

**LOGCOMP  
RESEARCH  
REPORT**



**INSTITUT FÜR LOGIC AND COMPUTATION  
FACHBEREICH WISSENSBASIERTE SYSTEME**

**ABSTRACTION FOR NON-GROUND ANSWER  
SET PROGRAMS**

**ZEYNEP G. SARIBATUR      THOMAS EITER      PETER SCHÜLLER**

**LOGCOMP RESEARCH REPORT 19-01**

**DECEMBER 2019**

Institut für Logic and Computation  
FB Wissensbasierte Systeme  
Technische Universität Wien  
Favoritenstraße 9-11  
A-1040 Wien, Austria  
Tel: +43-1-58801-18405  
Fax: +43-1-58801-18493  
sek@kr.tuwien.ac.at  
www.kr.tuwien.ac.at





## ABSTRACTION FOR NON-GROUND ANSWER SET PROGRAMS

Zeynep G. Saribatur<sup>1</sup>      Thomas Eiter<sup>1</sup>      Peter Schüller<sup>1</sup>

**Abstract.** Abstraction is an important technique utilized by humans in model building and problem solving, in order to figure out key elements and relevant details of a world of interest. This naturally has led to investigations of using abstraction in AI and Computer Science to simplify problems, especially in the design of intelligent agents and automated problem solving. By omitting details, scenarios are reduced to ones that are easier to deal with and to understand, where further details are added back only when they matter. Despite the fact that abstraction is a powerful technique, it has not been considered much in the context of nonmonotonic knowledge representation and reasoning, and specifically not in Answer Set Programming (ASP), apart from some related simplification methods. In this work, we introduce a notion for abstracting from the domain of an ASP program such that the domain size shrinks while the set of answer sets (i.e., models) of the program is over-approximated. To achieve the latter, the program is transformed into an abstract program over the abstract domain while preserving the structure of the rules. We show in elaboration how this can be also achieved for single or multiple sub-domains (sorts) of a domain, and in case of structured domains like grid environments in which structure should be preserved. Furthermore, we introduce an abstraction-&-refinement methodology that makes it possible to start with an initial abstraction and to achieve automatically an abstraction with an associated abstract answer set that matches an answer set of the original program, provided that the program is satisfiable. Experiments based on prototypical implementations reveal the potential of the approach for problem analysis, by its ability to focus on the parts of the program that cause unsatisfiability and by achieving concrete abstract answer sets that merely reflect relevant details. This makes domain abstraction an interesting topic of research whose further use in important areas like Explainable AI remains to be explored.

---

<sup>1</sup>Institute of Logic and Computation, TU Wien; email: (zeynep | eiter | ps)@kr.tuwien.ac.at.

**Acknowledgements:** This work has been partially supported by the Austrian Science Fund (FWF) grant W-1255 and by the EU's H2020 research and innovation programme under grant agreement 825619 (AI4EU).

Some of the results in this article were presented in preliminary form at JELIA 2019 (Saribatur et al., 2019) and XAI 2019 (Eiter et al., 2019).

Copyright © 2019 by the authors

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Background</b>	<b>8</b>
2.1	Answer Set Programming . . . . .	8
2.2	Abstraction in ASP . . . . .	11
2.2.1	Abstraction Refinement Methodology . . . . .	12
<b>3</b>	<b>Domain Abstraction</b>	<b>14</b>
3.1	Towards an Abstract Program . . . . .	15
3.2	Lifted Built-in Relations . . . . .	17
3.3	Abstract Program Construction . . . . .	18
3.3.1	Restricted Case . . . . .	18
3.3.2	Abstract Program (General Case) . . . . .	20
3.4	Syntactic Extensions and Further Considerations . . . . .	23
3.4.1	Treating Choice Rules . . . . .	23
3.4.2	Other Forms of Relations . . . . .	24
3.4.3	Concreteness with Projection . . . . .	24
3.5	Properties of Domain Abstraction . . . . .	25
3.5.1	Abstraction over Sorts . . . . .	27
3.5.2	Cartesian Abstraction . . . . .	28
3.6	Computational Complexity . . . . .	28
<b>4</b>	<b>Refinement by Debugging Non-Ground Spuriousness</b>	<b>30</b>
4.1	Non-Ground Debugging using Tagging . . . . .	32
4.2	Deciding on a Refinement . . . . .	36
<b>5</b>	<b>Multi-Dimensional Domain Abstraction</b>	<b>37</b>
5.1	Existential Abstraction on Relations . . . . .	38
5.1.1	Abstraction procedure . . . . .	40
5.1.2	Computing Joint Abstract Relation Types . . . . .	42
5.2	Quad-tree Abstraction . . . . .	43
<b>6</b>	<b>Implementation and Evaluation</b>	<b>45</b>
6.1	Overall Methodology . . . . .	45
6.2	Implementation . . . . .	46
6.2.1	DASPAR . . . . .	46
6.2.2	mDASPAR . . . . .	47
6.2.3	Implementation Aspects of mDASPAR . . . . .	49
6.3	Evaluation: Obtaining Abstract Solutions . . . . .	50
6.3.1	Experiments . . . . .	50
6.3.2	Results . . . . .	51
6.4	Evaluation: Unsolvable Problem Instances in Grid-Cells . . . . .	54
6.4.1	Effects of Different Debugging Approaches . . . . .	55
6.4.2	User Study on Unsatisfiability Explanations . . . . .	57

<b>7</b>	<b>Discussion</b>	<b>59</b>
7.1	Predicate Abstraction . . . . .	59
7.2	Use Case: Abstraction in ASP Planning . . . . .	61
7.2.1	Abstracting over Irrelevant Details . . . . .	62
7.2.2	Computing Abstract Plans . . . . .	64
7.3	Related Work . . . . .	65
7.3.1	Equivalence-based Rewriting and Program Transformations in ASP . . . . .	66
7.3.2	Abstraction in Planning and Agent Verification . . . . .	67
7.3.3	Generalized Planning . . . . .	68
<b>8</b>	<b>Conclusion</b>	<b>68</b>
8.1	Outlook . . . . .	69
<b>A</b>	<b>Proofs</b>	<b>77</b>
<b>B</b>	<b>Further Details</b>	<b>86</b>
B.1	Correctness Checking Failure for Non-tight Programs . . . . .	86
B.2	Grid-Cell Problem Encodings . . . . .	86
B.3	Example run of mDASPAR . . . . .	88
B.4	Modular Concreteness Checking . . . . .	90
<b>C</b>	<b>Use Case: Abstraction for Policy Refutation</b>	<b>92</b>

# 1 Introduction

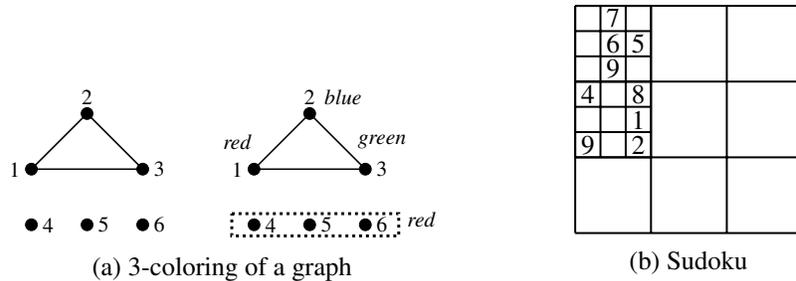
Abstraction is a technique applied in human reasoning and understanding, by reasoning over the models of the world that are built mentally ( Craik, 1952; Johnson-Laird, 1983). Although its meaning comes from “to draw away”, there is no precise definition that is capable of covering all meanings that abstraction has in its utilizations. There is a variety of interpretations in different disciplines such as Philosophy, Cognitive Science, Art, Mathematics and Artificial Intelligence, with the shared consensus of the aim to “distill the essential” (Saitta and Zucker, 2013). Among them is the capability of abstract thinking, which is achieved by removing irrelevant details and identifying the “essence” of a problem (Kramer, 2007). The notion of relevance is especially important in problem solving, as a problem may become too complex to solve if every detail is taken into account. A good strategy to solve a complex problem is to start with a coarse solution and then refine it by adding back more details. When planning a trip, for instance, one may first pick the destination and determine a coarse travel plan; fleshing out the precise details of the travel, such as taking the subway to the airport, comes later. This may be done in a hierarchy of levels of abstraction, with the lowest level containing all of the details. Another view of abstraction is the generalization aspect, which is singling out the relevant features and properties shared by objects. For example, features of an airplane such as color and cargo capacity with their possible differences may be irrelevant to the travel plan; we are (mostly) only interested in the fact that there is an airplane that takes us from Vienna to New York, say. Overall, the general aim of abstraction is to simplify the problem at hand to one that is easier to understand and deal with.

For solving combinatorial problems and figuring out the key elements, humans arguably employ abstraction. In Artificial Intelligence, such problems vary from planning problems like in which order to move blocks to achieve a final configuration, to solving constraint problems such as finding an admissible coloring of the nodes of a given graph. In the latter problem, for instance, isolated nodes can be viewed as a single node and colored the same without thinking about the specific details (Figure 1a). If a given graph is non-colorable, then we may try to find some subgraph (e.g., a clique) which causes the unsolvability, and we would not care about other nodes in the graph. Similarly with the blocks: if the labels are not important, we would disregard them when figuring out the actions. If the goal configuration can not be achieved from the initial one, we would aim to find out the particular blocks that cause this.

Notably, such disregard of detail also occurs for problems with multi-dimensional structures such as grid-cells in the well-known Sudoku problem, where a partially filled  $9 \times 9$  board must be completed by filling in numbers 1..9 into the empty cells under constraints. If an instance is unsolvable, the reason can only be meaningfully grasped by a human by focusing on the relevant sub-regions, as looking at the whole grid is too complex. For illustration, Figure 1b shows the sub-regions of an instance that contain the reason why no solution exists: as 6 and 7 occur in the middle column, they must appear in the sub-region below in the left column, which is unfeasible as there is only one empty cell. All these examples demonstrate abstraction abilities of humans that come naturally.

Due to its important role in knowledge representation and in reasoning, abstraction has been explored in AI research early on as a useful tool for problem solving: solve a problem at hand first in an abstracted space, and then use the abstract solution as a heuristics to guide the search for a solution in the original space (Newell and Simon, 1972; Sacerdoti, 1974; Knoblock, 1994). This approach was used in planning for speeding up the solving (Holte et al., 1996) and especially for computing heuristic functions to guide the plan search in the state space. Several abstraction methods were introduced towards this direction, especially to automatically compute abstractions that give a good heuristic (Edelkamp, 2001; Helmert et al., 2007; Seipp and Helmert, 2013). However, it is well known that the success in solving a problem relies on how “good” the abstraction is. For this, theoretical approaches for defining abstractions with desired properties

Figure 1: Use of abstraction



have been investigated (Hobbs, 1990; Giunchiglia and Walsh, 1992; Nayak and Levy, 1995). Apart from gaining efficiency (which however may not always materialize (Backstrom and Jonsson, 1995; Hoffmann et al., 2006)), abstraction forms a basis to obtain high-level explanations and an understanding of a problem.

Abstraction has been studied in other areas of AI and Computer Science as well, among them model-based diagnosis (Mozetič, 1991; Chittaro and Ranon, 2004), constraint satisfaction (Freuder, 1991; Bistarelli et al., 2002), theorem proving (Plaisted, 1981), to name a few. Particularly fruitful were applications in model checking, which is a highly successful approach to computer aided verification (Clarke et al., 2018), to tackle the state explosion problem by property preserving abstractions (Clarke et al., 1994; Loiseaux et al., 1995; Dams et al., 1997). Furthermore, the seminal counterexample guided abstraction refinement (CEGAR) method (Clarke et al., 2003) allows for automatic generation of such abstractions, by starting from an initial abstraction that over-approximates the behavior of a system to verify, and then stepwise refining the abstraction as long as needed, i.e., as long as spurious (false) counterexamples exist.

**Abstraction for Answer Set Programming.** Answer Set Programming (ASP) (Lifschitz, 2008b; Brewka et al., 2011) is a declarative problem solving paradigm that is rooted in knowledge representation, logic programming, and nonmonotonic reasoning. Akin to SAT solving, a problem is represented by a non-monotonic logic program whose models (called “answer sets” (Gelfond and Lifschitz, 1991)) correspond to the solutions of the problem. Thanks to the availability of efficient solvers and the expressiveness of the formalism, ASP has been gaining popularity for applications in many areas of AI and beyond, cf. (Erdem et al., 2016; Erdem and Patoglu, 2018; Falkner et al., 2018) and references therein, from combinatorial search problems (e.g. configuration, diagnosis, planning) over system modeling (e.g., behavior of dynamic systems, beliefs and actions of agents) to knowledge-intensive applications (e.g., query answering, explanation generation), to name a few (Erdem et al., 2016). The declarative nature of ASP enables a flexible use for solving different reasoning problems, and it provides a useful basis for investigating ways to help in understanding a problem with its key elements. Studies in understanding how ASP programs find a solution (or none) to a problem have been conducted, which mainly focus on debugging answer sets (Brain et al., 2007; Gebser et al., 2008; Oetsch et al., 2010) or finding justifications (Pontelli et al., 2009; Schulz and Toni, 2013; Cabalar et al., 2014). These approaches could be used to aid in understanding the problem at hand; however, as noted in (Fandinno and Schulz, 2019), the explanations offered may contain a high number of details which prevent one from seeing the crucial parts. This is where abstraction would come in handy and could be very fruitfully used.

Somewhat surprisingly, abstraction has not been considered much in the context of nonmonotonic knowledge representation and reasoning, including ASP as a premier formalism in this area. Simplification methods such as equivalence-based rewriting (Gebser et al., 2008; Pearce, 2004), partial evaluation (Brass and

Dix, 1997; Janhunen et al., 2006), or forgetting (see (Leite, 2017) for a recent survey), have been extensively studied; however, they strive for preserving semantics, while abstraction may lead to an over-approximation of the answer sets of a logic program, in a modified language. Spurious answer sets, i.e., abstract answer sets that do not correspond to some answer set of the original program, may be discarded, while cautious reasoning from all abstract answer sets preserves soundness; moreover, in case no spurious answer set exists we also have completeness (in particular, if no abstract answer set exists). This makes abstraction an interesting topic for research.

In a recent work (Saribatur and Eiter, 2018), a notion of abstraction for ASP was introduced that focuses on omission of atoms from the vocabulary and ensures over-approximation by rewriting the rules of a given program. That approach is propositional in nature related to forgetting, with the difference of over-approximation vs preserving the answer sets. We follow an orthogonal approach and introduce in this work a notion of abstraction for ASP *on the first-order level* that is concerned with collapsing (i.e., clustering) objects in the (Herbrand) domain of a program. It is that in this way, multiplicity is removed in the spirit of Occam’s razor.<sup>1</sup> If the graph in Figure 1a is represented by facts  $node(1), \dots, node(6)$  and  $edge(1,2), edge(1,3), edge(1,3)$ , then collapsing the nodes 4, 5, and 6 into an abstract node  $a$  would not affect 3-colorability of the graph; we thus expect that an abstract version of an ASP program that encodes all 3-colorings of the graph would yield answer sets from which these 3-colorings can be recovered. However, if there were an edge between 4 and 5, then over-approximation would yield that this abstract program will have spurious answer sets, as the abstract node  $a$  may have a single color in some abstract answer set  $I'$  while 4 and 5 must have in every original answer set  $I$  different colors, thus  $I$  can not be mapped to  $I'$ . Similarly, if nodes 1 and 2 were collapsed, the graph would become 2-colorable (where 1 and 2 share the same color), while the original graph is not 2-colorable. These simple examples shows that a naive use of domain abstraction – just replace individuals by a cluster of elements – does not work; and yet more subtle effects may surface when programs have recursive definitions such as reachability.

In fact, suitable domain abstraction for ASP is non-trivial and has several challenges. First, the abstract program should be automatically constructed, while the structure of the original rules should be preserved if this is feasible. Second, abstraction refinement, i.e., unclustering of objects for eliminating spurious answer sets, should be automated as well. This is non-trivial, given a large space of possible refinements and that objects might be related among each other in multiple ways, e.g. in temporal or spatial relationships as in reasoning about actions, for instance. And third, the capability of dealing with structure and to support hierarchical abstraction that can handle objects of different granularities at different levels of abstraction is needed.

**Contributions.** We address the issues above in this work, whose main contributions can be summarized as follows.

(1) We formally introduce the notion of domain abstraction for ASP programs. To this end, we define abstraction mappings  $m$  from the original (concrete) domain  $D$  of the program to an abstract domain  $D'$ , and construct an abstract program  $\Pi^m$  over  $D'$  such that each answer set  $I$  of  $\Pi$  maps to an abstract answer set  $I'$  of  $\Pi^m$ . The construction works modularly on the syntactic level and transforms each rule in  $\Pi$  into a set of abstract rules with a similar structure. In the transformation, built-in relations and in particular equality, whose treatment provides the backbone of the method, are lifted to the abstract level, and uncertainty caused by the abstracted domain  $D'$  is carefully respected. Our notion of abstraction can be used for different applications such as obtaining abstract solutions from ASP programs or showing reasons of unsatisfiability in

---

<sup>1</sup>Often referred to as “Entia non sunt multiplicanda praeter necessitatem.” that is, entities should not be multiplied beyond necessity.

case no answer set exists. We illustrate this on various problems expressed in ASP, among them problems from combinatorial search, planning, and agent behavior assessment.

(2) We present a method which, in case an abstract answer set  $I'$  of a program  $\Pi$  w.r.t. mapping  $m$  is spurious, computes a refinement  $m'$  of the abstraction in order to eliminate  $I'$ . To this end, we reduce the test for spuriousness to unsatisfiability of a non-ground ASP program  $\Pi'$  constructed from  $I'$ ,  $\Pi$ , and  $m$ . As unsatisfiability of  $\Pi'$  as such leaves one clueless, we use a debugging technique for ASP in order to obtain useful information for computing a promising refinement  $m'$ . To this end, we lift the SPOCK approach (Brain et al., 2007) for tight programs to the non-ground level such that decisions on refinements can be based on special atoms computed during the debugging. Intuitively, these atoms single out changes for a spurious answer set  $I'$  towards a corresponding answer set  $I$  of the original program  $\Pi$ , where always some such atoms will be found; based on heuristics, we introduce different refinement strategies to eliminate spurious answer sets. These strategies are employed in a CEGAR-style (Clarke et al., 2003) methodology of iterative abstraction and refinement, which starts with a highly coarse abstraction and automatically searches for and outputs an abstraction with a non-spurious (concrete) answer set if one exists.

(3) We introduce the possibility of multi-dimensional abstraction mappings over a domain. While the abstraction method from above can deal with sorts, we have to modify it to form an abstraction over the relations that is akin to existential abstraction (Clarke et al., 2003), in order to enable that elements at mixed levels of abstraction can be handled properly. This in fact is needed to relate in Figure 1b cells like the top-left corner with abstract cells such as the mid-left  $3 \times 3$  sub-region, and to express their abstract locations such as being above, left-of etc. We extend the abstraction-&-refinement methodology with handling the structural aspects of grid-cells by using a quad-tree abstraction, and we consider more sophisticated decision making for the refinement to observe its effects on the resulting abstractions.

(4) We analyze semantic and computational properties of the abstraction approach, which can be exploited for modeling and for guiding the design of suitable implementations. Among other results, we establish that abstractions for sequences of refinements can be built incrementally (Proposition 3.6), and that abstractions of independent sorts can be naturally composed (Proposition 3.7). Furthermore, the two variants of domain abstractions we consider are semantically equivalent, which we demonstrate here for the basic case (Theorem 5.1), but have features making them attractive in different contexts. As regards complexity, we show that checking whether an abstract answer set is spurious is coNEXP-complete and that deciding whether an abstraction mapping is faithful, i.e., has no spurious abstract answer sets, is coNEXP<sup>NP</sup>-complete. Thus, the worst case complexity of these problems is roughly by an exponential higher than the one of answer set checking resp. unsatisfiability testing of non-ground ASP programs (Eiter et al., 2007). Furthermore, if predicate arities are bounded by a constant, the problems are  $\Pi_2^p$ -complete and  $\Pi_p^3$ -complete, respectively, and thus one level higher up in the Polynomial Hierarchy than answer set checking and unsatisfiability testing, respectively. Reducing spurious checking to unsatisfiability testing as we consider is thus worst-case optimal in both settings.

(5) We have implemented the abstraction-&-refinement approach in prototypical tools DASPARE and mDASPARE for plain and multi-dimensional abstraction, respectively. They take as input a non-ground program  $\Pi$  and an initial coarse domain mapping  $m$  (which by default is the trivial mapping that clusters all elements), and output a refinement mapping  $m'$  and an abstract answer set  $\hat{I}$  for the abstract program  $\Pi^{m'}$  that is non-spurious, if one exists; otherwise, i.e., in case  $\Pi$  is unsatisfiable, they provide a refinement mapping  $m'$  that is faithful, i.e., such that  $\Pi^{m'}$  has no abstract answer sets. The implementations include different refinement strategies and support independent sorts 2-dimensional abstractions with a quadtree-style refinement process. Based on these tools, we have conducted an experimental evaluation of the approach,

where one set of experiments focused on finding non-trivial abstractions for problems expressed by well-known ASP programs (graph coloring, scheduling), while another one consisted in detecting the unsolvability of several benchmark problems involving grid-cells. In the experiments, different debugging strategies were considered and a measure for assessing the quality of multi-dimensional abstraction was developed. The results show the potential of the approach, where in particular a small user-study for a natural grid-cell problem indicates its capability of putting a human-like focus on the grid to show unsatisfiability.

Summarizing, our work on domain abstraction for ASP opens an intriguing line of research which aims at making it possible to identify the gist of a program’s domain that is responsible for matters such as inconsistency or certain solutions of interest. The approach that we present provides the ability to adjust the granularity of abstraction towards the details relevant for a problem in an (semi-)automated way. The experimental results indicate the value of domain abstraction for program analysis, whose further use in important areas like Explainable AI remains to be explored.

**Organisation.** The remainder of this article is organized as follows. The next section recalls notions from ASP as needed for this work, and reviews the seminal approach to abstraction over (in essence) propositional ASP programs introduced in (Saribatur and Eiter, 2018). After that, we turn in Section 3 to domain abstraction for non-ground ASP programs, where we present an abstraction procedure, consider various extensions, and study semantic and computational properties of the approach. In the subsequent Section 4, we present our refinement methodology that is based on debugging of ASP programs. This is followed by discussing in Section 5 multi-dimensional abstraction and an alternative abstraction method, based on existential abstraction, that is needed for it. Implementation and evaluation of the approach is considered in Section 6, while in Section 7 we discuss further notions of abstraction and possible use cases, as well as related work. The final Section 8 provides a summary and outlines issues for future research.

In order not to distract from the flow of reading, longer proofs and further details have been moved to the Appendix (A and B), which also provides a further use case in agent behavior assessment (C).

## 2 Background

### 2.1 Answer Set Programming

In this section, we recall the concepts and notions of Answer Set Programming (ASP) that we need for this article. We refer to (Brewka et al., 2011, 2016; Schaub and Woltran, 2018) for more background and references. We start with syntax and semantics of ASP programs and then recall some notions that are useful for this work.

**Syntax.** We consider a first order vocabulary  $\mathcal{L} = (\mathcal{P}, \mathcal{C})$  consisting of non-empty finite sets  $\mathcal{P}$  of predicates and  $\mathcal{C}$  of constants. Let  $\mathcal{V}$  represent the set of variable symbols. A *term* is either a constant from  $\mathcal{C}$  or a variable from  $\mathcal{V}$ . An *atom* is of the form  $p(t_1, \dots, t_n)$  where  $p \in \mathcal{P}$  and each  $t_i$  is a term;  $n$  is the *arity* of  $p$ . Atoms are called *propositional* if  $k = 0$  and *ground* if they do not contain variables. A *literal*  $l$  is an atom  $a$  or a strongly (“classically”) negated atom  $\neg a$ ,

A *rule*  $r$  is an expression of the form

$$\alpha_0 \leftarrow \alpha_1, \dots, \alpha_m, \text{not } \alpha_{m+1}, \dots, \text{not } \alpha_n, \quad 0 \leq m \leq n, \quad (1)$$

where each  $\alpha_i$  is a literal and *not* is negation as failure alias strong negation. Intuitively, a *default literal* *not*  $\alpha_i$  is true, if  $\alpha$  cannot be derived using rules, and false otherwise. We refer to  $\alpha_0$  as the *head* of  $r$ , and  $\alpha_1, \dots, \alpha_m, \text{not } \alpha_{m+1}, \dots, \text{not } \alpha_n$  as the *body* or  $r$ . We also write  $r$  as  $\alpha_0 \leftarrow B(r)$ , such that  $H(r) = \alpha_0$  denotes

the head and  $B(r)$  denotes the set of all the body literals  $B^+(r) \cup B^-(r)$ , where  $B^+(r) = \{\alpha_1, \dots, \alpha_m\}$  is the *positive body* and  $B^-(r) = \{\alpha_{m+1}, \dots, \alpha_n\}$  is the *negative body* of  $r$ ; thus, we have  $H(r) \leftarrow B^+(r), \text{not } B^-(r)$ . We may omit  $r$  from  $B(r), B^+(r)$  etc. if  $r$  is clear.

A rule  $r$  is a *constraint*, if  $\alpha_0$  is falsity ( $\perp$ , then omitted) and a *fact*, if  $n=0$  and no variable occurs in  $r$ . Furthermore,  $r$  is *positive* if  $n = m$ , and *normal*, if no *not* occurs in  $r$ . A rule is *ground*, if all literals occurring in it are ground.

A *program*  $\Pi$  is a finite set of rules. It is *positive*, *normal*, *ground* if the rules in it have the respective property; it is *safe*, if every variable that occurs in a rule also occurs in the same rule in some positive body literal. It is *domain-restricted*, if this body literal is over a positive *domain predicate*, which are predicates not defined via negative recursion or using choice rules.

**Semantics.** The answer set semantics is defined via ground programs. For a program  $\Pi$ , we define its *ground instantiation* as follows.

Given a program  $\Pi$ , its *Herbrand universe*, denoted by  $HU_\Pi$ , is the set of all constant symbols  $C \subseteq \mathcal{C}$  appearing in  $\Pi$ ; in case there is no constant symbol, then  $HU_\Pi = \{c\}$  for some arbitrary constant symbol. The *Herbrand base* of a program  $\Pi$ , denoted by  $HB_\Pi$ , is the set of all ground literals constructed using predicates from  $\mathcal{P}$  and constants from  $\mathcal{C}$ .

The *ground instances of a rule*  $r \in \Pi$ , denoted by  $grd(r)$ , is obtained by replacing all variables in  $r$  with constant symbols in  $HU_\Pi$ . The *grounding* of a program  $\Pi$  then becomes  $grd(\Pi) = \bigcup_{r \in \Pi} grd(r)$ . To group the rules in  $grd(\Pi)$  with the same head  $q$ , we use  $def(q, \Pi) = \{r \in \Pi \mid H(r) = q\}$ .

Let  $\Pi$  be a ground program. A set  $L \subseteq HB_\Pi$  of literals is *consistent*, if  $p, \neg p \notin L$  for every atom  $p \in HB_\Pi$ . An *interpretation*  $I$  is a consistent subset of  $HB_\Pi$ . An interpretation  $I$  *satisfies* a rule  $r \in \Pi$ , denoted by  $I \models r$ , if  $H(r) \subseteq I$  whenever  $B^+(r) \subseteq I$  and  $B^-(r) \cap I = \emptyset$ . An interpretation is a *model* of  $\Pi$ , denoted by  $I \models \Pi$ , if  $I \models r$  for all  $r \in \Pi$ . A model  $I$  is *minimal*, if there is no model  $J$  of  $\Pi$  such that  $J \subset I$ .

**Example 2.1.** Consider the program  $\Pi$  below and the interpretation  $I = \{a, b, d\}$ .

$c \leftarrow \text{not } d.$   
 $d \leftarrow \text{not } c.$   
 $a \leftarrow \text{not } b, c.$   
 $b \leftarrow d.$

$I$  is a model of  $\Pi$ , but it is not minimal, since the interpretation  $I' = \{b, d\}$  is also a model of  $\Pi$ .

**Definition 2.1** (GL-reduct). *The Gelfond-Lifschitz (GL-)reduct of a program  $\Pi$  relative to an interpretation  $I \subseteq HB_\Pi$ , denoted by  $\Pi^I$ , is the ground positive program obtained from  $grd(\Pi)$  when each rule  $H(r) \leftarrow B^+(r), \text{not } B^-(r)$*

- (i) *with  $B^-(r) \cap I \neq \emptyset$  is deleted, and*
- (ii) *is replaced by  $H(r) \leftarrow B^+(r)$ , otherwise.*

Informally, the first step is to remove the rules where  $I$  contradicts a default negated literal, and from the remaining rules, the second step removes their negative body. An interpretation  $I$  is an *answer set* of  $\Pi$ , if it is the minimal model of the GL-reduct  $\Pi^I$ . Apart from the GL-reduct which is considered to be the standard definition for stable models (i.e., answer sets), a collection of other definitions can be found in (Lifschitz, 2008a).

The set of answer sets of a program  $\Pi$  is denoted as  $AS(\Pi)$ . A program  $\Pi$  is *unsatisfiable*, if  $AS(\Pi) = \emptyset$ .

**Example 2.2** (ctd).  $\Pi$  has two answer sets, viz.  $I_1 = \{c, a\}$  and  $I_2 = \{d, b\}$ ; indeed,

- $\Pi^{I_1} = \{c \leftarrow \text{not } d.; a \leftarrow \text{not } b, c.\}$  and  $I_1$  is a minimal model of  $\Pi^{I_1}$ ; similarly,
- $\Pi^{I_2} = \{d \leftarrow \text{not } c.; b \leftarrow d.\}$  has  $I_2$  among its minimal models.

**language extensions.** *Choice rules* are a syntactic extension of the form  $\{\alpha\} \leftarrow B$ , which stands for the rules  $\alpha \leftarrow B, \text{not } \bar{\alpha}$  and  $\bar{\alpha} \leftarrow B, \text{not } \alpha$ , where  $\bar{\alpha}$  is a new atom. Cardinality constraints and conditional literals are further common syntactic extensions (Simons et al., 2002); in particular,  $i_\ell \{a(X) : b(X)\} i_u$  is true whenever at least  $i_\ell$  and at most  $i_u$  instances of  $a(X)$  subject to  $b(X)$  are true.

**Example 2.3.** Consider the following rules.

$$1\{a(X, Y, Z) : b(X), c(Y)\} 1 \leftarrow d(Z). \quad (2)$$

$$\{a(X) : b(X); c(X, Y) : b(X), d(Y)\} 1. \quad (3)$$

Rule (2) states that exactly one instance of  $a(X, Y, Z)$  subject to  $b(X)$  and  $c(Y)$  has to be true, where the value of  $Z$  depends on the instance of  $d(Z)$  that is true, and rule (3) states that at most one instance of either  $a(X)$  subject to  $b(X)$  or  $c(X, Y)$  subject to  $b(X)$  and  $d(Y)$  has to be true.

A *weak constraint* (Leone et al., 2006) is of the form

$$:\sim \alpha_1, \dots, \alpha_m, \text{not } \alpha_{m+1}, \dots, \text{not } \alpha_n. [w : l]$$

where  $w$  (the *weight*) and  $l$  (the *level*) are positive integer constants or variables. The costs of all violated (instances of) weak constraints (grouped by levels of priorities  $l$ ) are added up and assigned to the answer set. Among all answer sets, those whose cost vector is lexicographically smallest are chosen as *optimal answer sets*. Using weak constraints is a convenient way of performing optimizations.

**Dependencies.** The *dependency graph* of a ground program  $\Pi$  is a directed graph  $G_\Pi = (V, E)$ , where the vertices  $V$  equals  $HB_\Pi$ , and the edges  $E = E^+ \cup E^-$  consist of positive edges  $E^+$  from any  $q = H(r)$  to any  $p_1 \in B^+(r)$  and negative edges  $E^-$  from any  $q = H(r)$  to any  $p_2 \in B^-(r)$ , for all  $r \in \text{grd}(\Pi)$ .

**Example 2.4** (ctd).  $G_\Pi$  has positive edges  $a \rightarrow c$  and  $b \rightarrow d$  and negative edges  $c \rightarrow d$ ,  $d \rightarrow c$  and  $a \rightarrow b$ .

A non-empty set  $L$  of ground literals describes an *odd loop* of  $\Pi$  if for each pair  $p, q \in L$  there is a path  $\tau$  from  $p$  to  $q$  in  $G_\Pi$  with an odd number of negative edges. Constraints are viewed as simple odd loops. As well-known,  $\Pi$  is satisfiable, if it contains no odd loop, cf. (Papadimitriou and Yannakakis, 1992).

**Example 2.5** (ctd). The program  $\Pi$  has no odd loop, and thus has some answer set.

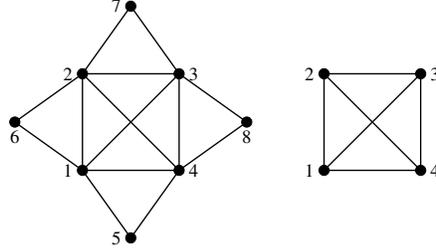
The *positive dependency graph* is the dependency graph containing only the positive edges, denoted by  $G_\Pi^+$ . A program  $\Pi$  is *tight* if  $G_\Pi^+$  is acyclic. A non-empty set  $L$  of ground literals describes a *positive loop* of  $\Pi$  if for each pair  $p, q \in L$  there is a path  $\tau$  from  $p$  to  $q$  in  $G_\Pi^+$  such that each literal in  $\tau$  is in  $L$ .

As we consider non-ground programs, we need to take care of cyclic dependencies of literals at the non-ground level. A *negative dependency cycle* of length  $n \geq 2$  is of the form

$$a_1(\bar{x}_1) \rightarrow a_2(\bar{x}_2) \rightarrow \dots \rightarrow a_n(\bar{x}_n) \rightarrow a_{n+1}(\bar{x}_{n+1})$$

with  $a_{n+1}(\bar{x}_{n+1}) = a_1(\bar{x}_1)$ , where  $a_i(\bar{x}_i) \rightarrow a_{i+1}(\bar{x}_{i+1})$  denotes that  $a_i(\bar{x}_i)$  is in the head of a rule  $r_i$ , and  $r_i$  has in its negative body some literal  $a_{i+1}(\bar{x}'_{i+1})$  that unifies with  $a_{i+1}(\bar{x}_{i+1})$ . For example, a choice rule

Figure 2: Non-3-colorable graph



consists of a negative cyclic dependency chain of length 2, and a rule of form  $p(X) \leftarrow \text{not } p(X), q(X)$  contains a cycle of length 1. In this article, we focus on the predicates  $a_i$  of the literals to determine the dependency, and we thus consider negative cyclic dependency chains of form  $a_1 \rightarrow a_2 \rightarrow \dots \rightarrow a_n \rightarrow a_{n+1}$  where  $a_{n+1} = a_1$ . Accordingly, we will consider a set  $L_c$  of literals as cyclic, if for each pair  $l_1, l_2 \in L_c$  a chain  $\text{pred}(l_1) \rightarrow \dots \rightarrow \text{pred}(l_2)$  exists.

ASP solvers first generate a grounding of the given program, and then a search for an answer set is conducted over the ground program.

## 2.2 Abstraction in ASP

Abstraction aims at discarding some details of a problem to obtain a more high-level view of a solution. This can be done in an over-approximation, which means that each original solution has some corresponding solution in the abstraction but one may encounter abstract solutions which do not have a corresponding original solution.

In (Saribatur and Eiter, 2018), the authors introduced such a notion of abstraction in ASP for over-approximating a given program  $\Pi$  on vocabulary  $\mathcal{A}$ , by constructing a program  $\Pi'$  on a smaller vocabulary  $\mathcal{A}'$ , i.e.,  $|\mathcal{A}| \geq |\mathcal{A}'|$ . More formally, abstraction was defined at the semantic level as follows.

**Definition 2.2** ((Saribatur and Eiter, 2018)). *Given two programs  $\Pi$  and  $\Pi'$  with  $|\mathcal{A}| \geq |\mathcal{A}'|$ , where  $\mathcal{A}, \mathcal{A}'$  are sets of ground atoms of  $\Pi$  and  $\Pi'$ , respectively,  $\Pi'$  is an abstraction of  $\Pi$  if there exists a mapping  $m: \mathcal{A} \rightarrow \mathcal{A}' \cup \{\top\}$  such that for any answer set  $I$  of  $\Pi$ ,  $I' = \{m(\alpha) \mid \alpha \in I\}$  is an answer set of  $\Pi'$ .*

We refer to  $m$  as an *abstraction mapping*. This notion of abstraction gives us the possibility to do clustering over atoms of the program. The (abstract) program  $\Pi'$  on the smaller vocabulary  $\mathcal{A}'$  serves to represent abstract answer sets. While the reduced vocabulary eases the search for an (abstract) answer set  $I'$ , an additional check is needed whether the original program  $\Pi$  has some answer that maps to  $I'$ . In (Saribatur and Eiter, 2018), the focus was on abstraction by omitting atoms from a program, i.e., by clustering them into  $\top$ .

**Example 2.6.** Consider the program that describes the graph 3-coloring problem below (adapted from the

coloring encoding in the ASP Competition 2013) and the graphs shown in Figure 1.

$$\begin{aligned}
& \text{color}(\text{red}). \text{color}(\text{green}). \text{color}(\text{blue}). \\
& \{\text{chosenColor}(N, C)\} \leftarrow \text{node}(N), \text{color}(C). \\
& \text{colored}(N) \leftarrow \text{chosenColor}(N, C). \\
& \perp \leftarrow \text{not colored}(N), \text{node}(N). \\
& \perp \leftarrow \text{chosenColor}(N, C_1), \text{chosenColor}(N, C_2), C_1 \neq C_2. \\
& \perp \leftarrow \text{chosenColor}(N_1, C), \text{chosenColor}(N_2, C), \text{edge}(N_1, N_2).
\end{aligned} \tag{4}$$

If we omit in the (ground version) of the encoding (4) with the instance shown in Figure 1a all atoms involving the nodes 4,5,6, the resulting abstract program will have answer sets which all correspond to some answer set of the original program, as the omitted nodes can be colored arbitrarily without destroying 3-colorability.

We could likewise map all atoms  $a$  involving 4,5,6 to atoms  $a'$  in which these nodes are replaced by a new node  $k$ ; e.g.,  $\text{node}(4)$  would become  $\text{node}(k)$ ,  $\text{node}(5)$  becomes  $\text{node}(k)$  etc. The abstract answer sets correspond then again to original answer sets, as the coloring of 4,5,6 does not matter. On the other hand, if we consider the graph in Figure 2 and omit all atoms that involve nodes 5,6,7,8, then the resulting abstract program has no an answer set, as the remaining clique 1-2-3-4 is not 3-colorable; also the original program has no answer set.

The latter observations is not by accident but in fact a useful property.

**Proposition 2.1.** *Let  $\Pi'$  be an abstraction of  $\Pi$ . If  $AS(\Pi') = \emptyset$ , then we have  $AS(\Pi) = \emptyset$ .*

In general, over-approximation can cause abstract answer sets that have no corresponding original answer set.

**Definition 2.3** (Spurious & concrete answer sets). *Let  $\Pi'$  be an abstraction of  $\Pi$  for the mapping  $m$ . The answer set  $I' \in AS(\Pi')$  is concrete if there exists an answer set  $I \in AS(\Pi)$  such that  $m(I) = I'$ ; otherwise, it is spurious.*

In these terms, the abstract answer set of the first abstract program constructed in Example 2.6 are all concrete, while if we drop all all atoms involving nodes 1,2,3,4 from the graph in Figure 2, the abstract answer sets of the resulting abstract program are all spurious.

### 2.2.1 Abstraction Refinement Methodology

To get rid of spurious abstract answer sets, the abstraction mapping  $m$  needs to be *refined* to a more fine-grained abstraction; in case of omission abstraction, the refinement would be to add back some of the omitted atoms.

We consider a CEGAR-style (Clarke et al., 2003) abstraction refinement approach which refines an initial abstraction repeatedly until a concrete solution is found or inconsistency (i.e., absence of solutions) is detected.

Before describing the general methodology, we first illustrate the idea with the graph coloring example.

**Example 2.7** (ctd). Figure 3 shows the abstraction of omitting 7 of the 8 nodes and their edges. Deciding on a color for the remaining node is easy. However, in the original domain, no coloring can match node 1 colored *red* as the graph is uncolorable. A refinement of this abstraction would be to add back some of the nodes and the knowledge about the edges. Until an abstraction is achieved where the four nodes causing the uncolorability is distinguished, a spurious coloring always occurs.

Figure 3: Abstraction refinement upon spurious graph colorings

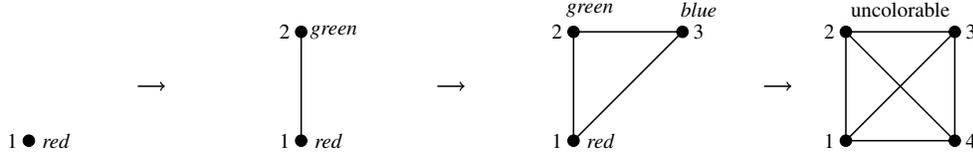


Figure 4: Abstraction & Refinement Methodology

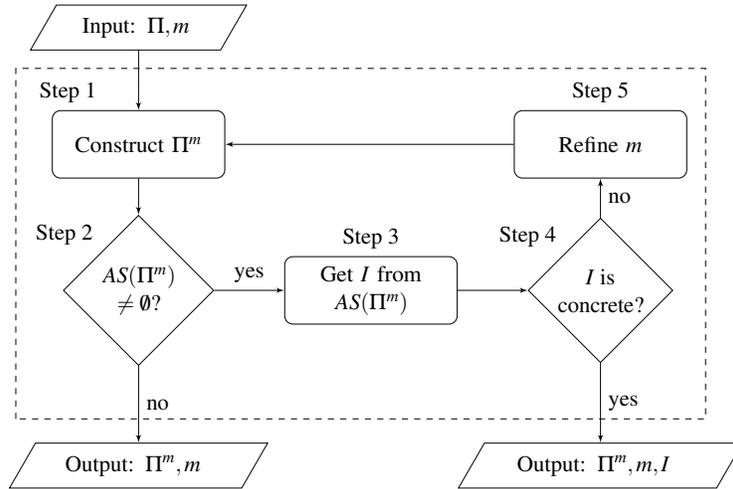


Figure 4 depicts the methodology introduced in (Saribatur and Eiter, 2018). For a program  $\Pi$ , we start with an initial abstraction mapping  $m$  to construct an abstract program  $\hat{\Pi}_m$  (Step 1) that over-approximates the original program  $\Pi$  and then compute the abstract answer sets. Over-approximation guarantees that if  $\Pi$  has a answer set  $I$ , then a corresponding abstract answer set  $m(I)$  of the abstract program  $\hat{\Pi}_m$  exists. If in turn the abstract program  $\hat{\Pi}_m$  has no answer set (Step 2), by Proposition 2.1  $\Pi$  is unsatisfiable. In this case, the abstract program  $\hat{\Pi}_m$  and the mapping  $m$  are returned. When we pick an abstract answer set  $I' \in AS(\hat{\Pi}_m)$  (Step 3), we check for concreteness (Step 4). If  $I'$  is concrete, it shows a solution to  $\Pi$ ; in this case, the abstract program  $\hat{\Pi}_m$ , the mapping  $m$  and the concrete abstract answer set  $I$  are returned. If  $I'$  is spurious, we refine the abstraction mapping  $m$  to  $m'$  (Step 5) and loop back to Step 1. This loop continues until either a picked abstract answer set is concrete, or the abstract program has no answer sets. Termination is guaranteed as in the extreme case  $m$  is refined to the trivial identity mapping, i.e., each element of the original domain is mapped to itself;  $\hat{\Pi}_m$  will coincide with  $\Pi$ . Thus, if  $\Pi$  is unsatisfiable, the procedure will stop at Step 2, otherwise at Step 4.

In the next section, we introduce the new abstraction method to be used in Step 1. We investigate an abstraction over non-ground ASP programs given a mapping over their domain (i.e., the Herbrand universe) that singles out the domain elements. The correctness checking of an abstract answer set (Step 4) and then deciding on a refinement (Step 5) is done using a debugging approach which is introduced in Section 4.

### 3 Domain Abstraction

The omission-based abstraction approach in (Saribatur and Eiter, 2018) is propositional in nature and does not account for the fact that in ASP, non-ground rules talk about a domain of discourse, where for the (non)existence of an answer set, the precise set of elements may not matter, but rather how certain elements are related. For example, the graph coloring encoding (4) expresses that each node should be colored differently from its neighbors. The names of the neighbor nodes are not relevant to the color determination, rather the relation of having a neighbor with a certain chosen color.

In this section, we tackle the issue of automatically constructing and evaluating a suitable abstract program  $\Pi'$  for a given non-ground ASP program  $\Pi$  with an abstraction over its domain.

To illustrate the abstraction and its various challenges, we use the following example.

