

Preference-Based Inconsistency Management in Multi-Context Systems

Thomas Eiter

*Technische Universität Wien
Institut für Informationssysteme
Abteilung für Wissensbasierte Systeme
Favoritenstr. 9-11, 184/3
1040 Vienna, Austria*

EITER@KR.TUWIEN.AC.AT

Antonius Weinzierl

*Aalto University
Department of Computer Science
Konemiehentie 2
02150 Espoo, Finland*

ANTONIUS.WEINZIERL@AALTO.FI

Abstract

Multi-Context Systems (MCS) are a powerful framework for interlinking possibly heterogeneous, autonomous knowledge bases, where information can be exchanged among knowledge bases by designated bridge rules with negation as failure. An acknowledged issue with MCS is inconsistency that arises due to the information exchange. To remedy this problem, inconsistency removal has been proposed in terms of repairs, which modify bridge rules based on suitable notions for diagnosis of inconsistency. In general, multiple diagnoses and repairs do exist; this leaves the user, who arguably may oversee the inconsistency removal, with the task of selecting some repair among all possible ones. To aid in this regard, we extend the MCS framework with preference information for diagnoses, such that undesired diagnoses are filtered out and diagnoses that are most preferred according to a preference ordering are selected. We consider preference information at a generic level and develop meta-reasoning techniques on diagnoses in MCS that can be exploited to reduce preference-based selection of diagnoses to computing ordinary subset-minimal diagnoses in an extended MCS. We describe two meta-reasoning encodings for preference orders: the first is conceptually simple but may incur an exponential blowup. The second is increasing only linearly in size and based on duplicating the original MCS. The latter requires nondeterministic guessing if a subset-minimal among all most preferred diagnoses should be computed. However, a complexity analysis of diagnoses shows that this is worst-case optimal, and that in general, preferred diagnoses have the same complexity as subset-minimal ordinary diagnoses. Furthermore, (subset-minimal) filtered diagnoses and (subset-minimal) ordinary diagnoses also have the same complexity.

1. Introduction

At the dawn of an age with growing information connectivity, the issue of interlinking and combining information from various knowledge sources is of increasing importance, posing a challenge to Artificial Intelligence and to Knowledge Representation and Reasoning in particular. Indeed, with the rise of the internet, sharing information has become as easy as never before, and a wealth of knowledge and information sources is available that can be accessed via communicating devices. Multi-Context Systems (Giunchiglia & Serafini, 1994; Roelofsen & Serafini, 2005; Brewka

& Eiter, 2007; Bikakis & Antoniou, 2010) are a well-known approach to address the challenge of sharing information, where individual knowledge bases, called *contexts*, are interlinked with special *bridge rules* which govern the information exchange, such that a global semantics of the system emerges from the local semantics of the constituent knowledge bases. Some practical applications of MCS are defeasible reasoning in ambient intelligence (Bikakis & Antoniou, 2010), cooperation in distributed information systems (Caire & Bikakis, 2011), and the METIS system for maritime situation awareness support (Velikova et al., 2014).

Rooted in the seminal work of McCarthy (1993), which proposed an explicit representation of context where combining different views may give a holistic picture of a situation, the Trento School around Giunchiglia and Serafini developed a notion of multi-context system that is geared to inter-link possibly non-monotonic knowledge bases and can be utilized for query answering (Giunchiglia & Serafini, 1994; Ghidini & Giunchiglia, 2001; Roelofsen & Serafini, 2005; Brewka, Roelofsen, & Serafini, 2007). Brewka and Eiter (2007) generalized this to an abstract framework in which contexts can have heterogeneous knowledge bases that are described using a very abstract notion of logic; Context Knowledge Repositories (Serafini & Homola, 2012) evolved MCS in a different direction for the Semantic Web, where meta and object knowledge can be intermingled. For a more detailed overview of MCS, see the work of Brewka et al. (2011a).

As the contexts of an MCS are typically autonomous and host knowledge bases that are inherited legacy systems, it may happen that the information exchange leads to unforeseen conclusions and in particular to inconsistency; to anticipate and handle all such situations at design time is difficult if not impossible, especially if sufficient details about the knowledge bases are lacking. Inconsistency of an MCS means that it has no model (called *equilibrium*) where a global model is composed of a local model for each context's knowledge base such that all bridge rules governing the information exchange are satisfied; thus, the whole MCS becomes useless.

