. M. Pereira. Minerva - a dynamic logic programming agent architecture. In *Intelligent Agents VIII*, number 2002 in *LNAI*, pages 141–157. Springer, 2002.
- [30] V. Lifschitz. Answer set programming and plan generation. *Journal of Artificial Intelligence*, 138(1-2):39–54, 2002.
- [31] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.
- [32] Protégé. <http://protege.stanford.edu/>.
- [33] C. Sakama and K. Inoue. Coordination between logical agents. In *Proceedings of the 5th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA-V)*, volume 3487 of *LNAI*, pages 161–177. Springer, 2005.
- [34] P. Simons. smodels homepage. <http://www.tcs.hut.fi/Software/smodels/>.
- [35] SOCS: <http://lia.deis.unibo.it/research/socs/>.
- [36] T. Soinen and I. Niemelä. Developing a declarative rule language for applications in product configuration. In *Proceedings of the First International Workshop on Practical Aspects of Declarative Languages (PADL '99)*, LNCS, San Antonio, Texas, 1999. Springer.
- [37] L. Stockmeyer and A. Meyer. Word problems requiring exponential time. In *Proceedings of the 5th ACM Symposium on Theory of Computing (STOC '73)*, pages 1–9, 1973.
- [38] D. Van Nieuwenborgh, M. De Vos, S. Heymans, and D. Vermeir. Hierarchical decision making in multi-agent systems using answer set programming. Technical report, Vrije Universiteit Brussel, Dept. of Computer Science, 2005.
- [39] D. Van Nieuwenborgh, S. Heymans, and D. Vermeir. Approximating extended answer sets. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*. IOS Press.
- [40] D. Van Nieuwenborgh, S. Heymans, and D. Vermeir. On programs with linearly ordered multiple preferences. In *Proceedings of 20th International Conference on Logic Programming (ICLP 2004)*, number 3132 in *LNCS*, pages 180–194. Springer, 2004.
- [41] D. Van Nieuwenborgh and D. Vermeir. Preferred answer sets for ordered logic programs. In *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA 2002)*, volume 2424 of *LNAI*, pages 432–443. Springer, 2002.
- [42] D. Van Nieuwenborgh and D. Vermeir. Order and negation as failure. In *Proceedings of the 19th International Conference on Logic Programming*, volume 2916 of *LNCS*, pages 194–208. Springer, 2003.
- [43] D. Van Nieuwenborgh and D. Vermeir. Ordered diagnosis. In *Proceedings of the 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR2003)*, volume 2850 of *LNAI*, pages 244–258. Springer, 2003.
- [44] J. Vázquez-Salceda, J. Padget, U. Cortés, A. López-Navidad, and F. Caballero. Formalizing an electronic institution for the distribution of human tissues. *Artificial Intelligence in Medicine*, 27(3):233–258, 2003. published by Elsevier.
- [45] T. Wakaki, K. Inoue, C. Sakama, and K. Nitta. Computing preferred answer sets in answer set programming. In *Proceedings of the 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR2003)*, volume 2850 of *LNAI*, pages 259–273. Springer, 2003.