**Example 3.1** (Running example). Consider the following example program  $\Pi$  with domain predicate  $int/1$  for an integer domain  $D = \{1, \dots, 5\}$ .

$$c(X) \leftarrow not\ d(X), X < 5, int(X). \quad (5)$$

$$d(X) \leftarrow not\ c(X), int(X). \quad (6)$$

$$b(X, Y) \leftarrow a(X), d(Y), int(X), int(Y). \quad (7)$$

$$e(X) \leftarrow c(X), a(Y), X \leq Y, int(X), int(Y). \quad (8)$$

$$\perp \leftarrow b(X, Y), e(X), int(X), int(Y). \quad (9)$$

We furthermore have facts  $a(1), a(3), int(1), \dots, int(5)$ . Note that  $\Pi$  has the following answer sets (with facts omitted):

$$I_1 = \{c(2), c(4), d(1), d(3), d(5), e(2)\} \cup S_b,$$

$$I_2 = \{c(2), d(1), d(3), d(4), d(5), e(2), b(1, 4), b(3, 4)\} \cup S_b,$$

$$I_3 = \{c(4), d(1), d(2), d(3), d(5), b(1, 2), b(3, 2)\} \cup S_b,$$

$$I_4 = \{d(1), d(2), d(3), d(4), d(5), b(1, 2), b(1, 4), b(3, 2), b(3, 4)\} \cup S_b,$$

where  $S_b = \{b(1, 1), b(1, 3), b(1, 5), b(3, 5), b(3, 1), b(3, 3)\}$ .

We take a first-order view in which  $\mathcal{A}$  is the Herbrand base of  $\Pi$ , which results from the available predicate symbols and the constant symbols (the domain  $D$  of discourse, i.e., the Herbrand universe), which are by default those occurring in  $\Pi$ . *Domain abstraction* induces domain abstraction mappings in which constants are merged.

**Definition 3.1.** Given a domain  $D$  of  $\Pi$ , a (domain abstraction) mapping is a function  $m: D \rightarrow \hat{D}$  for a set  $\hat{D}$  (the abstracted domain) with  $|\hat{D}| \leq |D|$ .

Thus, a domain abstraction mapping divides  $D$  into *clusters* of elements  $\{d \in D \mid m(d) = \hat{d}\}$ , where  $\hat{d} \in \hat{D}$ , seen as equal; if unambiguous, we also write  $\hat{d}$  for its cluster  $m^{-1}(\hat{d})$ .

**Example 3.2** (ctd). A possible mapping for  $\Pi$  with  $\hat{D}_1 = \{k_1, k_2, k_3\}$  clusters 1, 2, 3 to the element  $k_1$  and 4, 5 to singleton clusters, i.e.,  $m_1 = \{\{1, 2, 3\}/k_1, \{4\}/k_2, \{5\}/k_3\}$ . A naive mapping is  $m_2 = \{\{1, \dots, 5\}/k\}$  with  $\hat{D}_2 = \{k\}$ .

Each domain abstraction mapping  $m$  naturally extends to ground atoms  $a = p(v_1, \dots, v_n)$  by

$$m(a) = p(m(v_1), \dots, m(v_n)).$$

To obtain for a program  $\Pi$  and a Herbrand base  $\mathcal{A}$  an induced mapping  $m : \mathcal{A} \rightarrow \mathcal{A}'$  where  $\mathcal{A}' = m(\mathcal{A}) = \{m(a) \mid a \in \mathcal{A}\}$ , we need an abstract program  $\Pi'$  as in Definition 2.2. However, simply applying  $m$  to  $\Pi$  does not work in general. Moreover, we want domain abstraction for non-ground  $\Pi$  that results in a non-ground program  $\Pi'$ . Building a suitable  $\Pi'$  turns out to be challenging and needs to solve several issues, which we discuss in the next section.

### 3.1 Towards an Abstract Program

Given a mapping  $m$  that describes an abstraction over the domain of a program  $\Pi$ , we start with the intuition of applying  $m$  to each rule, i.e., each atom in a rule is modified according to  $m$ , in order to obtain an abstraction  $\Pi'$  of  $\Pi$ .

One may think that simply lifting the rules in  $\Pi$  to the abstract level and considering the abstract domain for evaluation should be enough to achieve an over-approximation. In some cases, this works.

**Example 3.3** (ctd). Lifting the rule (7) with the abstraction means considering the abstract domain predicate  $\widehat{int}$ :

$$b(X, Y) \leftarrow a(X), d(Y), \widehat{int}(X), \widehat{int}(Y). \quad (10)$$

For the mapping  $m_2$ , we obtain the abstract domain fact  $\widehat{int}(k)$ , and whenever  $a(k)$  and  $d(k)$  holds in an abstract answer set,  $b(k, k)$  would also hold. For any original answer set  $I \in AS(\Pi)$ , the mapped answer set  $m_2(I)$  always contains  $\{a(k), d(k), b(k, k)\}$ . Thus the lifted rule (10) preserves the behavior of the original answer sets. This similarly can be observed for the mapping  $m_1$ .

However, if a rule contains built-in relations or negation, the behavior of the lifted rule on the abstract domain might change on the abstraction of the original answer sets. In particular, it could spoil having an over-approximation. By looking at the cases where simply applying the lifting fails we gradually present our approach addressing this issue.

**Handling built-ins and (in)equalities.** Original rules may rely on certain *built-in relations* involving variables, such as  $<$ ,  $\leq$  in (5) and (8), or  $=$  and  $\neq$ . Simply lifting the rule in the abstraction would mean to apply these relations in the abstract domain. However, due to the domain clustering, some relations might not preserve the behavior of the relation in the original domain, which results in an *uncertainty* that needs to be dealt with.

**Example 3.4.** Lifting rule (8) simply to

$$e(X) \leftarrow c(X), a(Y), X \leq Y, \widehat{int}(X), \widehat{int}(Y).$$

does not work, as  $X \leq Y$  behaves differently over the cluster  $k$ . As  $k \leq k$  under  $m_2$ , the lifted rule derives  $e(k)$  if  $c(k)$  and  $a(k)$  hold. This applies, e.g., to the abstraction of  $I = \{a(1), a(3), c(4), d(0), \dots, d(3)\}$ , where (8) derives no  $e$ -atom as  $4 \not\leq 3$  and  $4 \not\leq 1$ . However,  $I$  is an answer set of  $\Pi$  and must not be lost in the abstraction. Thus, when a cluster causes uncertainties over built-ins, we permit  $e(k)$  to be false even if  $c(k)$  and  $a(k)$  hold by creating the following rule:

$$\{e(X)\} \leftarrow c(X), a(Y), X \leq Y, \widehat{int}(X), \widehat{int}(Y).$$

**Negation.** A naive abstraction approach then could be to turn all rule heads into choices. However, negative literals or certain built-ins (e.g.,  $\neq$ ,  $<$ ) may cause a loss of original answer sets in the abstraction.

**Example 3.5 (ctd).** Consider an alteration of the rule (8) as

$$\{e(X)\} \leftarrow c(X), a(Y), X \neq Y, \widehat{int}(X), \widehat{int}(Y).$$

As  $k = k$ , the abstract body is never satisfied and  $e(k)$  is never derived. However,  $\Pi$  has answer sets containing  $c(2)$ ,  $a(3)$ , and thus also  $e(2)$ , as  $2 \neq 3$ ; they are all lost. Adding a choice rule with a flipped relation,  $X = Y$ , catches such cases.

Similarly, let us change  $a(Y)$  in (8) to *not*  $a(Y)$ . When the rule is lifted to

$$\{e(X)\} \leftarrow c(X), \text{not } a(Y), X \leq Y, \widehat{int}(X), \widehat{int}(Y),$$

$e(k)$  is not derived as  $a(k)$  holds and originally  $a$  holds only for 1 and 3. Thus, original answer sets  $I$  may contain  $e(2)$  or  $e(4)$  but they are lost in the abstraction. Such cases are caught by additional rules with reversed negation for  $a(Y)$ :

$$\{e(X)\} \leftarrow c(X), a(Y), X \leq Y, \widehat{int}(X), \widehat{int}(Y).$$

**Standardizing apart.** The shared use of a variable in a positive rule body must be treated before abstraction is applied.

**Example 3.6 (ctd).** The constraint (9), i.e.,  $\perp \leftarrow b(X, Y), e(X), int(X), int(Y)$ , has a shared use of the variable  $X$ . If this rule is lifted with no change by following the intuition, then  $b(k, k)$  and  $e(k)$  would never occur in the abstract answer sets, while in the original program, the answer set  $I_2$  contains  $e(2)$  and  $b(1, 3)$ . Thus,  $I_2$  would not be mapped to an abstract answer set. If the variables in the rule are standardized apart as

$$\perp \leftarrow b(X, Y), e(X_1), X = X_1, int(X), int(X_1), int(Y), \quad (11)$$

then the focus of the abstraction can be directed towards the relation, i.e.,  $X = X_1$ , in the rule.

Furthermore, the occurrence of constants in the arguments of atom needs to be represented using variables: If a literal  $l(t_1, \dots, t_n)$  in a rule has a constant as one of its arguments, i.e.,  $t_i = c$ , then the rule is modified by having  $l(t_1, \dots, t_{i-1}, X, t_{i+1}, \dots, t_n), X = c$ .

**Example 3.7 (ctd).** If the constraint (9) is of form  $\perp \leftarrow b(2, 3), e(1)$ , then it needs to be changed to the form

$$\perp \leftarrow b(X, X_1), e(X_2), X = 2, X_1 = 3, X_2 = 1, int(X), int(X_1), int(X_2).$$

In summary, we require the input program to have no variables shared among non-domain predicates and no constants in non-domain predicates. This way, the effect of the abstraction over the domain can be treated more easily through the built-in relations in the rules.

**Constraints.** Finally, constraints also have to be treated specially in the abstraction for the cases of uncertainty. Naively lifting the constraints to abstract rules results in losing answer sets for the non-singleton domain clusters.

**Example 3.8 (ctd).** After standardizing apart, if the constraint (9) is lifted with no change with the lifted relation as in (11), this would result in losing the original answer sets  $I_1, I_2$  in which e.g.  $b(1, 3)$  and  $e(2)$  occur together.

In conclusion, only creating choices is not enough to preserve all original answer sets; we need a fine-grained systematic approach to deal with uncertainties.

Table 1: Cases for lifting a binary relation  $rel$

$$\begin{aligned}
\tau_I^{rel}(\hat{d}_1, \hat{d}_2): & \quad rel(\hat{d}_1, \hat{d}_2) \wedge \forall x_1 \in \hat{d}_1, \forall x_2 \in \hat{d}_2. rel(x_1, x_2) \\
\tau_{II}^{rel}(\hat{d}_1, \hat{d}_2): & \quad \neg rel(\hat{d}_1, \hat{d}_2) \wedge \forall x_1 \in \hat{d}_1, \forall x_2 \in \hat{d}_2. \neg rel(x_1, x_2) \\
\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2): & \quad rel(\hat{d}_1, \hat{d}_2) \wedge \exists x_1 \in \hat{d}_1, \exists x_2 \in \hat{d}_2. \neg rel(x_1, x_2) \\
\tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2): & \quad \neg rel(\hat{d}_1, \hat{d}_2) \wedge \exists x_1 \in \hat{d}_1, \exists x_2 \in \hat{d}_2. rel(x_1, x_2)
\end{aligned}$$

### 3.2 Lifted Built-in Relations

As shown above, in lifting the rules for abstraction the built-in relations need special treatment, and so do shared usages of variables in a rule. To unify both issues, we focus on rules of the form

$$r : l \leftarrow B(r), \Gamma_{rel}(r)$$

where the variable occurrences in  $B(r)$  are standardized apart and  $\Gamma_{rel}$  consists of built-in atoms that constrain the variables in  $B(r)$ .

**Example 3.9** (ctd). The rule (7) has  $\Gamma_{rel}(r) = \top$  while the rule (9) must be standardized apart into  $\perp \leftarrow b(X, Y), e(X_1), \Gamma_{rel}$  with  $\Gamma_{rel} = (X = X_1)$ .

The uncertainty that arises during the abstraction is caused by relation restrictions over non-singleton clusters (i.e.,  $|\hat{d}| > 1$ ) or by negative literals mapped to non-singleton abstract literals. In order to address the uncertainty due to relation restrictions in the rules, we consider a notion of *relation types* with respect to the abstraction. For simplicity, we first focus on binary built-ins, e.g.,  $=, <, \leq, \neq$ , and a built-in part  $\Gamma_{rel}(r)$  of the form  $rel(X, c)$  or  $rel(X, Y)$ . Later in Section 3.4, we show how other forms of relations can be addressed.

**Lifted relation types.** When the relation  $rel$  is lifted to the abstract domain, four cases  $\tau_I$ – $\tau_{IV}$  for  $rel(\hat{d}_1, \hat{d}_2)$  can occur in a mapping, as shown in Table 1.

**Example 3.10.** Consider a mapping  $m = \{\{1\}/k_1, \{2, 3\}/k_2, \{4, 5\}/k_3\}$ . For the relation “ $=$ ”,  $k_1 = k_1$  holds and for any  $x_1, x_2 \in k_1 = \{1\}$ ,  $x_1 = x_2$  holds and type I applies. In contrast,  $k_2 = k_2$  holds while  $2, 3 \in k_2$  and  $2 \neq 3$ ; so type III applies. Further,  $k_2 < k_3$  holds and for any  $x \in k_2 = \{2, 3\}$  and  $y \in k_3 = \{4, 5\}$ , we have  $x < y$  and so type I applies. The relation  $k_2 \neq k_2$  does not hold while  $k_2 = \{2, 3\}$  has different elements  $2 \neq 3$  (type IV). Moreover,  $k_1 = k_2$  does not hold in  $\hat{D}$  nor does  $x = y$  for every  $x \in k_1 = \{1\}$  and  $y \in k_2 = \{2, 3\}$  (type II).

If  $rel(\hat{d}_1, \hat{d}_2)$  holds for some  $\hat{d}_1, \hat{d}_2 \in \hat{D}$ , type III is more common in practice in domain abstractions with clusters due to the standardization and as  $=$  often occurs, while type I occurs for singleton mappings (i.e.,  $|\hat{d}_1| = |\hat{d}_2| = 1$ ) or for relations such as  $\neq$  or  $<$  when  $\hat{d}_1 \neq \hat{d}_2$ . If  $rel(\hat{d}_1, \hat{d}_2)$  does not hold for some  $\hat{d}_1, \hat{d}_2 \in \hat{D}$ , type II is common (again due to often having the relation  $=$ ), e.g.,  $=, \leq$ , whereas type IV may occur for  $\neq$  or  $<$ .

**Definition 3.2** ( $\mathcal{T}_m$ ). For an abstraction  $m$ , we let  $\mathcal{T}_m$  be the set of all atoms  $\tau_\iota^{rel}(\hat{d}_1, \hat{d}_2)$  where  $\iota \in \{I, \dots, IV\}$  is the type of the built-in instance  $rel(\hat{d}_1, \hat{d}_2)$  for  $m$

We remark that  $\mathcal{T}_m$  is easily computed.

**Respecting the order relation.** Notice that if the original domain  $D$  contains an order relation among its elements, i.e.,  $rel(x_1, x_2)$  where  $rel \in \{<, \leq\}$ , then in order to be able to talk about the relation  $rel(\hat{d}_1, \hat{d}_2)$  for the abstract elements  $\hat{d}_1, \hat{d}_2$ , in the abstract domain  $\hat{D}$  the relation  $rel$  should be defined. Furthermore, the abstraction mapping should respect the order relation among the elements to avoid unnecessary uncertainty.

**Example 3.11** (ctd). In Example 3.10, the abstract elements  $k_1, k_2, k_3$  were assumed to be ordered  $k_1 < k_2 < k_3$ . If the mapping  $m$  were to arbitrary abstract elements  $a, b, c$ , the relation types could not have been determined as  $i < j$  for  $i, j \in \{a, b, c\}$  is undefined.

Now consider the mapping  $m' = \{\{4\}/k_1, \{1, 5\}/k_2, \{2, 3\}/k_3\}$ , which does not respect the order relation of the elements in  $D$  for the abstraction. The relation types could still be defined for  $m'$ , however for the relation  $<$  the relation types will mostly be of type III and IV, resulting in many uncertainties.

### 3.3 Abstract Program Construction

By our analysis in Section 3.1, the basic idea to construct an abstract program  $\Pi'$  for a program  $\Pi$  with a domain mapping  $m$  is as follows. We either just abstract each atom in a rule, or in case of uncertainty due to domain abstraction, we guess rule heads to catch possible cases, or we treat negated literals by shifting their polarity depending on the abstract domain clusters.

For a ground literal  $l$ , we say that  $l$  is *mapped to a non-singleton cluster* if  $|m^{-1}(m(l))| > 1$ , and it is *mapped to a singleton cluster* otherwise. We use auxiliary facts  $isCluster(\hat{d})$  (resp.  $isSingleton(\hat{d})$ ) for the abstract domain elements  $\hat{d} \in \widehat{D}$  to denote  $|m^{-1}(\hat{d})| > 1$  (and  $|m^{-1}(\hat{d})| = 1$ ). These atoms can also be used to represent whether an abstract literal is a singleton or non-singleton cluster. For the abstract literal  $m(l)$ , if there exists some term  $t \in arg(m(l))$  for which  $isCluster(t)$  holds, then this means that  $m(l)$  is a non-singleton cluster; otherwise, it is a singleton cluster.

**Example 3.12** (Example 3.1 ctd). Consider the domain mapping  $m = \{\{1\}/k_1, \{2, 3\}/k_2, \{4, 5\}/k_3\}$ . For the abstract domain, we have  $isSingleton(k_1), isCluster(k_2), isCluster(k_3)$ . For the literals, the singleton clusters are  $a(k_1), c(k_1), d(k_1), e(k_1)$  and  $b(k_1, k_1)$ , while the remaining literals are non-singletons.

We remark that due to their definition, if either  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  or  $\tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$  holds true for some  $\hat{d}_1, \hat{d}_2 \in \widehat{D}$ , this means that either  $\hat{d}_1$  or  $\hat{d}_2$  is mapped to a cluster, i.e.,  $isCluster(\hat{d}_i)$  for some  $i \in \{1, 2\}$ .

#### 3.3.1 Restricted Case

For ease of presentation, we first consider programs  $\Pi$  with rules having

- (I) at most one negative body literal,
- (II) a single, binary built-in literal, and
- (III) no cyclic dependencies between non-ground literals.

**Example 3.13** (ctd). The program  $\Pi$  only adheres to the restrictions (i)-(ii) and not to (iii) because  $d(X)$  and  $c(X)$  negatively depend on each another.

We next define how to construct an abstract rule  $r^m$ .

**Definition 3.3** (rule abstraction). *Given a rule  $r : l \leftarrow B(r), rel(t_1, t_2)$  as above and a domain mapping  $m$ , the set  $r^m$  contains the following rules:*

- (a)  $m(l) \leftarrow m(B(r)), rel(\hat{t}_1, \hat{t}_2), \tau_I^{rel}(\hat{t}_1, \hat{t}_2).$
- (b)  $\{m(l)\} \leftarrow m(B(r)), rel(\hat{t}_1, \hat{t}_2), \tau_{III}^{rel}(\hat{t}_1, \hat{t}_2).$
- (c)  $\{m(l)\} \leftarrow m(B(r)), \overline{rel}(\hat{t}_1, \hat{t}_2), \tau_{IV}^{rel}(\hat{t}_1, \hat{t}_2).$
- (d) For  $l_i \in B^-(r)$ :

$$(i) \bigcup_{j \in arg(l_i)} \left\{ \{m(l)\} \leftarrow m(B_{l_i}^{sh}(r)), rel(\hat{t}_1, \hat{t}_2), isCluster(\hat{j}). \right\}$$

$$(ii) \bigcup_{j \in arg(l_i)} \left\{ \{m(l)\} \leftarrow m(B_{l_i}^{sh}(r)), \overline{rel}(\hat{t}_1, \hat{t}_2), \tau_{IV}^{rel}(\hat{t}_1, \hat{t}_2), isCluster(\hat{j}). \right\}$$

where

- $B_{l_i}^{sh}(r) = B^+(r) \cup \{l_i\}$ , not  $B^-(r) \setminus \{l_i\}$ ,
- $\overline{rel}$  denotes the complement of  $rel$ , and
- for  $k \in \{1, 2\}$ , if  $t_k$  is a constant then  $\hat{t}_k = m(t_k)$ , else  $\hat{t}_k = t_k$ , i.e., variables are not mapped; similarly, if  $j \in arg(l_i)$  is a constant then  $\hat{j} = m(\hat{j})$ , else  $\hat{j} = j$ .

In step (a), the case of having no uncertainty due to abstraction is applied. Steps (b) and (c) are for the cases of uncertainty. The head becomes a choice, and for case IV, we flip the relation,  $\overline{rel}$ , to catch the case of the relation holding true (which is causing the uncertainty). No rules are added for case II, since the body of the rule will never be satisfied due to the relation not holding true in the abstract domain (similar as in the original domain). As for constraints (e.g., (9)), we note that  $m(\perp) = \perp$ . Consequently, in (a) the head is unchanged; as an optimization, all other steps (b)-(d) can be omitted, since the choice  $\{\perp\}$  is ineffective (we always can choose that  $\perp$  is false).

**Example 3.14** (ctd). Consider the domain mapping  $m = \{\{1\}/k_1, \{2, 3\}/k_2, \{4, 5\}/k_3\}$ . We have  $\tau_I^{\leq}(x, y)$  true for  $(x, y) \in \{(k_1, k_1), (k_1, k_2), (k_1, k_3), (k_2, k_3)\}$ , and  $\tau_{III}^{\leq}(x, y)$  true for  $(x, y) \in \{(k_2, k_2), (k_3, k_3)\}$ , and only type II for all other tuples  $(x, y)$ . The abstract rules for (8) are:

$$e(X) \leftarrow c(X), a(Y), X \leq Y, \tau_I^{\leq}(X, Y), \widehat{int}(X), \widehat{int}(Y).$$

$$\{e(X)\} \leftarrow c(X), a(Y), X \leq Y, \tau_{III}^{\leq}(X, Y), \widehat{int}(X), \widehat{int}(Y).$$

In step (d) of Definition 3.3, we grasp the uncertainty arising from negation by adding rules that shift the negative literal only if it shares arguments that are mapped to a non-singleton cluster.

**Example 3.15** (ctd). Rule (5) has a negative literal, *not*  $d(X)$ , and the relation  $X < 5$  with shared argument  $X$ . When it is lifted to  $X < k_3$ , it has  $\tau_{II}^<(a, b)$  true for  $(a, b) \in \{(k_3, k_1), (k_3, k_2)\}$ ,  $\tau_{IV}^<(k_3, k_3)$ , and type I for all other tuples  $(a, b)$ .

By case (a), it is abstracted without change for  $\tau_I$  abstract values, and by case (c) the relation is flipped for  $\tau_{IV}$ . Furthermore, a shift on the polarity of the negative literal is made:

$$c(X) \leftarrow not\ d(X), X < k_3, \tau_I^<(X, k_3), \widehat{int}(X).$$

$$\{c(X)\} \leftarrow not\ d(X), X \geq k_3, \tau_{IV}^<(X, k_3), \widehat{int}(X).$$

$$\{c(X)\} \leftarrow d(X), X < k_3, isCluster(X), \widehat{int}(X).$$

$$\{c(X)\} \leftarrow d(X), X \geq k_3, \tau_{IV}^<(X, k_3), isCluster(X), \widehat{int}(X).$$

Notice that the case of having  $rel = \top$  is covered by step (a). For this case, the lifted relation will only have type I, thus the rules from the remaining steps need not be added (see Example 3.3).

Semantically, the rules added in steps (a)-(b) are to ensure that  $m(I)$  is a model of  $\Pi^m$ , as either the original rule is kept or it is changed to a choice rule. Steps (c)-(d) serve to catch the cases that may violate the minimality of the model due to a negative literal or a relation over non-singleton clusters. The abstract program is now as follows.

**Definition 3.4** (Abstract program  $\Pi^m$ ). *Given a program  $\Pi$  as above and a domain abstraction  $m$ , the abstract program for  $m$  consists of the rules*

$$\Pi^m = \{r^m \mid r : l \leftarrow B(r), \text{rel}(t_1, t_2) \in \Pi\} \cup \{x \mid x \in \mathcal{T}_m\} \cup \{m(p(\vec{c})) \mid p(\vec{c}) \in \Pi\}.$$

Notably, the construction of  $\Pi^m$  is modular, rule by rule. The following result states that this abstraction works.

**Theorem 3.1** (restricted program abstraction). *Let  $m$  be a domain mapping of a program  $\Pi$  under the above assumptions (i)–(iii). Then for every  $I \in AS(\Pi)$ ,  $m(I) \cup \mathcal{T}_m \in AS(\Pi^m)$ .*

### 3.3.2 Abstract Program (General Case)

We now describe how to remove the restrictions (i)–(iii) on programs from above.

**(G-I) Multiple negative literals.** If rule  $r$  has multiple negative literals, i.e.,  $|B^-(r)| > 1$ , we shift each negative literal that shares arguments that are mapped to a non-singleton cluster. Thus, instead of shifting one literal  $l_i \in B^-(r)$ , we consider shifting multiple literals  $L \subseteq B^-(r)$  at a time and all combinations of (non-)shifting the literals in  $B^-(r)$ .

**Definition 3.5** (Treating multiple negative literals). *Step (d) of Definition 3.3 is modified as*

(d) For  $\emptyset \subset L \subseteq B^-(r)$ :

- (i)  $\bigcup_{j \in \text{arg}(l_i), l_i \in L} \{ \{m(l)\} \leftarrow m(B_L^{sh}(r)), \text{rel}(\hat{t}_1, \hat{t}_2), \text{isCluster}(\hat{j}) \}$
- (ii)  $\bigcup_{j \in \text{arg}(l_i), l_i \in L} \{ \{m(l)\} \leftarrow m(B_L^{sh}(r)), \overline{\text{rel}}(\hat{t}_1, \hat{t}_2), \tau_{\text{V}}^{\text{rel}}(\hat{t}_1, \hat{t}_2), \text{isCluster}(\hat{j}) \}$

where  $B_L^{sh}(r) = B^+(r) \cup L$ , not  $B^-(r) \setminus L$ .

This definition allows us to discard the restriction that all negative literals must share a variable with the relation atom, by shifting them to ensure that the case of having a non-singleton mapping is considered. Step (d-i) coincides with steps (d-i) and (d-iii) of Definition 3.3 and step (d-ii) coincides with step (d-ii) of Definition 3.3.

**Example 3.16.** Consider the rule

$$d(X) \leftarrow \text{not } c(X), \text{not } a(X), \text{int}(X).$$

The constructed non-ground abstract rules following step (d-i) of Definition 3.5 will be

- $\{d(X)\} \leftarrow c(X), \text{not } a(X), \text{isCluster}(X), \widehat{\text{int}}(X), \widehat{\text{int}}(X_1).$
- $\{d(X)\} \leftarrow \text{not } c(X), a(X), \text{isCluster}(X), \widehat{\text{int}}(X), \widehat{\text{int}}(X_1).$
- $\{d(X)\} \leftarrow c(X), a(X), \text{isCluster}(X), \widehat{\text{int}}(X), \widehat{\text{int}}(X_1).$

Step (d-ii) is similarly applied.

Table 2: Cases for lifting an  $n$ -ary relation  $rel'$

$$\begin{aligned}
\tau_I^{rel'}(\hat{d}_1, \dots, \hat{d}_n): & \quad rel'(\hat{d}_1, \dots, \hat{d}_n) \wedge \forall x_1 \in \hat{d}_1, \dots, \forall x_n \in \hat{d}_n. rel'(x_1, \dots, x_n) \\
\tau_{II}^{rel'}(\hat{d}_1, \dots, \hat{d}_n): & \quad \neg rel'(\hat{d}_1, \dots, \hat{d}_n) \wedge \forall x_1 \in \hat{d}_1, \dots, \forall x_n \in \hat{d}_n. \neg rel'(x_1, \dots, x_n) \\
\tau_{III}^{rel'}(\hat{d}_1, \dots, \hat{d}_n): & \quad rel'(\hat{d}_1, \dots, \hat{d}_n) \wedge \exists x_1 \in \hat{d}_1, \dots, \exists x_n \in \hat{d}_n. \neg rel'(x_1, \dots, x_n) \\
\tau_{IV}^{rel'}(\hat{d}_1, \dots, \hat{d}_n): & \quad \neg rel'(\hat{d}_1, \dots, \hat{d}_n) \wedge \exists x_1 \in \hat{d}_1, \dots, \exists x_n \in \hat{d}_n. rel'(x_1, \dots, x_n)
\end{aligned}$$


---

**(G-II) Multiple relation literals.** A simple approach to handle a built-in part with multiple literals, i.e.,

$$\Gamma_{rel} = rel_1(t_{1,1}, t_{2,1}), \dots, rel_k(t_{1,k}, t_{2,k}), k > 1,$$

is to view it as a literal of an  $2k$ -ary built-in relation  $rel'(X_{1,1}, X_{2,1}, \dots, X_{1,k}, X_{2,k})$ . The abstract version of such  $rel'$  and the cases I-IV are then lifted from  $x_1, x_2$  to  $x_1, \dots, x_n$  as in Table 2.

**Example 3.17.** For  $\Gamma_{rel} = (X_1 = X_2, X_3 = X_4)$ , we use a new relation  $rel'(X_1, X_2, X_3, X_4)$ . If it holds for abstract values  $\hat{d}_1, \dots, \hat{d}_4$  such that  $\hat{d}_1 = \hat{d}_2 \wedge \hat{d}_3 = \hat{d}_4$  holds, then we have type  $\tau_I$  if all  $\hat{d}_i$  are singleton clusters and type  $\tau_{III}$  if some  $\hat{d}_i$  is non-singleton; otherwise (i.e.,  $\overline{rel}'(\hat{d}_1, \hat{d}_2, \hat{d}_3, \hat{d}_4)$  holds) type  $\tau_{II}$  applies.

**(G-III) Cyclic dependencies.** Rules which are involved in a negative cyclic dependency need special consideration.

**Example 3.18 (ctd).** For the rules (5)-(6) in Example 3.1 and the mapping  $\{\{1, \dots, 5\} / k\}$ , the abstract rules are

$$\{c(X)\} \leftarrow not\ d(X), X \geq k, \tau_{IV}^<(X, k), \widehat{int}(X). \quad (12)$$

$$\{c(X)\} \leftarrow d(X), X < k, isCluster(X), \widehat{int}(X). \quad (13)$$

$$\{c(X)\} \leftarrow d(X), X \geq k, \tau_{IV}^<(X, k), isCluster(X), \widehat{int}(X). \quad (14)$$

$$\{d(X)\} \leftarrow c(X), isCluster(X), \widehat{int}(X). \quad (15)$$

in addition to the abstracted rules due to step (a). Consider the answer set  $I = \{c(0), d(1), c(2), d(3), c(4), d(5)\}$  of  $\Pi$ . We have  $\hat{I} = m(I) = \{c(k), d(k)\}$ . Although  $\hat{I}$  is a model of  $(\Pi^m)^I$ , either  $c(k)$  or  $d(k)$  is unfounded; hence  $\hat{I}$  is not minimal, i.e., not an answer set of  $\Pi^m$ . The reason is that the negative cyclic dependency (i.e., “choice”) of  $c$ - and  $d$ -atoms does not occur for  $c(k)$  and  $d(k)$  in the constructed  $\Pi^m$ .

To resolve this, the literals of  $\Pi$  that are involved in a negative loop are treated specially.

**Definition 3.6** (Treating cyclic dependency). *Given a set  $L_c$  of literals involved in a negative cyclic dependency, Definition 3.3 is modified by redefining  $B_{l_i}^{sh}(r)$  as*

$$B_{l_i, L_c}^{sh}(r) = \begin{cases} B^+(r) \cup \{l_i\}, not\ B^-(r) \setminus \{l_i\} & \text{if } l_i \notin L_c, \\ B^+(r), not\ B^-(r) \setminus \{l_i\} & \text{if } l_i \in L_c \end{cases}$$

In step (d) of Definition 3.3, the newly defined  $B_{l_i, L_c}^{sh}(r)$  eliminates the literals  $l_i$  that are involved in a loop from the body instead of shifting their polarity.

**Example 3.19** (ctd). For the program  $\Pi$  in (5)-(9) with the mapping  $m = \{\{1, \dots, 5\}/k\}$ , the program  $\Pi^m$  is

$$c(X) \leftarrow \text{not } d(X), X < k, \tau_1^<(X, k), \widehat{\text{int}}(X). \quad (16)$$

$$\{c(X)\} \leftarrow \text{not } d(X), X \geq k, \tau_{\text{IV}}^<(X, k), \widehat{\text{int}}(X). \quad (17)$$

$$\{c(X)\} \leftarrow X \geq k, \tau_{\text{IV}}^<(X, k), \text{isCluster}(X), \widehat{\text{int}}(X). \quad (18)$$

$$\{c(X)\} \leftarrow X < k, \text{isCluster}(X), \widehat{\text{int}}(X). \quad (19)$$

$$d(X) \leftarrow \text{not } c(X), \widehat{\text{int}}(X). \quad (20)$$

$$\{d(X)\} \leftarrow \text{isCluster}(X), \widehat{\text{int}}(X). \quad (21)$$

$$b(X, Y) \leftarrow a(X), d(Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y). \quad (22)$$

$$e(X) \leftarrow c(X), a(Y), X \leq Y, \tau_1^<(X, Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y). \quad (23)$$

$$\{e(X)\} \leftarrow c(X), a(Y), X \leq Y, \tau_{\text{III}}^<(X, Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y). \quad (24)$$

$$\perp \leftarrow b(X, Y), e(X_1), X = X_1, \tau_1^-(X, X_1), \widehat{\text{int}}(X), \widehat{\text{int}}(X_1), \widehat{\text{int}}(Y). \quad (25)$$

The abstract facts for  $a(1)$ ,  $a(3)$ ,  $\text{int}(1)$ ,  $\dots$ ,  $\text{int}(5)$ , are  $a(k)$ ,  $\widehat{\text{int}}(k)$ , and the type facts are  $\mathcal{T}_m = \{\tau_{\text{III}}^<(k, k), \tau_{\text{III}}^-(k, k), \tau_{\text{IV}}^<(k, k)\}$ . Notice that when the rules are grounded to the relation type facts  $\mathcal{T}_m$ , only the rules (17)-(22) and (24) remain to be used for the answer set computation.

Cyclic dependency with multiple negative literals can then be treated by modifying the shifting procedure  $B_L^{\text{sh}}(r)$  in Definition 3.5 with  $B_{l_i, L_c}^{\text{sh}}(r)$  of Definition 3.6 as

$$B_{L, L_c}^{\text{sh}}(r) = B^+(r) \cup (L \setminus L_c), \text{not } B^-(r) \setminus L \quad (26)$$

where the negative literals in  $L$  get their polarity shifted if they do not occur in  $L_c$ , otherwise they are omitted.

Multiple cycles  $L_{c_1}, \dots, L_{c_n}$  can be handled by defining the set  $L_c$  in Definition 3.6 as the union of all the cycles, i.e.,  $L_c = \bigcup L_{c_i}$ . This way  $B_{l_i, L_c}^{\text{sh}}(r)$  will eliminate all the literals involved in some loop.

Let  $\Pi^m$  denote the program obtained from a general program  $\Pi$  with the generalized abstraction procedure. Then:

**Theorem 3.2** (general program abstraction). *Let  $m$  be a domain mapping of a program  $\Pi$ . Then for every  $I \in \text{AS}(\Pi)$ , the abstract interpretation  $\widehat{I} = m(I) \cup \mathcal{T}_m$  is an answer set of  $\Pi^m$ .*

**Example 3.20** (ctd). The constructed abstract program  $\Pi^m$  has the answer sets (with abstract facts omitted)

$$\hat{I}_1 = \{d(k), b(k, k)\}, \hat{I}_2 = \{c(k)\}, \hat{I}_3 = \{c(k), d(k), b(k, k)\}, \hat{I}_4 = \{c(k), e(k)\}, \hat{I}_5 = \{c(k), d(k), e(k), b(k, k)\},$$

where  $m(I_1) = m(I_2) = \hat{I}_5$ ,  $m(I_3) = \hat{I}_3$ ,  $m(I_4) = \hat{I}_1$ .

The abstraction yields in general an over-approximation of the answer sets of a program. The notion of spurious and concrete answer sets amounts to the following.

**Definition 3.7** (cf. Definition 2.3). *An abstract answer set  $\widehat{I} \in \text{AS}(\Pi^m)$  is concrete, if there exists an answer set  $I \in \text{AS}(\Pi)$  such that  $\widehat{I} = m(I) \cup \mathcal{T}_m$ , else it is spurious.*

A spurious abstract answer set has no corresponding concrete answer set.

**Example 3.21** (ctd). The abstract answer sets  $\hat{I}_2 = \{c(k)\}$  and  $\hat{I}_4 = \{c(k), e(k)\}$  are spurious.

## 3.4 Syntactic Extensions and Further Considerations

### 3.4.1 Treating Choice Rules

Choice rules are treated by ensuring that the body is abstracted and the choice over the abstracted head is kept.

**Definition 3.8.** *Given a choice rule  $r : \{l\} \leftarrow B(r), rel(t_1, t_2)$  and a domain mapping  $m$ , the set  $r^m$  contains the rules of Definition 3.3 for steps (b)-(d), and for step (a), it contains*

$$\{m(l)\} \leftarrow m(B(r)), rel(\hat{t}_1, \hat{t}_2), \tau_1^{rel}(\hat{t}_1, \hat{t}_2).$$

More sophisticated choice rules that involve cardinality constraints, e.g.,  $n_1 \leq \{l\} \leq n_2 \leftarrow B(r)$ , can not immediately be treated similarly. Lifting the cardinality constraints analogously to the abstract rule causes to force the occurrence of abstract atoms for ensuring the lower bound.

**Example 3.22 (ctd).** Consider instead of (7) the rule

$$2 \leq \{b(X, Y) : d(Y)\} \leq 4 \leftarrow a(X), dom(X).$$

which gets lifted to the same abstract rule. However, for the mapping  $m = \{\{1, 2, 3, 4, 5\}/k\}$ , if  $a(k)$  and  $d(k)$  holds true, this would cause to have  $b(k, k)$  hold true and no other atoms with the same predicate. Thus, the lower bound can not be satisfied, causing the abstract program to become unsatisfiable.

The issue arises from the fact that if the atom in the choice head is involved in some non-singleton cluster, then multiple original atom may be mapped to it, thus still satisfying the lower bound constraint in the original program. Such choice rules can be treated by modifying the lower bounds in the abstract program and adding a constraint to ensure that the original lower bound is met if the atom is only involved with singleton clusters.

**Definition 3.9.** *Given a rule  $r : n_1 \leq \{l\} \leq n_2 \leftarrow B(r), rel(t_1, t_2)$ , in the abstraction procedure the choice head is changed to  $\{m(l)\} \leq n_2$ , and an additional constraint of the following form is added.*

$$\perp \leftarrow \{m(l) : isSingleton(\hat{t}_1), \dots, isSingleton(\hat{t}_n)\} < n_1, \quad (27)$$

$$\{m(l) : isCluster(\hat{t}_1); \dots; m(l) : isCluster(\hat{t}_n)\} < 1. \quad (28)$$

where  $arg(l) = \{t_1, \dots, t_n\}$  and  $l = p(t_1, \dots, t_n)$  for some predicate  $m$ .

The idea with the additional constraint is to ensure that if the lower bound  $n_1$  is not satisfied through literals mapped to singleton clusters (27), then some literal with a non-singleton cluster (28) should also occur.

**Example 3.23 (ctd).** Instead of lifting the choice rule as in Example 3.22, we add the following abstract rules:

$$\{b(X, Y) : d(Y)\} \leq 4 \leftarrow a(X), dom(X).$$

$$\perp \leftarrow \{b(X, Y) : isSingleton(X), isSingleton(Y)\} < 2,$$

$$\{b(X, Y) : isCluster(X); b(X, Y) : isCluster(Y)\} < 1.$$

This way there is no lower bound on the number of occurrences of  $b(X, Y)$  that causes unsatisfiability at the abstract program. Furthermore, for the mapping  $m = \{\{1\}/k_1, \{2, 3, 4, 5\}/k_2\}$ , for an answer set containing  $b(k_1, k_1)$  the constraint ensures it also contains some  $b(\hat{d}_1, \hat{d}_2)$ , where  $\hat{d}_1$  or  $\hat{d}_2$  is a cluster, so that the original lower bound is met.

### 3.4.2 Other Forms of Relations

A program can also contain relations that are not binary, such as addition or multiplication. They can be treated as follows: (1) rewrite the relations by adding instead auxiliary atoms to represent them, (2) standardize apart the auxiliary atom arguments similarly as the remaining atoms, and (3) add to the original program facts of the auxiliary atom to show for which domain elements the relation holds true. In the abstraction procedure, the facts added will be lifted to the abstract domain, and the abstraction is handled over relations for the arguments which were standardized apart.

**Example 3.24.** Consider the rule

$$b(X, Y) \leftarrow a(X), d(Y), X + 1 = Y, \text{int}(X), \text{int}(Y).$$

For the addition relation, an auxiliary atom  $\text{plusOne}(X, Y)$  is introduced by adding the set  $\{\text{plusOne}(1, 2), \text{plusOne}(2, 3), \text{plusOne}(3, 4), \text{plusOne}(4, 5)\}$  of facts to  $\Pi$  to show on which domain elements this relation holds.

The respective rule gets standardized apart into the following form.

$$b(X, Y) \leftarrow a(X), d(Y), \text{plusOne}(X_1, Y_1), X = X_1, Y = Y_1, \text{int}(X), \text{int}(Y), \text{int}(X_1), \text{int}(Y_1).$$

For being able to lift the relations to the abstract domain, the built-in relations must also be defined over the abstract domain. If e.g. an ordered domain  $(\{1, 2, 3, 4, 5\}, <)$  is mapped to a domain  $\{a, b, c\}$  where  $<$  is undefined, then lifting  $<$  to the abstract domain will not be feasible; here the above approach of introducing auxiliary atoms must be taken.

### 3.4.3 Concreteness with Projection

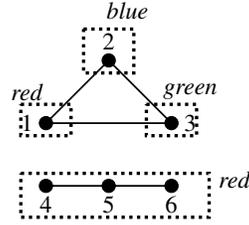
Usually the problem encodings contain auxiliary atoms that are insignificant for solutions. When constructing the abstract program, such auxiliary atoms are treated the same, by introducing choices whenever there is an uncertainty. However, this causes many spurious guesses over the auxiliary atoms, and making sure that the abstract answer set is concrete w.r.t. all of these atoms becomes too ambitious, as encountering a concrete abstract answer set among many spurious ones is more difficult. For this reason, we consider a projected notion of determining concreteness of an abstract answer set by only focusing on a certain set of atoms.

**Definition 3.10.** For a set  $A$  of atoms, an abstract answer set  $\hat{I} \in AS(\Pi^m)$  is concrete w.r.t.  $A$ , if  $\hat{I}|_{\hat{A}} = m(I|_A) \cup \mathcal{F}_m$  for an answer set  $I \in AS(\Pi)$ , where  $\hat{A} = m(A)$ .

**Example 3.25.** Consider a modified instance of graph coloring where the isolated nodes are connected as shown in Figure 5. For the abstraction, the abstract coloring is spurious as the nodes in the cluster  $\{4, 5, 6\}$  cannot all be colored to *red* in the original graph due to the edges. However, the abstract coloring is concrete w.r.t. the nodes  $\{1, 2, 3\}$ .

Such a notion of concreteness becomes useful when abstraction is applied to analyze problems, as one can focus on the atoms deemed to be important. For this, the user should have an idea of the atoms that matter for determining a valid solution. E.g. for planning problems, this notion can help in focusing on the actions and directly affected objects, which serve to describe a solution. One then obtains abstract answer sets that have concrete truth assignments of these atoms, while the auxiliary atoms and their concrete truth assignments become irrelevant.

Figure 5: Concreteness w.r.t. projection over nodes  $\{1, 2, 3\}$



### 3.5 Properties of Domain Abstraction

We now consider some basic semantic properties of our formulation of program abstraction. (Non-)existing spurious answer sets allow us to infer properties of the original program.

**Proposition 3.3.** *For any program  $\Pi$ ,*

- (i)  $AS(\Pi^{m_{id}}) = \{I \cup \mathcal{T}_{m_{id}} \mid I \in AS(\Pi)\}$  for the identity mapping  $m_{id} = \{\{x\}/x \mid x \in D\}$ .
- (ii)  $AS(\Pi^m) = \emptyset$  implies that  $AS(\Pi) = \emptyset$ .
- (iii)  $AS(\Pi) = \emptyset$  iff some  $\Pi^m$  has only spurious answer sets.

*Proof.* (i) Having the identity mapping  $id$  causes to only have singleton clusters in the abstract domain, thus resulting in only  $\tau_I$  and  $\tau_{\Pi}$  type facts in  $\mathcal{T}_{m_{id}}$ . This causes for only the rules of step (a) in Definitions 3.3 and 3.8 to remain when the rules are grounded to the relation types. Hence, the same answer sets are obtained.

(ii) Corollary of Theorem 3.2.