To repair an inconsistent MCS, basic notions for inconsistency management have been developed by Eiter, Fink, Schüller, and Weinzierl (2010, 2014). Most notably, the notion of diagnosis formalizes the removal of an inconsistency by modifying the information exchange, that is, the bridge rules for the information flow between the contexts. However, while an arbitrary diagnosis restores consistency, the modified information exchange that it affects may have serious consequences, as shown in the following example.

Example 1. Consider an MCS employed in a hospital, which interconnects three systems: (1) a patient knowledge base storing information e.g. about illnesses, insurance companies, and potential allergies of patients; (2) an expert system suggesting proper treatments to illnesses; and (3) a system billing the insurance company of patients for the administered treatments (a formal account is given later, cf. Example 3, Figure 1). The expert system only recommends treatments to which patients are not allergic, while the billing system only allows administered treatments that are covered by the insurance companies. Now suppose a patient with specific allergies can be cured only with a drug that is not covered by his/her insurance; this makes the whole MCS inconsistent and hence no treatment for any patient can be soundly inferred. It is easy to repair this inconsistency, e.g. by modifying the information flow such that either the illness or allergy of the patient is ignored, which results in either not treating the patient or causing an allergic reaction. An alternative repair is to not inform the billing system about the uncovered administration of the drug, so the patient is correctly treated at potential financial loss of the hospital.

Fully-automated, unreflected inconsistency removal may ignore vital information and lead to dangerous results. It is thus desirable – or even necessary – to keep a human operator in the loop while selecting a suitable diagnosis for repair. However, in realistic scenarios often a large (even exponential) number of diagnoses exists, which makes careful manual selection a very time consuming task if not infeasible under time and cognitive constraints. The risk of choosing an improper diagnosis or ending up with no (approved) diagnosis clearly is an obstacle to the deployment of MCS to a broader range of application domains. Also note that inconsistency can arise in all applications: some context may consider certain data to be unrealistic, e.g. in ambient intelligence when the reading from a light sensor is regarded as too high after the janitor installed new light bulbs in the rooms; or some context may not expect to receive a certain combination of data from other contexts, e.g. in case of maritime control when a ship is reported to be both faster and bigger than in the record of the knowledge base. In the former case, removing the actual data from the light sensor restores consistency but it might be more preferred to replace the too-high value with an acceptable one. In the latter case there are scenarios where it is more useful to know the speed of a ship correctly and underestimate its size while for other scenarios it can be the other way round, e.g. to ensure that an inspection crew is large enough for the ship’s size.

The goal of this work thus is to develop some machinery for automatic identification of preferred diagnoses and pruning of those that are unwanted, in order to only require from the human operator to select from a much smaller set (of most preferred diagnoses) the best diagnosis manually. What constitutes a preferred or best diagnosis, cannot be decided in general since it will be different for each MCS and depends on the environment into which an MCS is embedded. In the above example, the health of patients may be considered paramount, but from an economic perspective billing correctly may be considered of highest importance. In any case, we observe that such a decision is up to the person or institution employing the MCS.

Automatic selection of the preferred diagnoses according to some preference requires in turn a formalism for expressing and evaluating preferences. Many such formalisms are available; a prominent and important one are *ceteris paribus* preferences (Doyle et al., 1991), CP-nets (Boutilier et al., 2004; Domshlak et al., 2001; Goldsmith et al., 2008), or utility functions (Von Neumann & Morgenstern, 1944) widely used in economics. As there is no one-fits-all preference formalism that suits every use case, it is a challenge to accommodate any preference formalism that a user deems to fit for selecting most preferred diagnoses. Our approach is based on the idea that a user-customized preference on diagnoses, specified in a formalism chosen by the user, can be seen as a knowledge-base or context of an MCS. This context must be enabled to “see” the diagnoses of the MCS, which is technically challenging. Furthermore, the selection of most preferred diagnoses according to the preference context turns out to be computationally harder than originally thought. As we show, this complexity increase is not due to our meta-reasoning approach, but is intrinsic to the problem, and our approach is worst-case optimal.

Our main contributions are briefly summarized as follows:

- We propose two basic methods for selection of preferred diagnoses: one allows to filter out diagnoses that fail some properties (similar to hard constraints); the other method compares diagnoses with each other in a binary relation and identifies the most preferred one(s). We call the functionalities of these methods *filters* and *preference orders*, respectively. Both are general concepts that can capture many concrete instances to express unwanted or preferred diagnoses. In the flexible and open spirit of the MCS framework, we do not commit to a particular formalism in which filters and

preference orders are specified, but remain at an abstract level and leave the choice of a particular formalism to the user. As illustrative sample instantiations, we consider here CP-nets.