(iii) If  $AS(\Pi) = \emptyset$ , then no  $\hat{I} \in AS(\Pi^m)$  for any  $m$  has a concrete answer set in  $\Pi$ ; thus, all abstract answer sets of  $\Pi^m$  are spurious. Now assume the latter holds but  $AS(\Pi) \neq \emptyset$ . Then  $\Pi$  has some answer set  $I$ , and by Theorem 3.2  $m(I) \cup \mathcal{T}_m \in AS(\Pi^m)$ , which would contradict that  $\Pi^m$  has only spurious answer sets.  $\square$

The abstract program is built by a syntactic transformation. The abstraction over the domain can also be done incrementally which in the end amounts to the overall abstraction.

**Lemma 3.4.** *For any program  $\Pi$  and mappings  $m, m_1, m_2$  such that  $m_2(m_1(D)) = m(D)$ , we have  $grd_{\mathcal{T}_{m_2, m_1}}((\Pi^{m_1})^{m_2}) = grd_{\mathcal{T}_m}(\Pi^m)$ , where  $grd_{\mathcal{T}}$  denotes the grounding of the program to the relation type facts  $\mathcal{T}$ .*

For proving Lemma 3.4, we use the following result.

**Lemma 3.5.** *For a relation  $rel(d_1, d_2)$  and mappings  $m, m_1, m_2$  such that  $m_2(m_1(D)) = m(D)$ , we have  $\mathcal{T}_{m_2, m_1}^{rel} = \mathcal{T}_m^{rel}$ .*

*Proof.* The relation type computation  $\mathcal{T}_{m_1}^{rel}$  is done for  $rel(m_1(d_1), m_1(d_2))$  and then the relation type computation  $\mathcal{T}_{m_2, m_1}^{rel}$  for  $rel(m_2(m_1(d_1)), m_2(m_1(d_2))) = rel(m(d_1), m(d_2))$ , resulting in the same relation type facts of  $\mathcal{T}_m^{rel}$ .  $\square$

Figure 6: Abstract programs of Example 3.26 with  $m_1 = \{\{1, 2\}/k_0, \{3, 4\}/k_1, \{5\}/k_2\}$  and  $m_2 = \{\{k_0, k_1\}/a_0, \{k_2\}/a_1\}$

$$\begin{array}{ll}
c(k_0) \leftarrow k_0 < k_2, \text{not } d(k_0). & \\
c(k_1) \leftarrow k_1 < k_2, \text{not } d(k_1). & \\
\{c(X)\} \leftarrow X < k_2, \text{dom}(X), \text{isCluster}(X). & \\
d(X) \leftarrow \text{dom}(X), \text{not } c(X). & \\
\{d(X)\} \leftarrow \text{dom}(X), \text{isCluster}(X). & c(a_0) \leftarrow a_0 < a_1, \text{not } d(a_0). \\
b(X, Y) \leftarrow a(X), d(Y), \text{dom}(X), \text{dom}(Y). & \{c(X)\} \leftarrow X < a_1, \text{dom}(X), \text{isCluster}(X). \\
\{e(k_0)\} \leftarrow c(k_0), a(k_0), k_0 \leq k_0. & d(X) \leftarrow \text{dom}(X), \text{not } c(X). \\
\{e(k_1)\} \leftarrow c(k_1), a(k_1), k_1 \leq k_1. & \{d(X)\} \leftarrow \text{dom}(X), \text{isCluster}(X). \\
e(k_0) \leftarrow c(k_0), a(k_1), k_0 \leq k_1. & b(X, Y) \leftarrow a(X), d(Y), \text{dom}(X), \text{dom}(Y). \\
e(k_0) \leftarrow c(k_0), a(k_2), k_0 \leq k_2. & \{e(a_0)\} \leftarrow c(a_0), a(a_0), a_0 \leq a_0. \\
e(k_1) \leftarrow c(k_1), a(k_2), k_1 \leq k_2. & e(a_0) \leftarrow c(a_0), a(a_1), a_0 \leq a_1. \\
e(k_2) \leftarrow c(k_2), a(k_2), k_2 \leq k_2. & e(a_1) \leftarrow c(a_1), a(a_1), a_1 \leq a_1. \\
\perp \leftarrow b(k_2, k_2), e(k_2), k_2 = k_2. & \perp \leftarrow b(a_1, a_1), e(a_1), a_1 = a_1.
\end{array}$$

(a)  $\text{grd}_{\mathcal{F}_{m_1}}(\Pi^{m_1})$  (b)  $\text{grd}_{\mathcal{F}_{m_2, m_1}}((\Pi^{m_1})^{m_2})$

*Proof of Lemma 3.4.* From the rules of  $\Pi^{m_1}$ , the rules for  $(\Pi^{m_1})^{m_2}$  will be constructed according to Definitions 3.3 and 3.8. Consider a rule  $r$  with body  $B(r), \text{rel}(t_1, t_2)$  in  $\Pi$ . The set  $r^{m_1} \in \Pi^{m_1}$  contains rules with body  $m_1(B(r)), \text{rel}(\hat{t}_1, \hat{t}_2), \tau_i^{\text{rel}}(\hat{t}_1, \hat{t}_2)$  where  $\hat{t}_k = m_1(t_k)$  if  $t_k$  is a constant;  $\hat{t}_k = t_k$  otherwise.

For the set  $r^{m_1}$  of rules, a new set  $(r^{m_1})^{m_2}$  will be constructed. Let  $r' \in r^{m_1}$ , its body will be abstracted to

$$m_2(B(r')), \text{rel}(\hat{\hat{t}}_1, \hat{\hat{t}}_2), \tau_j^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2) \quad (29)$$

where  $m_2(B(r')) = m_2(m_1(B(r))), m_2(\tau_i^{\text{rel}}(\hat{t}_1, \hat{t}_2))$  and  $\hat{\hat{t}}_j = m_2(m_1(t_k))$  if  $t_k$  is a constant;  $\hat{\hat{t}}_k = t_k$  otherwise. Since  $m_2(\tau_i^{\text{rel}}(\hat{t}_1, \hat{t}_2)) = \tau_i^{\text{rel}}(m_2(\hat{t}_1), m_2(\hat{t}_2)) = \tau_i^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2)$ , (29) will take the form

$$m(B(r)), \tau_i^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2), \text{rel}(\hat{\hat{t}}_1, \hat{\hat{t}}_2), \tau_j^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2).$$

where  $\hat{\hat{t}}_k = m(t_k)$  if  $t_k$  is a constant;  $\hat{\hat{t}}_k = t_k$  otherwise.

The rules in  $(r^{m_1})^{m_2}$  where types of the relation differ, i.e.,  $i \neq j$  for  $\tau_i^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2), \tau_j^{\text{rel}}(\hat{\hat{t}}_1, \hat{\hat{t}}_2)$ , are insignificant as the atoms can not both hold true in  $\mathcal{F}_{m_2, m_1}$ , i.e., they do not appear in  $\text{grd}_{\mathcal{F}_{m_2, m_1}}((r^{m_1})^{m_2})$ . As for the remaining rules in  $(r^{m_1})^{m_2}$ , they correspond to the rules in  $r^m$ . Thus, by Lemma 3.5 and  $\{m_2(m_1(p(\vec{c}))) \mid p(\vec{c}) \in \Pi\} = \{m(p(\vec{c})) \mid p(\vec{c}) \in \Pi\}$ , we obtain  $\text{grd}_{\mathcal{F}_{m_2, m_1}}(\Pi^{m_1})^{m_2} = \text{grd}_{\mathcal{F}_m}(\Pi^m)$ .  $\square$

**Example 3.26** (Example 3.1 ctd). Applying first the mapping  $m_1 = \{\{1, 2\}/k_0, \{3, 4\}/k_1, \{5\}/k_2\}$  and then the mapping  $m_2 = \{\{k_0, k_1\}/a_0, \{k_2\}/a_1\}$  yields the mapping  $m = \{\{1, 2, 3, 4\}/a_0, \{5\}/a_1\}$ . Figure 6 shows the constructed abstract programs. Notice that the program in Figure 6b is the same as the non-ground program in Example 3.19 updated for the mapping  $m$ , i.e.,  $k$  is replaced with  $a_1$ , when it is grounded to  $\mathcal{F}_{m_2, m_1}$ .

An easy induction argument shows then the possibility of doing abstraction sequentially, by having abstract mappings defined over previously abstracted domains.

**Proposition 3.6.** *For any program  $\Pi$  and mapping  $m, m_1, \dots, m_n$  such that  $m_n(\dots(m_1(D))) = m(D)$ , we have  $\text{grd}_{\mathcal{F}_m}(\Pi^m) = \text{grd}_{\mathcal{F}_{m_n, \dots, m_1}}(((\Pi^{m_1}) \dots)^{m_n})$ .*

In Section 5.2 below, we demonstrate further uses of having a hierarchy of abstractions.

We remark that general properties of spurious answer set from over-approximation apply to domain abstraction as an instance of it. Examples of such properties, mentioned for omission abstraction in (Saribatur and Eiter, 2018), are *non-reoccurrence after elimination*, i.e., if a spurious answer  $\hat{I}$  set of a program  $\Pi$  w.r.t. an abstraction mapping  $m$  has no corresponding (can not be mapped to some) answer set  $\hat{I}'$  in a refinement  $m'$  of  $m$ , then no refinement  $m''$  of  $m'$  will have an answer set  $\hat{I}''$  corresponding to  $\hat{I}$  either, and *convexity*, i.e., if on the contrary  $\hat{I}$  has some corresponding answer set  $\hat{I}'$  under  $m'$ , then every refinement  $m''$  in between  $m'$  and  $m$  admits a spurious answer set  $\hat{I}''$  of  $\Pi$  w.r.t.  $m''$  that corresponds with  $\hat{I}$ .

### 3.5.1 Abstraction over Sorts

Applications of ASP usually contain *sorts* that form subdomains of the Herbrand universe. For example, in graph coloring there are sorts for nodes and colors. We define an abstraction over a sort as follows.

**Definition 3.11.** *An abstraction is limited to a sort  $D_i \subseteq D$ , if all elements  $x \in D \setminus D_i$  form singleton clusters  $\{x\}/x$ .*

**Example 3.27.** In graph coloring, we have sorts *node* and *color* in the domain  $\{1, \dots, 6, \text{red}, \text{green}, \text{blue}\}$  for the instance in Figure 1a. An abstraction mapping  $m$  limited to the sort *node* means  $m(x) = \{x\}$  for  $x \in \{\text{red}, \text{blue}, \text{green}\}$ .

In order to obtain much coarser abstractions, applying abstraction over multiple sorts is also possible, given that the individual sorts fulfill the following property.

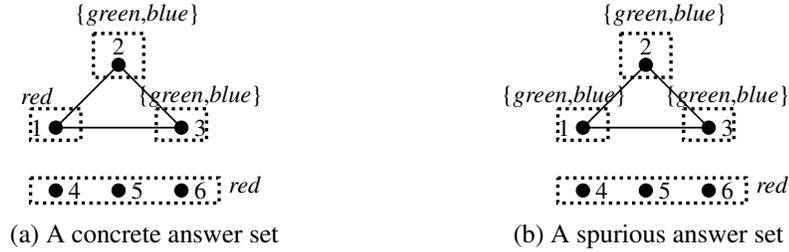
**Definition 3.12** (Sort independence). *For a program  $\Pi$  and domain  $D$ , subdomains  $D_1, \dots, D_n \subseteq D$  are independent, if  $D_i \cap D_j = \emptyset$  for all  $i \neq j$ .*

For independent sorts, abstractions can be composed.

**Proposition 3.7.** *For every domain mappings  $m_1$  and  $m_2$  over independent subdomains  $D_1$  and  $D_2$ , it holds that  $\text{grd}_{\mathcal{F}_{m_1, m_2}}((\Pi^{m_2})^{m_1}) = \text{grd}_{\mathcal{F}_{m_2, m_1}}((\Pi^{m_1})^{m_2})$ .*

*Proof.* The mapping  $m_i : D \mapsto \hat{D}$  is of form  $\{\{x\}/x \mid x \in D \setminus D_i\} \cup m_{D_i}$ ,  $i \in \{1, 2\}$ , where  $m_{D_i}$  describes the mapping over  $D_i$  to the abstract domain  $\hat{D}_i$ . We know that  $m_i(D \setminus D_i) = D \setminus D_i$ , and since  $D_1$  and  $D_2$  are independent, we have  $D_1 \subseteq m_2(D)$  and  $D_2 \subseteq m_1(D)$ . Consequently, we can apply the mappings independently from each other as  $m_2(m_1(D)) = m_1(m_2(D))$  to achieve an abstract domain  $\hat{D} = (D \setminus (D_1 \cup D_2)) \cup \hat{D}_1 \cup \hat{D}_2$ . Another mapping  $m$  can then be defined to map  $D$  to  $\hat{D}$ . By Lemma 3.4 we get the result.  $\square$

Figure 7: Abstraction over the set of nodes and the set of colors



### 3.5.2 Cartesian Abstraction

Given domain mappings  $m_1, \dots, m_n$  limited to subdomains  $D_1, \dots, D_n$ , respectively, a *cartesian abstraction* of the mappings corresponds to the abstract domain  $m(D_1) \times \dots \times m(D_n)$ . Assuming that the subdomains  $D_1, \dots, D_n$  are independent, Definition 3.3 can be altered to be applied over a rule of the form

$$r : l \leftarrow B(r), rel^{D_1}(t_1, t_2), \dots, rel^{D_n}(t_1, t_2)$$

by considering all possible combinations of  $\tau_j^{rel^{D_j}}(\hat{t}_1, \hat{t}_2), j = 1, \dots, n$ . Alternatively, we can define cartesian abstraction by applying abstraction over each subdomain one step at a time, by extending Prop. 3.7 to multiple sorts.

**Proposition 3.8.** *For domain mappings  $m_1, \dots, m_n$  over independent domains  $D_1, \dots, D_n$ , it holds that  $\Pi^{m_1 \times \dots \times m_n} = ((\Pi^{m_{\pi(1)}}) \dots)^{m_{\pi(n)}}$  where  $\pi$  is any permutation of  $\{1, \dots, n\}$ .*

**Example 3.28** (Example 2.6 ctd). In the graph coloring instance of Figure 1a, consider the mappings  $m_n = \{\{4, 5, 6\}/\hat{4}\}$  and  $m_c = \{\{red\}/\hat{r}, \{green, blue\}/\hat{g}\}$  over the sorts nodes and colors, respectively. The abstract program  $(\Pi^{m_n})^{m_c}$  has the concrete answer set

$$\{chosenColor(1, \hat{r}), chosenColor(2, \hat{g}), chosenColor(3, \hat{g}), chosenColor(\hat{4}, \hat{r})\}$$

(shown in Figure 7a) that chooses the color cluster  $\hat{g}$  for nodes 2 and 3, which matches the intuition of coloring the neighbor nodes of node 1 to some color different than its own color.

Notably,  $(\Pi^{m_n})^{m_c}$  also has the spurious answer set (shown in Figure 7b)

$$\{chosenColor(1, \hat{g}), chosenColor(2, \hat{g}), chosenColor(3, \hat{g}), chosenColor(\hat{4}, \hat{r})\}$$

due to the guesses introduced for the uncertainty.

In Section 7.2 we demonstrate further uses of such a multi-step abstraction over the subdomains.

## 3.6 Computational Complexity

In this section, we turn to the computational complexity of reasoning tasks that are associated with program abstraction. We build on the complexity results in (Dantsin et al., 2001; Eiter et al., 2007), which cover the basic reasoning tasks for arbitrary non-ground programs and for non-ground programs with bounded predicate arities, i.e., the maximum arity of a predicate occurring in the program is bounded by a constant.

**Lemma 3.9.** *Given an arbitrary non-ground program  $\Pi$ , a mapping  $m$ , and an abstract interpretation  $I$ , checking whether  $I \in AS(\Pi^m)$  holds is feasible in  $\Delta_2^P$ .*

Intuitively, this holds because we can nondeterministically generate each rule  $r$  in  $\Pi^m$  in polynomial time, and if  $I$  is not a model of the reduct  $(\Pi^m)^I$  also an instance of  $r$  witnessing this fact. The minimality of  $I$  for  $(\Pi^m)^I$  can be shown by a polynomial-size proof tree that can be guessed and checked.

Armed with this lemma, we consider the problem of identifying concrete abstract answer sets.

**Theorem 3.10.** *Given a program  $\Pi$ , a domain mapping  $m$ , and an abstract interpretation  $\hat{I}$ , deciding whether  $\hat{I}$  is a concrete abstract answer set of  $\Pi^m$  is NEXP-complete in general and  $\Sigma_2^P$ -complete for bounded predicate arities. Furthermore, the complexity remains unchanged if  $\hat{I} \in AS(\Pi^m)$  is asserted.*

That is, the worst case complexity is the one of answer set existence for non-ground programs (Dantsin et al., 2001; Eiter et al., 2007); the two problems can be reduced to each other in polynomial time. Intuitively, in general an abstract atom in  $\hat{I}$  may be mapped back to exponentially many atoms in an answer set  $I$  of the original program  $\Pi$  that witnesses the concreteness of  $\hat{I}$ ; such an  $I$  can be guessed and checked in nondeterministic exponential time. Accordingly, the complexity drops to  $\Sigma_2^P$  if the domain size  $|D|$  is polynomial in the abstracted domain size  $|\hat{D}|$  and interpretations are represented as bitmaps (as customary); e.g., it drops if each abstract cluster is small (and multiple clusters exist). Under bounded predicate arities, each abstract atom maps back to polynomially many original atoms, such that the guess  $I$  has polynomial size and checking  $I$  can be done with an NP oracle in polynomial time (cf. Lemma 3.9). The matching lower bounds are shown by reductions from deciding whether a given non-ground program has some answer set.

As an immediate consequence of Theorem 3.10, obtain the following result for spuriousness checking.

**Corollary 3.11.** *Given a program  $\Pi$ , a domain mapping  $m$ , and an abstract interpretation  $\hat{I}$ , deciding whether  $\hat{I}$  is a spurious abstract answer set of  $\Pi^m$  is NEXP-complete in general and  $\Sigma_2^P$ -complete for bounded predicate arities. Furthermore, the complexity remains unchanged if  $\hat{I} \in AS(\Pi^m)$  is asserted.*

Next we consider deciding whether the abstract program has some spurious answer set. This problem turns out to have higher complexity.

**Theorem 3.12.** *Given a program  $\Pi$  and a domain mapping  $m$ , deciding whether some  $\hat{I} \in AS(\Pi^m)$  exists that is spurious is  $\text{coNEXP}^{\text{NP}}$ -complete in general and  $\Pi_3^P$ -complete for programs with bounded predicate arities.*

Intuitively, compared to the previous problem we first must make a guess for  $\hat{I}$  such that it is an abstract answer set of  $\Pi^m$  but not concrete; the size of  $\hat{I}$  may be exponential in the input of the problem, and relative to this testing concreteness is feasible in nondeterministic polynomial time, i.e., with an NP oracle. The matching hardness is shown by reductions from evaluating second-order logic formulas of a suitable form over finite relational structures.

**Faithful abstraction.** An abstract program that does not have a spurious answer set is a faithful abstraction of the original program.

**Example 3.29** (Example 2.6 ctd). In the graph coloring instance of Figure 1a, the mapping  $m = \{\{4, 5, 6\}/\hat{4}\}$ , which maps nodes 1,2,3 to singleton clusters, yields an abstract program  $\Pi^m$  that has 42 answer sets, which are the combinations of 6 possible correct colorings of the nodes 1-3 with 7 possible colorings  $\{\{red\}, \{blue\}, \{green\}, \{red, blue\}, \{red, green\}, \{green, blue\}, \{red, green, blue\}\}$  of the node cluster  $\hat{4}$ , thus resulting in a faithful abstraction.

Ideally, we have faithfulness, but this is hard to achieve in general. From Theorem 3.12, we immediately obtain:

**Corollary 3.13.** *Given a program  $\Pi$  and a domain mapping  $m$ , deciding whether  $\Pi^m$  is faithful is  $\text{coNEXP}^{\text{NP}}$ -complete in general and  $\Pi_3^P$ -complete for bounded predicate arities.*

## 4 Refinement by Debugging Non-Ground Spuriousness

Over-approximation of an answer set program unavoidably introduces spurious answer sets. Once a spurious abstract answer set is encountered, one can either continue searching for a concrete abstract answer set, or *refine* the abstraction to reach one where less spurious answer sets occur.

**Definition 4.1.** *Given a domain mapping  $m : D \rightarrow D'$ , a mapping  $m' : D \rightarrow D''$  is a refinement of  $m$  if for all  $x \in D$ ,  $m'^{-1}(m'(x)) \subseteq m^{-1}(m(x))$ .*

That is, refinement is on dividing the abstract clusters to a finer grained domain.

**Example 4.1** (ctd). The mapping  $m' = \{\{1\}/k_1, \{2, 3, 4\}/k_2, \{5\}/k_3\}$  is a refinement of the mapping  $m$ . Furthermore,  $\Pi^{m'}$  has no answer set  $I'$  such that  $m(m'^{-1}(I')) = \hat{I}_2$ ; hence the spurious answer set  $\hat{I}_2$  of  $\Pi^m$  is eliminated.

In the CEGAR methodology (Clarke et al., 2003), the decision in a refinement step depends on the correctness checking of the spurious abstract solution, through which the problematic part of the abstraction is detected. Inspired by this, we develop an alternative for checking the correctness of abstract answer sets that can be used to determine how the refinement should be made.

**Correctness checking using constraints.** That an abstract answer set  $\hat{I}$  is spurious means the original program  $\Pi$  has no answer set matching  $\hat{I}$ . In other words, querying  $\Pi$  for a match to an abstract answer set  $\hat{I}$  would return no result exactly if  $\hat{I}$  is spurious.

**Definition 4.2** (Query of an answer set). *Given an abstract answer set  $\hat{I}$  and a mapping  $m$ , a query  $Q_{\hat{I}}^m$  for an answer set that matches  $\hat{I}$  is described by the following constraints.*

$$\perp \leftarrow \{ \alpha \mid m(\alpha) = \hat{\alpha} \} \leq 0. \quad \hat{\alpha} \in \hat{I} \setminus \mathcal{T}_m \quad (30)$$

$$\perp \leftarrow \alpha. \quad \hat{\alpha} \notin \hat{I} \setminus \mathcal{T}_m, m(\alpha) = \hat{\alpha} \quad (31)$$

Here (30) ensures that a witnessing answer set  $I$  of  $\Pi$  (i.e.,  $m(I) = \hat{I}$ ) contains for every non- $\tau_i$  abstract atom in  $\hat{I}$  some atom that is mapped to it, while (31) ensures that no atom in  $I$  is mapped to an abstract atom not in  $\hat{I}$ . The following is then easy to establish.

**Proposition 4.1.** *Suppose  $m$  is a domain abstraction mapping for a program  $\Pi$ , then an abstract answer set  $\hat{I} \in AS(\Pi^m)$  is spurious iff  $\Pi \cup Q_{\hat{I}}^m$  is unsatisfiable.*

*Proof.* As  $\hat{I}$  is spurious, there exists no  $I \in AS(\Pi)$  such that  $m(I) = \hat{I} \setminus \mathcal{T}_m$ , i.e., there is no match of an original answer set  $I$  for  $\hat{I}$  where the atoms in  $I$  can be mapped to the abstract atoms contained in  $\hat{I} \setminus \mathcal{T}_m$  and the atoms not in  $I$  can be mapped to the abstract atoms not contained in  $\hat{I} \setminus \mathcal{T}_m$ .  $Q_{\hat{I}}^m$  enforces such a match, thus returns unsatisfiability.

Having no match for  $\hat{I}$  means that no original answer set can be mapped to it, thus  $\hat{I}$  is spurious.  $\square$

**Correctness checking with debugging.** We will employ an ASP debugging approach to debug the inconsistency of the original program  $\Pi$  caused by checking a spurious answer set  $\hat{I}$ , referred to as *inconsistency of  $\Pi$  w.r.t.  $\hat{I}$* , in order to get hints for refining the abstraction. Different from a usual ASP program debugging approach, we need to shift the focus from “debugging the original program” to “debugging the inconsistency caused by the spurious answer set”. Unfortunately an immediate application of the available ASP debugging tools is not possible. For our purposes, we make use of the meta-level debugging language in (Brain et al.,

2007), which is based on a tagging technique that allows one to control the building of answer sets and to manipulate the program evaluation.

The meta-program constructed by `spock` (Brain et al., 2007) introduces *tags* to control answer set building. Given a program  $\Pi$  over  $\mathcal{A}$  and a set  $\mathcal{N}$  of names for all rules in  $\Pi$ , it creates an enriched alphabet  $\mathcal{A}^+$  obtained from  $\mathcal{A}$  by adding atoms such as  $ap(n_r), bl(n_r), ok(n_r), ko(n_r)$  where  $n_r \in \mathcal{N}$  for each  $r \in \Pi$ . The atoms  $ap(n_r), bl(n_r)$  express whether a rule  $r$  is applicable or blocked, respectively, while  $ok(n_r), ko(n_r)$  are used for manipulating the application of  $r$ .

For domain abstraction, debugging the non-ground program has its own difficulties. The approach in (Brain et al., 2007) is on the propositional level, thus can not be immediately applied. Also debugging non-ground programs is not as straightforward as in the propositional case, as there is the additional need to debug the checking for an original answer set that can be mapped to the given abstract answer set. However, non-ground program debugging approaches such as (Oetsch et al., 2010; Dodaro et al., 2015) are not easily adjustable due to the need for shifting the focus towards debugging the correctness checking.

**Using available debuggers.** Debugging non-ground ASP programs through a meta-programming (Gebser et al., 2008) approach has been studied by (Oetsch et al., 2010), with the drawback of considering all possible explanations why a given interpretation  $I$  is not an answer set of the program  $\Pi$ . For the given input  $I$ , in order to prove that  $I$  is not an answer set of  $\Pi$ , the debugging considers many possible guesses of variable assignments that matches  $I$  with a faulty behavior. In our case, the input  $I$  is an abstract answer set stating that there should be some original answer set  $I'$  of  $\Pi$  such that each atom in  $I'$  can be mapped to some abstract atom  $\hat{\alpha}$  in  $I$ . This adds a guess of some original atom that could be mapped to  $\hat{\alpha}$ . However, as the debugging aims at showing that  $I$  is not an answer set of  $\Pi$ , when this additional guessing comes into play, it guesses original atoms to create some faulty behavior for  $I$  even if these atoms do not occur in an original answer set. Thus, an immediate application of the meta-programming approach is infeasible.

In order to use the available non-ground debugging tools off-the-shelf, one possibility is to first guess all possible combinations of the original atoms to match the abstract answer set  $\hat{I}$ , and then separately debug each of them. If  $\hat{I}$  is in fact spurious, this will be caught as each possible guess would return some inconsistency. If  $\hat{I}$  is concrete, then at some point some guess will correspond to an original answer set, with no inconsistency. However, this approach is too cumbersome, as there can be many possible concrete guesses for an abstract  $I$  and checking each of them one by one until a concrete one is found (if exists) is highly inefficient.

**Our approach to debugging.** As existing non-ground debugging tools are not readily applicable, we debug the unsatisfiability of  $\Pi \cup Q_I^m$  for a spurious abstract answer set  $\hat{I}$  following the debugging approach based on (Brain et al., 2007) from above. In previous work on domain abstraction (Saribatur et al., 2019), we considered a simplified debugging approach inspired from the *ko* atoms of (Brain et al., 2007), which is based on detecting the rules that must be deactivated in order to keep the satisfiability while checking the concreteness of an abstract answer set  $\hat{I}$ , in case it is spurious. As the naive debugging can not address all debugging cases, in this work we show an extension of the refinement method by lifting the `spock` (Brain et al., 2007) debugging approach to the non-ground case, confining to tight programs (i.e., we omit unfounded loop checking).

When demonstrating the different debugging approaches, we use a non-ground version of  $Q_I^m$ .

**Definition 4.3** (Non-ground query). *Given an abstract answer set  $\hat{I}$  and a mapping  $m$  expressed as a set of facts of form  $m(x, a)$  (where  $m(x) = a$ ), a (non-ground) query for an answer set that matches  $\hat{I}$  is described*

as follows:

$$\perp \leftarrow \text{in}(\hat{\alpha}), \{\alpha : m(X_1, \hat{X}_1), \dots, m(X_k, \hat{X}_k)\} \leq 0. \quad (32)$$

$$\perp \leftarrow \alpha, \text{not in}(\hat{\alpha}), m(X_1, \hat{X}_1), \dots, m(X_k, \hat{X}_k). \quad (33)$$

where  $\alpha = p(X_1, \dots, X_k)$  and  $\hat{\alpha} = p(\hat{X}_1, \dots, \hat{X}_k)$ , and  $m(X_i, \hat{X}_i)$  expresses the abstract mapping, plus the facts

$$\text{in}(\hat{\alpha}), \quad \hat{\alpha} \in \hat{I} \setminus \mathcal{T}_m. \quad (34)$$

**Example 4.2** (Example 3.1 ctd). For the program  $\Pi$  and the mapping  $m = \{\{1, 2, 3, 4, 5\}/k\}$  given as facts  $m(1, k), m(2, k), m(3, k), m(4, k), m(5, k)$ , the abstract program  $\Pi^m$  has an answer set  $\hat{I} = \{a(k), c(k)\}$ . The query  $Q_{\hat{I}}^m$  is

$$\begin{aligned} \perp &\leftarrow \text{not in}(d(A_1)), d(X_1), m(X_1, A_1). & \perp &\leftarrow \text{in}(d(A_1)), \{d(X_1) : m(X_1, A_1)\} \leq 0. \\ \perp &\leftarrow \text{not in}(c(A_1)), c(X_1), m(X_1, A_1). & \perp &\leftarrow \text{in}(c(A_1)), \{c(X_1) : m(X_1, A_1)\} \leq 0. \\ \perp &\leftarrow \text{not in}(a(A_1)), a(X_1), m(X_1, A_1). & \perp &\leftarrow \text{in}(a(A_1)), \{a(X_1) : m(X_1, A_1)\} \leq 0. \\ \perp &\leftarrow \text{not in}(e(A_1)), e(X_1), m(X_1, A_1). & \perp &\leftarrow \text{in}(e(A_1)), \{e(X_1) : m(X_1, A_1)\} \leq 0. \\ \perp &\leftarrow \text{not in}(b(A_1, A_2)), b(X_1, X_2), m(X_1, A_1), m(X_2, A_2). \\ \perp &\leftarrow \text{in}(b(A_1, A_2)), \{b(X_1, X_2) : m(X_1, A_1), m(X_2, A_2)\} \leq 0. \\ \text{in}(a(k)). & & \text{in}(c(k)). & \end{aligned}$$

## 4.1 Non-Ground Debugging using Tagging

We extend the refinement method described in (Saribatur et al., 2019) by lifting the “tagging” approach of `spock` (Brain et al., 2007) to the non-ground case, confining to tight programs (i.e., we omit unfounded loop checking). Given  $\Pi$ , we construct the meta program  $\mathcal{T}_{meta}[\Pi]$  similar to `spock` (Brain et al., 2007), but with an extension of having arguments in the  $ap_{n_r}, bl_{n_r}$  atoms to have information for which constants the rules are applicable and blocked.

**Definition 4.4.** Given a non-ground program  $\Pi$ , the program  $\mathcal{T}_{meta}[\Pi]$  consists of the following rules for  $r \in \Pi$  with  $\text{arg}(H(r)) = \{c_1, \dots, c_n\}$  and  $\text{arg}(B(r)) = \{d_1, \dots, d_m\}$ :

$$\begin{aligned} \text{If } B(r) = \emptyset : & & r \\ & & H(r) \leftarrow ap_{n_r}(c_1, \dots, c_n), \text{not } ko_{n_r}. \\ \text{If } H(r) \neq \perp \wedge n > 0 : & & ap_{n_r}(c_1, \dots, c_n) \leftarrow B(r). \\ & & bl_{n_r}(c_1, \dots, c_n) \leftarrow \text{not } ap_{n_r}(c_1, \dots, c_n). \\ & & H(r) \leftarrow ap_{n_r}(d_1, \dots, d_m), \text{not } ko_{n_r}. \\ \text{If } H(r) = \perp \vee n = 0 : & & ap_{n_r}(d_1, \dots, d_m) \leftarrow B(r). \\ & & ap_{n_r} \leftarrow ap_{n_r}(d_1, \dots, d_m). \\ & & bl_{n_r} \leftarrow \text{not } ap_{n_r}. \end{aligned}$$

In case the head of rule  $r$  is  $\perp$  or does not contain arguments in the atom, we use the arguments from the body to know whether  $r$  is applicable.

We have abnormality atoms to indicate the actions that are required to avoid the inconsistency:

- $ab\_deact_{n_r}$  signals that  $r$  was applicable under some interpretation, but had to be deactivated;
- similarly for  $ab\_deactCons_{n_r}$  which only talks about the constraints; and
- $ab\_act(\alpha)$  says that atom  $\alpha$  must be made true while no rule deriving  $\alpha$  was applicable (i.e.,  $\alpha$  is *unsupported*<sup>2</sup>).

**Definition 4.5.** Given a non-ground program  $\Pi$  over  $\mathcal{A}$ , the following additional meta-programs are constructed:

1. *Rule Deactivation:*  $\mathcal{T}_{deact}[\Pi]$ : for all  $r \in \Pi$  with  $B(r) \neq \emptyset$  and  $H(r) \neq \perp$ :

$$\begin{aligned} & ko_{n_r}. \\ & \{H(r)\} \leftarrow ap_{n_r}(c_1, \dots, c_n). \\ & ab\_deact_{n_r}(c_1, \dots, c_n) \leftarrow ap_{n_r}(c_1, \dots, c_n), not H(r). \end{aligned}$$

2. *Constraint Deactivation:*  $\mathcal{T}_{deactCons}[\Pi]$ : for all  $r \in \Pi$  with  $H(r) = \perp$ :

$$\begin{aligned} & \{ko_{n_r}\}. \\ & ab\_deactCons_{n_r}(c_1, \dots, c_n) \leftarrow ap_{n_r}(c_1, \dots, c_n), ko_{n_r}. \end{aligned}$$

3. *Rule Head Activation:*  $\mathcal{T}_{act}[\Pi, \mathcal{A}]$ : for all  $\alpha \in \mathcal{A}$ , with  $def(\alpha, \Pi) = \{r_1, \dots, r_k\}$ :

$$\begin{aligned} & \{\alpha\} \leftarrow bl_{n_{r_1}}(c_1, \dots, c_n), \dots, bl_{n_{r_k}}(c_1, \dots, c_n). \\ & ab\_act(\alpha) \leftarrow \alpha, bl_{n_{r_1}}(c_1, \dots, c_n), \dots, bl_{n_{r_k}}(c_1, \dots, c_n). \end{aligned}$$

The arguments of  $ab\_deact$  only contain the ones from the head of the rule. This is a representation choice, to avoid dealing with many variables involved in the body while only few of them are used in the head of the rule. For the definition of  $ab\_deactCons$  however, the variables of the body must be used. Having a different representation for the deactivation of the constraints will allow to steer the debugging towards the constraints by assigning different costs for their occurrence when computing the answer sets with the smallest number of  $ab$  atoms.

**Definition 4.6.** For a program  $\Pi$  over atoms  $\mathcal{A}$ , we let the program be defined by

$$\Pi_{debug} = \mathcal{T}_{meta}[\Pi] \cup \mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}].$$

Furthermore, we denote by  $\mathcal{A}^* = \mathcal{A}^+ \cup AB(\Pi)$  the vocabulary of  $\Pi_{debug}$ .

We use  $\Pi_{debug}$  for checking the correctness of an abstract answer set and then deciding on the refinement. Adding weak constraints over the abnormality atoms yields an answer set with fewest  $ab$  atoms.

**Example 4.3** (Example 4.2 ctd). The programs  $\mathcal{T}_{meta}[\Pi]$  and  $\mathcal{T}_{deact}[\Pi]$  with additional weak constraints on abnormality atoms are shown in Figs. 8a and 8b, respectively. The minimal answer set of  $\mathcal{T}_{meta}[\Pi] \cup \mathcal{T}_{deact}[\Pi] \cup Q_I^m$  is then

$$\{ab\_deact_{r_2}(5), ab\_deact_{r_4}(1), ab\_deact_{r_4}(3), ab\_deact_{r_4}(2)\}.$$

<sup>2</sup>A literal  $q$  is *unsupported* by an interpretation  $I$  if for each  $r \in def(q, \Pi)$ ,  $B^+(r) \not\subseteq I$  or  $B^-(r) \cap I \neq \emptyset$  (Van Gelder et al., 1991).

Figure 8: Debugging programs for Example 4.2

$c(X) \quad \leftarrow ap_{r1}(X), not\ ko_{r1}.$ $ap_{r1}(X) \quad \leftarrow X < 5, not\ d(X).$ $bl_{r1}(X) \quad \leftarrow not\ ap_{r1}(X).$	$ko_{r1}.$ $\{c(X)\} \quad \leftarrow ap_{r1}(X).$ $ab\_deact_{r1}(X) \quad \leftarrow ap_{r1}(X), not\ c(X).$ $\perp \quad \quad \quad : \sim ab\_deact_{r1}(X).$
$d(X) \quad \leftarrow ap_{r2}(X), not\ ko_{r2}.$ $ap_{r2}(X) \quad \leftarrow not\ c(X).$ $bl_{r2}(X) \quad \leftarrow not\ ap_{r2}(X).$	$ko_{r2}.$ $\{d(X)\} \quad \leftarrow ap_{r2}(X).$ $ab\_deact_{r2}(X) \quad \leftarrow ap_{r2}(X), not\ d(X).$ $\perp \quad \quad \quad : \sim ab\_deact_{r2}(X).$
$b(X, Y) \quad \leftarrow ap_{r3}(X, Y), not\ ko_{r3}.$ $ap_{r3}(X, Y) \quad \leftarrow a(X), d(Y).$ $bl_{r3}(X, Y) \quad \leftarrow not\ ap_{r3}(X, Y).$	$ko_{r3}.$ $\{b(X, Y)\} \quad \leftarrow ap_{r3}(X, Y).$ $ab\_deact_{r3}(X, Y) \quad \leftarrow ap_{r3}(X, Y), not\ b(X, Y).$ $\perp \quad \quad \quad : \sim ab\_deact_{r3}(X, Y).$
$e(X) \quad \leftarrow ap_{r4}(X), not\ ko_{r4}.$ $ap_{r4}(X) \quad \leftarrow c(X), a(Y), X \leq Y.$ $bl_{r4}(X) \quad \leftarrow not\ ap_{r4}(X).$	$ko_{r4}.$ $\{e(X)\} \quad \leftarrow ap_{r4}(X).$ $ab\_deact_{r4}(X) \quad \leftarrow ap_{r4}(X), not\ e(X).$ $\perp \quad \quad \quad : \sim ab\_deact_{r4}(X).$
$falsum \quad \leftarrow ap_{r5}(X, Y), not\ ko_{r5}.$ $ap_{r5}(X, Y) \quad \leftarrow b(X, Y), e(X).$ $ap_{r5} \quad \leftarrow ap_{r5}(X, Y).$ $bl_{r5} \quad \leftarrow not\ ap_{r5}.$ $\perp \quad \quad \quad \leftarrow falsum.$	$ko_{r5}.$ $\{falsum\} \quad \leftarrow falsum.$ $ab\_deact_{r5} \quad \leftarrow falsum.$ $\perp \quad \quad \quad : \sim ab\_deact_{r5}.$
<p>(a) <math>\mathcal{T}_{meta}[\Pi]</math></p>	<p>(b) <math>\mathcal{T}_{deact}[\Pi]</math> with weak constraints</p>

This debugging approach is also able to handle the shortcomings of the naive approach (Saribatur et al., 2019), as  $\mathcal{T}_{act}[\Pi, \mathcal{A}]$  is used to activate original atoms if it is necessary for achieving satisfiability for  $\Pi_{debug} \cup Q_{\hat{I}}^m$ .

As a first property, we show that  $\Pi_{debug} \cup Q_{\hat{I}}^m$  always has an answer set, i.e., no abstract answer set  $\hat{I}$  is dismissed.

**Proposition 4.2.** *Given a tight program  $\Pi$  and a mapping  $m$ , for each answer set  $\hat{I} \in AS(\Pi^m)$ ,  $\Pi_{debug} \cup Q_{\hat{I}}^m$  has an answer set.*

The next result now shows that we can use  $\Pi_{debug} \cup Q_{\hat{I}}^m$  to obtain hints for the spuriousness reason of  $\hat{I}$ .

**Proposition 4.3.** *Given a tight program  $\Pi$  and a mapping  $m$ , if an answer set  $\hat{I} \in AS(\Pi^m)$  is spurious, then for every answer set  $S \in AS(\Pi_{debug} \cup Q_{\hat{I}}^m)$  either (i)  $ab\_deact_{n_r}(c_1, \dots, c_n) \in S$  or  $ab\_deactCons_{n_r}(c_1, \dots, c_n) \in S$  for some  $r \in \Pi$ , or (ii)  $ab\_act(\alpha(c_1, \dots, c_n)) \in S$ .*

*Proof.* If  $\hat{I}$  is spurious, by Proposition 4.1 the program  $\Pi \cup Q_{\hat{I}}^m$  is unsatisfiable. We focus on debugging the cause of inconsistency introduced by  $Q_{\hat{I}}^m$ . Since  $\Pi$  is tight, this inconsistency can either be due to (i) an unsatisfied rule or (ii) an unsupported atom.

- (i) Let  $r \in \Pi$  be an unsatisfied rule w.r.t.  $S$ . This means that the constraints in  $Q_{\hat{I}}^m$  is causing  $H(r)$  to be false while  $B(r)$  is satisfied. By the program  $\mathcal{T}_{meta}[\Pi]$ , depending on  $r$ , either  $ap_{n_r}(c_1, \dots, c_n)$  or  $ap_{n_r}(d_1, \dots, d_m)$  is true. By  $\mathcal{T}_{deact}[\Pi]$ , we get  $ab\_deact(r, c_1, \dots, c_n) \in S$ . If  $H(r) = \perp$ , then by  $\mathcal{T}_{deactCons}[\Pi]$ , we have  $ko_{n_r} \in S$  (else  $\perp \in S$  by  $\mathcal{T}_{meta}[\Pi]$ ), and we get  $ab\_deactCons(r, c_1, \dots, c_n) \in S$ .
- (ii) Let  $\alpha = \alpha(c_1, \dots, c_n) \in S$  be an unsupported atom in  $\Pi$  w.r.t  $S$  for the domain elements  $c_1, \dots, c_n$ . Then, for each rule instance  $r$  deriving  $\alpha$ , we have  $bl_{n_r}(c_1, \dots, c_n) \in S$  and by  $\mathcal{T}_{act}[\Pi, \mathcal{A}]$ , we have  $ab\_act(\alpha) \in S$ .  $\square$

Less suprisingly, for non-tight programs, debugging the correctness checking could result in unsatisfiability (see Appendix B.1 for an example). To avoid this, unfounded loop checking can be handled by introducing an additional abnormality atom, say  $ab_{loop}$  as in (Brain et al., 2007), and lifting it to the non-ground setting. However, this solution causes further guessing rules involved in the non-ground debugging. Also the existence of  $ab_{loop}(\alpha)$  sometimes does not even indicate that a loop formula is violated and just makes the search more difficult due to considering many possibilities of the guesses. Therefore, we choose to focus only on tight programs and concentrate on the determination of a refinement.

The obtained debugging atoms during a correctness check gives hints on which domain elements should not be involved in a cluster.

**Definition 4.7.** *The refinement-hint gathering program  $\Pi_{hint}(AB)$  contains the following rules:*

- For  $c_i \in arg(ab\_deact_{n_r}(c_1, \dots, c_n))$ :  

$$refine(c_1, \dots, c_n) \leftarrow ab\_deact_{n_r}(c_1, \dots, c_n), m(c_i, a_i), isCluster(a_i).$$
- For  $c_i \in arg(ab\_deactCons_{n_r}(c_1, \dots, c_n))$ :  

$$refine(c_1, \dots, c_n) \leftarrow ab\_deactCons_{n_r}(c_1, \dots, c_n), m(c_i, a_i), isCluster(a_i).$$
- For  $c_i \in arg(\alpha(c_1, \dots, c_n))$ :  

$$refine(c_1, \dots, c_n) \leftarrow ab\_act(\alpha(c_1, \dots, c_n)), m(c_i, a_i), isCluster(a_i).$$

From  $\Pi_{hint}$  we get as hints the domain elements that are mapped to abstract cluster elements and cause to obtain  $ab$  atoms in the debugging.

## 4.2 Deciding on a Refinement

The introduced debugging approach finds a set of abnormality atoms in case the abstract answer set is spurious. We consider two ways of using the obtained debugging output for deciding on a refinement.

- (v1) The smallest number of *ab* atoms occurring in an answer set is the cost of the corresponding mapping.
- (v2) The inferred *refine* atoms are used to decide on a refinement of the abstraction.

In (v1), the cost is used for a local search among the possible refinements of an abstraction, where the one with the minimum cost is picked. Approach (v2) is closer to the CEGAR-like approach (Clarke et al., 2003), where a refinement is determined from the spuriousness check. We now describe the approaches in more detail and report on a comparison in Section 6.3.

**(v1) Local Refinement Search.** The idea is to search among possible refinements of a mapping for deciding on a refinement. To single out close refinements, we measure the distance  $dist(m, m')$  between a mapping  $m: D \rightarrow \hat{D}$  and a refinement  $m': D \rightarrow \hat{D}'$  of it by the number of additional clusters, i.e.,  $dist(m, m') = |\hat{D}'| - |\hat{D}|$ . In case  $dist(m, m') = 1$ , we call  $m'$  a *1-distance refinement* of  $m$ .

**Example 4.4.** Each mapping  $m' \in \bigcup_{C \in \mathcal{C}} \{C/k_1, \{1, \dots, 5\} \setminus C/k_2\}$ , where  $\mathcal{C} = \{\{1\}, \{1, 2\}, \{1, 2, 3\}, \{1, 2, 3, 4\}\}$ , is a 1-distance refinement of  $m = \{\{1, 2, 3, 4, 5\}/k\}$ .

Algorithm 1 shows the procedure of deciding on a refinement for a given mapping  $m$ , by doing a distance-based search among all possible refinements of the mapping and picking the one with the least cost. All 1-distance refinements of  $m$  are computed, and then the cost of each of them is determined, by calling *getCostOfMapping*. This function constructs the abstract program  $\Pi'$  according to the mapping and picks an abstract answer set  $I$ . It then finds the answer set with smallest number  $s$  of *ab*-atoms of the program  $\Pi_{debug} \cup Q_m(I)$  and returns  $s$ . If some refinement  $m'$  has cost 0, it is returned. Otherwise, all the refinements and their were collected. In Line 11 the refinements with minimum cost are gathered, and then a random pick is made over them. If the given mapping contains only singleton clusters, this means the original domain has been reached.

**(v2) Abstraction Refinement Using Hints.** The abstract answer set correctness checking returns *ab*-atoms that contain the domain elements involved in the debugging of the unsatisfiability. The latter can be used as hints on which part of the mapping to refine. The idea of the *refine* atoms is to get hints about which domain elements should not be involved in a cluster. Given a hint atom  $refine(c_1, \dots, c_n)$ , we consider two actions to describe a refinement  $m'$  of  $m$ :

- (1) For  $refine(c_1, \dots, c_n)$  and  $i \in \{1, \dots, n\}$  such that  $m^{-1}(m(c_i)) > 1$ , the refinement  $m'$  satisfies  $m'^{-1}(m'(c_i)) = 1$ .
- (2) For  $refine(c_1, \dots, c_n)$  and  $c_i \neq c_j \in \{c_1, \dots, c_n\}$  such that  $m(c_i) = m(c_j)$ , the refinement  $m'$  satisfies  $m'(c_i) \neq m'(c_j)$ .

Applying refinement action (1) means to refine the abstraction by mapping all elements occurring in some *refine* atom to singletons, while the refinement action (2) should ensure that distinct elements occurring in the same *refine* atom are no longer mapped to the same cluster. Note that obtaining some *refine* atom during the correctness checking it is not guaranteed whenever  $\hat{I}$  is spurious, as the *ab* atoms may contain only domain elements that are mapped to singleton clusters. In this case, another abstract answer set  $\hat{I}'$  of  $\hat{\Pi}^m$  can be picked for the correctness checking.

---

**Algorithm 1:** *decideRefinement* with Search

---

**Input:** program  $\Pi$ , domain mapping  $m$ **Output:** refinement  $m'$  of  $m$ 

```
1 if  $m$  has non-singleton clusters then
2    $refinecosts = []$ ;  $allrefs = []$ ;
   /* compute all 1-distance refinements of  $m$  */
3    $refs = computeRefinements(m, I)$ 
4   forall  $m' \in refs$  do
5      $c = getCostOfMapping(\Pi, m')$ ;
6     if  $c = 0$  then /* found a concrete abstract answer set */
7       return  $m'$ ;
8     else
9        $allrefs.append(m')$ ;
10       $refinecosts.append(c)$ 
11   $minrefs = getRefsMinCost(refinecost, allrefs)$ 
12   $m = pickRandomRef(minrefs)$ 
13 return  $m$ 
14
15 def  $getCostOfMapping(\Pi, m)$ 
16    $\Pi' = constructAbsProg(\Pi, m)$ ;
17    $\Pi_{debug} = constructDebugProg(\Pi)$ ;
18   Pick some  $I \in AS(\Pi')$ 
19   Find optimum answer set  $I'$  of  $\Pi_{debug} \cup Q_m(I)$  /* with smallest number  $s$  of
      $ab$ -atoms */
20  return  $|I'|_{ab}$  /*  $s = |I'|_{ab}$  */
```

---

## 5 Multi-Dimensional Domain Abstraction

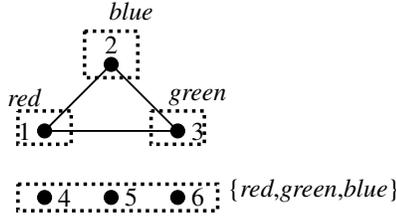
As we have seen from Proposition 3.8, it is possible to construct an abstract program over multiple sorts in the manner of cartesian abstraction. This is achieved by doing abstraction over the sorts one at a time, which means one can not take certain interdependencies among the sorts into account. However, in some scenarios this may be needed.

**Example 5.1** (Example 3.28 ctd). An interesting abstraction would be to assign a color cluster  $\hat{rgb}$  only for the nodes  $\{4, 5, 6\}$ , which are clustered to a node  $\hat{4}$ , while for nodes  $\{1, 2, 3\}$  the original colors are considered (see Figure 9). Such an abstraction can not be achieved with a cartesian style abstraction, since the color cluster  $\hat{rgb}$  is only meant to be considered for the node cluster  $\hat{4}$ . Thus, the desired abstraction can only be defined with a multi-dimensional mapping  $m : D_n \times D_c \rightarrow \hat{D}_n \times \hat{D}_c$  as follows:

$$m(i, j) = \begin{cases} (i, j) & i \in \{1, 2, 3\}, j \in \{red, green, blue\} \\ (\hat{4}, \hat{rgb}) & i \in \{4, 5, 6\}, j \in \{red, green, blue\} \end{cases} .$$

To further motivate the need for multi-dimensionality, we consider grid-cell domains, which are commonly used.

Figure 9: Joint abstraction of nodes and colors



**Example 5.2** (Grid-Cell Domains). Usually the grid-cells are represented by using two sorts *row* and *column*. The following rules show the part of a Sudoku encoding that guesses an assignment of symbols to the cells and ensures that each cell has a number.

$$\begin{aligned}
 \{sol(X, Y, N)\} &\leftarrow not\ occupied(X, Y), num(N), row(X), column(Y). \\
 hasNum(X, Y) &\leftarrow sol(X, Y, N), row(X), column(Y). \\
 \perp &\leftarrow not\ hasNum(X, Y), row(X), column(Y).
 \end{aligned} \tag{35}$$

Further constraints ensure that cells in the same column (36) or same row (37) do not contain the same symbol.

$$\perp \leftarrow sol(X, Y_1, M), sol(X, Y_2, M), Y_1 < Y_2. \tag{36}$$

$$\perp \leftarrow sol(X_1, Y, M), sol(X_2, Y, M), X_1 < X_2. \tag{37}$$

A further more involved constraint (cf. Appendix B.2) ensures that the cells in the same sub-region also satisfies this.

An abstraction over the grid-cells would be to cluster the rows and columns together in order to define an abstract grid-cell. Although abstraction over the sorts one at a time can achieve certain abstract cell structures, to obtain more sophisticated abstractions these sorts must be *jointly* abstracted. Consider for example the abstractions in Figure 10. Those in Figures 10a-10b can be achieved by independent mappings over the rows and columns such as  $m_{row} = m_{col} = \{\{123\}/a_{123}, \{4\}/a_4\}$  and  $m_{row} = m_{col} = \{\{12\}/a_{12}, \{34\}/a_{34}\}$ . For a given program  $\Pi$ , one can construct the abstract program  $(\Pi^{m_{row}})^{m_{col}}$ . However to achieve Figure 10c, rows and columns must be *jointly* abstracted. While the cells  $(a_i, b_j), 1 \leq i, j \leq 2$  are singletons mapped from  $(i, j)$ , the other abstract regions are only given by

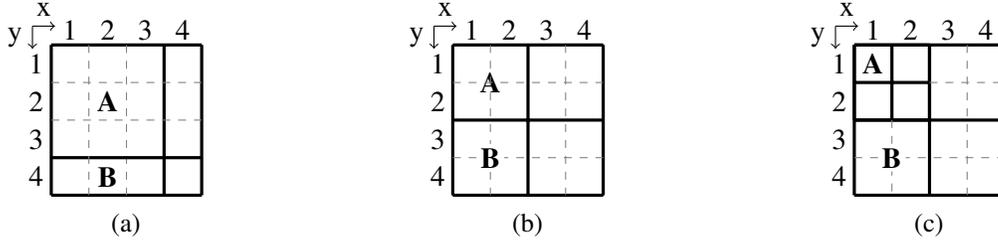
$$m_{row.col}(x, y) = \begin{cases} (a_{12}, b_{34}) & x \in \{1, 2\}, y \in \{3, 4\} \\ (a_{34}, b_{12}) & x \in \{3, 4\}, y \in \{1, 2\} \\ (a_{34}, b_{34}) & x \in \{3, 4\}, y \in \{3, 4\} \end{cases} \tag{38}$$

Observe that the abstract row  $a_{12}$  describes a cluster that abstracts over the individual abstract rows  $a_1, a_2$ . The original rows  $\{1, 2\}$  are mapped to  $\{a_{12}\}$  only in combination with columns  $\{3, 4\}$ , otherwise they are mapped to  $\{a_1, a_2\}$ .

## 5.1 Existential Abstraction on Relations

The abstraction method described in Section 3.2 aims at keeping the built-ins in the abstract program and finds a way to handle their different behavior in the abstract domain. However, this approach can not be used

Figure 10: Abstractions over grid-cells



to achieve the above mentioned multi-dimensional abstraction. Consider the rule (36) standardized apart over rows and columns, thus having relations  $X_1 = X_2$  and  $Y_1 < Y_2$ . If for the mapping  $m_{row,col}$  in Figure 10c these relations are lifted following Section 3.2, while the relation over the y-axis is still defined (as A is located above of B), i.e.,  $A_Y \leq B_Y$ , the relation  $A_X = B_X$  is unclear as the abstract clusters for X-values are different due to different levels of abstraction.

To tackle this issue, an alternative abstraction method is needed that also abstracts over the built-in relations and reasons over the abstracted relation (in the abstract domain). This leads us to a notion of abstraction that is similar in spirit to so called existential abstraction (Clarke et al., 2003) and allows us to introduce domain mappings over multiple sorts such as

$$m: D_1 \times \dots \times D_n \rightarrow \hat{D}_1 \times \dots \times \hat{D}_n,$$

and to handle relations over different levels of abstraction.

For this, we introduce an abstract relation  $\widehat{rel}$  for a  $k$ -ary relation  $rel$  as follows:

$$(\forall_{i=1}^k \hat{d}_i \in \hat{D}) \widehat{rel}(\hat{d}_1, \dots, \hat{d}_k) \Leftrightarrow \exists_{i=1}^k x_i \in m^{-1}(\hat{d}_i). rel(x_1, \dots, x_k). \quad (39)$$

$$(\forall_{i=1}^k \hat{d}_i \in \hat{D}) neg\_rel(\hat{d}_1, \dots, \hat{d}_k) \Leftrightarrow \exists_{i=1}^k x_i \in m^{-1}(\hat{d}_i). \neg rel(x_1, \dots, x_k). \quad (40)$$

I.e.,  $\widehat{rel}(\hat{d}_1, \dots, \hat{d}_k)$  is true if for some corresponding original values the original relation holds; the negation of  $\widehat{rel}(\hat{d}_1, \dots, \hat{d}_k)$  is true otherwise. Notably, both versions may hold simultaneously, depending on the domain clusters.

**Example 5.3** (Example 3.1 ctd). For the mapping  $\{1, \dots, 5\}/k$ , the abstract relation  $k \widehat{\leq} k$  holds true, as  $X \leq Y$  for all  $X, Y$  mapped to  $k$ . Both  $k \widehat{=} k$  and  $k \neg \widehat{=} k$  hold true, as  $X_1 = X_2$  holds only for some  $X_1, X_2$  values mapped to  $k$ .

Notice that having both  $rel$  and  $neg\_rel$  hold means an uncertainty on the truth value of the relation in the abstract clusters. This brings us to determining the types of the relations over the abstract clusters, similar as before.

**Abstract relation types.** The following cases  $\tau_I - \tau_{III}$  occur in a mapping for the abstract relation predicates  $\widehat{rel}(\hat{d}_1, \dots, \hat{d}_k)$  and  $neg\_rel(\hat{d}_1, \dots, \hat{d}_k)$ :

$$\begin{aligned} \tau_I^{rel}(\hat{d}_1, \dots, \hat{d}_k): & \quad \widehat{rel}(\hat{d}_1, \dots, \hat{d}_k) \wedge not \ neg\_rel(\hat{d}_1, \dots, \hat{d}_k) \\ \tau_{II}^{rel}(\hat{d}_1, \dots, \hat{d}_k): & \quad neg\_rel(\hat{d}_1, \dots, \hat{d}_k) \wedge not \quad \widehat{rel}(\hat{d}_1, \dots, \hat{d}_k) \\ \tau_{III}^{rel}(\hat{d}_1, \dots, \hat{d}_k): & \quad \widehat{rel}(\hat{d}_1, \dots, \hat{d}_k) \wedge \quad neg\_rel(\hat{d}_1, \dots, \hat{d}_k) \end{aligned} \quad (41)$$

Type I is the case where the abstraction does not cause uncertainty for the relation, thus the rules that contain  $\widehat{rel}$  with type I can remain the same in the abstract program. Type II shows the cases where  $\widehat{rel}$  does not hold in the abstract domain. Type III is the uncertainty case, which needs to be dealt with when creating the abstract rules. To ensure that an over-approximation is achieved, the head of the respective rule will be changed into a choice.

For an abstraction  $m$ , we compute the set  $\mathcal{T}_{m_{\exists}}$  of all atoms  $\tau_i^{\widehat{rel}}(\hat{d}_1, \dots, \hat{d}_k)$  where  $i \in \{I, II, III\}$  is the type of  $\widehat{rel}(\hat{d}_1, \dots, \hat{d}_k)$  for  $m$ .

### 5.1.1 Abstraction procedure

For simplicity and ease of presentation, we consider programs with rules having (i) a single relation atom; and (ii) no cyclic dependencies between non-ground literals.

**Definition 5.1** (rule abstraction). *Given a rule  $r : l \leftarrow B(r), rel(t_1, \dots, t_k)$  and a domain mapping  $m$ , the set  $r_{\exists}^m$  contains the following rules.*

- (a)  $m(l) \leftarrow m(B(r)), \tau_I^{\widehat{rel}}(\hat{t}_1, \dots, \hat{t}_k)$ .
- (b)  $\{m(l)\} \leftarrow m(B(r)), \tau_{III}^{\widehat{rel}}(\hat{t}_1, \dots, \hat{t}_k)$ .
- (c) For all  $L \subseteq B^-(r)$ :
 
$$\begin{aligned} \{m(l)\} &\leftarrow m(B_L^{sh}(r)), \tau_I^{\widehat{rel}}(\hat{t}_1, \dots, \hat{t}_k), isCluster(\hat{j}). \\ \{m(l)\} &\leftarrow m(B_L^{sh}(r)), \tau_{III}^{\widehat{rel}}(\hat{t}_1, \dots, \hat{t}_k), isCluster(\hat{j}) \end{aligned} \quad j \in arg(l_i), l_i \in L$$

where  $B_L^{sh}(r) = B^+(r) \cup L, not B^-(r) \setminus L$ .

The idea is to introduce guesses when there is an uncertainty over the relation holding in the abstract domain (b), or over the negated atoms due to the abstract clusters (c) (by considering all combinations of the negative literals), and otherwise just abstracting the rule (a).

The abstraction procedure introduced in Definition 5.1 obtains semantically the same abstract program as in Definition 3.3 for rules of form

$$l \leftarrow B(r), rel(t_1, t_2).$$

with binary relations  $=, \neq, \leq, <$ . Formally,

**Definition 5.2** (Existential abstract program  $\Pi_{\exists}^m$ ). *Given a program  $\Pi$  and a domain mapping  $m$ , we denote by*

$$\Pi_E^m = \{r^m \mid r : l \leftarrow B(r), rel(t_1, t_2) \in \Pi\} \cup \{x \mid x \in \mathcal{T}_{m_{\exists}}\} \cup \{m(p(\vec{c})) \mid p(\vec{c}) \in \Pi\}. \quad (42)$$

the program obtained from  $\Pi$  under existential abstraction using  $m$ .

We then have

**Theorem 5.1.** *For any domain mapping  $m$  of a program  $\Pi$ ,  $AS(\Pi^m)$  and  $AS(\Pi_{\exists}^m)$  coincide (modulo auxiliary atoms).*

A generalization to multiple relation atoms and handling cyclic dependencies by removing the restrictions (i)-(ii) can be done similarly as in cases (G-II) and (G-III) of Section 3.3.

**Example 5.4** (Example 3.1 ctd). For the program  $\Pi$  in (5)-(9) and  $m = \{\{1, \dots, 5\}/k\}$ , the program  $\Pi_{\exists}^m$  is:

$$\begin{aligned} c(X) &\leftarrow \text{not } d(X), \tau_1^{\hat{<}}(X, k), \widehat{\text{int}}(X). \\ \{c(X)\} &\leftarrow \text{not } d(X), \tau_{\text{III}}^{\hat{<}}(X, k), \widehat{\text{int}}(X). \\ \{c(X)\} &\leftarrow \tau_1^{\hat{<}}(X, k), \text{isCluster}(X), \widehat{\text{int}}(X). \end{aligned} \tag{43}$$

$$\{c(X)\} \leftarrow \tau_{\text{III}}^{\hat{<}}(X, k), \text{isCluster}(X), \widehat{\text{int}}(X). \tag{44}$$

$$d(X) \leftarrow \text{not } c(X), \widehat{\text{int}}(X).$$

$$\{d(X)\} \leftarrow \text{isCluster}(X), \widehat{\text{int}}(X).$$

$$b(X, Y) \leftarrow a(X), d(Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y).$$

$$e(X) \leftarrow c(X), a(Y), \tau_1^{\hat{<}}(X, Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y).$$

$$\{e(X)\} \leftarrow c(X), a(Y), \tau_{\text{III}}^{\hat{<}}(X, Y), \widehat{\text{int}}(X), \widehat{\text{int}}(Y).$$

$$\perp \leftarrow b(X, Y), e(X_1), \tau_1^{\hat{<}}(X, X_1), \widehat{\text{int}}(X), \widehat{\text{int}}(X_1), \widehat{\text{int}}(Y).$$

and the abstract facts are  $a(k), \widehat{\text{int}}(k)$ ; furthermore, we have  $\mathcal{F}_{m_{\exists}} = \{\tau_{\text{III}}^{\hat{<}}(k, k), \tau_{\text{III}}^{\hat{>}}(k, k), \tau_{\text{III}}^{\hat{<}}(k, k)\}$ . Note that the atom  $d(X)$  is omitted in (43)-(44) instead of shifting the polarity.

The abstract program is similar to the one constructed by lifting the relations in Example 3.19. As can be easily checked, the programs have modulo the auxiliary atoms the same abstract answer sets.

We observe that for treating  $n$ -ary relations where  $n > 2$ , we can modify Definition 5.1 to create finer abstractions.

**Example 5.5.** Consider the argument  $Z$  of the following rule involving addition:

$$r : e(Z) \leftarrow c(X), a(Y), Z = X + Y. \tag{45}$$

We denote  $Z = X + Y$  with the relation  $\text{plus}(X, Y, Z)$ . Regarding the arguments, we have  $\text{arg}(e(Z)) \cap \text{arg}(\text{plus}(X, Y, Z)) = \{Z\} \neq \emptyset$  while  $\text{arg}(e(Z)) \cap \{X, Y\} = \emptyset$ , where  $X, Y$  are the shared arguments of the body literals with the relation  $\text{plus}$ , i.e.,  $\text{arg}(B(r)) \cap \text{arg}(\text{plus}(X, Y, Z)) = \{X, Y\}$ . Consider the mapping  $m : \{1\} \mapsto a_1, \{2, 3\} \mapsto a_{23}, \{4, 5\} \mapsto a_{45}$  and  $X = a_1, Y = a_1$ . For the abstract relation  $\hat{\text{plus}}$ , both  $\hat{\text{plus}}(a_1, a_1, a_{23})$  and  $\text{neg\_}\hat{\text{plus}}(a_1, a_1, a_{23})$  hold true, due to  $1 + 1 = 2$  and  $1 + 1 \neq 3$ . As  $Z$  is not used in the body literals, it does not cause uncertainties for applying the rule in the abstraction, which is caught by

$$e(Z) \leftarrow c(X), a(Y), \tau_{\text{III}}^{\widehat{\text{rel}}(X, Y, Z)}, \text{isSingleton}(X), \text{isSingleton}(Y).$$

In general, by adding in Definition 5.1 the rule

- $m(l) \leftarrow m(B(r)), \tau_{\text{III}}^{\widehat{\text{rel}}}(\hat{t}_1, \dots, \hat{t}_k), \bigwedge_{\hat{t}_i \in \text{arg}_i(\text{rel}) \setminus \text{arg}(l)} \text{isSingleton}(\hat{t}_i)$   
if  $\text{arg}(l) \cap \text{arg}(\text{rel}) \neq \emptyset$  and  $\text{arg}(l) \cap (\text{arg}(B(r)) \cap \text{arg}(\text{rel})) = \emptyset$ .

the guess in (b) can be avoided, if all arguments of  $\text{rel}$  not involved in the head  $l$  are singleton clusters.

The use of abstract relations opens a wide range of possible applications, as it simplifies the use of a given program without preprocessing it to match the restrictions over the forms of the relations for the previous abstraction method.

### 5.1.2 Computing Joint Abstract Relation Types

Abstract relations can be easily employed with abstraction mappings over several sorts in the domain as  $m : D_1 \times \dots \times D_n \rightarrow \hat{D}_1 \times \dots \times \hat{D}_n$ . If a rule has relations over the sorts, a joint abstract relation combining them must be computed. We show an example of grid-cell abstraction for illustration.

**Example 5.6** (Abstracting grid-cells). Consider the relations  $rel_1(X_1, X_2): X_1 = X_2$  and  $rel_2(Y_1, Y_2): Y_1 < Y_2$  for  $X_1, X_2 \in row, Y_1, Y_2 \in column$ , from standardizing apart the variables in (36). The rules to compute the types  $\tau_I^{rel}, \tau_{III}^{rel}$ , where  $\widehat{rel}$  combines  $rel_1$  and  $rel_2$ , are as follows:

1. Define the abstract relations. This step corresponds to the existential abstraction (39).

$$\widehat{rel}_1((\hat{X}_1, \hat{Y}_1), (\hat{X}_2, \hat{Y}_2)) \leftarrow rel_1(X_1, X_2), m((X_1, Y_1), (\hat{X}_1, \hat{Y}_1)), m((X_2, Y_2), (\hat{X}_2, \hat{Y}_2)).$$

$$\widehat{rel}_2((\hat{X}_1, \hat{Y}_1), (\hat{X}_2, \hat{Y}_2)) \leftarrow rel_2(Y_1, Y_2), m((X_1, Y_1), (\hat{X}_1, \hat{Y}_1)), m((X_2, Y_2), (\hat{X}_2, \hat{Y}_2)).$$

The negations  $\neg\widehat{rel}_1, \neg\widehat{rel}_2$  are computed similarly as (40).

$$\neg\widehat{rel}_1((\hat{X}_1, \hat{Y}_1), (\hat{X}_2, \hat{Y}_2)) \leftarrow \neg rel_1(X_1, X_2), m((X_1, Y_1), (\hat{X}_1, \hat{Y}_1)), m((X_2, Y_2), (\hat{X}_2, \hat{Y}_2)).$$

$$\neg\widehat{rel}_2((\hat{X}_1, \hat{Y}_1), (\hat{X}_2, \hat{Y}_2)) \leftarrow \neg rel_2(Y_1, Y_2), m((X_1, Y_1), (\hat{X}_1, \hat{Y}_1)), m((X_2, Y_2), (\hat{X}_2, \hat{Y}_2)).$$

2. Compute the types of each abstract relation  $\widehat{rel}_i, i \in \{1, 2\}$  with the objects  $\hat{C}_i = (\hat{X}_i, \hat{Y}_i), i \in \{1, 2\}$  as (41).

$$\tau_I^{rel_i}(\hat{C}_1, \hat{C}_2) \leftarrow \widehat{rel}_i(\hat{C}_1, \hat{C}_2), not \neg\widehat{rel}_i(\hat{C}_1, \hat{C}_2).$$

$$\tau_{II}^{rel_i}(\hat{C}_1, \hat{C}_2) \leftarrow not \widehat{rel}_i(\hat{C}_1, \hat{C}_2), \neg\widehat{rel}_i(\hat{C}_1, \hat{C}_2).$$

$$\tau_{III}^{rel_i}(\hat{C}_1, \hat{C}_2) \leftarrow \widehat{rel}_i(\hat{C}_1, \hat{C}_2), \neg\widehat{rel}_i(\hat{C}_1, \hat{C}_2).$$

3. Compute the types of the joint abstract relation  $\widehat{rel}$  over  $\widehat{rel}_i, i \in \{1, 2\}$ :

$$\tau_I^{rel}(\hat{C}_1, \hat{C}_2) \leftarrow \tau_I^{rel1}(\hat{C}_1, \hat{C}_2), \tau_I^{rel2}(\hat{C}_1, \hat{C}_2).$$

$$\tau_{III}^{rel}(\hat{C}_1, \hat{C}_2) \leftarrow not \tau_{II}^{rel1}(\hat{C}_1, \hat{C}_2), \tau_{III}^{rel2}(\hat{C}_1, \hat{C}_2).$$

$$\tau_{III}^{rel}(\hat{C}_1, \hat{C}_2) \leftarrow \tau_{III}^{rel1}(\hat{C}_1, \hat{C}_2), not \tau_{II}^{rel2}(\hat{C}_1, \hat{C}_2).$$

The mapping (38) shown in Figure 10c gives the types  $\tau_I^{rel}((a_1, b_1), (a_1, b_2)), \tau_I^{rel}((a_2, b_1), (a_2, b_2))$  and  $\tau_{III}^{rel}$  for the remaining abstract pairs.

Figure 11 presents the multi-dimensional case, that is computing abstract  $k$ -tuple relations for given relations  $rel_1, \dots, rel_{n_1}$  over variables from  $\hat{D}_1, \dots, \hat{D}_{n_2}$ . We assume for simplicity a uniform arity  $k$ .

Note that for the joint abstract relation  $\widehat{rel}$ , type  $\tau_{II}^{rel}$  computation is not needed, as the abstract rule construction only deals with types I and III. To emphasize the abstracted relations, we may denote  $\widehat{rel}$  in  $\tau^{rel}$  with the combination of the relations the abstract relation is built on; e.g., for the joint relation type of  $rel_1(X_1, X_2): X_1 = X_2$  and  $rel_2(Y_1, Y_2): Y_1 < Y_2$  we write  $\tau_I^{\widehat{rel}_1, <}, \tau_{III}^{\widehat{rel}_1, <}$ .

The multi-dimensional abstraction constructs an abstract structure, i.e., *object*, over the abstracted sorts where not all combinations of the abstract sorts yields a valid object. For example, in Example 5.1 the color cluster  $\hat{r}\hat{g}\hat{b}$  can only be considered with the node cluster  $\hat{4}$ . This also needs to be taken into account when constructing the abstract program.

Figure 11: Computation of multi-dimensional relation types ( $\hat{c}_j = (d_j^1, \dots, d_j^{n_2}), 1 \leq j \leq k$ )

$$\begin{aligned}
\widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow rel_i(d_1^i, \dots, d_k^i), m((d_1^1, \dots, d_1^{n_2}), \hat{c}_1), \dots, m((d_k^1, \dots, d_k^{n_2}), \hat{c}_k). & i = 1, \dots, n_1 \\
\tau_I^{rel_i}(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k), not \neg \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k). & i = 1, \dots, n_1 \\
\tau_{II}^{rel_i}(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow not \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k), \neg \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k). & i = 1, \dots, n_1 \\
\tau_{III}^{rel_i}(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k), \neg \widehat{rel}_i(\hat{c}_1, \dots, \hat{c}_k). & i = 1, \dots, n_1 \\
\tau_I^{rel_1, \dots, rel_{n_1}}(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow \tau_I^{rel_1}(\hat{c}_1, \dots, \hat{c}_k), \dots, \tau_I^{rel_{n_1}}(\hat{c}_1, \dots, \hat{c}_k). \\
\tau_{III}^{rel_1, \dots, rel_{n_1}}(\hat{c}_1, \dots, \hat{c}_k) &\leftarrow \tau_{III}^{rel_i}(\hat{c}_1, \dots, \hat{c}_k), \bigwedge_{j=1: j \neq i}^{n_2} not \tau_{II}^{rel_j}(\hat{c}_1, \dots, \hat{c}_k). & i = 1, \dots, n_1
\end{aligned}$$

**Example 5.7.** The abstract program for Sudoku (35)-(37), where the occurrences of  $row(X)$ ,  $column(Y)$  are replaced by  $cell(X, Y)$ , is as follows.

$$\begin{aligned}
hasNum(X, Y) &\leftarrow sol(X, Y, N), cell(X, Y). \\
\{sol(X, Y, N)\} &\leftarrow not occupied(X, Y), num(N), cell(X, Y). \\
\{sol(X, Y, N)\} &\leftarrow occupied(X, Y), num(N), isCluster(X), cell(X, Y). \\
\{sol(X, Y, N)\} &\leftarrow occupied(X, Y), num(N), isCluster(Y), cell(X, Y). \\
\perp &\leftarrow not hasNum(X, Y), cell(X, Y). \\
\perp &\leftarrow sol(X_1, Y_1, M), sol(X_2, Y_2, M), \tau_1^{<}(X_1, Y_1, X_2, Y_2), cell(X, Y_1), cell(X_2, Y_2). \\
\perp &\leftarrow sol(X_1, Y_1, M), sol(X_2, Y_2, M), \tau_1^{<=} (X_1, Y_1, X_2, Y_2), cell(X_1, Y), cell(X_2, Y_2).
\end{aligned}$$

## 5.2 Quad-tree Abstraction

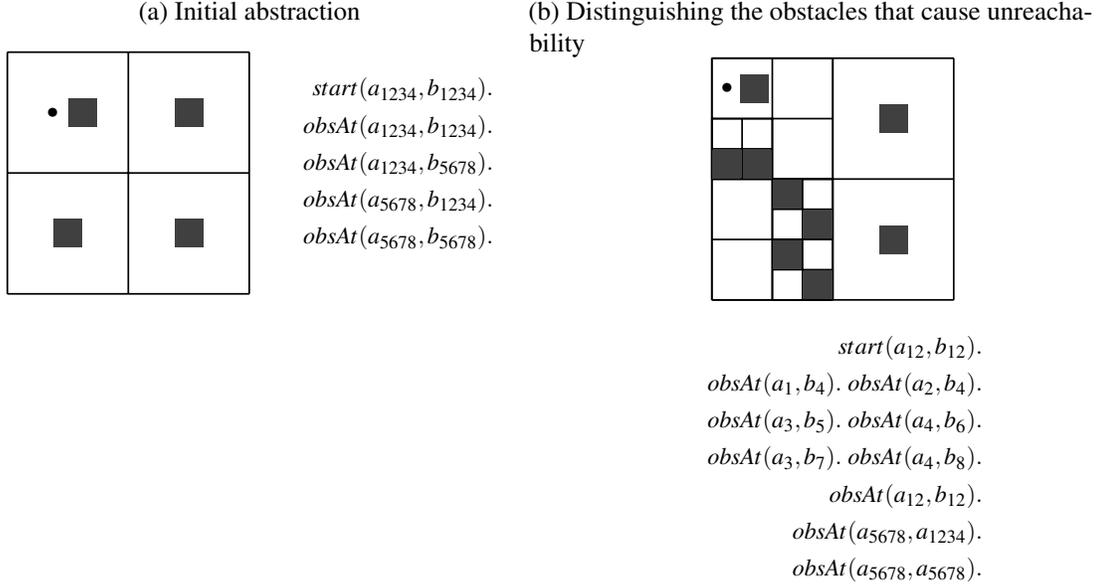
Grid-cell environments are a particular type of environment which describes a structure. For problems over grid-cells, it is often the case that certain parts of the environment are crucial to finding a solution. In order to obtain an abstraction over a grid-cell that allows to adjust its granularity, multi-dimensionality must be considered. Multi-dimensional abstraction allows us to express abstractions where one domain (e.g., an X coordinate) is abstracted depending on its context, i.e., depending on a second domain it occurs with (e.g., a Y coordinate).

For a systematic refinement of abstractions on grid-cell environments, we consider a generic quad-tree representation (Figure 12), which is a concept used, e.g., in path planning (Kambhampati and Davis, 1986). Initially, an environment may be abstracted to four regions of  $n/2 \times n/2$  grid-cells each. This amounts to a tree with four leaf nodes that correspond to the main regions. Each region then contains 4 leaves of smaller regions. The leaves of the quad-tree are then the original cells of the grid-cell at level 0. A refinement of a region amounts to dividing it into four subregions, i.e., sprouting the respective node to four children. Given the original X and Y coordinates  $a_1, \dots, a_n$  and  $b_1, \dots, b_n$ , respectively, we represent the coordinates of an abstract region with level  $\log_2(k+1)$ , for  $0 \leq k < n$ , defined over the cells within the coordinates  $a_i, \dots, a_{i+k}$  and  $b_j, \dots, b_{j+k}$ , respectively, by the shorthand notation  $(a_{i..i+k}, b_{j..j+k})$ .

Starting with an initial abstraction of level  $\log_2(n)$ , using quad-tree split operations as abstraction refinement operations, we can automatically search for suitable quad-tree-structured abstractions in grids (see Section 6.2.2). Importantly, multi-dimensional abstraction refinement is *structure aware*: refining one of



Figure 14: Abstractions over the grid-cell domain with obstacles in Figure 13




---

**Algorithm 2:** *Domain-Abs&Ref*

---

**Input:** program  $\Pi$ , domain mapping  $m_{init}$ , *mode* (lifted or existential)

**Output:**  $\Pi^m$ , a mapping  $m$  refining  $m_{init}$ , and an abstract answer set  $\hat{I}$  of  $\Pi^m$  resp. *null* if  $\Pi^m$  is unsatisfiable.

```

1  $m = m_{init};$ 
2  $[\Pi^m, \mathcal{T}_m] = constructAbsProg(\Pi, m, mode);$ 
3  $\Pi_{debug} = constructDebugProg(\Pi);$ 
4 while  $AS(\Pi^m \cup \mathcal{T}_m) \neq \emptyset$  do
5   Get  $I \in AS(\Pi^m \cup \mathcal{T}_m);$ 
6    $C = getASWithMinAbAtoms(\Pi_{debug} \cup Q_I^m);$ 
7   if  $C|_{ab} = \emptyset$  then  $/* I$  concrete */
8     return  $\Pi^m, m, I$ 
9    $m = decideRefinement(m, C|_{ab});$ 
10   $[\Pi^m, \mathcal{T}_m] = constructAbsProg(\Pi, m, mode);$ 
11 return  $\Pi^m, m, \emptyset$ 

```

---

## 6 Implementation and Evaluation

### 6.1 Overall Methodology

The abstraction and refinement method for domain abstraction is shown in Algorithm 2. In the initialization phase, the abstract program is constructed for the mapping  $m$  depending on the *mode* of abstraction (lifted or existential) on the built-ins (Line 2); the relation types are also computed in this step. After constructing the debugging program (Line 3), some abstract answer set is computed (Lines 5), and its concreteness checked

with a call to the function *getASWithMinAbAtoms* on the non-ground query  $Q_i^m$  (Definition 4.3) for  $m$  and  $I$  (Line 6), which returns an answer set of  $\Pi_{debug} \cup Q_i^m$  with minimal number of *ab* atoms. If the latter is positive (i.e.,  $I$  is spurious),  $m$  is refined given the *ab* atoms found in the check (Line 9) and the abstract program and the relation types are recomputed before looping back to evaluation. Among possible variations of the algorithm, the following are particularly noticeable.

**Abstraction over subdomains/sorts** For abstraction over a particular subdomain resp. sort, Algorithm 2 can be extended with a predicate name  $s$  for the sort in the input and the computation of the abstract program, the relation types and the debugging program focused on the domain elements related with  $s$ .

**Correctness checking for relevant atoms** The correctness checking of abstract answer sets can be confined to the *relevant* atoms describing the solution by constructing  $Q_i^m$  (Line 6) only for these atoms. An abstract answer set will then pass as concrete as long as it describes a concrete solution with respect to them.

**Diverse abstract answer sets** The refinement decision may be made by looking at multiple abstracts answer sets rather than a single one. Lines 5-8 are changed to collect the checking results  $C_1, \dots, C_n$  for abstract answer sets  $I_1, \dots, I_n$ ; then *decideRefinement* can use the collected results  $C_1|_{ab}, \dots, C_n|_{ab}$  (all assumed to be nonempty) for deciding on a refinement. For refinement approach (v1), the one with the least cost (aggregated from the collection) is picked, while refinement approach (v2) decides by choosing the most occurring *refine* atom.

## 6.2 Implementation

The methodology in Section 6.1 has been implemented in tools called DASPAP and mDASPAP based on clingo 5.2.2, Python and the Ouroboros debugging tool (Oetsch et al., 2010), whose MetaTranslator is exploited to obtain a reified program for which then the debugging program is constructed (cf. Section 4); more details are provided in the next subsections. The implementations are online available at <http://www.kr.tuwien.ac.at/research/systems/abstraction/>.

### 6.2.1 DASPAP

The program DASPAP supports abstraction from sorts, but the input program  $\Pi$  must adhere to certain restrictions: each variable in a rule must be guarded by a domain predicate; if the abstraction should be on some subset  $S$  of sorts, then the variables referring to the sorts in  $S$  must be standardized apart. For example, a rule of form

$$a(X) \leftarrow b(X, X_1), c(X_2), d(X_2), X \leq X_1.$$

needs to be converted into

$$a(X) \leftarrow b(X, X_1), c(X_2), d(X_3), X \leq X_1, X_2 = X_3, dom(X), dom(X_1), dom(X_2), dom(X_3).$$

with domain predicate *dom*. In order to support the case of having more than one relation, a syntactic change on the rule has to be made. These relations need to be combined into an auxiliary relation atom which represents the combination of the relations. The above rule needs to be converted into

$$a(X) \leftarrow b(X, X_1), c(X_2), d(X_3), leqEqu4(X, X_1, X_2, X_3).$$

where *leqEqu4*( $X, X_1, X_2, X_3$ ) is an auxiliary atom which holds true whenever the respective relation holds true for its arguments. A basic set of auxiliary relation combinations are built in to the tool, and more can easily be added.

DASPAR is invoked as follows.

```
python daspar.py prog mapping pred ref_type <focus_atoms>
```

Here `prog` contains the original program in the input format and the `mapping` the abstraction mapping information. DASPAR supports abstraction on one sort (see Section 6.2.2 for multi-dimensional tool mDASPAR), thus `pred` should be the name of the sort to be abstracted. The parameter `ref_type` allows to specify whether the refinement should respect an order (1) or not (0). If 1 is given, the refinement step considers only splitting the domain, while when 0 is given the refinement step is unrestricted. The parameter `<focus_atoms>` is an optional input for projection in the correctness check (see Section 6.1).

DASPAR has different settings for picking abstract answer sets and for deciding on a refinement. For the former, by default the first computed answer set is picked. This can be changed to considering a diverse set of abstract answer sets. For deciding on a refinement, the two forms mentioned in Section 4 are implemented. Later, we evaluate the effects of having these different settings in the methodology on the achieved resulting abstractions.

For practical purposes, sorts can be overlapping, provided that all occurrences of a sort are guarded by domain predicates. E.g., the blocksworld has sorts *block* and *time* which both can use integers. Note that this restriction is to aid the machine knowing about the relations of the arguments, which the user implicitly knows when encoding the problem. With this guidance, it becomes clear over which arguments in the rule the abstraction should focus on.

## 6.2.2 mDASPAR

The program mDASPAR extends DASPAR to multi-dimensional domain abstraction. Currently, it handles 2-dimensional abstractions with a quad-tree style refinement process, and it can be applied to problems described over grid-cells. We discuss some challenges of multi-dimensional abstractions that are tackled in the system.

**Abstract objects** A multi-dimensional abstraction creates abstract objects; not all combinations of the abstracted sorts, e.g., *row* and *column*, correspond to a valid object. To avoid such combinations, the constructed abstract program should comply to only using the abstract objects in the rules. For this, mDASPAR post-processes the abstract program and replaces the occurrence of the abstracted sorts with a new object name.

Note that for “grouping” objects automatically and correctly, the system needs some guidance. For a given encoding, humans are capable of detecting the cells implicitly, whereas machine can not do this readily. The user must provide it with some guidelines to recognize the objects, by adjusting the encoding so that the grids are explicitly shown. For this, we impose some syntactic restrictions on the input program, on which the post-processing technique relies.

Given two sorts  $s_1, s_2$  for a 2-dimensional abstraction, the input program should adhere to the following restrictions in order to achieve a correct object naming:

- (1) The rules should have atoms that contain pairs  $X, Y$  of variables where  $X \in s_1, Y \in s_2$ , and
- (2) the domain predicates for sorts  $s_1, s_2$  should be written in the order of the pairs.

If these restrictions are satisfied, then mDASPAR can correctly convert the sort names to the abstract object name *cell*. For example,  $row(X_1), column(Y_1)$  is changed to  $cell(X_1, Y_1)$ .

Figure 15: Input program with the rules (46)-(51)

```

point(X,Y) :- not obsAt(X,Y), row(X), column(Y).
reachable(X,Y) :- start(X,Y), row(X), column(Y).
reachable(X1,Y1) :- reachable(X,Y), point(X1,Y1), neighbor(X2,Y2,X3,Y3),
                    equEqu4(X,X2,X1,X3), equEqu4(Y,Y2,Y1,Y3),
                    row(X), column(Y), row(X1), column(Y1),
                    row(X2), column(Y2), row(X3), column(Y3).
                    :- point(X,Y), not reachable(X1,Y1), X=X1, Y=Y1,
                    row(X), column(Y), row(X1), column(Y1).
neighbor(X,Y,X1,Y1) :- dist1(X,X1), Y=Y1, row(X), column(Y), row(X1), column(Y1).
neighbor(X,Y,X1,Y1) :- X=X1, dist1(Y,Y1), row(X), column(Y), row(X1), column(Y1).

```

**Example 6.1** (ctd). The rule (48) will be standardized apart into

$$reachable(X_1, Y_1) \leftarrow reachable(X, Y), point(X_1, Y_1), neighbor(X_2, Y_2, X_3, Y_3), \\ X = X_2, X_1 = X_3, Y = Y_2, Y_1 = Y_3.$$

Then the multiple relations related with a sort to be converted into an auxiliary relation atom:

$$reachable(X_1, Y_1) \leftarrow reachable(X, Y), point(X_1, Y_1), neighbor(X_2, Y_2, X_3, Y_3), \\ equEqu4(X, X_2, X_1, X_3), equEqu4(Y, Y_2, Y_1, Y_3).$$

The domain predicates for the rule above also need to be written in a format where the pairs  $(X, Y)$ ,  $(X_1, Y_1)$ ,  $(X_2, Y_2)$ , and  $(X_3, Y_3)$  appear together. Figure 15 shows the resulting rules in the input program.

**Relation type computation** When abstracting a rule, mDASPAR gathers the relations in it related with the abstracted sorts and creates an abstract relation atom following the description in Section 5. The relation type facts ( $\tau$ ) are computed using auxiliary programs.

The program mDASPAR is invoked similarly to DASPAR, but with an additional parameter `size` which is the number  $n = 2^k$ . The next example shows the input format of mDASPAR and the created abstract program.

**Example 6.2** (ctd). Figure 15 shows the input program for mDASPAR with the rules (46)-(51) where variables are standardized apart. The program constructed for abstracting over the sorts `row`, `column` is shown in Figure 16, where the occurrence of the sorts are renamed with a new object `cell`. The rules of the original program are numbered  $(r_1, \dots, r_6)$  and the relation atoms in the abstract program are named w.r.t. the rule number. For example, the constraint (51) is numbered  $r_4$ , and the standardization creates the relations  $X = X_1$  and  $Y = Y_1$ ; in the abstraction the joint relation type atom becomes  $\tau_I^{\text{---}}(X, Y, X_1, Y_1)$  for type I. The abstracted constraint containing the Type III relation atom, i.e.,  $\tau_{III}^{\text{---}}(X, Y, X_1, Y_1)$ , in its body is not important as gets omitted in the abstraction.

Furthermore, standardizing apart the variables of the negative literal in (51) relaxes its aim of ensuring that all points are reachable to hold only when the abstraction is refined enough to satisfy the relation atom `relr4(X, Y, X1, Y1, i)`. Having (51) without standardization would ensure it is satisfied in coarser abstractions. We standardized apart the variables of the negative literal as well to obtain more fine-grained abstractions that distinguish the original cells to reach a concrete solution. This makes it easier to visualize the resulting abstractions and understand solutions obtained.