- To realize the selection of diagnoses in such an open way, we develop three transformations to enable meta-reasoning about diagnoses in MCS, i.e., given an MCS and a filter or preference order, a transformed MCS is constructed such that the diagnoses of the original MCS also occur in the transformed MCS, but an additional context is able to observe these diagnoses and apply custom (preference) reasoning. Since the observer context is not restricted to any particular formalism, this allows one to express filters and preference orders in any formalism that can be couched into a context of an MCS.

- For the selection of (most) preferred diagnoses, three extensions of the notion of diagnosis are introduced, namely *protected-minimal*, *prioritized-minimal*, and *subset-minimal prioritized minimal diagnosis*. We investigate the computational complexity of these notions and show by polynomial-time reductions that the first two are of the same complexity as checking whether a pair of sets of bridge rules constitutes a subset-minimal diagnosis. For the third notion, we provide a novel non-deterministic refutation algorithm that works in polynomial time with the help of an oracle for one of the other notions. Still the algorithm is worst-case optimal in a number of settings, as it matches the complexity of the underlying problem. A byproduct of these results are concrete algorithms that can exploit an existing implementation of inconsistency explanation (Bögl, Eiter, Fink, & Schüller, 2010).

The results of this work may be applied to concrete instances of MCS, and the basic notions and results may be carried over to generalizations and recent extensions of MCS, as we shall discuss; furthermore, they may be of use for formalisms that can be modeled using (extensions or variants of) MCS, such as hybrid MKNF knowledge bases (Knorr, Slota, Leite, & Homola, 2014), knowledge base networks (Eiter & Šimkus, 2015), or Boolean networks (Kauffman, 1969, 1993, cf. Inoue, 2011), to mention a few.

1.1 Organization

The remainder of this article is structured as follows. After recalling preliminary notions and fixing notation in Section 2, we introduce in Section 3 filters and preference orders and consider some sample instantiations. In Section 4 we investigate how an (extended) MCS can be enabled to select diagnoses of the original MCS. In Section 5 we show how diagnoses may be selected according to a filter or a preference order and prove the correctness of these realizations, while in Section 6 their computational complexity is investigated. Section 7 discusses related work and in Section 8 we conclude with a summary and an outlook. Proofs of theorems and propositions as well as some detailed examples are in the appendix. This article is strongly based on the work of Weinzierl (2014), which in turn is a significant extension and revision of work by Eiter, Fink, and Weinzierl (2010).

2. Preliminaries

In this section we recall the framework of Multi-Context Systems (MCS) by Brewka and Eiter (2007) and notions for inconsistency management in MCS from Eiter et al. (2014). The MCS framework is based on three basic concepts: abstract logics to capture knowledge-representation formalisms, contexts which represent concrete instances of knowledge bases, and bridge rules to

specify the information exchange; an MCS then simply is a collection of such contexts and their respective bridge rules. Finally, the semantics of an MCS is given in terms of equilibria.

To capture all kinds of knowledge-representation formalisms, the concept of an abstract logic is used, which reduces it to the set-theoretic level.

Definition 1 (Brewka & Eiter, 2007). *An abstract “logic” L , is a triple $L = (\mathbf{KB}, \mathbf{BS}, \mathbf{ACC})$ where:*

- **KB** is the set of knowledge bases of L , where each knowledge base $kb \in \mathbf{KB}$ is a set of elements called “formulas”.
- **BS** is the set of possible belief sets, where each $S \in \mathbf{BS}$ is a set of elements called “beliefs”.
- **ACC** : $\mathbf{KB} \rightarrow 2^{\mathbf{BS}}$ is a function describing the “semantics” of the logic, by assigning to each knowledge base a set of acceptable belief sets.

Intuitively, each knowledge base $kb \in \mathbf{KB}$ is a set of well-formed formulas while each belief set $bs \in \mathbf{BS}$ is a set of beliefs (statements) that a reasoner may jointly hold. The acceptability function $\mathbf{ACC}(kb)$ singles out, given a knowledge base $kb \in \mathbf{KB}$, those sets of beliefs that are acceptable according to some reasoning method for kb . \mathbf{ACC} is a multi-valued function in order to capture also nonmonotonic formalisms, where a knowledge base may have multiple acceptable belief sets (as e.g. for Answer-Set Programming, see Gelfond & Lifschitz, 1991; Default Logic, see Reiter, 1980; or in Abstract Argumentation, see Dung, 1995).