Figure 16: Non-ground abstract program constructed by mDASPAR

```

    point(X,Y) :- cell(X,Y), not obsAt(X,Y).
  { point(X,Y) } :- cell(X,Y), obsAt(X,Y), isCluster(X).
  { point(X,Y) } :- cell(X,Y), obsAt(X,Y), isCluster(Y).
  reachable(X,Y) :- start(X,Y), cell(X,Y).
  reachable(X1,Y1) :- reachable(X,Y), point(X1,Y1), neighbor(X2,Y2,X3,Y3),
    cell(X2,Y2), cell(X3,Y3), cell(X,Y), cell(X1,Y1),
    relr3(X,Y,X2,Y2,X1,Y1,X3,Y3,i).
  { reachable(X1,Y1) } :- reachable(X,Y), point(X1,Y1), neighbor(X2,Y2,X3,Y3),
    cell(X2,Y2), cell(X3,Y3), cell(X,Y), cell(X1,Y1),
    relr3(X,Y,X2,Y2,X1,Y1,X3,Y3,iii).
    :- point(X,Y), cell(X,Y), cell(X1,Y1),
    not reachable(X1,Y1), relr4(X,Y,X1,Y1,i).
  neighbor(X,Y,X1,Y1) :- dist1(X,X1), cell(X,Y), cell(X1,Y1), relr5(Y,Y1,i).
  { neighbor(X,Y,X1,Y1) } :- dist1(X,X1), cell(X,Y), cell(X1,Y1), relr5(Y,Y1,iii).
  neighbor(X,Y,X1,Y1) :- dist1(Y,Y1), cell(X,Y), cell(X1,Y1), relr6(X,X1,i).
  { neighbor(X,Y,X1,Y1) } :- dist1(Y,Y1), cell(X,Y), cell(X1,Y1), relr6(X,X1,iii).

```

### 6.2.3 Implementation Aspects of mDASPAR

**Two-phase debugging.** The multi-dimensionality of the domain mapping gives rise to many possible causes of spuriousness. Debugging the non-ground spuriousness by searching for the answer set with smallest number of *ab* atoms can become more difficult. To handle this, we implemented a two-phase debugging approach. In phase 1, the debugging program  $\Pi_{debug}$  is created by modifying the debugging atoms *ab\_deact*, *ab\_deactCons* of Definition 4.5 to have only the rule name as arguments. We denote this program by  $\Pi_{debug_0}$ . This then results in an easier computation of an answer set with minimal *ab* atoms. In phase 2 a new program  $\Pi_{debug}$  is created according to the original definition, but the *ab* atoms are only created for the rule names or atoms occurring in the *ab* atoms of  $I$ . This way, the search for an optimal answer set focuses on the trouble-making rules/atoms.

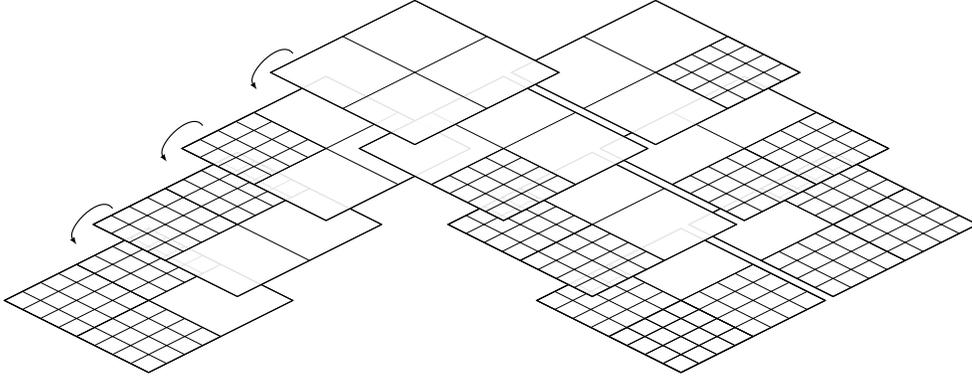
**Steer debugging towards constraints.** In the problems we focus on, the constraints in the program cause to have unsatisfiability or to obtain a particular solution for a given instance. In order to help with reaching abstractions where the relevant constraints are distinguished, we assign less cost to obtaining answer sets with *ab\_deactCons* atoms. in the optimal answer set search during debugging.

**Getting hints.** Since the refinement of a region means to split it into four subregions, we only need to get the hint of which region to refine. This is different from the hints obtained for DASPAR, as there a decision for refinement relies on the domain elements occurring in the debugging atoms. We alter the *refine* atoms to get the information of which abstract domain occurs as a reason for spuriousness.

**Modular concreteness checking.** In some cases, even the two-phase checking may not help with easily finding the optimal answer set during the debugging step as the original domain is large or many atoms cause to consider many possible concretizations We thus considered two orthogonal approaches:

- (1) For programs that are modular and contain a clear order on the atoms (e.g. in a plan), the checking is done incrementally over the approach, similar in spirit to (Gebser et al., 2008) which builds on the concept of modules (Oikarinen and Janhunen, 2006).

Figure 17: Step-wise partial concretization of a grid-cell abstraction



- (2) Using a hierarchy of abstractions (as possible by Proposition 3.6), the checking is done via incrementally concretizing the abstract domain, following an iterative deepening style (Figure 17).

The aim of Approach (1) is to avoid checking the whole ordered sequence of atoms (e.g., a plan), and catching the spuriousness in some prefix. Approach (2) is applied to avoid making the concreteness check directly at the original domain. If the abstract answer set is spurious, this may be detected in the partially concretized domain. We then check correctness of  $I$  on the abstract level  $m_i$  using  $\Pi^{m_i}$ . If  $I$  is concrete w.r.t. the partially concretized abstraction, the concretization is increased for redoing the check. However, if spuriousness is detected, the mapping is refined and the partial concretization continues from the updated mapping.

Further details on the implementation can be found in Appendix B.4.

### 6.3 Evaluation: Obtaining Abstract Solutions

The main aim of our evaluation was to see whether the domain abstraction and refinement method from above can find automatically non-trivial domain abstractions that yield concrete answer sets. We also wanted to observe the effect of variants of picking abstract answer sets (Section 6.1) or making refinement decisions (Section 4).

#### 6.3.1 Experiments

We used DASPARE v0.2 for the experiments, which employs the sophisticated debugging program for concreteness checking.<sup>3</sup> The variations we considered are as follows.

- When computing abstract answer sets, we either **(s)** pick a single abstract answer set or **(div)** pick a (diverse) set of answer sets w.r.t. the focus atoms.
- Deciding on a refinement is either done **(v1)** by assigning costs to possible refinements and picking the one with smallest cost or **(v2)** by using the hints obtained from the debugging atoms while checking. For **(v2)**, it is ensured that the distinct domain elements in the picked debugging atom do not occur in the same cluster.

<sup>3</sup>Results for version v0.1 were reported in (Saribatur et al., 2019).

We conducted experiments on two benchmark problems from ASP competitions, viz. graph coloring and disjunctive scheduling. For graph coloring, we randomly generated 20 graphs on 10 nodes with edge probability 0.1, 0.2, ..., 0.5 each; out of the 100 graphs, 74 were 3-colorable. We used two different graph coloring encodings shown in Figure 18, to see their effect in the resulting abstractions. In the first encoding  $GC_{enc1}$  (Figure 18a), a color assignment to each node is guessed as (52)-(54) with the common approach of using default negation, the auxiliary atom  $hasEdgeTo(X, C)$  shows which colors  $C$  the node  $X$  has as its neighbors. The second encoding  $GC_{enc2}$  (Figure 18b) uses a choice rule (59) to guess an assignment and then ensures with (60) that a node is not assigned more than one color. The rules (61)-(62) are an alternative way of writing the rule  $\perp \leftarrow chosenColor(X_1, C), chosenColor(X_2, C), edge(X_1, X_2), X_1 < X_2$ . so that when the variables are standardized apart for the sort *node*, fewer relation atoms occur in one rule. Also notice that  $GC_{enc2}$  imposes an order relation among the nodes, to reduce duplications of the constraints.

For disjunctive scheduling, for each  $t \in \{10, 20, 30\}$ , we generated 20 instances with 5 tasks over time  $\{1, \dots, t\}$ . We used the encoding<sup>4</sup> from ASP Competition 2011 and precomputed the deterministic part (i.e., not involved in unstratified negation resp. guesses) of the program, so that they are lifted to the abstract program without introducing (unnecessary) nondeterminism (see Appendix B.2). The initial abstraction mapping is the single-cluster abstraction, i.e., clustering all nodes into one for graph coloring and all time points into one for disjunctive scheduling.

In the experiments, we use the lifted relation approach. While in existential abstraction on the relations (Section 5.1) there are fewer relation types to consider, we observed no improvement in the computation effort, since type III for abstract relations is a combination of type III and IV of lifted relations. Thus, it does not make a difference on which relation type approach is used in the experiments.

### 6.3.2 Results

We report the average results over 10 runs for each variation. To ease presentation, we discuss the results for each benchmark separately by concentrating on the different observations made throughout the experimental evaluation.

**Graph coloring.** The evaluation results of the obtained abstractions are presented in Table 3. The first two rows show the average number of refinement steps and the average domain size (i.e., the number of clusters) of the resulting abstractions. The best abstraction (i.e., with smallest domain size) found for each instance in the runs is further checked for faithfulness, to observe whether the corresponding abstract program only contains concrete answer sets. The domain size of the faithful abstractions is shown in the third row. The frequencies of the abstractions that are trivial (thus faithful), non-trivial and faithful, and on-faithful are shown in the last three rows.

The left column shows the results of full concreteness checking with different variations. We can observe that deciding on a refinement based on single abstract answer set (**s**) results in finer abstractions (i.e., with larger domain size) than on diverse set of abstract answer sets (**div**). The number of trivial abstractions obtained is also smaller for (**div**) (better decisions are made) and the chance of encountering a concrete abstract answer set is larger. The latter causes to obtain more non-faithful abstractions, as then no refinement to an abstraction with less spurious answer sets is made. As for using (**v2**), i.e., to decide on refinements, we can observe that this is not better for obtaining coarser abstractions than the minimal cost method (**v1**) in general; it also yields more trivial abstractions than (**v2**), as splitting the domain repeatedly to break up clusters of certain abstract elements quickly ends up with the original domain.

<sup>4</sup>[www.mat.unical.it/aspcomp2011/files/DisjunctiveScheduling/disjunctive\\_scheduling.enc.asp](http://www.mat.unical.it/aspcomp2011/files/DisjunctiveScheduling/disjunctive_scheduling.enc.asp)

Figure 18: Two encodings of the Graph Coloring problem:  $GC_{enc1}$  and  $GC_{enc2}$

$$chosenColor(X, r) \leftarrow not\ chosenColor(X, g), not\ chosenColor(X, y), node(X). \quad (52)$$

$$chosenColor(X, g) \leftarrow not\ chosenColor(X, r), not\ chosenColor(X, y), node(X). \quad (53)$$

$$chosenColor(X, y) \leftarrow not\ chosenColor(X, g), not\ chosenColor(X, r), node(X). \quad (54)$$

$$hasEdgeTo(X, C) \leftarrow edge(X, Y), chosenColor(Y, C). \quad (55)$$

$$\perp \leftarrow hasEdgeTo(X, C), chosenColor(X, C). \quad (56)$$

$$colored(X) \leftarrow chosenColor(X, C). \quad (57)$$

$$\perp \leftarrow node(X), not\ colored(X). \quad (58)$$

(a)  $GC_{enc1}$

$$\{chosenColor(X, C)\} \leftarrow node(X), color(C). \quad (59)$$

$$\perp \leftarrow chosenColor(X, C_1), chosenColor(X, C_2), C_1 \neq C_2. \quad (60)$$

$$adj(X, Y) \leftarrow edge(X, Y), X < Y. \quad (61)$$

$$\perp \leftarrow adj(X, Y), chosenColor(X, C), chosenColor(Y, C). \quad (62)$$

$$colored(X) \leftarrow chosenColor(X, C). \quad (63)$$

$$\perp \leftarrow node(X), not\ colored(X). \quad (64)$$

(b)  $GC_{enc2}$

Table 3: Experimental results for graph coloring

		full				projected			
		(s)		(div)		(s)		(div)	
		(v1)	(v2)	(v1)	(v2)	(v1)	(v2)	(v1)	(v2)
$GC_{enc1}$	number of steps	7.38	7.83	7.04	7.69	5.24	6.48	4.83	6.14
	abstraction domain size	8.38	8.84	8.04	8.69	6.24	7.48	5.83	7.14
	faithful abstraction domain size	6.84	8.04	6.12	7.51	6.02	5.71	5.65	5.82
	trivial abstractions ( <i>id</i> )	13%	23%	4%	12%	2%	1%	2%	2%
	faithful & non- <i>id</i> abstractions	30%	32%	29%	27%	56%	61%	50%	47%
	non-faithful abstractions	57%	45%	67%	61%	42%	38%	48%	51%
$GC_{enc2}$	number of steps	7.01	6.40	6.56	6.37	3.53	3.76	3.40	3.52
	abstraction domain size	8.01	8.64	7.56	8.29	4.53	6.73	4.40	6.36
	faithful abstraction domain size	8.88	8.62	7.97	8.66	4.86	5.44	4.75	5.72
	trivial abstractions ( <i>id</i> )	19%	13%	5%	13%	3%	2%	3%	2%
	faithful & non- <i>id</i> abstractions	22%	24%	25%	22%	54%	59%	54%	48%
	non-faithful abstractions	59%	63%	70%	65%	43%	39%	43%	50%

Table 4: Experimental results for scheduling

time		(s)			(div)		
		(v1)	(v2)	(v2')	(v1)	(v2)	(v2')
$t = 10$	number of steps	7.22	4.81	3.56	6.04	4.81	3.54
	abstract domain size	8.22	8.48	8.46	7.04	8.38	8.35
	calls abstract program	41.35	5.81	4.56	40.72	5.81	4.54
	calls debugging program	40.90	5.36	4.11	56.30	6.87	5.44
$t = 20$	number of steps	14.71	7.65	5.47	12.00	7.53	5.33
	abstract domain size	15.71	14.16	14.12	13.00	14.16	13.81
	calls abstract program	168.48	8.65	6.47	157.41	8.53	6.33
	calls debugging program	168.28	8.45	6.27	244.45	12.08	8.74
$t = 30$	number of steps	22.82	9.57	7.76	20.57	9.56	7.68
	abstract domain size	23.82	19.02	19.12	21.57	19.07	18.68
	calls abstract program	391.88	10.57	8.76	366.09	10.56	8.68
	calls debugging program	391.43	10.12	8.31	580.23	14.59	12.24

The right column shows the results for a projected notion of concreteness that limits checking to a set of relevant atoms; for this, we picked the nodes 1,2,3 and their assigned colors. As expected, a concrete abstract answer set is encountered in much coarser abstractions, as the colors assigned to the other nodes do not matter. In case of projection, the trivial abstraction is reached much less often; moreover, more non-trivial faithful abstractions are reached. This is beneficial, as the computed abstractions can be used to obtain all (concrete) solutions over the nodes focused on.

The main difference of the various encodings is the size of the achieved abstract domains.  $GC_{enc2}$  requires fewer refinement steps to achieve an abstraction with a concrete solution than  $GC_{enc1}$ , as the need to preserve the node ordering leaves fewer refinement possibilities. On average, the resulting abstractions are coarser than by  $GC_{enc1}$  while the domain size of the faithful abstractions are larger. This may be due to the choice rule in  $GC_{enc2}$ , causing spurious answer sets that must be treated by further refinement steps.

**Disjunctive scheduling.** We compared the effects of the variations for the resulting abstractions and the calls to the ASP solver to obtain an abstract answer set respectively to check concreteness with debugging; Table 4 shows the collected results. For the refinement search, we considered besides (v1) and (v2) the variant (v2') of (v2) where each abstract element in the obtained debugging atom is mapped to a singleton cluster in the refinement.

As expected, the minimal cost method (v1) causes much more calls to the ASP solver, as the cost for each possible refinement must be computed. While it achieves coarser abstractions in half of the cases, the large number of calls is a clear disadvantage. For example, for the case  $t = 20$ , (v1) achieves with (div) on average an abstract domain of 13.00 clusters with 400 calls to the ASP solver, while (v2') achieves an average of 13.81 clusters with only around 15 calls.

For the instances with  $t = 20$ , refinement through hints (v2) achieves coarser abstractions than (v1) when single abstract answer sets are picked. Here hints guide the refinement much better than the cost from a single abstract answer set. For the cost from a diverse set of abstract answer sets, significantly coarser abstractions are achieved. Looking at  $t = 30$ , we can observe that the cost approach (v1) results in much finer abstractions than the hint based approaches (v2) and (v2'), which provide better guidance. This shows that a local search over the 1-step refinements does not always yield the best outcome, and it is moreover also more expensive.

We can also observe that (v2') mostly achieves much coarser abstractions than (v2); immediately singling out the domain elements connected with the spuriousness helps. It also needs the smallest number of

refinement steps compared to other approaches, as it reaches a concrete solution much faster with the refinement decisions.

The results show that with larger domains, the effect of the abstraction can be seen much better; e.g., the best abstract domain size reached for  $t = 10$  on average is 70.4% ( $= 7.04/10$ ) of the original domain size, while for  $t = 30$  it shrinks to 62% ( $= 18.68/30$ ).

**Summary.** The results show that with domain abstraction it is possible to achieve concrete solutions while abstracting over some of the details of the program. Reaching faithful abstractions is desired; however it does not occur often, unless a projected concreteness check is considered that only distinguishes the details relevant for a solution of the problem. Obtaining hints from a set of abstract spurious answer sets instead from a single such answer set results in better decisions, and thus coarser abstractions. direction.

## 6.4 Evaluation: Unsolvable Problem Instances in Grid-Cells

We investigated obtaining explanations of unsatisfiable grid-cell problems by achieving an abstraction over the instance to focus on the troubling area. We considered the following benchmark problems:

- *Reachability (R)*: This problem needs the neighboring cell information and can be encoded without introducing guesses. We check whether every cell is reachable; unsatisfiability is due to the layout of the obstacles.
- *Knight's Tour (KT)*: This problem is on finding a tour on which a knight visits each square of a board once and returns to the starting point. It is commonly used in ASP Competitions, with possible addition of forbidden cells. Unsolvability is due to forbidden cells that prevent the knight from moving. In ASP competitions, this problem is encoded by guessing a set of  $move(X_1, Y_1, X_2, Y_2)$  atoms and ensuring that each cell has only one incoming and one outgoing movement.<sup>5</sup> There is no *time* sort (as in planning) which would describe an order.
- *Visitall*: We extended the planning problem of visiting every cell (without revisiting a cell) with obstacles. This problem needs the neighboring cell information and can be encoded in two forms;

(**V**) as a planning problem, in order to find a sequence of actions that visits every cell, or

(**V<sub>KT</sub>**) as a combinatorial problem similar to the Knight's Tour encoding.

To allow for shorter plans, we encoded (**V**) using  $go(X, Y, T)$  actions that can move horizontally/vertically to a cell  $X, Y$  (without passing through obstacles) and the passed cells become visited; we set a limit of 30 time steps.

- *Sudoku (S)*: This problem has also been used in ASP competitions.<sup>6</sup> Its encoding consists of a guess of numbers in the cells combined with simple constraints such as one symbol per column, one symbol per subregion etc. The unsolvability occurs due to violation of these constraints.

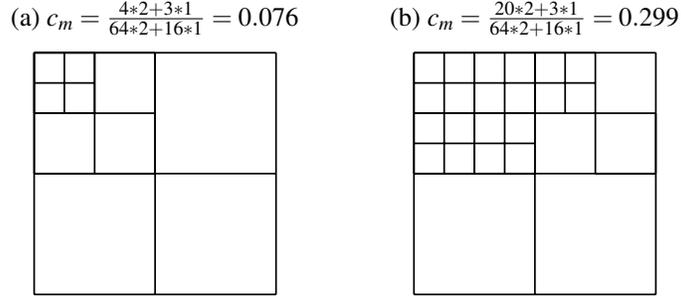
We generated 10 unsatisfiable instances complying to the following properties so that the unsolvability can be explained by focusing on a trouble making area:<sup>7</sup>

<sup>5</sup>[www.mat.unical.it/aspcomp2013/KnightTour](http://www.mat.unical.it/aspcomp2013/KnightTour)

<sup>6</sup>[dtai.cs.kuleuven.be/events/ASP-competition/Benchmarks/Sudoku.shtml](http://dtai.cs.kuleuven.be/events/ASP-competition/Benchmarks/Sudoku.shtml)

<sup>7</sup>All benchmark instances and encodings as well as user explanations for Visitall and Reachability are available at [www.kr.tuwien.ac.at/research/systems/abstraction/mdasparmaterial.zip](http://www.kr.tuwien.ac.at/research/systems/abstraction/mdasparmaterial.zip)

Figure 19: Measure for quality of a quadtree abstraction



- In Reachability instances, a group of neighboring cells is unreachable due to the obstacles surrounding them.
- For Knight’s Tour instances, one or two cells are picked to have only one valid movement to an obstacle-free cell. This way, these cells and the obstacles that do not allow the valid movements become a reason for unsolvability.
- The Visitall instances consist of either two dead-end cells or areas with only one cell passage, so that one is forced to pass some cells more than once, which is not allowed.
- For Sudoku, we generated a layout of numbers that force to violate the constraints when solving the problem.

**Measuring abstraction quality** We consider a *quality measure* of the quad-tree abstraction by normalizing the number of abstract regions of certain size and their level in the quadtree. The cost of a mapping  $m$  over an  $n \times n$  grid is

$$c(m) = \sum_{i=0}^{\ell} r_{2^i}(m)(\ell - i) / \sum_{i=0}^{\ell} n^2 2^{-i^2} (\ell - i),$$

where  $\ell = \log_2(n) - 1$  is the level,  $r_{2^i}(m)$  is the number of abstract regions of size  $2^i \times 2^i$  in  $m$ , and  $n^2 2^{-i^2}$  is the number of abstract regions of size  $2^i \times 2^i$  in the  $n \times n$ -sized cell. The factor  $\ell - i$  is a weight that gives higher cost to abstractions with more low-level regions. An abstraction mapping with the smaller cost, i.e., intuitively smaller level of detail, is considered to be of *better quality*.

Figure 19 shows measures of two abstraction mappings. The abstraction in Figure 19a is coarser than the one in Figure 19b, and this is reflected in the computed measures. Assigning more weight to having coarser regions would stress the importance of having a coarse abstraction even more. The computation of the measure is purely structural and domain-independent. Other measures can be defined that are dependent on the domain which considers further aspects, e.g., such that an abstraction that singles out smallest number of cells with obstacles is preferred.

#### 6.4.1 Effects of Different Debugging Approaches

We compared different debugging approaches from Section 6.2.3 to observe their effects on the resulting abstractions and the taken refinement steps. Due to their encodings and constraints, the Knight’s Tour and Visitall problems are the challenging ones. To observe whether an incremental checking could help in deciding on a refinement and achieve better abstractions, we applied for KT and  $V_{KT}$  partial concretization

Table 5: Evaluation results of applying different debugging approaches

	debugging type	average		minimum		best	
		steps	cost	steps	cost	steps	cost
R	default	5.4	0.227	5.4	0.227	5.0	0.208
	2-phase	5.5	0.233	5.3	0.222		
S	default	6.5	0.696	5.1	0.550	3.2	0.371
	2-phase	4.3	0.476	3.4	0.391		
KT	2-phase	14.3	0.643	10.4	0.460	5.6	0.245
	grid-inc	10.1	0.442	6.3	0.277		
V	2-phase <sup>8</sup>	16.2	0.708	13.9	0.608	8.7	0.360
	time-inc	16.3	0.712	13.5	0.569		
V <sub>KT</sub>	2-phase	15.7	0.693	13.0	0.572	7.6	0.317
	grid-inc	13.0	0.569	10.3	0.449		

and for V incremental *time* checking. To evaluate how far a resulting abstraction from the *best possible* abstraction showings unsolvability, we also checked whether a coarser abstraction with this property exists.

Table 5 shows the main evaluation results. We compare different debugging approaches in terms of the average refinement steps and average costs of the resulting abstractions over 10 runs, and also on the best outcome obtained (with minimum refinement steps and minimum mapping cost). The two right-most columns concern the existence of a coarser abstraction for best outcome obtained. The time to find an optimal solution in the debugging step was limited by 50 seconds. If exceeded, the refinement is decided on the basis of suboptimal analyses by considering the optimal debugging solution that could be computed within the time limit.

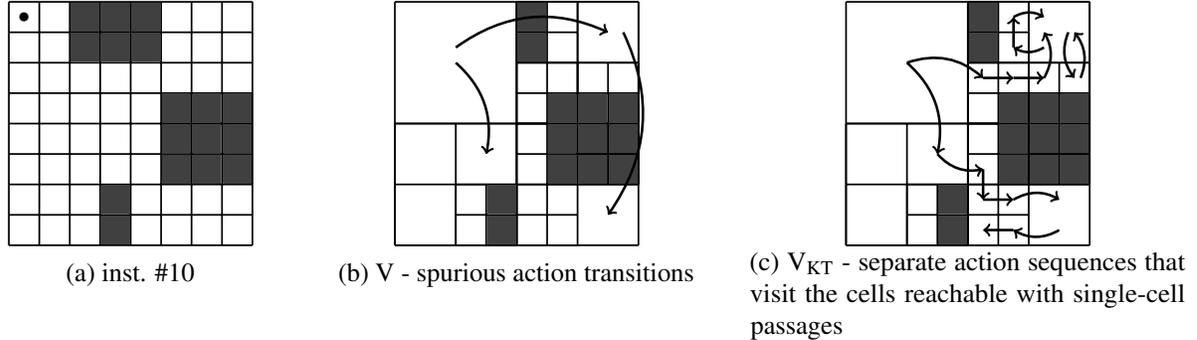
For Reachability and Sudoku, we observe that abstractions close to the best possible ones can be obtained. Better abstractions were obtained with 2-phase debugging in these cases (the majority with a clear margin for S), as after the first step the focus was on the right part of the abstraction. For Knight’s Tour and Visitall, we observe that incremental checking can obtain better abstractions. This is because for 2-phase debugging, the programs mostly had due to timeouts to decide on suboptimal concreteness checking outputs. Moreover, for the V encoding 2-phase debugging caused memory outages (limit 500 MB) on some runs for some instances, thus not all 10 runs could be completed.

We can also see a difference of the resulting abstractions for the two encodings of Visitall. The planning encoding causes to achieve unsatisfiability with less coarse abstractions. Guesses of spurious sequences of actions in the abstraction cause the debugging to decide on refinements that avoid these sequences. The focus moves towards the unsolvability when the abstract action sequence is not executable due to an obstacle. In some instances where the reason for unsolvability is not easily caught by having two dead-ends, focusing on the existence of obstacles does not achieve unsatisfiability: the abstract encoding manages to find a plan passing through different sized regions by avoiding the constraints due to uncertainty. For these instances, the abstraction needs to be fine-grained enough to get rid of most of the uncertainty.

Figure 20a shows such an example of such an instance. An abstraction that distinguishes the one-passage-entries and the obstacles that surround the cells cannot achieve unsatisfiability for encoding V. Figure 20b shows some spurious action transitions that are determined in a plan found with the V encoding among

<sup>8</sup>A total of 16 runs could not be completed due to memory errors. The results are computed among the runs that have been completed.

Figure 20: Spurious plans in abstractions that distinguish the single-cell passages



regions by avoiding the constraints due to uncertainty. Unsatisfiability cannot be achieved for  $V_{KT}$  as well. This is due to guess of *move* atoms which achieve that every cell is visited, but lack a corresponding original order of movements. Figure 20c shows the spurious order of movements that is split to visit each cell that is only reachable through a one-passage-entry. If the abstraction is refined to distinguish the cells in the respective corners, then unsatisfiability is realized.

## 6.4.2 User Study on Unsatisfiability Explanations

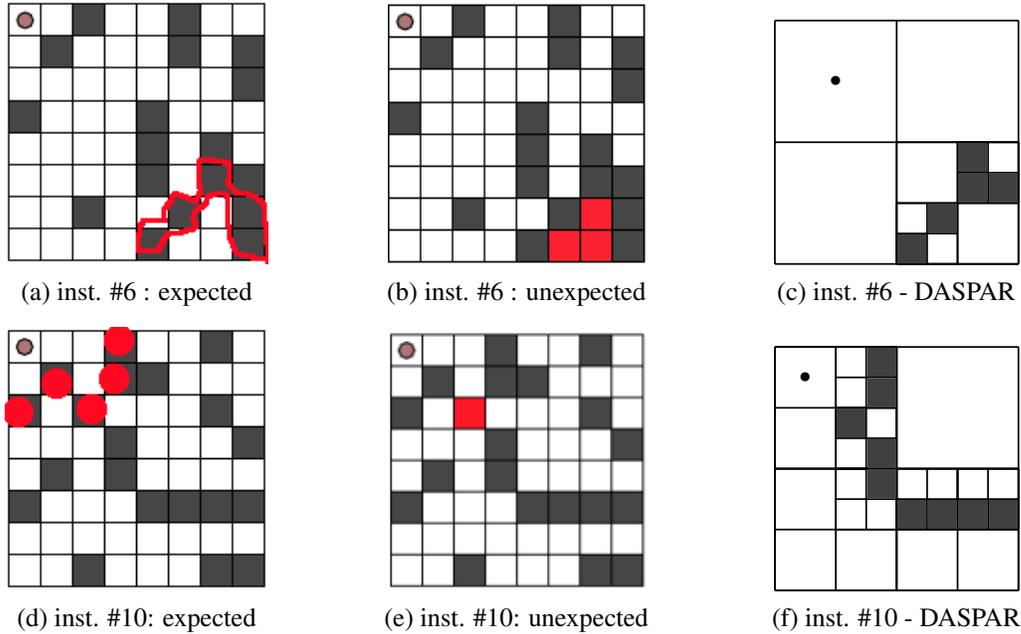
We were interested in checking whether the obtained abstractions match the intuition behind a human explanation. For Reachability and Visitall, finding the reason for unsolvability of an instance is possible by looking at the obstacle layout. Thus, we conducted a user study for these problems in order to obtain the regions that humans focus on to realize the unsolvability of the problem instance.

As participants, we had ten PhD students of Computer Science at TU Wien. We asked them to mark the area which shows the reason for having unreachable cells in the Reachability instances and the reason for not finding a solution that visits all the cells in the Visitall instances; multiple reasons are to mark with different colors. Explanations for 10 instances of each problem were collected.<sup>7</sup> We discuss the results for both problems by showing two of the responses (expected and unexpected) and the best abstraction obtained from mDASPAR when starting with the initial mapping.

**Reachability.** The expected explanations (e.g., Figures 21a and 21d) focus on the obstacles that surround the unreachable cells, as they prevent them from being reachable. When their respective abstraction mappings are given to mDASPAR, the constructed abstract program is also unsatisfiable. The explanation in Fig. 21b puts the focus on the unreachable cells themselves, and Fig. 21e distinguishes a particular obstacle as a reason. When the respective abstraction mappings are given to mDASPAR, it needs to refine further to distinguish more obstacles and achieve unsatisfiability. The mark in Figure 21e is a possible solution to the unreachability of the cells, since removing the marked obstacle makes all the cells reachable.

In ASP, checking whether all cells are reachable is straightforward, without introducing guesses. This is also observed to be helpful for mDASPAR, as most of the resulting abstractions were similar to the gathered answers. Since in the initial abstraction, the abstract program only knows that the agent is located in the upper-left abstract region, in instance #10, mDASPAR follows a different path in refining the abstraction, and reaches the abstraction shown in Figure 21f. Although different from the one by them, by the users, it also shows a reason for having unreachable cells. Humans use the implicit knowledge that the agent is located in the upper-left corner in order to determine the reason for unreachability of the cells, and thus focus on a

Figure 21: Explanatory abstractions for unsatisfiable Reachability instances



different area than mDASPAR. Such an abstraction can also be achieved with the method, by influencing the refinement decisions towards singling out the initial location of the agent.

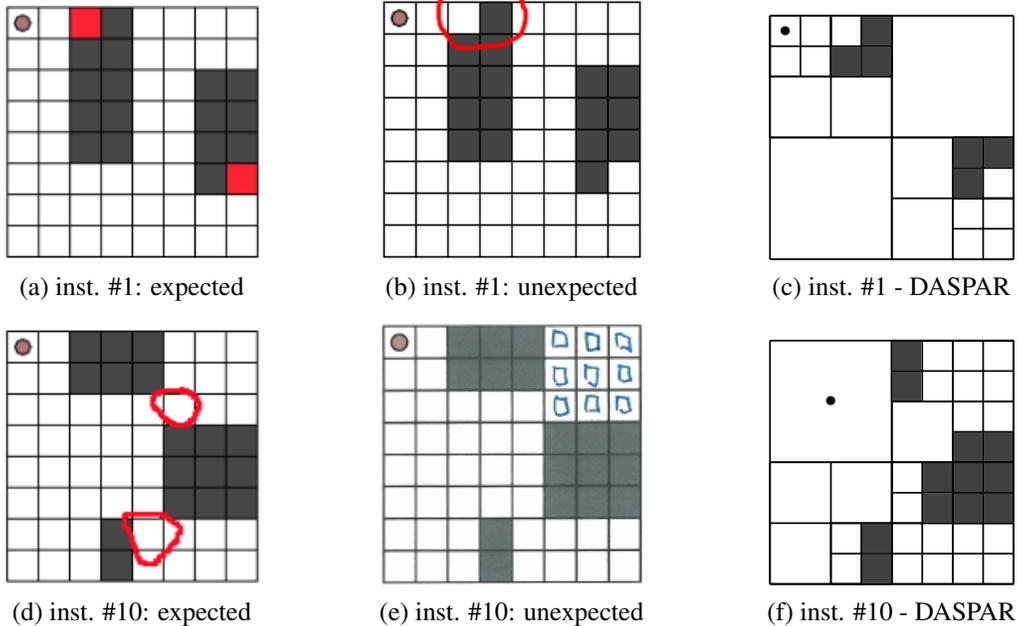
The abstractions achieved by mDASPAR are more general as the precise initial location of the agent is immaterial to distinguish the unreachable cells: it can be in any of the cells mapped to the respective abstract region. The precise obstacle layout in the abstracted regions also plays no role in determining the unreachability of the distinguished cells.

**Visittal.** Most of the users picked two dead-end cells in the instances (if such occur) as an explanation for unsatisfiability, instead of the obstacles surrounding these cells (see Fig. 22a), which are the actual cause for them being dead-end cells. Even with abstraction mappings that also distinguish the surrounding obstacles of these dead-end cells, the corresponding abstract program still remains satisfiable. Then mDASPAR needs to further refinement to distinguish the neighboring cells (as in Fig. 22c) and to realize that it can only pass through one grid-cell when reaching the dead-end cells, and thus achieve unsatisfiability. Unexpectedly, some users marked only one of the dead-end cells as an explanation (Fig. 22b), which is actually focusing on a possible solution to the unsolvability, since if the marked area was not a dead-end, all cells could have been visited.

Some instances do not contain two dead-end cells, but single-cell passages to some regions. Fig. 22d shows an entry that distinguishes these passages, while again focusing only on the cells themselves. For these instances, the results of mDASPAR are quite different. We discussed this already for Figure 20 in Section 6.4.1. Figure 22f shows the best abstraction achieved for  $V_{KT}$ . It distinguishes all cells in the one-passage-entry regions to realize that a desired action sequence can not be found.

The generality of the achieved abstractions is also here observable: the precise agent position and obstacle layout in the abstracted areas do not change the unsatisfiability result as not all cells in the distinguished parts can be visited.

Figure 22: Explanatory abstractions for unsatisfiable Visital instances



**Observations** The abstraction method can demonstrate the capability of human-like focus on certain parts of the grid to show the unsolvability reason. However, humans implicitly also use their background knowledge and do not need to explicitly state the relations among the objects. Empowering machines with such capabilities remains a challenge. The study also showed a difference in understanding the meaning of “explanation”. For some study participants, showing how to get rid of unsolvability was also seen as an explanation. This discrepancy shows that one needs to clearly specify what is wanted (e.g., “mark only the obstacles that cause to have unreachable cells”), to achieve less variety of results.

## 7 Discussion

### 7.1 Predicate Abstraction

Predicate abstraction in ASP would introduce literals involving new predicates that describe an abstraction of original literals, and rewriting the program to mention only the new literals. Naively replacing literals with the abstract ones would not always achieve an over-approximation.

**Example 7.1** (Example 3.1 ctd). Consider a predicate abstraction that maps the atoms  $a(X)$  and  $d(X)$  to  $ad(X)$ , i.e., if  $a(d)$  or  $d(d)$  holds true for some  $d \in D$ , then  $ad(d)$  holds true. When we replace the  $a$ - and

$d$ -atoms in Example 3.1, we obtain the following program:

$$\begin{aligned}
c(X) &\leftarrow \text{not } ad(X), X < 5, \text{int}(X). \\
ad(X) &\leftarrow \text{not } c(X), \text{int}(X). \\
b(X, Y) &\leftarrow ad(X), ad(Y), \text{int}(X), \text{int}(Y). \\
e(X) &\leftarrow c(X), ad(Y), X \leq Y, \text{int}(X), \text{int}(Y). \\
\perp &\leftarrow b(X, Y), e(X), \text{int}(X), \text{int}(Y). \\
ad(1). \quad ad(3). \quad \text{int}(0), \dots, \text{int}(5).
\end{aligned}$$

However, each answer set of this program contains  $b(5, 1), b(5, 3)$  and it matches no original answer set.

A simple way to achieve predicate abstraction is via domain abstraction after reification of predicates of the original program. For example,  $p(X, Y)$  is written to  $x(p, X, Y)$  and a sort of predicate names (viewed as constants) is introduced. By standardizing apart the variables, predicate names can be clustered via the built-in relations.

**Example 7.2 (ctd).** Rewriting the atoms  $a(X)$  and  $d(X)$  to  $x(a, X)$  and  $x(d, X)$ , respectively. yields the program

$$\begin{aligned}
c(X) &\leftarrow \text{not } x(d, X), X < 5, \text{int}(X). \\
x(d, X) &\leftarrow \text{not } c(X), \text{int}(X). \\
b(X, Y) &\leftarrow x(P_1, X), x(P_2, Y), P_1 = a, P_2 = d, \text{pred}(P_1), \text{pred}(P_2), \text{int}(X), \text{int}(Y). \\
e(X) &\leftarrow c(X), x(P, Y), P = a, \text{pred}(P), X \leq Y, \text{int}(X), \text{int}(Y). \\
\perp &\leftarrow b(X, Y), e(X), \text{int}(X), \text{int}(Y). \\
x(a, 1). \quad x(a, 3). \quad \text{int}(0), \dots, \text{int}(5). \\
\text{pred}(a). \quad \text{pred}(d).
\end{aligned}$$

Then an abstraction  $m$  over the sort  $\text{pred}$  such as  $\{a, d\}/ad$  can be applied.

This approach works for predicate abstraction where the corresponding literals have arguments from the same sort in the same argument position. In case a literal has fewer arguments, dummy values can be used to fill in the remaining argument positions.

Another way to achieve predicate abstraction is by following the motivation behind existential abstraction of the relations. The idea is to introduce a new set of predicates along with their relation types according to the abstraction; then the abstract rules will be formed for all combinations of the abstraction types in the bodies, where choice is added to the head unless all are type I.

**Example 7.3 (ctd).** Similar to Example 7.1, the abstract predicate name  $ad$  is introduced with the relation type  $\tau_{\text{III}}^{ad}$  and the additional fact  $isCluster^{ad}$ . Note that the arguments of the literal are not important, as the

abstraction is over the predicate name and not over domain elements. The abstract program is then as follows:

$$\begin{aligned}
c(X) &\leftarrow \text{not } ad(X), X < 5, \text{int}(X). \\
\{c(X)\} &\leftarrow ad(X), \text{isCluster}^{ad}, X < 5, \text{int}(X). \\
ad(X) &\leftarrow \text{not } c(X), \text{int}(X). \\
\{b(X, Y)\} &\leftarrow ad(X), ad(Y), \tau_{\text{III}}^{ad}, \text{int}(X), \text{int}(Y). \\
\{e(X)\} &\leftarrow c(X), ad(Y), \tau_{\text{III}}^{ad}, X \leq Y, \text{int}(X), \text{int}(Y). \\
&\perp \leftarrow b(X, Y), e(X), \text{int}(X), \text{int}(Y). \\
&ad(1). \quad ad(3). \quad \text{int}(0), \dots, \text{int}(5).
\end{aligned}$$

We remark that this approach is similar to using the rewriting of the original program with reification of predicates, and applying existential abstraction on the relations.

## 7.2 Use Case: Abstraction in ASP Planning

Domain abstraction gives us the possibility to adjust the granularity of a problem towards the relevant details. By achieving abstract answer sets that are concrete and thus catch all the relevant details, it also allows for problem solving over abstract notions, which can be useful in a wide range of applications. We discuss here the possible use of domain abstraction in ASP planning, in particular, in understanding planning problems expressed in ASP by abstracting over the unnecessary details. Another use case about policy refutation is described in Appendix C.

Planning problems in ASP are represented by using a `time` sort to describe the sequence of states and the changes according to actions taken (Lifschitz, 1999). There are usually two types of objects, represented with different sort types:

- objects on which the actions have a direct effect, e.g., the blocks in the blocksworld which can be moved, and
- other objects unaffected by the actions but involved in the decision making, e.g., the table in the blocksworld onto which blocks can be moved.

We discuss how to use domain abstraction over these objects such that we can talk about abstract states and plans.

**Describing Actions in ASP** First, we emphasize on the blocksworld example the different ways of expressing planning problems in ASP. The effects of moving a block on top of another block are expressible by the following rules:

$$onB(B, B_1, T + 1) \leftarrow moveToBlock(B, B_1, T). \quad (65)$$

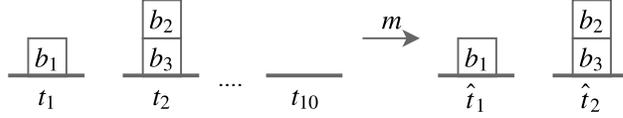
$$\neg onB(B, B_2, T) \leftarrow onB(B, B_1, T), B_1 \neq B_2. \quad (66)$$

where (65) models the direct and (66) the indirect effect. Alternatively, all effects are expressed as direct effects by

$$onB(B, B_1, T + 1) \leftarrow moveToBlock(B, B_1, T). \quad (67)$$

$$\neg onB(B, B_2, T) \leftarrow moveToBlock(B, B_1, T), B_1 \neq B_2. \quad (68)$$

Figure 23: Initial state of a blocksworld with multiple tables (concrete  $\xrightarrow{m}$  abstract).



The preconditions of an action can either be described through constraints, or as a condition for an action to become applicable. For example, that a block cannot sit on a block with a smaller name can be expressed as a constraint

$$\perp \leftarrow onB(B, B_1, T), B_1 \leq B.$$

Alternatively, the respective action can be forbidden if the condition is not satisfied using the following rules:

$$\begin{aligned} \perp &\leftarrow moveToBlock(B, B_1, T), not\ precondmtb(B, B_1, T). \\ precondmtb(B, B_1, T) &\leftarrow B < B_1, block(B), block(B_1). \end{aligned}$$

Note that the alternative version is much closer to the PDDL-style encoding. The law of inertia is described by the rule

$$onB(B, B_1, T+1) \leftarrow onB(B, B_1, T), not\ \neg onB(B, B_1, T+1).$$

### 7.2.1 Abstracting over Irrelevant Details

We first show the possibility of abstraction over the details of the objects that are indirectly affected by the actions. For demonstration, we consider two extensions of well-known planning domains.

- *Multi-table blocksworld (MTB)*: here blocks can be moved onto one of multiple tables (where each table can hold multiple blocks); a plan is needed that piles the blocks up on a given specific table.
- *Package delivery with checkpoints (PDC)*: packages must be carried from an initial to a goal location, while passing through a checkpoint reachable from the initial location.

**Multi-table blocksworld.** Figure 23 illustrates an instance of MTB where the blocks must be piled up on table  $t_1$  such that  $b_1$  is above  $b_2$  and  $b_2$  is above  $b_3$ . Here reaching the goal state does not depend on the concrete tables to which blocks are moved before moving them to the goal table. However, when computing a plan based on the original program, the planner has to consider all possible movements.

Figure 24 shows a (natural) encoding of the problem with the actions  $moveToT(B, Ta, T)$  and  $moveToB(B, B', T)$  for moving block  $B$  onto table  $Ta$  and onto block  $B'$ , resp., at time  $T$ . Consider the initial state shown in Figure 23:

$$onT(b_1, t_1), onB(b_2, b_3), onT(b_3, t_2), chosenTable(t_1).$$

After ensuring that all variables are guarded by domain predicates and those related with the *table* sort are standardized apart, we run DASPAR with the initial mapping  $\{\{t_1, \dots, t_n\}/\hat{t}\}$ . The abstraction obtained is shown in Figure 23; it singles out the chosen table  $\hat{t}_1$  and clusters all others into  $\hat{t}_2$ . We then can compute a concrete abstract answer set

$$\{moveToT(b_2, \hat{t}_2, 0), moveToT(b_3, \hat{t}_1, 1), moveToB(b_2, b_3, 2), moveToB(b_1, b_2, 3)\}.$$

Figure 24: Encoding for Multi-table Blocksworld

```

% action choice
{moveToB(B,B1,T) : bl(B),bl(B1);moveToT(B,L,T) : bl(B),tbl(L)} ≤ 1 ← T < tmax.

% no gaps between moves
done(T) ← moveToB(B,B1,T).
done(T) ← moveToT(B,L,T).
⊥ ← done(T+1),not done(T).
% preconditions
⊥ ← moveToB(B,B2,T),onB(B1,B,T).
⊥ ← moveToB(B1,B2,T),onB(B1,B2,T).
⊥ ← moveToT(B,L,T),onB(B1,B,T).
⊥ ← moveToT(B,L,T),onT(B,L,T).
% effects
onB(B,B1,T+1) ← moveToB(B,B1,T),T < tmax.
onT(B,L,T+1) ← moveToT(B,L,T),T < tmax.
¬onB(B,B2,T) ← onB(B,B1,T),B1 ≠ B2.
¬onT(B,L,T) ← onB(B,B1,T).
¬onB(B,B1,T) ← onT(B,L,T).
¬onT(B,L1,T) ← onT(B,L2,T),L1 ≠ L2.

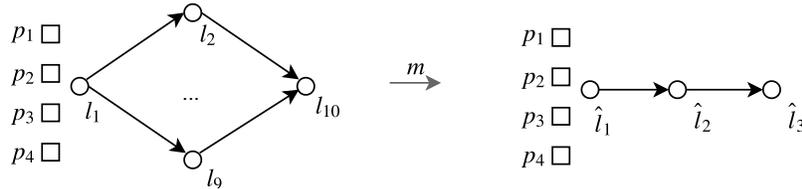
% inertia
onB(B,B1,T+1) ← onB(B,B1,T),
not ¬onB(B,B1,T+1),T < tmax.
onT(B,L,T+1) ← onT(B,L,T),
not ¬onT(B,L,T+1),T < tmax.

% state constraints
⊥ ← onB(B1,B,T),onB(B2,B,T),B1 ≠ B2.
⊥ ← onB(B,B1,T),onB(B,B2,T),B1 ≠ B2.
⊥ ← onB(B,B1,T),onT(B,L,T).
⊥ ← onB(B,B1,T),B1 ≤ B.

% goal constraints
notblockgoal(T) ← onT(B,L,T),onT(B1,L1,T),B ≠ B1.
⊥ ← notblockgoal(T),maxTime(T).
⊥ ← not notblockgoal(T),onT(B,L,T),not chosenTable(L).

```

Figure 25: Initial state of a package delivery with checkpoints (concrete  $\xrightarrow{m}$  abstract).



that describes a plan without going into detail on which table the blocks are moved. The abstraction shows that for solving the problem, it is essential to distinguish the picked table from all others; the number of tables is irrelevant. Furthermore, this abstraction is faithful for projection to the actions *moveToB*, *moveToT*.

**Package delivery with checkpoints.** Figure 25 illustrates an instance of PDC where the packages in location  $l_1$  must be carried to location  $l_{10}$ . As these locations are not directed connected, the truck has to pass through a middle point; through which point the truck passes is immaterial for reaching the goal state.

For this problem, we used the Nomystery encoding from ASPCOMP2015 and altered it to have no fuel computation. Furthermore, for a *drive*( $T, L_1, L_2, S$ ) action to be possible we added an additional condition that the locations  $L_1$  and  $L_2$  should be connected by an edge,  $edge(L_1, L_2)$ . Consider the initial state shown in Figure 25:  $\{atT(t, l_1, 0), atP(p_1, l_1, 0), atP(p_2, l_1, 0), atP(p_3, l_1, 0), atP(p_4, l_1, 0), goal(p_1, l_{10}), goal(p_2, l_{10}), goal(p_3, l_{10}), goal(p_4, l_{10})\}$  with the depicted *edge* facts. Running DASPARG with the mapping  $\{\{l_1, \dots, l_{10}\}/\hat{l}\}$  over the sort *location* results in the abstraction mapping  $\{\{l_1\}/\hat{l}_1, \{l_2, \dots, l_9\}/\hat{l}_2, l_{10}/\hat{l}_3\}$

Figure 26: Encoding for Package Delivery

```

% action choice
{unload(P,T,L,S) : package(P),truck(T),loc(L);
 load(P,T,L,S) : package(P),truck(T),loc(L);
 drive(T,L1,L2,S) : edge(L1,L2),loc(L1),loc(L2),truck(T) } ≤ 1 ← step(S),S > 0.

% no gaps between moves
done(S) ← unload(P,T,L,S).
done(S) ← load(P,T,L,S).
done(S) ← drive(T,L1,L2,S).
⊥ ← done(S+1),not done(S).
% effects
atP(P,L,S) ← unload(P,T,L,S).
¬in(P,T,S) ← unload(P,T,L,S).
¬atP(P,L,S) ← load(P,T,L,S).
in(P,T,S) ← load(P,T,L,S).
¬atT(T,L1,S) ← drive(T,L1,L2,S).
atT(T,L2,S) ← drive(T,L1,L2,S).

% precondition check
⊥ ← unload(P,T,L,S),not precondn(P,T,L,S).
precondn(P,T,L,S) ← atT(T,L,S-1),
                    in(P,T,S-1).
⊥ ← load(P,T,L,S),not precondl(P,T,L,S).
precondl(P,T,L,S) ← atT(T,L,S-1),
                    atP(P,L,S-1).
⊥ ← drive(T,L1,L2,S),not precondn(T,L1,L2,S).
precondn(T,L1,L2,S) ← atT(T,L1,S-1).
% inertia
atT(T,L,S) ← atT(T,L,S-1),not ¬atT(T,L,S).
atP(P,L,S) ← atP(P,L,S-1),not ¬atP(P,L,S).
in(P,T,S) ← in(P,T,S-1),not ¬in(P,T,S).
% goal check
⊥ ← goal(P,L),not atP(P,L,S),maxstep(S).

```

(shown in Figure 25). With this abstraction, the following concrete abstract answer set is computed:

$$\{load(p_4,t,\hat{l}_1,1),load(p_3,t,\hat{l}_1,2),load(p_1,t,\hat{l}_1,3),load(p_2,t,\hat{l}_1,4),$$

$$drive(t,\hat{l}_2,\hat{l}_1,5),drive(t,\hat{l}_1,\hat{l}_3,6),$$

$$unload(p_3,t,\hat{l}_3,7),unload(p_1,t,\hat{l}_3,8),unload(p_4,t,\hat{l}_3,9),unload(p_2,t,\hat{l}_3,10)\}.$$

It describes a plan that loads all the packages, moves to the middle cluster location, moves to the goal location, and unloads the packages; the resulting abstraction is faithful for the projection to the actions *load,unload,drive*.

Domain abstraction simplified details that are unimportant for the essence of whether the plan is feasible. The faithful abstraction gives an understanding of the problem by realizing its neuralgic points. If however there are further constraints over details needed to construct a plan, then faithfulness might not be achievable in a non-trivial abstraction.

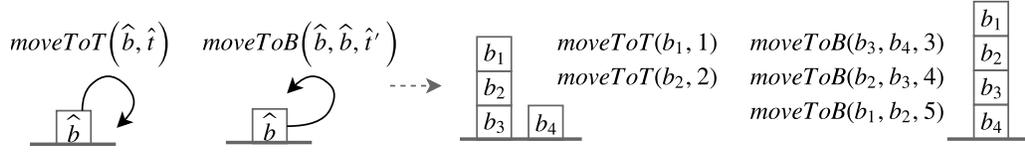
## 7.2.2 Computing Abstract Plans

Abstracting over the objects directly affected by the actions would empower us to talk about abstract plans. However, in ASP-style encodings, abstracting over the object sort only causes the abstract program to compute plans with the original *time* sort. For example, say in the Package Delivery problem (with no checkpoints and two locations  $l_1, l_2$ ) we cluster the packages into one abstract package,  $\hat{p}_0$ . Then, the abstract program will have the abstract actions  $load(\hat{p}_0,t,l,s), unload(\hat{p}_0,t,l,s)$  which then lead to a plan

$$load(\hat{p}_0,t,l_1,1),drive(t,l_1,l_2,2),unload(\hat{p}_0,t,l_2,3).$$

However, this plan is clearly spurious as it is no original action can match  $load(\hat{p}_0,t,l_1,1)$ , which loads all packages in one step; thus many spurious answer sets will result. In order to avoid this, also abstraction over

Figure 27: Abstract and concrete plan of Example 7.4



the *time* sort is necessary. By doing this, we can talk about abstract instances of actions and abstract from the concrete order of their application. Given that the sorts (i.e., blocks, time respectively packages,time) are independent, multiple calls of DASPARG to abstract over each sort one-by-one achieves the desired abstract program.

For the Package Delivery problem, consider two abstraction mappings  $m_{package} = \{\{p_1, p_2, p_3, p_4\}/\hat{p}\}$  and  $m_{time} = \{\{1, 2, 3, 4\}/\hat{t}_1, \{5\}/\hat{t}_2, \{6, 7, 8, 9\}/\hat{t}_3\}$ . The constructed abstract program yields the abstract plan

$$load(\hat{p}, t, l_0, \hat{t}_1), drive(t, l_1, l_2, \hat{t}_2), unload(\hat{p}, t, l_2, \hat{t}_3)$$

which abstracts over the order of package (un)loading and includes abstract actions over time clusters.

Unfortunately, finding a suitable abstraction over multiple sorts, especially if one is over the *time* domain, is non-trivial. The abstraction over time via time clusters steers the plan computation and the action ordering. For example, for the time mapping  $\{\{1\}/\hat{t}_1, \{2, 3\}/\hat{t}_2, \{4, 5, 6, 7, 8, 9\}/\hat{t}_3\}$  the abstract plan from above is spurious. A policy may help in deciding the next sequence of actions, which we demonstrate in the next example.

**Example 7.4** (Blocksworld with a Policy). Consider the blocksworld problem in Figure 27 modify the encoding in Figure 24 by dropping the table argument from *onT* and by standardizing apart with respect to sorts block and time. Suppose the rules realize a 2-phase policy that first puts all blocks on the table and then piles them up:

$$\begin{aligned} existsOnBlock(T) &\leftarrow onB(B, B_1, T). \\ allOnTable(T) &\leftarrow not\ existsOnBlock(T), time(T). \\ atPhase2(T_1) &\leftarrow allOnTable(T), T < T_1. \\ 1\{moveToT(B, T) : onB(B, B_1, T)\} &\leftarrow T < t_{max}, not\ atPhase2(T), not\ allOnTable(T). \\ 1\{moveToB(B, B_1, T) : onT(B, T), block(B_1)\} &\leftarrow T < t_{max}, allOnTable(T). \\ 1\{moveToB(B, B_1, T) : onT(B, T), onB(B_1, B_2, T)\} &\leftarrow T < t_{max}, atPhase2(T). \end{aligned}$$

Given the initial state  $\{onT(b_4, 1), onT(b_3, 1), onB(b_2, b_3, 1), onB(b_1, b_2, 1)\}$  and time domain  $\{1, \dots, 6\}$ , we abstract using the block mapping  $\{\{b_1, \dots, b_4\}/\hat{b}\}$  and the time mapping  $\{\{1, 2\}/\hat{t}, \{3, \dots, 6\}/\hat{t}'\}$ . The abstract program has eight answer sets, among them

$$\{moveToT(\hat{b}, \hat{t}), onT(\hat{b}, \hat{t}), onB(\hat{b}, \hat{b}, \hat{t}), onT(\hat{b}, \hat{t}'), onB(\hat{b}, \hat{b}, \hat{t}'), moveToB(\hat{b}, \hat{b}, \hat{t}')\},$$

which contains two abstract actions, viz.  $moveToT(\hat{b}, \hat{t})$  and  $moveToB(\hat{b}, \hat{t}')$ .

### 7.3 Related Work

In the context of logic programming, abstraction has been considered many years back in the classic work of Cousot and Cousot (Cousot and Cousot, 1992). However, the focus of their studies was on the use of abstract interpretations and termination analysis of programs, and moreover stable semantics was not addressed.

The work most related to abstraction in ASP are the simplification methods that strive for preserving the semantics. Such methods have been extensively studied over the years; we give here an overview of some notions. Notice that, different from these simplification methods, abstraction may lead to an over-approximation of the models (answer sets) of a program, which changes the semantics, in a modified language.

Over-approximation by abstraction reduces the vocabulary which makes it different from *relaxation* methods (Lin and Zhao, 2004; Giunchiglia et al., 2004). These methods translate a ground program into its completion (Clark, 1978) and search for an answer set over the relaxed model. As they focus only on ground programs, they can be compared with the abstraction that omits atoms from the program, which does not need account for loop formulas when searching for a concrete abstract answer set. However, finding the reason for spuriousness of an abstract answer set is trickier than finding the reason why a model of the program completion is not an answer set of the original program, since the abstract answer set contains fewer atoms and a search over the original program is needed to detect the reason why no matching answer set can be found.

### 7.3.1 Equivalence-based Rewriting and Program Transformations in ASP

Equivalence of logic programs is considered under answer set semantics as follows: a program  $\Pi_1$  is equivalent to a program  $\Pi_2$ , if  $AS(\Pi_1) = AS(\Pi_2)$ . *Strong equivalence* (Lifschitz et al., 2001) is a much stricter condition:  $\Pi_1$  and  $\Pi_2$  are strongly equivalent if, for any set  $R$  of rules,  $\Pi_1 \cup R$  and  $\Pi_2 \cup R$  are equivalent. This notion makes it possible to replace a part  $Q$  of a logic program by a strongly equivalent (simpler) program  $Q'$ , without looking at the rest; (Osorio et al., 2002; Turner, 2003; Eiter et al., 2004; Pearce, 2004) show ways of transforming programs by ensuring that the property holds. A more liberal notion is *uniform equivalence* (Maher, 1986; Sagiv, 1987) where  $R$  is restricted to a set of facts; that is  $Q$  and  $Q'$  are equivalent with respect to all factual inputs (Eiter and Fink, 2003).

In terms of abstraction, there is the abstraction mapping that needs to be taken into account, since the constructed program may contain a modified language and the mapping may relate it back to the original language. Thus, to define equivalence between the original program  $\Pi$  and its abstraction  $\widehat{\Pi}^m$  according to a mapping  $m$ , we need to compare  $m(AS(\Pi))$  with  $AS(\Pi^m)$ . The equivalence of  $\Pi$  and  $\widehat{\Pi}^m$  then becomes similar to the notion of faithfulness. However, as we have shown, even if the abstract program  $\widehat{\Pi}^m$  is faithful, refining  $m$  may lead to an abstract program having spurious answer sets. Thus, simply lifting the current notions of equivalence to abstraction may not achieve useful results.

Refinement-safe faithfulness, however, would allow one to use of  $\widehat{\Pi}^m$  instead of  $\Pi$ , as it preserves the answer sets. This property is achieved when abstract program is unsatisfiable (which then implies that the original program was unsatisfiable). However, for original programs that are consistent, reaching an abstraction that is refinement-safe faithful is not easy; dividing the domain cluster may immediately cause a guess that introduces spurious solutions.

Other transformation methods, especially to help with grounding and solving of ASP programs, were investigated. A preprocessing technique was considered in (Gebser et al., 2008) along with an assignment and a relation expressing equivalences among the parts of the program that could be assigned. Another form of preprocessing in (Morak and Woltran, 2012; Bichler et al., 2016) was applied to each rule of a program by computing a tree decomposition and then splitting the rule into multiple, smaller rules accordingly.

### 7.3.2 Abstraction in Planning and Agent Verification

Starting from the early years of AI planning, applications of abstraction to help with the search and planning for complex domains have received a lot of attention. One main research focus has been on hierarchical planning, which considers different abstraction levels over the problem space. A plan is searched at the abstract level and then the solution is refined successively to more detailed levels in the abstraction hierarchy, until a concrete plan is computed at the original level. Sacerdoti (Sacerdoti, 1974) showed an abstraction notion that keeps the “critical” preconditions of actions and ignores the rest. For example, Knoblock (Knoblock, 1994) proposed an ordered monotonicity property to ensure that solving the subproblems by refining certain parts of the plan does not change the remainder of the abstract plan. A similar property was considered by Bacchus and Yang (Bacchus and Yang, 1994), which states that if the original problem is solvable, then any abstract solution must have a refinement. Anderson and Farley (Anderson and Farley, 1988) constructed operator hierarchies by having classes of operators that share common effects and forming new abstract operators with the shared preconditions.

Another research focus has been on using abstractions to compute heuristics, which are estimates of the distances to the solution that guide the plan search. Pattern databases (Culberson and Schaeffer, 1998) are constructed from the results of projecting the state space to a set of variables of the planning task, called a pattern, which is to be solved optimally. The omission abstraction in (Saribatur and Eiter, 2018) matches the intuition behind this projection notion. Edelkamp (Edelkamp, 2001) was the first to apply this technique in planning. He showed that a pre-compiled look-up table with the costs of abstract solutions can help the heuristic search in finding optimal solutions. The merge & shrink abstraction method of Helmert et al. (Helmert et al., 2007) starts with a suite of single projections and then computes an abstraction by merging them and shrinking. A CEGAR-inspired method was proposed by Seipp and Helmert (Seipp and Helmert, 2013) based on cartesian abstractions, which form a general class of abstractions. The reason for the abstract plan being spurious is detected when trying to construct a concrete plan, and the abstraction is refined by splitting the states. Obtaining such a cartesian abstraction is also possible with domain abstraction introduced in Section 3; we further empower abstraction with a multi-dimensional handle in Section 5 that has the capability of representing a hierarchy of abstraction levels.

Although not investigated in detail, notions related to domain abstraction were also considered in heuristic-search planning. Hernádvolgyi and Holte (Hernádvolgyi and Holte, 1999) presented a domain abstraction notion over the states which are represented as fixed length vectors of labels; they also noted the possibility of encountering spurious states with some abstractions. Hoffman et al. (Hoffmann et al., 2006) considered variable domain abstraction by modifying the add and delete lists of the operators accordingly. They argued that obtaining efficient results from abstraction in planning mostly relies on the how much irrelevance is in the problem; this is an observation we similarly made in our experiments. To further investigate the structure of problems that can obtain good results, especially in the context of ASP, is an interesting research direction.

The notion of irrelevant information and its effects were analyzed for planning by Nebel et al. (Nebel et al., 1997), in which different heuristics were introduced to omit such information. Fox and Long (Fox and Long, 1999) described a method for detecting symmetries in a problem which are then treated as indistinguishable to help the planner.

Abstraction was studied for situation calculus action theories by Banihashemi et al. (Banihashemi et al., 2017), who imposed a bisimulation restriction on the abstraction in order to ensure that reasoning about the actions of an agent at the abstract level can be mapped to concrete reasoning. They later showed how this restricted notion of abstraction can be used in reasoning about a strategy for an agent to achieve a goal at the high level and then mapping it back into a low-level strategy (Banihashemi et al., 2018). However, their

focus was not on how such an abstraction can be found.

For verifying the behavior of multi-agent systems, the use of abstraction has been investigated by Lomuscio et al. for abstracting over each agent to construct an abstract system while preserving the properties expressed in a temporal-epistemic logic (Cohen et al., 2009) or alternating-time temporal logic (Lomuscio and Michaliszyn, 2014). In (Cohen et al., 2009) the focus is not on how such an abstraction can be built. In (Lomuscio and Michaliszyn, 2014), an abstraction over the states is made that have the same possible actions to execute and action abstraction keeps the actions of certain agents while omitting the rest.

They considered a three-valued logic and the abstraction also preserves the behavior of not satisfying a property. Spuriousness may occur for the case of achieving an “uncertain” result for checking a specification in the abstract level, which then forces one to refine the abstraction by splitting the states by investigating the subformulas of the specification. They later extended this work to infinite state models (Lomuscio and Michaliszyn, 2016) and abstracted them to finite models using predicate abstraction, and they presented an interpolant-based refinement method (Belardinelli et al., 2016).

In the context of ASP and action languages, Dix et al. (Dix et al., 2003) proposed a way of formulating and solving hierarchical planning under the ASP semantics, with a focus on *ordered task decomposition*, which is planning each step in the order it will later be executed. For a particular application of mobile robot planning, Zhang et al. (Zhang et al., 2015) performed hierarchical planning using the action language  $\mathcal{BC}$ .

### 7.3.3 Generalized Planning

Finding a plan that can achieve the goal for a class of problem instances can give an understanding of the details relevant for these problems. The plan can then be used for any particular problem instance without the need for further search. Note that, as discussed in Section 7.2, the plans that are computed with our domain abstraction method can also be seen as generalized plans as they work for any original problem instance that maps to the abstract instance.

Srivastava et al. (Srivastava et al., 2011) proposed an abstraction method for constructing generalized plans with loops, by focusing on classical planning; however, selecting a good abstraction was beyond their scope. Bonet and Geffner (Bonet and Geffner, 2015) considered a setting where uncertainty is represented by a set of states, by clustering the states that provide the same observations. indistinguishability notion we proposed in (Saribatur and Eiter, 2016). They studied the conditions for a policy (i.e., plan) to be general enough to work on other instances. Later they considered also trajectory constraints (Bonet et al., 2017).

Illanes and McIlraith (Illanes and McIlraith, 2016) studied abstraction for numeric planning problems by compiling them into classical planning. Recently, they used abstraction for problems with quantifiable objects (Illanes and McIlraith, 2019), e.g., a number of packages to deliver to points A and B, to find by abstracting from the quantification generalized plans that work for multiple instances. For this, they built a quantified planning problem by clustering indistinguishable objects using reformulation techniques (Riddle et al., 2016) to reduce symmetry, and then compute a general policy. While the quantifiability conditions of (Illanes and McIlraith, 2019) restrict applicability, our method has the potential drawback spurious answers.

## 8 Conclusion

Abstraction is an important aspect of Artificial Intelligence aiming at the omission of detail and fine-grained structure in problem solving, to reduce the cognitive and/or the computational complexity in order to better understand respectively effectively find solutions. In this spirit, we have introduced the notion of domain abstraction to Answer Set Programming, where the size of the domain of an ASP program shrinks while the

collection of its answer sets is over-approximated, i.e., every original answer set can be mapped to some abstract answer set. We have shown how this can be applied to single or multiple sorts of the domain, and how multi-dimensionality can be handled that enables a hierarchical view of abstraction, with quad-tree abstractions as a showcase for multi-granular abstraction over grids.

More specifically, we have introduced two structure-preserving approaches that apply abstraction to the rules. The first approach keeps built-in relations in rules, which then must be lifted to the abstract domain, while the second approach, existential abstraction, loses their original format at the benefit of the ability to handle different levels of abstractions among the abstract elements, as needed in hierarchical abstraction.

As over-approximation may result in spurious abstract answer sets that do not correspond to original answer sets, we have presented a method for refining abstractions which uses ASP-debugging techniques to obtain hints for refinements and a CEGAR-style methodology of iterated abstraction refinement (Clarke et al., 2003). The approach has been implemented in tools DASPAR and mDASPAR (for multi-dimensional domain mappings), which given an ASP program and an initial abstraction, automatically refine it until for the induced abstract program either a concrete answer set is encountered or unsatisfiability is detected (which proves that the original program is unsatisfiable).

Our experiments showed the potential of the approach for understanding the core parts of an ASP program. In case of satisfiability, abstract answer sets focus on relevant details, as in case of planning problems such as Blocksworld and Package Delivery; a justification technique in ASP (cf. (Pontelli et al., 2009; Cabalar et al., 2014)) can be used to understand why a particular abstract answer set is computed, and moreover, if the abstraction is faithful, to identify details which are irrelevant for finding a solution. In case of unsatisfiability, the automatic abstraction refinement was able to catch the unsatisfiability without refining back to the original program. Furthermore, in grid-cell problems, a multi-dimensional view of abstraction enables zooming in to the area of the grid-cell which shows the reason for unsolvability; compared to the results of a small user study, explanations of decent quality were achieved, which suggests to continue this line of research.

## 8.1 Outlook

This article has provided seminal concepts and notions for domain abstraction in ASP, an assessment of semantic and computational properties, and results for a prototypical evaluation. The work on domain abstraction can be continued and extended in several directions.

One direction is to obtain more general notions of abstraction, and to apply abstraction to larger classes of ASP programs. As for the former, we remark that domain abstraction can be combined with omission abstraction (Saribatur and Eiter, 2018) to obtain an abstraction that omits certain details and also abstracts over some part of the domain. This can be achieved with the current definitions, by first applying the desired domain abstraction to the program and then grounding the constructed non-ground abstract program to omit some of the atoms from it. The refinement decisions then need to take into account two causes for spuriousness; bad clustering of domain elements or bad omission of atoms. As regards larger classes of ASP programs, further language constructs like disjunction in rule heads, aggregates, or weak constraints as in the ASP Core-2 standard (Calimeri et al., 2019) are natural targets. Furthermore, extensions with nested rules, external atoms, or constraint solving are interesting other target languages.

Another research direction concerns abstraction refinement. Different methods can be explored to help with the decision making in the refinement step. On the one hand, further heuristics for deciding about a refinement from a collection of abstract answer sets may be considered, where the range for local search may be increased and in addition domain-specific knowledge is exploited. On the other hand, by using justification methods such as (Pontelli et al., 2009; Cabalar et al., 2014) we can obtain an explanation of how an abstract

answer set is built and check it on the original program. In case of failure, a reason for the spuriousness of the abstract justification may be distilled and exploited for abstraction refinement.

Related to this is the issue of abstraction assessment, where the question of what is a “good” abstraction needs to be further studied. Different criteria can be relevant in this respect, from technical ones like the degree of spuriousness (measured e.g. by the number of spurious answer sets) or the level of abstraction (measured e.g. by the granularity of clustering), to the cognitive appeal from a human user perspective in terms of understandability of the abstraction and the abstract program. Addressing the latter appears to be challenging and harder than developing measures for technical criteria, given that humans have implicit background knowledge about the domain.

A further research direction is advanced implementations beyond tuning the current prototypes. An apparent bottleneck is concreteness checking, which can be costly due to the need for grounding the original program in this process. Here one could explore the use of lazy-grounding, e.g. (Palù et al., 2009; Cat et al., 2015; Lefèvre et al., 2017; Weinzierl, 2017) and non-ground ASP solving, e.g. (Marple et al., 2017; Arias et al., 2018), or develop native techniques for concreteness checking. Further improvements may be using justification methods to explain abstract answer sets (as mentioned above), which as a further benefit mitigates the grounding issue and can yield substantial gains, or to embed the debugging program into the evaluation of the abstract program, such that hints can be obtained during the embedded checking; however, to achieve that this works efficiently is non-trivial.

Last but not least, the use of domain abstraction remains to be explored for applications. Different possibilities can be envisaged, with ASP program development as an obvious candidate. Different from common debugging techniques, domain abstraction aims to not just show the rules themselves that effect a certain behavior, but can moreover be used to identify the gist of the domain that is responsible for the latter and thus aids in gaining more insight into a program at hand. Another possible use of domain abstraction is as a solving technique to address scalability. While the state-of-the-art ASP solvers are quite efficient in solving problems, they may struggle with problems that create huge search spaces or require optimization. For such problems, abstraction could be useful. However, achieving a good abstraction that could help with solving is non-trivial, and advances in performance and in particular of concreteness checking would be necessary for a fruitful deployment. Finally, we believe that domain abstraction has potential for building systems that explain matters to a human end user, and thus can be a useful tool for realizing explainable AI. Our experiments with grid cell puzzles have nurtured this view, since the reasons for unsolvability obtained in an automated way by our tools are a good match with human intuition. However, this is just an initial step, and significant research efforts will have to be invested to make this view become reality.

## References

- ANDERSON, J. S. AND FARLEY, A. M. 1988. Plan abstraction based on operator generalization. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI 1988)*. 100–104.
- ARIAS, J., CARRO, M., SALAZAR, E., MARPLE, K., AND GUPTA, G. 2018. Constraint answer set programming without grounding. *TPLP 18*, 3-4, 337–354.
- BACCHUS, F. AND YANG, Q. 1994. Downward refinement and the efficiency of hierarchical problem solving. *Artificial Intelligence 71*, 1, 43–100.

- BACKSTROM, C. AND JONSSON, P. 1995. Planning with abstraction hierarchies can be exponentially less efficient. In *Proceedings of the 14th International Joint conference on Artificial intelligence (IJCAI 1995) - Volume 2*. Morgan Kaufmann Publishers Inc., 1599–1604.
- BANIHASHEMI, B., DE GIACOMO, G., AND LESPÉRANCE, Y. 2017. Abstraction in situation calculus action theories. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*. 1048–1055.
- BANIHASHEMI, B., DE GIACOMO, G., AND LESPÉRANCE, Y. 2018. Abstraction of agents executing online and their abilities in the situation calculus. In *Proceedings of the 27th International Joint conference on Artificial intelligence (IJCAI 2018)*. 1699–1706.
- BELARDINELLI, F., LOMUSCIO, A., AND MICHALISZYN, J. 2016. Agent-based refinement for predicate abstraction of multi-agent systems. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*. IOS Press, 286–294.
- BICHLER, M., MORAK, M., AND WOLTRAN, S. 2016. The power of non-ground rules in answer set programming. *Theory and Practice of Logic Programming* 16, 5-6, 552–569.
- BISTARELLI, S., CODOGNET, P., AND ROSSI, F. 2002. Abstracting soft constraints: Framework, properties, examples. *Artificial Intelligence* 139, 2, 175–211.
- BONET, B., DE GIACOMO, G., GEFFNER, H., AND RUBIN, S. 2017. Generalized planning: Non-deterministic abstractions and trajectory constraints. In *Proceedings of the 26th International Joint conference on Artificial intelligence (IJCAI 2017)*. 873–879.
- BONET, B. AND GEFFNER, H. 2015. Policies that generalize: Solving many planning problems with the same policy. In *Proceedings of the 24th International Joint conference on Artificial intelligence (IJCAI 2015)*. AAAI Press.
- BRAIN, M., GEBSER, M., PÜHRER, J., SCHAUB, T., TOMPITS, H., AND WOLTRAN, S. 2007. Debugging asp programs by means of asp. In *Proceedings of the 9th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2007)*. Springer, 31–43.
- BRASS, S. AND DIX, J. 1997. Characterizations of the disjunctive stable semantics by partial evaluation. *The Journal of Logic Programming* 32, 3, 207–228.
- BREWKA, G., EITER, T., AND TRUSZCZYŃSKI, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54, 12, 92–103.
- BREWKA, G., EITER, T., AND TRUSZCZYŃSKI, M., Eds. 2016. *Answer Set Programming [special issue]*. AI Magazine, vol. 37(3). AAAI Press.
- CABALAR, P., FANDINNO, J., AND FINK, M. 2014. Causal graph justifications of logic programs. *Theory and Practice of Logic Programming (TC)* 14, 4-5, 603–618.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., MARATEA, M., RICCA, F., AND SCHAUB, T. 2019. ASP-Core-2 Input Language Format.
- CAT, B. D., DENECKER, M., BRUYNOOGHE, M., AND STUCKEY, P. J. 2015. Lazy model expansion: Interleaving grounding with search. *J. Artif. Intell. Res.* 52, 235–286.

- CHITTARO, L. AND RANON, R. 2004. Hierarchical model-based diagnosis based on structural abstraction. *Artificial Intelligence* 155, 1-2, 147–182.
- CLARK, K. L. 1978. Negation as failure. In *Logic and data bases*. Springer, 293–322.
- CLARKE, E., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. 2003. Counterexample-guided abstraction refinement for symbolic model checking. *Journal of the ACM* 50, 5, 752–794.
- CLARKE, E. M., GRUMBERG, O., AND LONG, D. E. 1994. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 1512–1542.
- CLARKE, E. M., HENZINGER, T. A., VEITH, H., AND BLOEM, R. 2018. *Handbook of model checking*. Springer.
- COHEN, M., DAM, M., LOMUSCIO, A., AND RUSSO, F. 2009. Abstraction in model checking multi-agent systems. In *International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. 945–952.
- COUSOT, P. AND COUSOT, R. 1992. Abstract interpretation and application to logic programs. *The Journal of Logic Programming* 13, 2, 103–179.
- CRAIK, K. J. W. 1952. *The nature of explanation*. Vol. 445. CUP Archive.
- CULBERSON, J. C. AND SCHAEFFER, J. 1998. Pattern databases. *Computational Intelligence* 14, 3, 318–334.
- DAMS, D., GERTH, R., AND GRUMBERG, O. 1997. Abstract interpretation of reactive systems. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 19, 2, 253–291.
- DANTSIN, E., EITER, T., GOTTLÖB, G., AND VORONKOV, A. 2001. Complexity and expressive power of logic programming. *ACM Computing Surveys (CSUR)* 33, 3, 374–425.
- DIX, J., KUTER, U., AND NAU, D. 2003. Planning in answer set programming using ordered task decomposition. In *Annual Conference on Artificial Intelligence*. Springer, 490–504.
- DODARO, C., GASTEIGER, P., MUSITSCH, B., RICCA, F., AND SHCHEKOTYKHIN, K. 2015. Interactive debugging of non-ground asp programs. In *Logic Programming and Nonmonotonic Reasoning (LPNMR)*. Springer, 279–293.
- EDELKAMP, S. 2001. Planning with pattern databases. In *Proceedings of the 6th European Conference on Planning (ECP 2001)*. 13–24.
- EITER, T., FABER, W., FINK, M., AND WOLTRAN, S. 2007. Complexity results for answer set programming with bounded predicate arities and implications. *Annals of Mathematics and Artificial Intelligence* 51, 2-4, 123.
- EITER, T. AND FINK, M. 2003. Uniform equivalence of logic programs under the stable model semantics. In *International Conference on Logic Programming*. Springer, 224–238.
- EITER, T., FINK, M., TOMPITS, H., AND WOLTRAN, S. 2004. Simplifying logic programs under uniform and strong equivalence. In *Proceedings of the 7th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2004)*, V. Lifschitz and I. Niemelä, Eds. Springer, 87–99.

- EITER, T., GOTTLÖB, G., AND GUREVICH, Y. 1996. Normal forms for second-order logic over finite structures, and classification of NP optimization problems. *Ann. Pure Appl. Logic* 78, 1-3, 111–125.
- EITER, T., SARIBATUR, Z. G., AND SCHÜLLER, P. 2019. Abstraction for zooming-in to unsolvability reasons of grid-cell problems. In *Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence (XAI)*. 7–13. online available at [https://drive.google.com/file/d/1ma5wilaj31A0d5KC4I2fYaTC\\_Lqm\\_d9X/view](https://drive.google.com/file/d/1ma5wilaj31A0d5KC4I2fYaTC_Lqm_d9X/view) and <http://arxiv.org/abs/1909.04998>.
- ERDEM, E., GELFOND, M., AND LEONE, N. 2016. Applications of answer set programming. *AI Magazine* 37, 3, 53–68.
- ERDEM, E. AND PATOGLU, V. 2018. Applications of ASP in robotics. *KI* 32, 2-3, 143–149.
- FALKNER, A. A., FRIEDRICH, G., SCHEKOTIHIN, K., TAUPE, R., AND TEPPAN, E. C. 2018. Industrial applications of answer set programming. *KI* 32, 2-3, 165–176.
- FANDINNO, J. AND SCHULZ, C. 2019. Answering the “why” in answer set programming - A survey of explanation approaches. *Theory and Practice of Logic Programming* 19, 2, 114–203.
- FOX, M. AND LONG, D. 1999. The detection and exploitation of symmetry in planning problems. In *Proceedings of the 16th International Joint conference on Artificial intelligence (IJCAI 1999)*. Vol. 99. 956–961.
- FREUDER, E. C. 1991. Eliminating interchangeable values in constraint satisfaction problems. In *AAAI*. Vol. 91. 227–233.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T., AND THIELE, S. 2008. Engineering an incremental ASP solver. In *Proceedings of the 24th International Conference on Logic Programming (ICLP 2008)*. 190–205.
- GEBSER, M., KAUFMANN, B., NEUMANN, A., AND SCHAUB, T. 2008. Advanced preprocessing for answer set solving. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*. Vol. 178. IOS Press, 15–19.
- GEBSER, M., PÜHRER, J., SCHAUB, T., AND TOMPITS, H. 2008. A meta-programming technique for debugging answer-set programs. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI 2008)*. Vol. 8. 448–453.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9, 3, 365–385.
- GIUNCHIGLIA, E., LIERLER, Y., AND MARATEA, M. 2004. Sat-based answer set programming. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI 2004)*. 61–66.
- GIUNCHIGLIA, F. AND WALSH, T. 1992. A theory of abstraction. *Artificial Intelligence* 57, 2-3, 323–389.
- GOTTLÖB, G., LEONE, N., AND VEITH, H. 1999. Succinctness as a source of complexity in logical formalisms. *Ann. Pure Appl. Logic* 97, 1-3, 231–260.

- HELMERT, M., HASLUM, P., HOFFMANN, J., ET AL. 2007. Flexible abstraction heuristics for optimal sequential planning. In *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS 2007)*. 176–183.
- HERNÁDVÖLGYI, I. T. AND HOLTE, R. C. 1999. Psvn: A vector representation for production systems.
- HOBBS, J. R. 1990. Granularity. In *Readings in qualitative reasoning about physical systems*. Elsevier, 542–545.
- HOFFMANN, J., SABHARWAL, A., AND DOMSHLAK, C. 2006. Friends or Foes? an AI planning perspective on abstraction and search. In *Proceedings of the 16th International Conference on Automated Planning and Scheduling (ICAPS 2006)*. 294–303.
- HOLTE, R. C., MKADMI, T., ZIMMER, R. M., AND MACDONALD, A. J. 1996. Speeding up problem solving by abstraction: A graph oriented approach. *Artificial Intelligence* 85, 1-2, 321–361.
- ILLANES, L. AND MCILRAITH, S. A. 2016. Numeric planning via search space abstraction. In *Proceedings of the Workshop on Knowledge-based Techniques for Problem Solving and Reasoning*.
- ILLANES, L. AND MCILRAITH, S. A. 2019. Generalized planning via abstraction: Arbitrary numbers of objects. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019)*.
- JANHUNEN, T., NIEMELÄ, I., SEIPEL, D., SIMONS, P., AND YOU, J.-H. 2006. Unfolding partiality and disjunctions in stable model semantics. *ACM Transactions on Computational Logic* 7, 1, 1–37.
- JOHNSON-LAIRD, P. N. 1983. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press.
- KAMBHAMPATI, S. AND DAVIS, L. 1986. Multiresolution path planning for mobile robots. *IEEE Journal on Robotics and Automation* 2, 3, 135–145.
- KNOBLOCK, C. A. 1994. Automatically generating abstractions for planning. *Artificial Intelligence* 68, 2, 243–302.
- KRAMER, J. 2007. Is abstraction the key to computing? *Communications of the ACM* 50, 4, 36–42.
- LEFÈVRE, C., BÉATRIX, C., STÉPHAN, I., AND GARCIA, L. 2017. Asperix, a first-order forward chaining approach for answer set computing. *TPLP* 17, 3, 266–310.
- LEITE, J. 2017. A bird’s-eye view of forgetting in answer-set programming. In *Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2017)*. Springer, 10–22.
- LEONE, N., PFEIFER, G., FABER, W., EITER, T., GOTTLOB, G., PERRI, S., AND SCARCELLO, F. 2006. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic* 7, 3, 499–562.
- LIFSCHITZ, V. 1999. Answer set planning. In *International Conference on Logic Programming (ICLP)*. 23–37.

- LIFSCHITZ, V. 2008a. Twelve definitions of a stable model. In *Logic Programming*, M. Garcia de la Banda and E. Pontelli, Eds. Springer Berlin Heidelberg, 37–51.
- LIFSCHITZ, V. 2008b. What is answer set programming? In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI 2008)*. 1594–1597.
- LIFSCHITZ, V., PEARCE, D., AND VALVERDE, A. 2001. Strongly equivalent logic programs. *ACM Transactions on Computational Logic* 2, 4 (Oct.), 526–541.
- LIN, F. AND ZHAO, Y. 2004. Assat: Computing answer sets of a logic program by sat solvers. *Artificial Intelligence* 157, 1-2, 115–137.
- LOISEAUX, C., GRAF, S., SIFAKIS, J., BOUAJJANI, A., BENSLEM, S., AND PROBST, D. 1995. Property preserving abstractions for the verification of concurrent systems. *Formal methods in system design* 6, 1, 11–44.
- LOMUSCIO, A. AND MICHALISZYN, J. 2014. An abstraction technique for the verification of multi-agent systems against atl specifications. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR 2014)*. AAAI Press.
- LOMUSCIO, A. AND MICHALISZYN, J. 2016. Verification of multi-agent systems via predicate abstraction against ATLK specifications. In *Proceedings of AAMAS*. 662–670.
- MAHER, M. J. 1986. Equivalences of logic programs. In *Third International Conference on Logic Programming*, E. Shapiro, Ed. Springer Berlin Heidelberg, 410–424.
- MARPLE, K., SALAZAR, E., AND GUPTA, G. 2017. Computing stable models of normal logic programs without grounding.
- MORAK, M. AND WOLTRAN, S. 2012. Preprocessing of complex non-ground rules in answer set programming. In *Technical Communications of the 28th International Conference on Logic Programming*. 247.
- MOZETIČ, I. 1991. Hierarchical model-based diagnosis. *International Journal of Man-Machine Studies* 35, 3, 329–362.
- NAYAK, P. P. AND LEVY, A. Y. 1995. A semantic theory of abstractions. In *Proc. International Joint Conference on Artificial intelligence*. 196–203.
- NEBEL, B., DIMOPOULOS, Y., AND KOEHLER, J. 1997. Ignoring irrelevant facts and operators in plan generation. In *European Conference on Planning*. Springer, 338–350.
- NEWELL, A. AND SIMON, H. A. 1972. *Human Problem Solving*. Prentice-Hall.
- OETSCH, J., PÜHRER, J., AND TOMPITS, H. 2010. Catching the ouroboros: On debugging non-ground answer-set programs. *Theory and Practice of Logic Programming* 10, 4-6, 513–529.
- OIKARINEN, E. AND JANHUNEN, T. 2006. Modular equivalence for normal logic programs. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*. Vol. 6. IOS Press, 412–416.

- OSORIO, M., NAVARRO, J. A., AND ARRAZOLA, J. 2002. Equivalence in answer set programming. In *Logic Based Program Synthesis and Transformation*, A. Pettorossi, Ed. Springer Berlin Heidelberg, 57–75.
- PALÙ, A. D., DOVIER, A., PONTELLI, E., AND ROSSI, G. 2009. Gasp: Answer set programming with lazy grounding. *Fundam. Inform.* 96, 3, 297–322.
- PAPADIMITRIOU, C. H. AND YANNAKAKIS, M. 1992. Tie-breaking semantics and structural totality. In *Proc. ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. 16–22.
- PEARCE, D. 2004. Simplifying logic programs under answer set semantics. In *Logic Programming*. 210–224.
- PLAISTED, D. A. 1981. Theorem proving with abstraction. *Artificial Intelligence* 16, 1, 47–108.
- PONTELLI, E., SON, T. C., AND ELKHATIB, O. 2009. Justifications for logic programs under answer set semantics. *Theory and Practice of Logic Programming* 9, 1, 1–56.
- RIDDLE, P., DOUGLAS, J., BARLEY, M., AND FRANCO, S. 2016. Improving performance by reformulating PDDL into a bagged representation. In *Proceedings of the Workshop on Heuristics and Search for Domain-independent Planning (HSDIP 2016)*. 28–36.
- SACERDOTI, E. D. 1974. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence* 5, 2, 115–135.
- SAGIV, Y. 1987. Optimizing datalog programs. In *Proceedings of the 6th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. ACM, 349–362.
- SAITTA, L. AND ZUCKER, J.-D. 2013. *Abstraction in artificial intelligence and complex systems*. Vol. 456. Springer.
- SARIBATUR, Z. G. AND EITER, T. 2016. Reactive policies with planning for action languages. In *Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016)*. Lecture Notes in Computer Science, vol. 10021. Springer, 463–480.
- SARIBATUR, Z. G. AND EITER, T. 2018. Omission-based abstraction for answer set programs. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018)*. AAAI Press, 42–51.
- SARIBATUR, Z. G., SCHÜLLER, P., AND EITER, T. 2019. Abstraction for non-ground answer set programs. In *Proceedings of the 16th European Conference on Logics in Artificial Intelligence (JELIA 2019)*. Lecture Notes in Computer Science. Springer, 576–592.
- SCHAUB, T. AND WOLTRAN, S. 2018. Answer set programming unleashed! *KI* 32, 2-3, 105–108.
- SCHULZ, C. AND TONI, F. 2013. ABA-based answer set justification. *Theory and Practice of Logic Programming (TC)*, 4–5.
- SEIPP, J. AND HELMERT, M. 2013. Counterexample-guided cartesian abstraction refinement. In *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS 2013)*.
- SIMONS, P., NIEMELÄ, I., AND SOININEN, T. 2002. Extending and implementing the stable model semantics. *Artificial Intelligence* 138, 1-2, 181–234.

- SRIVASTAVA, S., IMMERMANN, N., AND ZILBERSTEIN, S. 2011. A new representation and associated algorithms for generalized planning. *Artificial Intelligence* 175, 2, 615–647.
- TURNER, H. 2003. Strong equivalence made easy: nested expressions and weight constraints. *Theory and Practice of Logic Programming* 3, 4–5, 609–622.
- VAN GELDER, A., ROSS, K. A., AND SCHLIPF, J. S. 1991. The well-founded semantics for general logic programs. *Journal of the ACM* 38, 3, 619–649.
- WEINZIERL, A. 2017. Blending lazy-grounding and CDNL search for answer-set solving. In *Logic Programming and Nonmonotonic Reasoning - 14th International Conference, LPNMR 2017, Espoo, Finland, July 3-6, 2017, Proceedings*. 191–204.
- ZHANG, S., YANG, F., KHANDELWAL, P., AND STONE, P. 2015. Mobile robot planning using action language BC with an abstraction hierarchy. In *Proc. International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*. Springer, 502–516.

## A Proofs

**Proof of Thm 3.1.** Let  $\hat{I}$  and  $\hat{\Pi}$  denote  $m(I)$  and  $\Pi^m$ , respectively. Towards a contradiction, assume that there exists some  $I \in AS(\Pi)$  s.t.  $\hat{I} \cup \mathcal{T}_m \notin AS(\hat{\Pi})$ . This can occur either because (i)  $\hat{I} \cup \mathcal{T}_m$  is not a model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$  or (ii)  $\hat{I} \cup \mathcal{T}_m$  is not a minimal model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ .

(i) Suppose  $\hat{I} \cup \mathcal{T}_m$  is not a model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ . Then there exists some rule  $\hat{r} \in \hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$  where  $\hat{I} \cup \mathcal{T}_m \models B(\hat{r})$  and  $\hat{I} \cup \mathcal{T}_m \not\models H(\hat{r})$ . By construction of  $\Pi$ ,  $\hat{r}$  is only obtained by step (a), otherwise  $\hat{r}$  would be a choice rule with head  $H(\hat{r}) = \{m(l)\}$ , and  $\hat{r}$  would be satisfied. Consequently  $\hat{r}$  is a rule from step (a) for  $r$  in  $\Pi$ .

Since  $\hat{I} \cup \mathcal{T}_m \models m(B(r)), rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ , we have  $\hat{I} \cup \mathcal{T}_m \models m(B(r))$ . If we have  $p(\hat{e}_1, \dots, \hat{e}_n) \in m(B^+(r))$ , some  $e_i \in \hat{e}_i$  exists such that  $p(e_1, \dots, e_n) \in I$  as all variables are standardized apart,  $I \models B^+(r)$  for this choice. As for  $p(\hat{e}_1, \dots, \hat{e}_n) \in m(B^-(r))$ , then  $p(e_1, \dots, e_n) \notin I$  for all  $e_i \in \hat{e}_i$ . So we can instantiate the abstract body  $m(B(r))$  to some original body  $B(r)$  where  $I \models B(r)$ . Also having  $\hat{I} \models rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  means  $I \models rel(d_1, d_2)$  for all  $d_i \in \hat{d}_i$ , thus we have  $I \models B(r), rel(d_1, d_2)$ . So  $r : l \leftarrow B(r), rel(d_1, d_2)$  is in  $\Pi^I$ . As  $I$  is a model, it follows that  $I \models l$ , which then means  $\hat{I} \models m(l)$ ; this is a contradiction.

(ii) Suppose there exists some  $\hat{J} \subset \hat{I}$  such that  $\hat{J} \cup \mathcal{T}_m$  is a model of  $\hat{\Pi}^{\hat{J} \cup \mathcal{T}_m}$ . We claim that  $J = m^{-1}(\hat{J}) \cap I$  is a model of  $\Pi^I$ ; as  $J \subset I$  holds, this would contradict that  $I \in AS(\Pi)$ . Assume  $J \not\models \Pi^I$ . Then  $J$  does not satisfy some rule  $r : l \leftarrow B(r), rel(d_1, d_2)$  in  $\Pi^I$ , i.e.,  $J \models B(r), rel(d_1, d_2)$  but  $J \not\models l$ . As  $J \subset I$  and  $I$  is a model of  $\Pi^I$ , we have  $I \models l$ , thus,  $l \in I \setminus J$ .

Now, we look at the cases for applying the mapping  $m$  to  $r$ , by considering the abstractions  $m(B(r))$  and  $rel(\hat{d}_1, \hat{d}_2)$ , and show that a contradiction is always achieved.

First, assume that  $\hat{I} \models m(B(r))$ . There are the following cases for  $m(J)$ : (1-1)  $m(J) \models m(B(r))$ , or (1-2)  $m(J) \not\models m(B(r))$ .

(1-1) As  $m(J) \models m(B(r))$ , we look at  $rel(\hat{d}_1, \hat{d}_2)$ . We know that  $J \models rel(d_1, d_2)$ .

- If  $rel(\hat{d}_1, \hat{d}_2)$  has the relation type  $\tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ , this means that we have  $m(J) \models rel(\hat{d}_1, \hat{d}_2)$ , and thus  $m(J) \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ . As  $\hat{J} = m(J)$  and  $\hat{J} \subset \hat{I}$ , we also get  $\hat{I} \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ , thus the non-ground rule created by step (a) has an instantiation  $m(l) \leftarrow$

$m(B(r)), rel(\hat{d}_1, \hat{d}_2), \tau_I^{rel}(\hat{d}_1, \hat{d}_2)$  in  $\widehat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ . As  $\hat{J}$  and  $\hat{I}$  are models of  $\widehat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ , we have  $\hat{J} \models m(l)$  and  $\hat{I} \models m(l)$ . Thus,  $l \in m^{-1}(\hat{J})$  and  $l \in I$ ; by definition of  $J$ , we get  $l \in J$  thus  $J \models l$ , which is a contradiction.

- If  $rel(\hat{d}_1, \hat{d}_2)$  has the relation type  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , this again means that we have  $m(J) \models rel(\hat{d}_1, \hat{d}_2)$ , and thus  $m(J) \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  and  $\hat{I} \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ . Thus, as  $m(l) \in \hat{I}$  the non-ground choice rule created by step (b) amounts to  $m(l) \leftarrow m(B(r)), rel(\hat{t}_1, \hat{t}_2), \tau_{III}^{rel}(\hat{t}_1, \hat{t}_2)$  in  $\widehat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ , which again achieves  $\hat{J} \models m(l)$ , thus  $J \models l$ , a contradiction.
- If  $rel(\hat{d}_1, \hat{d}_2)$  has the relation type  $\tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$ , then we have  $m(J) \not\models rel(\hat{d}_1, \hat{d}_2)$ , i.e.,  $m(J) \models \bar{rel}(\hat{d}_1, \hat{d}_2)$ , and thus  $m(J) \cup \mathcal{T}_m \models \bar{rel}(\hat{d}_1, \hat{d}_2), \tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$ . With similar reasoning as above, we find an instantiation of the non-ground rule created by step (c) and achieve  $J \models l$ , a contradiction.

(1-2) We show that there is no possibility to have  $\hat{J} \not\models m(B(r))$ , for  $\hat{J} = m(J)$ , while  $\hat{I} \models m(B(r))$ . In order to have  $\hat{J} \not\models m(B(r))$ , some positive literal  $\hat{l}_i \in m(B^+(r))$  must occur in  $\hat{I} \setminus \hat{J}$  so that  $\hat{J} \not\models m(B^+(r))$ . However, this contradicts with  $J \models B^+(r)$ .

Now, assume that  $\hat{I} \not\models m(B(r))$ . As  $I \models B(r)$ , we know that  $\hat{I} \models m(B^+(r))$  holds. So we have the rule  $l \leftarrow B^+(r), not\ l_i, rel(d_1, d_2)$  (according to restriction (i) on having at most one negative literal) where  $l_i \neq l$  and  $\hat{I} \not\models m(B(r))$  means  $\hat{I} \models m(l_i)$  for  $l_i \in B^-(r)$  while  $I \not\models l_i$ , i.e.,  $l_i \notin I$ . So we get  $\hat{I} \models m(B_i^{sh}(r))$ . Then there are the following cases for  $m(J)$ : (2-1)  $m(J) \models m(B_i^{sh}(r))$ , or (2-2)  $m(J) \not\models m(B_i^{sh}(r))$ .

(2-1) As we have  $m(J) \models m(B_i^{sh}(r))$ , we look at  $rel(\hat{d}_1, \hat{d}_2)$ . We know that  $J \models rel(d_1, d_2)$ .

- For cases  $\tau_I^{rel}(\hat{d}_1, \hat{d}_2)$  and  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , as we have  $J \models rel(d_1, d_2)$ , we get  $\hat{J} \models rel(\hat{d}_1, \hat{d}_2)$  and  $\hat{I} \models rel(\hat{d}_1, \hat{d}_2)$ . Notice that since  $m(l_i) \in \hat{I}$ , there must be some  $l'_i \in I$  such that  $m(l_i) = m(l'_i)$ , thus  $l_i$  is mapped to a non-singleton cluster  $m(l_i)$ . So the atom  $isCluster(\hat{j})$  holds true in  $\hat{J}$  and  $\hat{I}$  for some  $j \in arg(l_i)$  for which  $|m^{-1}(m(j))| > 1$ . Thus in  $\widehat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$  we get an instantiation  $m(l) \leftarrow m(B_i^{sh}(r)), rel(\hat{d}_1, \hat{d}_2), isCluster(\hat{j})$  of the non-ground rule created by (step d-i), and again achieve  $J \models l$ , which is a contradiction.
- For the case  $\tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$ , with similar reasoning as in (1-1), we find instantiations of the non-ground rules created by (step d-ii) and achieve  $J \models l$ , which is a contradiction.

(2-2) We show that there is no possibility to have  $m(J) \not\models m(B_i^{sh}(r))$ , while  $\hat{I} \models m(B_i^{sh}(r))$ . As  $J \models B(r)$ , we know that  $m(J) \models m(B^+(r))$  holds. So  $m(J) \not\models m(B_i^{sh}(r))$  means  $m(J) \not\models m(l_i)$  while  $\hat{I} \models m(l_i)$ . Now, we take a look at  $\Pi^I$ . As there must be some  $l'_i \in I$  (such that  $m(l_i) = m(l'_i)$ ), this means that there is some rule  $r' : l'_i \leftarrow B(r'), rel(d'_1, d'_2)$  in  $\Pi^I$ . We then take a look at the abstraction of  $r'$ . By doing the same case analysis of (1-1), (1-2) and (2-1), we achieve  $m(J) \models m(l'_i)$ , i.e.,  $m(J) \models m(l_i)$ , which yields a contradiction. As for (2-2), this means the rule  $r'$  is of form  $r' : l'_i \leftarrow B^+(r'), not\ l_{i_2}, rel(d'_1, d'_2)$ , where we want to claim  $m(J) \not\models m(l_{i_2})$ . For this, we take a look at another rule  $r''$  in  $\Pi^I$  of form  $r'' : l'_{i_2} \leftarrow B(r''), rel(d''_1, d''_2)$  with  $m(l'_{i_2}) = m(l_{i_2})$ . By restriction (iii) on no negative cyclic dependency among the literals, this recursive process eventually ends, say, after  $n$  steps, at some rule  $r^n : l'_{i_n} \leftarrow B(r^n), rel(d'_n, d''_n)$  where case (2-2) is not applicable, and  $m(J) \models l'_{i_n}$  is achieved. Then by tracing the rules back to  $r$  we get  $m(J) \models m(l_i)$ . Thus  $m(J) \not\models m(B_i^{sh}(r))$  is not possible.  $\square$

**Proof of Thm 3.2.** Similar to proof of Theorem 3.1, we assume towards a contradiction that there exists some  $I \in AS(\Pi)$  such that  $\hat{I} \cup \mathcal{T}_m \notin AS(\hat{\Pi})$ . This can occur either because (i)  $\hat{I} \cup \mathcal{T}_m$  is not a model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$  or (ii)  $\hat{I} \cup \mathcal{T}_m$  is not a minimal model.

(i) Let  $\hat{I} \cup \mathcal{T}_m$  be not a model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ , then there exists some rule  $\hat{r} \in \hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$  where  $\hat{I} \cup \mathcal{T}_m \models B(\hat{r}) \wedge \hat{I} \cup \mathcal{T}_m \not\models H(\hat{r})$ . For cases (G-i) and (G-iii), the contradiction is achieved similar to the proof of Theorem 3.1, since  $\hat{r}$  is a rule from step (a). As for case (G-ii), we will have  $\hat{I} \cup \mathcal{T}_m \models m(B(r)), rel'(\hat{d}), \tau_1^{rel}(\hat{d})$ , where  $\hat{d}$  is a shorthand for  $\hat{d}_{1,1}, \hat{d}_{2,1}, \dots, \hat{d}_{1,k}, \hat{d}_{2,k}$ ; then by definition of  $rel'$  this means  $I \models B(r), rel_1(d_{1,1}, d_{2,1}), \dots, rel_k(d_{1,k}, d_{2,k})$  of  $r$  in  $\Pi^I$ . This reaches a contradiction as  $I$  is a model and  $I \models l$ , which means  $\hat{I} \models m(l)$ .

(ii) Now let there be  $\hat{J} \subset \hat{I}$  such that  $\hat{J} \cup \mathcal{T}_m$  is a model of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ . We claim that  $J = m^{-1}(\hat{J}) \cap I \subset I$  is a model of  $\Pi^I$ ; which would contradict that  $I \in AS(\Pi)$ . Assume  $J \neq \Pi^I$ .  $J$  does not satisfy some rule  $r : l \leftarrow B(r), rel(d_1, d_2)$  in  $\Pi^I$ , i.e.,  $J \models B(r), rel_1(d_{1,1}, d_{2,1}), \dots, rel_k(d_{1,k}, d_{2,k})$  but  $J \not\models l$ , i.e.,  $l \notin J$ . As  $J \subset I$  and  $I$  is a model of  $\Pi^I$ , we have  $I \models l$ , i.e.,  $l \in I \setminus J$ . We consider the abstractions  $m(B(r))$  and  $rel_1(\hat{d}_{1,1}, \hat{d}_{2,1}), \dots, rel_k(\hat{d}_{1,k}, \hat{d}_{2,k})$ .

First, assume  $\hat{I} \models m(B(r))$ . There are the following cases for  $m(J)$ : (1-1)  $m(J) \models m(B(r))$ , or  $m(J) \not\models m(B(r))$ .

(1-1) As  $m(J) \models m(B(r))$ , we look at  $rel_1(\hat{d}_{1,1}, \hat{d}_{2,1}), \dots, rel_k(\hat{d}_{1,k}, \hat{d}_{2,k})$ . We know that  $J \models rel_1(d_{1,1}, d_{2,1}), \dots, rel_k(d_{1,k}, d_{2,k})$ .

(1-1-a) If all  $rel_i(\hat{d}_{1,i}, \hat{d}_{2,i})$  have the relation type  $\tau_1^{rel_i}(\hat{d}_{1,i}, \hat{d}_{2,i})$ , this means that we have  $m(J) \models rel_1(\hat{d}_{1,1}, \hat{d}_{2,1}), \dots, rel_k(\hat{d}_{1,k}, \hat{d}_{2,k})$ , and thus

$$m(J) \cup \mathcal{T}_m \models rel'(\hat{d}), \tau_1^{rel}(\hat{d}). \quad (69)$$

As  $\hat{J} = m(J)$  and  $\hat{J} \subset \hat{I}$ , we also get  $\hat{I} \cup \mathcal{T}_m \models rel'(\hat{d}), \tau_1^{rel}(\hat{d})$ , thus the non-ground rule created by step (a) has an instantiation  $m(l) \leftarrow m(B(r)), rel'(\hat{d}), \tau_1^{rel}(\hat{d})$  in  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ . As  $\hat{J}$  and  $\hat{I}$  are models of  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ , we have  $\hat{J} \models m(l)$  and  $\hat{I} \models m(l)$ . Thus, by definition of  $J$ , we get  $J \models l$ , which is a contradiction.

(1-1-b) If at least one  $rel_i(\hat{d}_{1,i}, \hat{d}_{2,i})$  has the relation type  $\tau_{III}^{rel_i}(\hat{d}_{1,i}, \hat{d}_{2,i})$ , while no  $rel(\hat{d}_{1,j}, \hat{d}_{2,j})$  has the relation type  $\tau_{IV}^{rel}(\hat{d}_{1,j}, \hat{d}_{2,j})$ , above (69) is achieved for  $\tau_{III}^{rel}(\hat{d})$  in place of  $\tau_1^{rel}(\hat{d})$ .

(1-1-c) If at least one  $rel_i(\hat{d}_{1,i}, \hat{d}_{2,i})$  has the relation type  $\tau_{IV}^{rel_i}(\hat{d}_{1,i}, \hat{d}_{2,i})$ , this means that we have  $m(J) \models \dots, \bar{rel}(\hat{d}_{1,i}, \hat{d}_{2,i}), \dots$ . Thus, for  $rel'(\hat{d})$  we have  $m(J) \not\models rel'(\hat{d})$  but  $m(J) \models \bar{rel}'(\hat{d})$ , and

$$m(J) \cup \mathcal{T}_m \models \bar{rel}'(\hat{d}), \tau_{IV}^{rel}(\hat{d}). \quad (70)$$

By the same reasoning in (1-1-a), we get that the non-ground choice rule created by step (c) amounts to  $m(l) \leftarrow m(B(r)), \bar{rel}'(\hat{d}), \tau_{IV}^{rel}(\hat{d})$  in  $\hat{\Pi}^{\hat{I} \cup \mathcal{T}_m}$ , and thus we reach a contradiction.

(1-2) This case is handled the same as in proof (1-2) of Theorem 3.1.

Now, we focus on the case  $I \not\models m(B(r))$ . As  $I \models B(r)$ , we know that  $\hat{I} \models m(B^+(r))$  should hold. Then  $\hat{I} \not\models m(B(r))$  means that for a non-empty set  $L \subseteq B^-(r)$  of negative literals  $\hat{I} \models l_i$  for each  $l_i \in L$ , while  $I \not\models l_i$ , i.e.,  $l_i \notin I$ . So we get  $\hat{I} \models m(B_L^{sh}(r))$ . Assume we further have a set  $L_c$  of literals involved in a negative loop. We also get  $\hat{I} \models m(B_{L, L_c}^{sh}(r))$  (26), since the set  $L' = L \cap L_c$  of literals gets omitted from  $B^-(r)$ , and we get  $\hat{I} \models \text{not } B^-(r) \setminus L'$ .

Then there are the following cases for  $m(J)$ : (2-1)  $m(J) \models m(B_{L, L_c}^{sh}(r))$ , or (2-2)  $m(J) \not\models m(B_{L, L_c}^{sh}(r))$ .

- (2-1) As  $m(J) \models m(B_{L,L_c}^{sh}(r))$ , similar to proof (1-1) above and (2-1) of Theorem 3.1, the abstraction  $rel_1(\hat{d}_{1,1}, \hat{d}_{2,1}), \dots, rel_k(\hat{d}_{1,k}, \hat{d}_{2,k})$  on relations is considered, and the contradiction  $J \models l$  is achieved.
- (2-1-a) By (1-1-a) and (1-1-b), we get the case (69) and same for  $\tau_{III}^{rel'}(\hat{d})$ . We know that as  $m(J) \models m(l_i)$  and as  $l_i \notin J$  (since  $J \subseteq I$  and  $l_i \notin I$ , this means that  $isCluster(\hat{j})$  holds true in  $m(J)$  and  $\hat{I}$  for some  $j \in arg(l_i)$ ). Thus we have  $m(J) \cup \mathcal{T}_m \models rel'(\hat{d}), isCluster(\hat{j})$ , which means that  $m(l) \in \hat{I}$ , thus  $l \in I$  and by definition of  $J$ ,  $l \in J$ , which is a contradiction.
- (2-1-b) By (1-1-c), we get the case (70) and by a similar reasoning as in (2-1-a) we also have  $m(J) \models isCluster(\hat{j})$ , hence  $m(J) \cup \mathcal{T}_m \models \overline{rel}'(\hat{d}), \tau_{IV}^{rel'}(\hat{d}), isCluster(\hat{j})$ . Thus we similarly achieve a contradiction.
- (2-2) We show that there is no possibility to have  $m(J) \not\models m(B_{L,L_c}^{sh}(r))$ , while  $\hat{I} \models m(B_{L,L_c}^{sh}(r))$ . As  $J \models B(r)$ , we know that  $m(J) \models m(B^+(r))$  should hold. So  $m(J) \not\models m(B_{L,L_c}^{sh}(r))$  means  $m(J) \not\models m(l_i)$  for some  $l_i \in L \setminus L_c$ , while  $\hat{I} \models m(l_i)$ . We do the same recursive reasoning as in proof (2-2) of Theorem 3.1 over the literals not in  $L_c$ . Thus, the process eventually ends and achieves that  $m(J) \models m(l_i)$  actually holds, and that  $m(J) \not\models m(B_{L,L_c}^{sh}(r))$  is not possible.  $\square$

**Proof of Lemma 3.9.** By definition, we need to check (1) that  $I$  is a model of  $(\Pi^m)^I$  and (2) that  $I$  is minimal, no  $J \subset I$  is a model of  $(\Pi^m)^I$ .

As for (1), we can refute the property by guessing a rule  $r \in \Pi^m$  and a variable substitution  $\theta$  and verifying that  $I$  does not satisfy  $(r\theta)^I$ , where  $r\theta$  denotes the ground instance of  $r$  obtained by applying  $\theta$  to its variables; note that in this case  $(r\theta)^I \in (\Pi^m)^I$  holds.

Each rule  $r \in \Pi^m$  has polynomial size in the input. Checking whether  $r \in \Pi^m$  holds is feasible in polynomial time, as computing the set of atoms that occur in negative cycles is feasible in polynomial time as well. Furthermore, checking whether  $r' = r\theta$  is in  $(\Pi^m)^I$  is feasible in polynomial time. Overall, refuting (1) is in NP.

As for (2),  $I$  is minimal if each atom  $a \in I$  has a proof, given by a sequence  $r_1, r_2, \dots, r_k$  of applications of rules from  $r_i \in (\Pi^m)^I$  where each positive body literal of  $r_i$  occurs in some head of  $r_j$ ,  $j < i$ . Note that w.l.o.g.  $I = \{a_1, \dots, a_k\}$  and  $a_i$  has as proof  $r_1, \dots, r_i$ ,  $i = 1, \dots, k$ . As the proof can be guessed and nondeterministically verified in polynomial time, it follows that (2) is in NP. Hence it follows that the problem is in  $\Delta_2^P$  (more precisely, in the class DP).  $\square$

**Proof of Thm 3.10.** To show that  $\hat{I}$  is a concrete abstract answer set of  $\Pi^m$ , we can guess an interpretation  $J$  of  $\Pi$  and check that (a)  $m(J) = \hat{I}$ , (b)  $m(J) \in AS(\Pi^m)$ , and (c)  $J \in AS(\Pi)$ . Testing (a) is clearly polynomial in the size of  $J$ , and by Lemma 3.9, (b) and (c) are feasible in  $\Delta_2^P$  in the size of  $J$  and  $\Pi$  (and thus in exponential time in the size of  $\hat{I}$  and  $\Pi$ ); consequently, deciding whether  $\hat{I}$  is a concrete abstract answer set of  $\Pi^m$  is in NEXP. For bounded predicate arities, the guess for  $J$  has polynomial size in the input, and we can check the conditions (b) and (c) by Lemma 3.9 with an NP oracle in polynomial time; this establishes  $\Sigma_2^P$  membership.

The matching lower bounds are shown by a reduction from deciding whether a given non-ground program  $\Pi$  has some answer set, which is NEXP-complete in the general case and  $\Sigma_2^P$ -complete for bounded predicate arities (Dantsin et al., 2001; Eiter et al., 2007).

Without loss of generality,  $\Pi$  involves a single predicate  $p$  (which can be achieved by reification and padding arguments) and contains some fact  $p(\vec{d})$ . The mapping we define is  $m = \{\{d_1, \dots, d_n\} / \hat{d}\}$  where  $d_1, \dots, d_n$  form the Herbrand domain. Then  $\hat{I} = \{p(\hat{d}, \dots, \hat{d})\}$  is a concrete abstract answer set of  $\Pi^m$  iff  $\Pi$

has some answer set. Note that actually  $\widehat{I} \in AS(\Pi^m)$  holds; thus the overall complexity does not change if this property is asserted. This proves the result.  $\square$

**Proof of Thm 3.12.** For the membership, one can guess an interpretation  $\widehat{I}$  of  $\Pi^m$  such that  $\widehat{I}$  is an answer set of  $\Pi^m$ , and then check whether  $\widehat{I}$  is spurious. By Theorem 3.10, the spuriousness check can be done with a coNEXP oracle in general and with a  $\Sigma_2^P$  oracle in the bounded predicate case. However, by applying standard padding techniques,<sup>9</sup> it follows that a coNP oracle is sufficient in the general case. This proves membership of the problem in coNEXP<sup>NP</sup> in the general case and in  $\Sigma_3^P$  in the bounded predicate case, respectively.

The coNEXP<sup>NP</sup>-hardness in the general case is shown by a reduction from evaluating second-order logic formulas of a suitable form over finite relational successor structures, i.e., relational structures  $S = (D, R^S)$  with a universe  $D$  and interpretations  $R_i^S$  for the relations  $R_i$  in  $R = R_1, \dots, R_k$ , which include the relations  $first(x)$ ,  $next(x, y)$  and  $last(x)$  associated with a linear ordering  $\leq$  of  $D$ .

**Lemma A.1.** *Given a second-order (SO) sentence of the form  $\Phi = \exists P \forall Q. \varphi$  where  $P = P_1, \dots, P_{m_1}$  and  $Q = Q_1, \dots, Q_{m_2}$  are predicate variables and  $\varphi = \bigvee_j \varphi_j$  is FO such that each  $\varphi_j$  is of the form  $\varphi_j = \exists x_1, \dots, x_n l_{j,1} \cdots \wedge \cdots \wedge l_{j,k}$  where each  $l_{i,j}$  is a FO-literal, and a finite relational successor structure  $S$ , deciding whether  $S \models \Phi$  is NEXP<sup>NP</sup>-complete.*

This lemma can be obtained from the facts that (1) evaluating SO-sentences of the form  $\Psi \exists P \forall Q. \varphi$ , where  $\Psi$  is a first-order formula, over finite relational successor structures is NEXP<sup>NP</sup>-complete, cf. (Gottlob et al., 1999), and (2) that  $\Psi$  can be transformed into some  $\Phi$  of the form described in polynomial time; the latter is possible using second-order skolemization and auxiliary predicates for quantifier elimination, cf. (Eiter et al., 1996) and for denoting subformulas, such that  $\varphi(\vec{x}) \equiv P_\varphi(\vec{x})$  and  $\varphi(\vec{x}) = \varphi_1(\vec{x}) \wedge \varphi_2(\vec{x})$  is expressed by  $P_\varphi(\vec{x}) \equiv P_{\varphi_1}(\vec{x}) \wedge P_{\varphi_2}(\vec{x})$  etc.

We first describe how to encode evaluating the sentence  $\Phi' = \exists P \exists Q \neg \varphi$  into an ordinary program  $\Pi_0$ , and then extend the encoding to prove the result. We define the rules of  $\Pi_0$  as follows:

$$P_{j,i}(X_1, \dots, X_n) \leftarrow \overline{P_{j,i}}(X_1, \dots, X_n), D(X_1), \dots, D(X_n). \quad \text{for each } P_{j,i} \in P \quad (71)$$

$$\overline{P_{j,i}}(X_1, \dots, X_n) \leftarrow \text{not } P_{j,i}(X_1, \dots, X_n), D(X_1), \dots, D(X_n). \quad \text{for each } P_{j,i} \in P \quad (72)$$

$$Q_{j,i}(Y_1, \dots, Y_n) \leftarrow \overline{Q_{j,i}}(Y_1, \dots, Y_n), D(Y_1), \dots, D(Y_n). \quad \text{for each } Q_{j,i} \in Q \quad (73)$$

$$\overline{Q_{j,i}}(Y_1, \dots, Y_n) \leftarrow \text{not } Q_{j,i}(Y_1, \dots, Y_n), D(Y_1), \dots, D(Y_n). \quad \text{for each } Q_{j,i} \in Q \quad (74)$$

$$\text{sat} \leftarrow l_{j,1}^{\neg/\text{not}} \wedge \cdots \wedge l_{j,k}^{\neg/\text{not}}. \quad (75)$$

$$\text{ok} \leftarrow \text{not } \text{ok}. \quad (76)$$

$$\text{ok} \leftarrow \text{not } \text{sat}. \quad (77)$$

where  $l^{\neg/\text{not}}$  denotes the replacement of  $\neg$  in  $l$  by  $\text{not}$ .<sup>10</sup>

Informally, the rules (71),(72) and (73),(74) guess extensions for the predicates in  $P$  and  $Q$ , respectively, while the rules (75) evaluate the formula  $\varphi$ . A guess for  $P$  and  $Q$  yields an answer set of  $\Pi_0$  augmented with  $S$  (provided as positive facts) iff  $\varphi$  evaluates over  $S$  to false; in this case, no rule (75) fires and thus  $\text{sat}$  can not be derived, which means in turn that  $\text{ok}$  can be derived by (77) and thus the constraint (76) is satisfied. On the other hand, deriving  $\text{ok}$  is necessary to have an answer set, which means that  $\text{sat}$  must not be derived from the guess for  $P$  and  $Q$ .

<sup>9</sup>The input  $x$  to the oracle is changed to  $(x, y)$ , where  $y$  is an (exponentially) long string  $y$ , and the oracle query considers  $x$  from the input only. This artificially lowers the time bound within the query (measured in the size of  $(x, y)$ ) can be answered.

<sup>10</sup>To make the rules safe, domain predicates  $D(X)$  can be added for unsafe variables  $X$ .

We extend the program  $\Pi_0$  now for spuriousness checking. To this end, we introduce for the domain  $D = \{x_1, \dots, x_n\}$  at hand a copy  $D' = \{y_1, \dots, y_n\}$  and link  $x_i$  to  $y_i$  via a predicate  $eq(x, y)$  that holds for  $x, y \in D \cup D'$  iff  $x = x_i \wedge y = y_i$  for some  $i = 1, \dots, n$ . The idea is to use  $D$  and  $D'$  in the predicates from  $P$  and  $Q$ , respectively, and to abstract  $D'$  into a single element, such that for every guess  $\chi$  for  $P$ , some abstract answer set  $\hat{I}_\chi$  of the abstract program  $\Pi^m$  will exist; and that, moreover,  $\hat{I}_\chi$  will be concrete if for some guess for  $Q$ , we have an answer set of  $\Pi$ , where the latter program is equivalent to  $\Pi_0$ ; thus  $\hat{I}_\chi$  will be spurious iff no guess for  $Q$  will yield an answer set of  $\Pi_0$ , which means that the formula  $\exists P \forall Q. \varphi$  evaluates to true.

We make the following adjustments.

1. First, we replace in (73) and (74) the predicate  $D$  with  $D'$ .
2. Next, for each rule  $r$  from (75) we add for each term  $t$  that occurs in the rule body a “typing” atom  $D(t)$ , we replace each term  $t$  that occurs in a  $Q$ -literal with a fresh variable  $X_t$  and add the atoms  $D'(X_t)$  and  $eq(t, X_t)$ .
3. To each rule  $r$  obtained from the previous step we add *not succ*( $y_1, y_1$ ) in the body (this literal evaluates to true with no abstraction).
4. We add facts  $eq(x_i, y_i)$ , for  $i = 1, \dots, n$ .
5. We add facts  $Q_{j,i}(y_0, \dots, y_0), \overline{Q_{j,i}}(y_0, \dots, y_0)$  for all  $Q_{j,i} \in Q$ , where  $y_0$  is a fresh constant.

It is not hard to establish that the answer sets  $I$  of the resulting program  $\Pi$  (over  $S$ ) correspond to the answer sets  $I_0$  of  $\Pi_0$  over  $S$ ; each  $I$  is obtained from some  $I_0$  by replacing in the  $Q_{j,i}$ - and  $\overline{Q_{j,i}}$ -atoms the constant  $x_i$  with the corresponding  $y_i$ , adding all facts  $Q_{j,i}(y_0, \dots, y_0), \overline{Q_{j,i}}(y_0, \dots, y_0)$  and stripping off the *eq*-atoms.

The mapping that we construct is  $m = \{\{x_1\}/x_1, \dots, \{x_n\}/x_n\} \cup \{\{y_0, y_1, \dots, y_n\}/\hat{y}\}$ . In the abstract program  $\Pi^m$ , the rules (71), (72) are carried over, while the modified rules (73), (74) are turned into rules to derive abstract atoms over  $Q_{j,i}$  resp.  $\overline{Q_{j,i}}$ . However, since  $\Pi^m$  contains the abstracted facts  $Q_{j,i}(\hat{y}, \dots, \hat{y}), \overline{Q_{j,i}}(\hat{y}, \dots, \hat{y})$ , these rules are redundant.

The modified rules (75) are turned into guessing rules for *sat*, while the other rules (76) and (77) remain unchanged. The abstract answer sets of  $\Pi^m$  correspond to guesses  $\chi$  for  $P$  to which *ok* and all  $Q_{j,i}(\hat{y}, \dots, \hat{y}), \overline{Q_{j,i}}(\hat{y}, \dots, \hat{y})$  are added (*sat* is guessed false); denote this answer set by  $I_\chi$ .

The answer set  $I_\chi$  is concrete, if there is some guess  $\mu$  for  $Q$  such that we obtain an answer set  $I$  of the program  $\Pi$  that is mapped to  $I_\chi$ , i.e.,  $m(I) = I_\chi$ ; this  $I$  corresponds to some answer set  $I_0$  as described above. Thus  $I_\chi$  is spurious, if no such guess  $\mu$  for  $Q$  exists.

Putting it all together, it holds that  $\Pi$  has with respect to the mapping  $m = \{\{x_1\}/x_1, \dots, \{x_n\}/x_n\} \cup \{\{y_0, y_1, \dots, y_n\}/\hat{y}\}$  some spurious answer iff the formula  $\Phi$  in Lemma A.1 evaluates over  $S$  to true. Since  $\Pi$  and  $m$  are constructable in polynomial time from  $\Phi$  and  $S$ , this proves  $\text{coNEXP}^{\text{NP}}$  hardness in the general case.

For the bounded predicate arities case, the evaluation of a formula  $\Phi$  as in Lemma A.1 is  $\Sigma_3^P$ -complete; furthermore, all steps in producing the program  $\Pi$  preserves bounded arities. Thus with the same argument, we obtain  $\Sigma_3^P$ -hardness for deciding whether some spurious answer set exists for bounded predicate arities. This proves the result.  $\square$

**Proof of Prop 4.2.** Let  $X$  be an interpretation over  $\mathcal{A}^*$ . We will show that with the help of the auxiliary rules/atoms, some interpretation  $X'$  which is a minimal model of  $(\Pi_{\text{debug}} \cup Q_{\hat{I}}^m)^{X'}$  can be reached starting from  $X$ . We have the cases (i)  $X \not\models (\Pi_{\text{debug}} \cup Q_{\hat{I}}^m)^X$ , and (ii)  $X \models (\Pi_{\text{debug}} \cup Q_{\hat{I}}^m)^X$ .

- (i) We show that  $X$  can be changed to some interpretation  $X'$  that avoids the conditions for  $X$  not satisfying a rule  $(\Pi_{debug} \cup Q_{\hat{I}}^m)^X$ , by doing a case analysis on where the rule may occur.
- (a) Let  $r$  be an unsatisfied rule in  $(\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}])^X$ . This means that  $X \models B(r)$  and  $X \not\models H(r)$ . The rule  $r$  can not be an instantiation of the choice rules in  $\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}]$ , as it would be instantiated for  $X$ , and hence be satisfied. Thus  $r$  can either (a-1) have  $H(r) \in AB(\Pi)$  or (a-2) have  $H(r) = ko_{n,r}$  for some  $r' \in \Pi$  and be in  $\mathcal{T}_{deact}[\Pi]^X$ . For both of these cases, we construct  $X' = X \cup \{H(r)\}$ , and the reduct  $(\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}])^{X'}$  will not have further rules, since the added  $H(r)$  does not occur in the body of any auxiliary rule. Note that adding  $ko_{n,r}$  also avoids having the corresponding meta-rule of  $r'$  in the reduct  $(\mathcal{T}_{meta}[\Pi])^{X'}$ . So no further rules will be added in  $(\Pi_{debug} \cup Q_{\hat{I}}^m)^{X'}$ .
- (b) Let  $r$  be an unsatisfied rule in  $(\mathcal{T}_{meta}[\Pi])^X$ .
- (b-1) If  $B(r) = \emptyset$ , we construct  $X' = X \cup H(r)$ . Thus, we get  $X' \models H(r)$ .
- (b-2) If the rule is of form  $H(r') \leftarrow ap_{n,r}(c_1, \dots, c_n), not\ ko_{n,r}$ , where  $H(r') \neq \perp$  for some  $r' \in \Pi$ , this means  $ko_{n,r} \notin X$ . However, rules for  $r'$  are added in  $\mathcal{T}_{deact}[\Pi]$  which uses the rule  $ko_{n,r}$  to deactivate the meta-rule in  $\mathcal{T}_{meta}[\Pi]$ , which is then also unsatisfied in the reduct  $(\mathcal{T}_{deact}[\Pi])^X$ . So we construct  $X' = X \setminus \{ko_{n,r}\}$ . Thus, the rule  $r$  does not appear in  $(\mathcal{T}_{meta}[\Pi])^{X'}$ .
- (b-3) Let the rule be of form  $H(r') \leftarrow ap_{n,r}(d_1, \dots, d_m), not\ ko_{n,r}$ , where  $H(r') = \perp$  for some  $r' \in \Pi$ . The rules in  $\mathcal{T}_{deactCons}[\Pi]$  are added for  $r'$ .  $X \models B(r)$  means that  $ko_{n,r} \notin X'$ . So the choice rule  $\{ko_{n,r}\}$  in  $\mathcal{T}_{deactCons}[\Pi]$  gets instantiated to  $ko'_{n,r}$  in  $(\mathcal{T}_{deactCons}[\Pi])^X$ , i.e.,  $ko'_{n,r} \in X$ . Now consider the interpretation  $X' = (X \setminus \{ko'_{n,r}\}) \cup \{ko_{n,r}\}$ . Since  $ko_{n,r} \in X'$ , the rule  $r$  does not appear in  $\mathcal{T}_{meta}[\Pi]^{X'}$ .
- (b-4) If  $r$  is of form  $ap_{n,r}(c_1, \dots, c_n) \leftarrow B(r')$  (or  $ap_{n,r}(d_1, \dots, d_m) \leftarrow B(r')$ ) for some  $r' \in \Pi$ , then we construct  $X' = X \cup \{ap_{n,r}(c_1, \dots, c_n)\}$  (or  $X' = X \cup \{ap_{n,r}(d_1, \dots, d_m)\}$ ). If  $r$  is of the remaining forms involving  $ab$  and  $bl$ , for some  $r' \in \Pi$ , we first check if  $X \models B(r')$ , then either add the respective  $ab$  atoms or the respective  $bl$  atom to  $X'$ .
- (c) Assume  $X \not\models (Q_{\hat{I}}^m)^X$ . As  $Q_{\hat{I}}^m$  has two forms of rules,  $X$  not satisfying a rule in  $(Q_{\hat{I}}^m)^X$  means that either (a) some  $\alpha \in X \cap \mathcal{A}$  exists while  $m(\alpha) \notin \hat{I}$ , or (b) for some  $\hat{\alpha} \in \hat{I}$ , no  $\alpha \in X \cap \mathcal{A}$  exists such that  $m(\alpha) = \hat{\alpha}$ .
- (c-1) The literal  $m(\alpha)$  not occurring in  $\hat{I}$  means that all rules in  $def(m(\alpha), \Pi^m)$  are not applicable for  $\hat{I}$ . This means that for each  $r \in def(\alpha, \Pi)$  we have  $B(r) \neq \emptyset$ , and for some  $r \in def(\alpha, \Pi)$  the choice rule  $\{\alpha\} \leftarrow ap_{n,r}(c_1, \dots, c_n)$  in  $\mathcal{T}_{deact}[\Pi]$  gets instantiated to  $\alpha \leftarrow ap_{n,r}(c_1, \dots, c_n)$  in  $\Pi_{debug}^X$ . Now consider the interpretation  $X' = (X \setminus \{\alpha\}) \cup \{\alpha'\}$ , for which the choice rule gets instantiated to  $\alpha' \leftarrow ap_{n,r}(c_1, \dots, c_n)$  in  $\Pi_{debug}^{X'}$ . Thus an interpretation  $X'$  is constructed where case (c-1) is avoided.
- (c-2) Similarly as above, depending on whether the original rules in  $def(\alpha, \Pi)$  for all  $\alpha \in m^{-1}(\hat{\alpha})$  are applicable or blocked for  $X$ , using the auxiliary choice rules in  $\mathcal{T}_{deact}[\Pi]$  and  $\mathcal{T}_{act}[\Pi, \mathcal{A}]$  some  $X'$  can be constructed that avoids the case of having no  $\alpha \in X'$  with  $m(\alpha) = \hat{\alpha}$ .

For the constructed  $X'$  case (a-1) may occur, and can be avoided as shown by adding  $ab$  atoms to the interpretation.

By iterating the reasoning on the constructed  $X'$ , since none of the cases will undo a construction step that is previously made, some interpretation  $X''$  that satisfies  $(\Pi_{debug} \cup Q_{\hat{I}}^m)^X$  will eventually be achieved.

(ii) If  $X$  is a minimal model, then  $X$  is an answer set of  $\Pi_{debug} \cup Q_I^m$ , which achieves the result. We assume this is not the case, and that there exists  $Y \subset X$  such that  $Y \models (\Pi_{debug} \cup Q_I^m)^X$ . As  $Y \models (Q_I^m)^X$ , we have  $m(Y \cap \mathcal{A}) = \hat{I}$ . Let  $\alpha \in X \setminus Y$ , then we have two cases: (a)  $\alpha \in \mathcal{A}$ , (b)  $\alpha \in \mathcal{A}^* \setminus \mathcal{A}$ , i.e.,  $\alpha$  is an auxiliary debugging atom.

(a) Let  $\alpha = p(c_1, \dots, c_n)$ . As  $\alpha \in X \setminus Y$  and  $\alpha \in \mathcal{A}$ , this means that there is some  $\hat{\alpha} = p(\hat{c}_1, \dots, \hat{c}_n)$  in  $\hat{I}$  and some  $\alpha_1 = p(c'_1, \dots, c'_n)$  in  $Y$  (and  $\alpha_1 \in X$ ) such that  $m(\alpha_1) = m(\alpha) = \hat{\alpha}$ . However, since  $\hat{\alpha} \in \hat{I}$  and  $\hat{I}$  is an answer set of  $\Pi^m$ , there exists a rule instantiation  $\hat{r} \in (\Pi^m)^{\hat{I}}$  such that  $H(\hat{r}) = \hat{\alpha}$ . Then, since  $\alpha, \alpha_1 \in X$  the abstract rule  $\hat{r}$  has corresponding rules  $r, r_1 \in \mathcal{T}_{meta}[\Pi]^X$  where  $r : p(c_1, \dots, c_n) \leftarrow ap_{n,r}(c_1, \dots, c_n), not\ ko_{n,r}$ . and  $r_1 : p(c'_1, \dots, c'_n) \leftarrow ap_{n,r}(c'_1, \dots, c'_n), not\ ko_{n,r}$  for some  $r', r'' \in \Pi$ .

Now, as we assume  $\alpha \notin Y$ , then  $ap_{n,r}(c_1, \dots, c_n) \notin Y$  should hold (to satisfy the rule  $r \in \mathcal{T}_{meta}[\Pi]^X$ ). Then in order to avoid not satisfying the rule  $ap_{n,r}(c_1, \dots, c_n) \leftarrow B(r')|_{c_1, \dots, c_n}$  in  $\mathcal{T}_{meta}[\Pi]^X$ , where  $B(r')|_{c_1, \dots, c_n}$  denotes that  $B(r')$  is instantiated over  $c_1, \dots, c_n$ , either (a-1) for some  $\beta \in B^+(r')|_{c_1, \dots, c_n}$  we must have  $\beta \notin Y$  or (a-2) for some  $\beta \in B^-(r')|_{c_1, \dots, c_n}$  we must have  $\beta \in Y$ . As case (a-2) contradicts  $Y \subset X$ , let us consider case (a-1). For  $\beta \notin Y$ , we do the same reasoning as above. By doing this recursive reasoning, since  $\Pi$  is tight, we eventually reach a rule in which case (a-1) is not applicable, and thus by case (a-2) we reach a contradiction.

(b) Let  $\alpha \in \mathcal{A}^* \setminus \mathcal{A}$ . If  $\alpha$  is in  $AB(\Pi)$ , then we pick  $Y$  as the interpretation, and the reduct  $(\Pi_{debug} \cup Q_I^m)^Y$  will not have further rules, since the  $ab$  atoms of  $AB(\Pi)$  do not occur in body of any auxiliary rule. If  $\alpha$  is in  $\mathcal{A}^+ \setminus \mathcal{A}$ , then we pick  $Y$  as the interpretation and apply the above reasoning for  $Y$ . When we recursively continue with this reasoning, eventually, the case (ii-b) will not be applicable, and thus we would construct a minimal model.  $\square$

Note that in the proof of Proposition 4.2, since the auxiliary programs  $\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}]$  can always have a model (as they do not contain constraints), one can start with an interpretation  $X$  which models  $(\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}])^X$ . However, since we construct another interpretation  $Y$  from  $X$ , a check on whether  $Y$  models  $(\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}])^Y$  then becomes necessary.

**Proof of Theorem 5.1.** For an assignment  $I$ , we need to show that  $I \cup \mathcal{T}_{m\exists}$  is a minimal model of  $(\Pi_{\exists}^m)^I$  if and only if  $I \cup \mathcal{T}_m$  is a minimal model of  $(\Pi^m)^I$ .

( $\Rightarrow$ ) Towards a contradiction, assume  $I \cup \mathcal{T}_{m\exists}$  is a minimal model of  $(\Pi_{\exists}^m)^I$  but  $I \cup \mathcal{T}_m$  is not a minimal model of  $(\Pi^m)^I$ . Then either (i)  $I \cup \mathcal{T}_m$  is not a model of  $(\Pi^m)^I$ , or (ii)  $I \cup \mathcal{T}_m$  is not a minimal model of  $(\Pi^m)^I$ .

(i) There is a rule  $\hat{r} \in (\Pi^m)^I$  such that  $I \cup \mathcal{T}_m \models B(\hat{r})$  but  $I \not\models H(\hat{r})$ . By construction of  $\Pi^m$ ,  $\hat{r}$  is only obtained by step (a) of Definition 3.3, otherwise  $\hat{r}$  would be a choice rule with head  $H(\hat{r}) = \{m(I)\}$ , and  $\hat{r}$  would be satisfied. Consequently  $\hat{r}$  is a rule from step (a) for  $r$  in  $\Pi$ . Thus, we have  $I \cup \mathcal{T}_m \models m(B(r)), rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ . Since the definitions of relation type I for lifted relations and abstract relations correspond to each other, we have  $\mathcal{T}_m \models \tau_1^{rel}(\hat{d}_1, \hat{d}_2) \iff \mathcal{T}_{m\exists} \models \tau_1^{\widehat{rel}}(\hat{d}_1, \hat{d}_2)$ . This means we get  $I \cup \mathcal{T}_{m\exists} \models m(B(r)), \tau_1^{\widehat{rel}}(\hat{d}_1, \hat{d}_2)$  which is the abstract rule of  $r$  constructed by step (a) of Definition 5.1. Since  $I \cup \mathcal{T}_{m\exists}$  is a minimal model of  $(\Pi_{\exists}^m)^I$ ,  $I \models H(\hat{r})$ . Hence, we reach a contradiction.

(ii) Let there be  $J \subset I$  such that  $J \cup \mathcal{T}_m$  is a model of  $(\Pi^m)^I$ . We claim that  $J \cup \mathcal{T}_{m\exists}$  is a model of  $(\Pi_{\exists}^m)^I$ , which would contradict  $I \cup \mathcal{T}_{m\exists} \in AS(\Pi_{\exists}^m)$ . Assume  $J \cup \mathcal{T}_{m\exists} \not\models (\Pi_{\exists}^m)^I$ . Then there is a rule  $\hat{r} \in (\Pi_{\exists}^m)^I$

such that  $J \cup \mathcal{T}_{m_{\exists}} \models B(\hat{r})$  but  $J \not\models H(\hat{r})$ , while  $I \models H(\hat{r})$ . We need to show that there is a corresponding rule in  $(\Pi^m)^I$  for  $\hat{r}$ , which would then achieve the contradiction that is  $J \models H(\hat{r})$ . Below, we denote by  $B(\hat{r}) \setminus \hat{\Gamma}_r$ , the abstract body excluding the abstracted relation (and its relation type atom).

- If  $\hat{r}$  contains  $\tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  (step (a) or (c) of Definition 5.1), then since we know  $\mathcal{T}_m \models \tau_1^{rel}(\hat{d}_1, \hat{d}_2) \iff \mathcal{T}_{m_{\exists}} \models \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$ , we achieve  $J \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2)$ ,  $\tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  (also  $J \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2)$ ). Since  $\hat{r} \in (\Pi_{\exists}^m)^I$ , we know that  $I \models B(\hat{r})$  and also  $I \models B(\hat{r}) \setminus \hat{\Gamma}_r$ . Thus we get  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  (also  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, rel(\hat{d}_1, \hat{d}_2)$ ) in  $(\Pi^m)^I$ . Since  $J \cup \mathcal{T}_m$  is a model of  $(\Pi^m)^I$ , we get  $J \models H(\hat{r})$ , which is a contradiction.
- If  $\hat{r}$  contains  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  (step (b) or (c) of Definition 5.1), then  $J \cup \mathcal{T}_{m_{\exists}} \models \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  means that there exist some  $d_{11}, d_{12} \in m^{-1}(\hat{d}_1)$ ,  $d_{21}, d_{22} \in m^{-1}(\hat{d}_2)$  and some  $J' \in m^{-1}(J)$  such that  $J' \models rel(d_{11}, d_{21})$  and  $J' \not\models rel(d_{12}, d_{22})$ . There are the following cases for  $rel(\hat{d}_1, \hat{d}_2)$ : (1)  $J \models rel(\hat{d}_1, \hat{d}_2)$ , or (2)  $J \not\models rel(\hat{d}_1, \hat{d}_2)$ .
  - (1) Since we know  $J' \not\models rel(d_{12}, d_{22})$ , this case obtains  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , thus  $J \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ . With similar reasoning as above on obtaining  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  in  $(\Pi^m)^I$  (also  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, rel(\hat{d}_1, \hat{d}_2)$  in  $(\Pi^m)^I$ ), we achieve  $J \models H(\hat{r})$ , a contradiction.
  - (2) Since we know  $J' \models rel(d_{11}, d_{21})$ , this case obtains  $\tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$ , thus  $J \cup \mathcal{T}_m \models \overline{rel}(\hat{d}_1, \hat{d}_2), \tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$ . With similar reasoning as above we reach a contradiction.

( $\Leftarrow$ ) Towards a contradiction, assume  $I \cup \mathcal{T}_m$  is a minimal model of  $(\Pi^m)^I$  but  $I \cup \mathcal{T}_{m_{\exists}}$  is not a minimal model of  $(\Pi_{\exists}^m)^I$ . Then either (i)  $I \cup \mathcal{T}_{m_{\exists}}$  is not a model of  $(\Pi_{\exists}^m)^I$ , or (ii)  $I \cup \mathcal{T}_{m_{\exists}}$  is not a minimal model of  $(\Pi_{\exists}^m)^I$ .

- (i) There is a rule  $\hat{r} \in (\Pi_{\exists}^m)^I$  such that  $I \cup \mathcal{T}_{m_{\exists}} \models B(\hat{r})$  but  $I \not\models H(\hat{r})$ . By construction of  $\Pi_{\exists}^m$ ,  $\hat{r}$  is only obtained by step (a) of Definition 5.1. With an analogous reasoning as above item (i), we achieve a contradiction.
- (ii) Let there be  $J \subset I$  such that  $J \cup \mathcal{T}_{m_{\exists}}$  is a model of  $(\Pi_{\exists}^m)^I$ . We claim that  $J \cup \mathcal{T}_m$  is a model of  $(\Pi^m)^I$ , which would contradict  $I \cup \mathcal{T}_m \in AS(\Pi^m)$ . Assume  $J \cup \mathcal{T}_m \not\models (\Pi^m)^I$ . Then there is a rule  $\hat{r} \in (\Pi^m)^I$  such that  $J \cup \mathcal{T}_m \models B(\hat{r})$  but  $J \not\models H(\hat{r})$ , while  $I \models H(\hat{r})$ . We need to show that there is a corresponding rule in  $(\Pi_{\exists}^m)^I$  for  $\hat{r}$ , which would then achieve the contradiction that  $J \models H(\hat{r})$ .

- If  $\hat{r}$  contains  $rel(\hat{d}_1, \hat{d}_2), \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  (step (a) of Definition 3.3), an analogous reasoning as above item (ii) obtains  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  in  $(\Pi_{\exists}^m)^I$  which achieves  $J \models H(\hat{r})$  a contradiction.
- If  $\hat{r}$  contains  $rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  (step (b) of Definition 3.3), then  $J \cup \mathcal{T}_m \models rel(\hat{d}_1, \hat{d}_2), \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  means that  $J \models rel(\hat{d}_1, \hat{d}_2)$  and there exist some  $d_1 \in m^{-1}(\hat{d}_1)$ ,  $d_2 \in m^{-1}(\hat{d}_2)$  and some  $J' \in m^{-1}(J)$  such that  $J' \not\models rel(d_1, d_2)$ . This obtains abstract relation type  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , thus  $J \cup \mathcal{T}_{m_{\exists}} \models \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ . Notice that also  $J \models isCluster(\hat{d}_i)$  holds for some  $i \in \{1, 2\}$ . With similar reasoning as above on obtaining  $H(\hat{r}) \leftarrow B(\hat{r}) \setminus \hat{\Gamma}_r, \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  in  $(\Pi_{\exists}^m)^I$ , we achieve  $J \models H(\hat{r})$ , a contradiction.
- If  $\hat{r}$  contains  $\overline{rel}(\hat{d}_1, \hat{d}_2), \tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$  (step (c) or (d-ii) of Definition 3.3), then  $J \cup \mathcal{T}_{m_0} \models \overline{rel}(\hat{d}_1, \hat{d}_2), \tau_{IV}^{rel}(\hat{d}_1, \hat{d}_2)$  means that  $J \not\models rel(\hat{d}_1, \hat{d}_2)$  and there exist some  $d_1 \in m^{-1}(\hat{d}_1)$ ,  $d_2 \in m^{-1}(\hat{d}_2)$  and some  $J' \in m^{-1}(J)$  such that  $J' \models rel(d_1, d_2)$ . This again obtains abstract relation type  $\tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , i.e.,  $J \cup \mathcal{T}_{m_{\exists}} \models \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$ , thus reaches a contradiction as above.
- If  $\hat{r}$  contains only  $rel(\hat{d}_1, \hat{d}_2)$  (step (d-i) of Definition 3.3), then this means either  $J \cup \mathcal{T}_{m_{\exists}} \models \tau_1^{rel}(\hat{d}_1, \hat{d}_2)$  or  $J \cup \mathcal{T}_{m_{\exists}} \models \tau_{III}^{rel}(\hat{d}_1, \hat{d}_2)$  holds. Also we know that  $J \models isCluster(\hat{d}_i)$  holds for some  $i \in \{1, 2\}$ . So similar as above, we achieve a contradiction.  $\square$

## B Further Details

### B.1 Correctness Checking Failure for Non-tight Programs

The following example shows that the procedure for correctness checking does not work for non-tight programs in general, i.e., in case of a positive loop.

**Example B.1.** Consider the program below which is unsatisfiable and also contains a positive loop.

$$\begin{aligned} a(X) &\leftarrow \text{not } a(X), \text{int}(X). \\ a(X) &\leftarrow a(X), \text{int}(X). \\ \text{int}(1). \text{int}(2). \text{int}(3). \end{aligned}$$

For the mapping  $m = \{\{1, 2, 3\}/k\}$ , the constructed abstract program  $\Pi^m$  is

$$\begin{aligned} a(X) &\leftarrow \text{not } a(X), \widehat{\text{int}}(X). \\ \{a(X)\} &\leftarrow \text{isCluster}(X), \widehat{\text{int}}(X). \\ a(X) &\leftarrow a(X), \widehat{\text{int}}(X). \\ \widehat{\text{int}}(k). \end{aligned}$$

which has the abstract answer set  $\hat{I} = \{a(k)\}$ . Checking the correctness using  $\Pi_{\text{debug}} \cup Q_I^m$  results in unsatisfiability, because it requires to have some  $a(c)$  for  $c \in m^{-1}(k)$  to hold true through a loop, which is not covered in the definition of  $\Pi_{\text{debug}}$ .

### B.2 Grid-Cell Problem Encodings

In this appendix, we provide details about the slight modifications made from the existing (or common) encodings, in order to use them in our experiments. The full encodings can be found in [www.kr.tuwien.ac.at/research/systems/abstraction/mdaspar.material.zip](http://www.kr.tuwien.ac.at/research/systems/abstraction/mdaspar.material.zip).

**Sudoku** We used the encoding from DLV group in ASPCOMP09 with slight modifications. The guessing of the assignment of numbers to the free cells is written as

$$\begin{aligned} \{sol(X, Y, N) : \text{num}(N)\} &\leftarrow \text{notoccupied}(X, Y), \text{row}(X), \text{column}(Y). \\ \text{hasNum}(X, Y) &\leftarrow sol(X, Y, N). \\ \perp &\leftarrow \text{not hasNum}(X, Y), \text{row}(X), \text{column}(Y). \end{aligned}$$

The constraints of assigning one symbol per column and one symbol per row are the same as in the original encoding, but with standardizing apart over the sorts *row* and *column*.

$$\begin{aligned} \perp &\leftarrow sol(X, Y_1, M), sol(X_2, Y_2, M), X = X_2, Y_1 < Y_2. \\ \perp &\leftarrow sol(X_1, Y, M), sol(X_2, Y_2, M), X_1 < X_2, Y = Y_2. \end{aligned}$$

For the constraint of assigning one symbol per subregion, standardizing apart the original rules caused to have relations with many argument, thus we converted them into the rules

$$\begin{aligned} \perp &\leftarrow sol(X_1, Y_1, M), sol(X_2, Y_2, M), \\ &\quad sameSubSquareLessThan(X_1, Y_1, X_2, Y_2). \\ sameSubSquareLessThan(X_1, Y_1, X_2, Y_2) &\leftarrow sameSubSquare(X_1, Y_1, X_2, Y_2), X_1 < X_2. \\ sameSubSquareLessThan(X_1, Y_1, X_2, Y_2) &\leftarrow sameSubSquare(X_1, Y_1, X_2, Y_2), Y_1 < Y_2. \\ sameSubSquare(X_1, Y_1, X_2, Y_2) &\leftarrow subrangeR(X_1, M), subrangeR(X_2, M), \\ &\quad subrangeC(Y_1, R), subrangeC(Y_2, R). \end{aligned}$$

with the hardcoded facts  $subrangeR(X, M)$  and  $subrangeC(Y, R)$  for subregions w.r.t. rows and columns, respectively.

**Knight's Tour** We used the encoding from ASPCOMP11<sup>11</sup> with slight modifications. At most one *move* atom is made for each *valid* movement among the cells.

$$\{move(X_1, Y_1, X_2, Y_2)\}1 \leftarrow valid(X_1, Y_1, X_2, Y_2).$$

In the original encoding, the valid cells computations were done using rules of the form

$$valid(X_1, Y_1, X_2, Y_2) \leftarrow point(X_1, Y_1), point(X_2, Y_2), X_1 = X_2 + 2, Y_1 = Y_2 + 1.$$

which are modified as

$$\begin{aligned} validcell(X_1, Y_1, X_2, Y_2) &\leftarrow dist1(X_1, X_2), dist2(Y_1, Y_2). \\ validcell(X_1, Y_1, X_2, Y_2) &\leftarrow dist2(X_1, X_2), dist1(Y_1, Y_2). \\ valid(X_1, Y_1, X_2, Y_2) &\leftarrow validcell(X_1, Y_1, X_2, Y_2), point(X_1, Y_1), point(X_2, Y_2). \end{aligned}$$

where the auxiliary facts  $dist1(X_1, X_2), dist2(X_1, X_2)$  represent the arithmetic operations  $X_1 = X_2 + 2, Y_1 = Y_2 + 1$ .

The constraints to ensure that exactly one entering/leaving movement is made for each cell is the same as the original encoding. Having each cell visited is ensured by the following rules

$$\begin{aligned} reached(X, Y) &\leftarrow move(X_1, Y_1, X, Y), start(X_1, Y_1). \\ reached(X_2, Y_2) &\leftarrow reached(X_1, Y_1), move(X_1, Y_1, X_2, Y_2). \\ \perp &\leftarrow point(X, Y), not\ reached(X, Y), row(X), column(Y). \end{aligned}$$

where the atom  $start(X, Y)$  is used to show the starting point, instead of having in the rule the atom  $move(1, 1, X, Y)$  as it is originally. This change makes treating the program more convenient, as the rules do not contain constants that need to be mapped to different abstract constants depending on the mapping.

<sup>11</sup>[www.mat.unical.it/aspcomp2011/files/KnightTour/knight\\_tour.enc.asp](http://www.mat.unical.it/aspcomp2011/files/KnightTour/knight_tour.enc.asp)

**Visitall** We encoded a planning problem along the guidelines in Section 7.2 on representing actions and change. We considered  $go(X, Y, T)$  actions that can move horizontally/vertically to a cell  $X, Y$ . For such an action, we must ensure that the action does not pass through an obstacle or a previously visited cell, and that all the passed cells become visited. A common way of encoding this is to have auxiliary atoms that keep track of the cells that are in between, such as

$$\begin{aligned}
aux\_passed(X, Y_2, T + 1) &\leftarrow rAt(X, Y, T), go(X, Y_1, T), Y < Y_2, Y_2 \leq Y_1. \\
aux\_passed(X, Y_2, T + 1) &\leftarrow rAt(X, Y, T), go(X, Y_1, T), Y_1 < Y_2, Y_2 \leq Y. \\
aux\_passed(X_2, Y, T + 1) &\leftarrow rAt(X, Y, T), go(X_1, Y, T), X < X_2, X_2 \leq X_1. \\
aux\_passed(X_2, Y, T + 1) &\leftarrow rAt(X, Y, T), go(X_1, Y, T), X_1 < X_2, X_2 \leq X. \\
passed(X, Y) &\leftarrow aux\_passed(X, Y, T).
\end{aligned}$$

which are then used to ensure the above conditions.

$$\begin{aligned}
\perp &\leftarrow passed(X, Y), obsAt(X, Y). \\
visited(X, Y, T) &\leftarrow aux\_passed(X, Y, T). \\
\perp &\leftarrow aux\_passed(X, Y, T + 1), visited(X, Y, T).
\end{aligned}$$

We follow the remark in Section 5 on handling different abstraction levels on variables in a rule. For example, for the first rule, in addition to the standardizing apart the rule as

$$aux\_passed(X, Y_2, T + 1) \leftarrow rAt(X, Y, T), go(X_1, Y_1, T), X = X_1, Y < Y_2, Y_2 \leq Y_1.$$

we add the additional rule

$$aux\_passed(X_1, Y_2, T + 1) \leftarrow rAt(X, Y, T), go(X_1, Y_1, T), X = X_1, Y < Y_2, Y_2 \leq Y_1.$$

We proceed similarly with the remaining rules.

Furthermore, knowing that the action  $go(X_1, Y_1, T)$  will only be picked in a horizontal (resp. vertical) direction of  $rAt(X, Y, T)$ , we drop the relation  $X = X_1$  (resp.  $Y = Y_1$ ) from the body to make it smaller.

### B.3 Example run of mDASPAR

To illustrate further considerations for debugging and refinement, we show an example run of mDASPAR.

**Example B.2** (ctd). We run mDASPAR with the input program (Figure 15) and the instance shown in Figure 13, with the initial mapping  $m$  of clustering the grid-cell into four regions (Figure 14a).

**step 1** After constructing the non-ground abstract program (Figure 16) and computing the relation types, mDASPAR computes an abstract answer set

$$\{reachable(a_{1234}, b_{1234}), reachable(a_{5678}, b_{1234}), reachable(a_{1234}, b_{5678})\}.$$

**step 2** Correctness checking first uses  $\Pi_{debug_0}$  where the  $ab$  atoms only contain rule names (Figure 28) to obtain the optimal answer set

$$\{ab\_deactCons_{r_3}, ab\_deact_{r_2}\}$$

Figure 28: Constructed debugging program  $\mathcal{T}_{deact}[\Pi] \cup \mathcal{T}_{deactCons}[\Pi] \cup \mathcal{T}_{act}[\Pi, \mathcal{A}]$

$$\begin{aligned}
& ko_{r1}. \\
& \{reachable(X, Y)\} \leftarrow ap_{r1}(X, Y). \\
& ab\_deact_{r1} \leftarrow ap_{r1}(X, Y), not\ reachable(X, Y). \\
& \perp : \sim ab\_deact_{r1}.[1] \\
& ko_{r2}. \\
& \{reachable(X_1, Y_1)\} \leftarrow ap_{r2}(X_1, Y_1). \\
& ab\_deact_{r2} \leftarrow ap_{r2}(X_1, Y_1), not\ reachable(X_1, Y_1). \\
& \perp : \sim ab\_deact_{r2}.[1] \\
& ko_{r3}. \\
& ab\_deactCons_{r3} \leftarrow ko_{r3}, ap_{r3}(X, Y, X_1, Y_1). \\
& \perp : \sim ab\_deactCons_{r3}.[1] \\
& ko_{r5}. \\
& \{neighbor(X, Y, X_1, Y_1)\} \leftarrow ap_{r5}(X, Y, X_1, Y_1). \\
& ab\_deact_{r5} \leftarrow ap_{r5}(X, Y, X_1, Y_1), not\ neighbor(X, Y, X_1, Y_1). \\
& \perp : \sim ab\_deact_{r5}.[1] \\
& ko_{r6}. \\
& \{neighbor(X, Y, X_1, Y_1)\} \leftarrow ap_{r6}(X, Y, X_1, Y_1). \\
& ab\_deact_{r6} \leftarrow ap_{r6}(X, Y, X_1, Y_1), not\ neighbor(X, Y, X_1, Y_1). \\
& \perp : \sim ab\_deact_{r6}.[1] \\
& \{neighbor(X, Y, X_1, Y_1)\} \leftarrow bl_{r5}(X, Y, X_1, Y_1), bl_{r6}(X, Y, X_1, Y_1). \\
& ab\_act(neighbor(X, Y, X_1, Y_1)) \leftarrow bl_{r5}(X, Y, X_1, Y_1), bl_{r6}(X, Y, X_1, Y_1), \\
& \quad neighbor(X, Y, X_1, Y_1). \\
& \perp : \sim ab\_act(neighbor(X, Y, X_1, Y_1)).[1, X, Y, X_1, Y_1] \\
& \{reachable(X, Y)\} \leftarrow bl_{r1}(X, Y), bl_{r2}(X, Y). \\
& ab\_act(reachable(X, Y)) \leftarrow bl_{r1}(X, Y), bl_{r2}(X, Y), reachable(X, Y). \\
& \perp : \sim ab\_act(reachable(X, Y)).[1, X, Y]
\end{aligned}$$

**step 3**  $\Pi_{debug}$  is constructed only for  $r2, r3$  now respecting variables in the rules, and by defining  $ab\_deact_{r2}(X_1, Y_1)$  as

$$\begin{aligned}
ab\_deact_{r2}(X_1, Y_1) &\leftarrow ap_{r2}(X_1, Y_1), not\ reachable(X_1, Y_1). \\
\perp &:\sim ab\_deact_{r2}(X_1, Y_1). [1, X_1, Y_1] \\
\perp &:\sim ab\_deact_{r2}(X_1, Y_1), mapTo(X_1, Y_1, A_1, B_1), \\
&\quad isSingleton(A_1), isSingleton(B_1). [20, X_1, Y_1] \\
refine(A_1, B_1) &\leftarrow ab\_deact_{r2}(X_1, Y_1), mapTo(X_1, Y_1, A_1, B_1), isCluster(A_1). \\
refine(A_1, B_1) &\leftarrow ab\_deact_{r2}(X_1, Y_1), mapTo(X_1, Y_1, A_1, B_1), isCluster(B_1).
\end{aligned}$$

and similarly  $ab\_deact_{r3}(X, Y, X_1, Y_1)$ . Correctness checking finds an optimal answer set with the atoms

$$refine(a_{1234}, b_{5678}), refine(a_{5678}, b_{5678}).$$

**step 4** The region  $(a_{1234}, b_{5678})$  is randomly picked to refine to  $\{(a_{12}, b_{56}), (a_{12}, b_{78}), (a_{34}, b_{56}), (a_{34}, b_{78})\}$  by updating the corresponding mapping  $m$ .

**step 5** Relation types according to the new mapping are computed and the loop goes back to step 1 to compute a new abstract answer set.

The loop continues until unsatisfiability is achieved. The abstraction shown in Figure 14b is one such abstraction where unsatisfiability is observed.

## B.4 Modular Concreteness Checking

**Ordered modularity** An incremental approach to ASP solving proposed by Gebser et al. (Gebser et al., 2008) builds on the concept of modules (Oikarinen and Janhunen, 2006) and gradually increases the bound to the solution size, represented by a parameter  $k$ , to help with both grounding and solving. they are searching for an answer set with minimum size over  $k$ , thus they increment the parameter until an answer set is computed. We use a similar idea to detect the spuriousness of an abstract answer set by gradually increasing the parameter. However, in our case, the increment is continued until the spuriousness is realized with debugging, i.e., an answer set with an abnormality atom is obtained. We take a simpler view by limiting the generated grounding of the program to the parameter.

Let  $\Pi$  be a program with the Herbrand base  $HB_{\Pi} = \mathcal{L}_B \cup \mathcal{L}_k$ , for parameter  $k$  ranging over the natural numbers, where  $\mathcal{L}_B$  represents the *static* literals with arguments independent of parameter  $k$ , and  $\mathcal{L}_k$  represents the *dynamic* literals which have an argument  $k$ . For a set  $X$  of literals, we denote by  $grd(\Pi)|_X = \{r \mid H(r) \cup B(r) \subseteq X\}$  the set of ground rules that contain only literals over  $X$ . Let  $X_i \subseteq HB_{\Pi}$  denote the set of literals until the parameter value  $i$ , i.e.,  $X_i = \mathcal{L}_B \cup \bigcup_{j=0}^i \mathcal{L}_{k/j}$ , where  $\mathcal{L}_{k/j}$  denotes the set of literals with the respective argument of value  $j$ . The rules of  $grd(\Pi)$  until parameter value  $i$  are then given by  $grd(\Pi)|_{X_i}$ , simply denoted  $grd(\Pi)|_i$ .

Let  $I_{\leq i}$  denote the projection of an interpretation  $I$  to the literals related with the parameter value  $i$ , i.e.,  $I_{\leq i} = I \cap (\mathcal{L}_B \cup \bigcup_{0 \leq j \leq i} \mathcal{L}_j)$ . We say that  $\Pi$  is *ordered modular*, if for each  $I \in AS(\Pi)$  it holds that  $I_{\leq i} \in AS(grd(\Pi)|_i)$  for all  $0 \leq i \leq k$ . We then know that determining the occurrence of a literal  $l$  in an answer set  $I_{\leq i}$  relies only on the decisions made until point  $i$ .

**Proposition B.1.** *Let  $\Pi$  be an ordered modular program,  $m$  a domain mapping for  $\Pi$ , and let  $\hat{I} \in AS(\Pi^m)$ . If  $\hat{I}_{\leq i} \subseteq \hat{I}$  is spurious for some  $i \leq n$ , then  $\hat{I}$  is spurious.*

*Proof.* Assume  $\hat{I}$  is concrete. This means that there exists some  $I \in \Pi$  such that  $m(I) = \hat{I}$ . As  $\Pi$  is ordered modular,  $I_{\leq i} \in AS(\text{grd}(\Pi)|_i)$ . Thus,  $m(I_{\leq i}) = \hat{I}_{\leq i}$  is concrete.  $\square$

We describe in detail the implementation of these approaches.

**Incrementing time.** Approach (1) is implemented in mDASPAR to handle planning problems with atoms having *time* arguments. By Proposition B.1, we know that if the first few actions of a potential plan described by the abstract answer set have no corresponding original plan, one can conclude that this plan is spurious.

A common description of the planning problem in ASP uses two sorts for time:  $timea = \{0, \dots, n-1\}$ , which is used for action atoms, and  $time = \{0, \dots, n\}$  which is used to for the fluents. For a given program  $\Pi$  with a description of a planning problem that contains facts for sort *time*, mDASPAR works as follows. We denote by  $\mathcal{T}_{meta}[\Pi]|_i$  the meta-program  $\mathcal{T}_{meta}[\Pi]$  which contains *time* facts (resp. *timea*) until domain element  $i$  (resp.  $i-1$ ), and by  $\Pi_{debug}|_i$  the analogous restricted version of  $\Pi_{debug}$ . For a computed abstract answer set  $\hat{I}$  encoding a plan  $\langle s_0, a_0, s_1, \dots, s_{n-1}, a_{n-1}, s_n \rangle$  we denote by  $\hat{I}|_i$  the part of the plan until time point  $i$ .

Starting with  $i = 1$ , we continue the iteration below while  $i \leq n$ .

**step (1)** Create  $\mathcal{T}_{meta}[\Pi]|_{i-1}$  and  $\Pi_{debug}|_i$ .

**step (2)** Check correctness of  $\hat{I}|_i$  with  $\Pi_{debug}|_i \cup \mathcal{T}_{meta}[\Pi]|_{i-1} \cup \mathcal{Q}_{\hat{I}|_i}^m$ .

**step (3)** If  $\hat{I}|_i$  is spurious, exit loop; otherwise, increase  $i$  by 1.

This way, we check the correctness of  $\hat{I}$  for the action taken at time  $i$ , by debugging only for time point  $i$  as the guessing for time points  $t < i$  is restricted by using  $\mathcal{T}_{meta}[\Pi]|_{i-1}$ . The time is increased incrementally while the partial solution yields a concrete partial plan. Once spuriousness is observed, the checking is stopped.

**Partial concretization.** For Approach (2), we use the possibility to have a hierarchy of abstractions mentioned in Proposition 3.6. The idea is to partially concretize the abstract domain, by fully concretizing certain regions and keeping the remaining ones abstract. Figure 17 shows the hierarchy of some partial concretizations of the initial mapping. For a given mapping  $m$ , we consider a set of possible partial concretizations. We then check the correctness of an abstract answer set  $I$  over the program with partially concretized domain. As the latter still describes an abstraction compared to the original domain, this check can not be immediately done over the original program. For that, we have must check correctness with debugging over the abstract program relative to the partial concretization.

The approach is implemented in mDASPAR as follows. For a given mapping  $m$ , starting with  $j = 1$ , the iteration focuses on concretizing  $j$  regions at a time, and checks the correctness in each such  $j$ -region combination. The iteration continues until spuriousness is detected or  $m = m_{id}$ :

**step (1)** Compute  $j$ -region concretizations of  $m$ , say  $m_1, \dots, m_n$ .

**step (2)** For every  $m_i \in \{m_1, \dots, m_n\}$ ;

1. Create  $\Pi^{m_i}$  with  $\mathcal{T}_{m_i}$  and the set  $\{m_i(p(c)).|p(c). \in \Pi\}$  of facts, and  $\Pi_{debug}^{m_i}$ .
2. Create the mapping  $m'$  such that  $m'(m_i(D)) = m(D)$ .
3. Check correctness of  $I$  with  $\Pi_{debug}^{m_i} \cup \mathcal{Q}_I^{m'}$ .
4. If spurious, exit loop with debug answer  $C$ .

**step (3)** If  $C \neq \emptyset$ , refine  $m$  according to  $C$  and go back to step (1); otherwise, increase  $j$  by 1, and go back to step (1)

We do correctness checking on the abstract level  $m_i$  using  $\Pi^{m_i}$ . If  $I$  is concrete w.r.t. the partially concretized abstraction, this does not guarantee that  $I$  is concrete; thus, the concretization is increased to redo the check. If spuriousness is detected, the mapping is refined and the partial concretization continues from the updated mapping.

## C Use Case: Abstraction for Policy Refutation

As a further example, consider checking whether an agent always manages to find a missing person with a given policy in a grid environment with obstacles. If the policy does not work, a counterexample trajectory over some part of the environment will reveal this; by inspecting the latter, one may guess why it fails. Depending on the problem, the focus points may have different nature. For the reachability example shown in Section 5.2, the focus area in the environment can remain local, while for the person search example the path of a trajectory needs to be distinguished.

For illustration, we use the following running example.

**Example C.1** (Example 5.8 cont'd). Consider the reachability problem described in the following encoding, where reachability (48-50) is redefined by prioritizing the east neighbor over the rest, and in case the east neighbor has an obstacle, choosing the south neighbor.

$$point(X, Y) \leftarrow not\ obsAt(X, Y), row(X), column(Y). \quad (78)$$

$$reachable(X, Y) \leftarrow start(X, Y). \quad (79)$$

$$neighborE(X, Y, X_1, Y) \leftarrow X_1 - X = 1, column(Y). \quad (80)$$

$$neighborS(X, Y, X, Y_1) \leftarrow Y_1 - Y = 1, row(X). \quad (81)$$

$$reachableE(X, Y, X_1, Y) \leftarrow reachable(X, Y), point(X_1, Y), \quad (82)$$

$$neighborE(X, Y, X_1, Y).$$

$$hasNeighborE(X, Y) \leftarrow reachableE(X, Y, X_1, Y_1). \quad (83)$$

$$reachable(X_1, Y_1) \leftarrow reachableE(X, Y, X_1, Y_1). \quad (84)$$

$$reachable(X, Y_1) \leftarrow reachable(X, Y), point(X, Y_1), \quad (85)$$

$$neighborS(X, Y, X, Y_1), not\ hasNeighborE(X, Y).$$

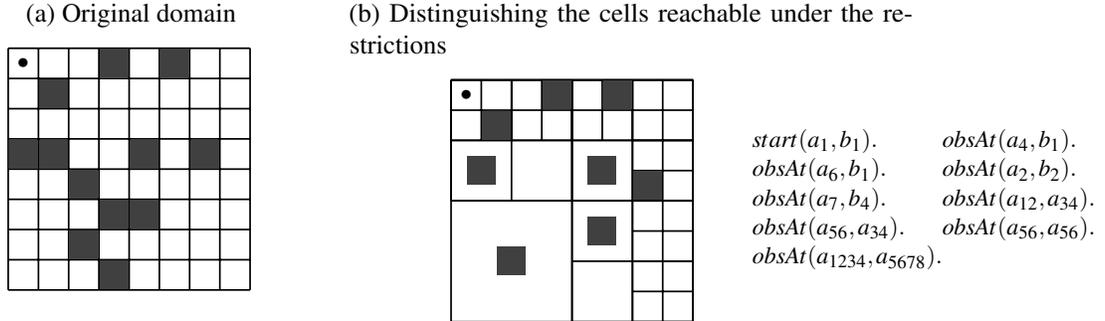
For the instance shown in Figure 29a the reachable cells are determined in the order  $\rightarrow^2 \downarrow \rightarrow^5 \downarrow^6$ . The abstraction shown in Figure 29b singles out the area that contains the cells that are reachable according to the restrictions.

We now focus on using the abstraction over grid-cells for the problem of checking policies on whether they manage to guide the agent towards the goal. We consider two versions of this problem and discuss the use of abstraction.

As a running example, we consider an agent trying to find its way in a maze towards a goal point (similar in spirit to the example of finding a missing person). For representing and generating the mazes, we use an altered version of the Maze Generation encoding from ASP Competition 2011.<sup>12</sup> A policy that may come to

<sup>12</sup><https://www.mat.unical.it/aspcomp2011/files/MazeGeneration/maze-generation.enc.asp>

Figure 29: Reachability abstraction under east neighbor prioritization



one’s mind when talking about mazes is the well-known “right-hand rule”, which is known to work in many maze instances, except when the goal is in the middle area and the agent is forced to loop due to the obstacle layout.

**Does the policy work on a given instance?** For fixed problem instances, this check is done by a search of a counterexample trajectory which follows the policy but does not reach the goal. If none is found (i.e., unsatisfiability is achieved), we conclude that the policy works for this instance. Abstraction can be used to focus on the part of the instance which is enough to show that the policy fails or works; notice that the latter case becomes similar to having unsatisfiable problems. The necessary granularity of the abstraction depends on the complexity of the policy. As demonstrated in Figure 31, for refuting the “right-hand rule” policy the abstraction must refine at least the outer area (if not more).

To observe how the policy type affects the resulting abstraction, we did some experiments. To help with the refinement decisions, the initial abstraction distinguishes the starting point of the agent and abstracts over the rest.

We consider the following two policies:

- (A) Right-hand rule: Follow the wall on the right-hand side.
- (B) Naive policy: Choose the direction to move to with the priority right > down > left > up.

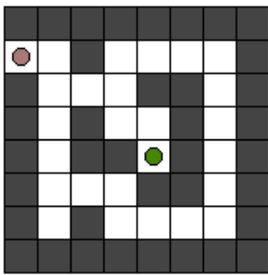
We generated 20 instances where on some of them both, one, or none of the policies work. For the debugging method we picked time increment, as we wanted it to focus on each step of the abstract trajectory starting from the beginning, and on whether or not the steps match the policy’s decisions in the corresponding original trajectory. Furthermore, the refinement decision is made only from one abstract answer set, to avoid that a concrete answer set is encountered among spurious ones; this would finalize the search and achieve a clearly non-faithful b abstraction.

Table 6 shows the results of using mDASPAR to achieve an abstraction with a concrete solution. Obtaining SAT means that the program found a concrete solution, i.e., a concrete counterexample trajectory, which shows that the policy does not work, while having UNSAT means that the policy works. As expected, the naive policy failed to work for most of the instances. Since the right-hand rule forces to traverse the environment more, mDASPAR required to have finer abstractions to figure out the concrete solution. In both cases, the obtained abstractions were not too distant from the best possible ones, although still sometimes the focus was shifted to the irrelevant parts of the grid. All of the obtained resulting abstractions were faithful, which means that they were able to show the actual behavior of the policy. Figure 30 shows the resulting abstractions for three of the instances.

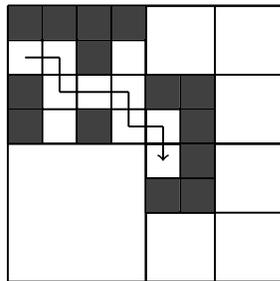
Table 6: Policy checking in maze instances

	sat/unsat	avg. step	avg. cost	avg. best step	avg. best cost
naive	16/4	7.2	0.391	6.5	0.362
right-hand	6/14	12.5	0.630	11.8	0.599

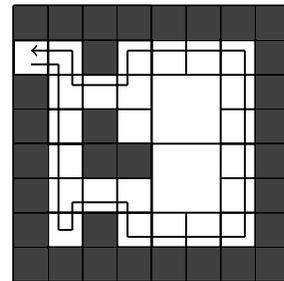
Figure 30: Abstractions on policy checking in maze instances (with the supporting/refuting paths)



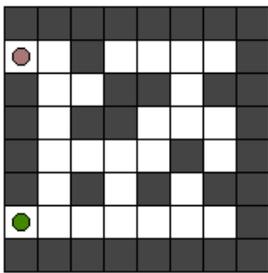
(a) naive: works, right-hand: fails



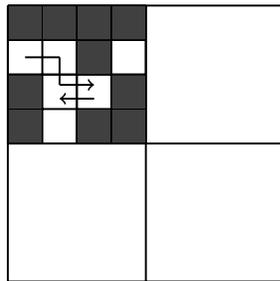
(b) Supporting the naive policy



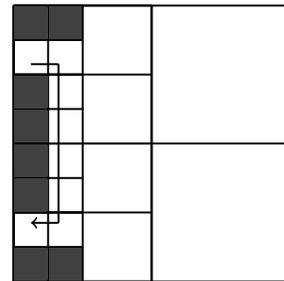
(c) Refuting the right-hand policy



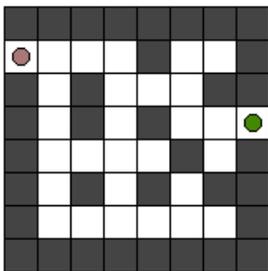
(d) naive: fails, right-hand: works



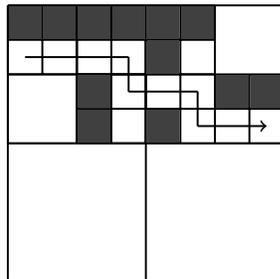
(e) Refuting the naive policy



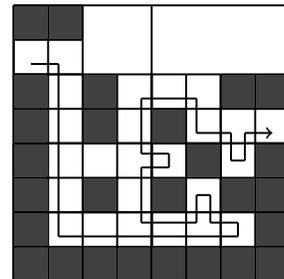
(f) Supporting the right-hand policy



(g) naive: works, right-hand: works

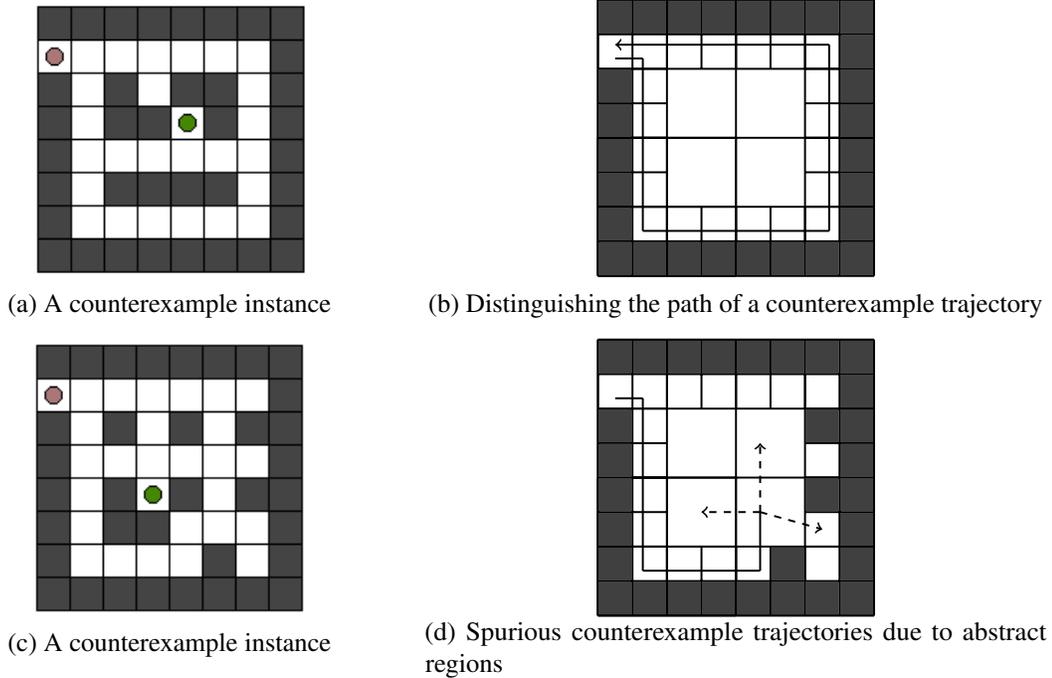


(h) Supporting the naive policy



(i) Supporting the right-hand policy

Figure 31: Can we refute the right-hand rule policy in all maze instances with one abstraction?



**Does the policy always work?** This is a more involved check, since a set of possible instances has to be considered and a search of a counterexample trajectory among each instance needs to be done. If the policy works, then all possible policy trajectories in all instances have to be checked to conclude this result. For this case, considering an abstraction that focuses on a certain part of the grid may not be useful, since depending on the structure of the instances different parts of the grid may need to be singled out. However, if the policy does not work, it is enough to find an instance in which a counterexample policy trajectory can be shown. Thus, an abstraction that focuses on a certain part of the grid where some instance can show a counterexample would be useful.

In ASP, such a check can be done by making two sets of guesses: (1) choose a valid instance, by guessing the layout of the environment and the position of the goal, and (2) determine a counterexample trajectory, by guessing the movements of the agent following the policy which do not achieve the goal in the instance. If the policy is deterministic (i.e., chooses exactly one action at a state), then the second guessing part becomes straightforward. However, for nondeterministic policies, a choice of possible actions to take exists, which adds to the complexity of the search.

The experiments showed that combining these guesses and those introduced in the syntactic transformation causes many spurious abstract answer sets, which sometimes force refinement decisions towards useless parts of the grid. For example, in general policy checking for the right-hand rule, mDASPAR must go back to the original domain to catch an instance with a counterexample trajectory, as the policy forces to traverse the environment and in the abstract encoding the guesses of the instance and the movements cause to create many spurious trajectories. As for the naive policy, mDASPAR can encounter a counterexample trajectory in few refinement steps: it is sufficient to realize that this policy fails by creating a partial instance where the agent enters a dead-end and has to leave by moving left; it then starts looping by moving right and left.

We remark that for a failing policy we may not expect to have one abstraction mapping that can be applied with any possible instance and catch a counterexample trajectory, the less a mapping that is faithful for any instance. Figure 31a shows an instance in which the right-hand policy is unable to reach the (green) goal point from the (red) entry point in the upper left corner. An abstraction such as Figure 31b is enough to realize that a loop occurs and a goal can not be reached (it is a faithful abstraction for this instance). However, this abstraction does not always distinguish the cells that force to obtain a counterexample trajectory in each possible refuting instance. For example, the instance in Figure 31c also forces the agent to loop; since with the same abstraction (Figure 31d) there is uncertainty among the abstract regions, it is still possible to create spurious counterexample trajectories. Thus, faithfulness can not be achieved. Here the identity abstraction would be the one that can be used to (faithfully) refute the policy in all possible instances.