Depending on the concrete situation, e.g. given an existing legacy system or a theorem prover for a specific logic, different formalizations for some logic might be used. There is no fixed mapping between a given logic and an abstract logic representing it, and the mapping may be adjusted to specific application needs. The approach allows one to capture flexibly, e.g., a knowledge-base, an expert system using a logic program, and a billing system using a description logic ontology as they might occur in the scenario described in Example 1. Let us consider two examples for abstract logics.

Example 2. *Classical propositional logic might be modeled as follows:*

- **KB** is the set of all (well-formed) formulas over a signature Σ built using $\wedge, \vee, \neg, \rightarrow$;
- **BS** is the set of deductively closed sets S of Σ -formulas (i.e., $S = Cn(S)$); and
- **ACC**(kb) is the singleton set $\{Cn(kb)\}$.

Disjunctive logic programs under answer set semantics over a function-free first order signature Σ may be modeled as follows:

- **KB** is the set of disjunctive logic programs over Σ , i.e., each $kb \in \mathbf{KB}$ is a set of rules r

$$a_1 \vee \dots \vee a_n \leftarrow b_1, \dots, b_i, \text{not } b_{i+1}, \dots, \text{not } b_m. \quad n + m > 0, \quad (1)$$

also written $H(r) \leftarrow B(r)$, where all a_i, b_j , are atoms over Σ and “not” is negation as failure; we further require that each variable in r occurs also in b_1, \dots, b_i (safety).

- **BS** is the set of Herbrand interpretations over Σ , i.e, each $S \in \mathbf{BS}$ is a set of ground (variable-free) atoms from Σ , and

- **ACC**(kb) is the set of answer sets of kb , i.e., consists of all $S \in \mathbf{BS}$ such that (i) S is a model of kb^S and (ii) no $S' \subset S$ is a model of kb^S (Faber, Leone, & Pfeifer, 2004), where $kb^S = \{r \in \text{grnd}(P) \mid S \models B(r)\}$ is the set of all ground instances r of rules in P whose body $B(r)$ is satisfied by S ; here for evaluation, “not” is treated like classical negation \neg .

We denote these modelings by L_Σ^{pl} and L_Σ^{asp} , respectively.

We remark that each rule r with $n = 0$ is a *constraint*; its heads $H(r)$ amounts to \perp , where \perp is a falsity. We view the latter as a special atom that is false in every Herbrand interpretation.

In the remainder of this work we often omit the explicit definition of the signature Σ for an abstract logic if it is clear from the context.

To specify information exchange between contexts, so-called bridge rules are used. Bridge rules are similar in form and behavior to rules in logic programming. They differ from each other by the fact that bridge rules are based on beliefs from (possibly) different abstract logics and corresponding contexts. Based on the presence (or absence) of beliefs at other contexts, a bridge rule can add information to a context.

Definition 2 (Brewka & Eiter, 2007). *Given a sequence $L = (L_1, \dots, L_n)$ of abstract logics $L_j = (\mathbf{KB}_j, \mathbf{BS}_j, \mathbf{ACC}_j)$, $1 \leq j \leq n$, an L^k -bridge rule over L , with $k \in \{1, \dots, n\}$ is of form:*

$$(k : s) \leftarrow (c_1 : p_1), \dots, (c_i : p_i), \mathbf{not} (c_{i+1} : p_{i+1}), \dots, \mathbf{not} (c_m : p_m). \quad (2)$$

where for each $1 \leq i \leq m$ we have that $c_i \in \{1, \dots, n\}$, $p_i \in \bigcup \mathbf{BS}_{c_i}$ is an element of some belief set of L_{c_i} , and $s \in \bigcup \mathbf{KB}_k$ is a knowledge base formula of L_k .

Each bridge rule in an MCS is associated to a certain context in such a way that all L^k bridge rules belong to the context with identifier k .

We denote by $\varphi(r)$ the formula s in the head of r and by $C_h(r)$ the context k where r belongs to. The full head of r is denoted by $\text{head}(r) = (k : s)$, thus $\text{head}(r) = (C_h(r) : \varphi(r))$. The literals in the body of r are referred to as $\text{body}^\pm(r)$, $\text{body}^+(r)$, $\text{body}^-(r)$, $\text{body}(r)$, which denote the sets $\{(c_1 : p_1), \dots, (c_m : p_m)\}$, $\{(c_1 : p_1), \dots, (c_j : p_j)\}$, $\{(c_{j+1} : p_{j+1}), \dots, (c_m : p_m)\}$, $\{(c_1 : p_1), \dots, (c_j : p_j), \mathbf{not} (c_{j+1} : p_{j+1}), \dots, \mathbf{not} (c_m : p_m)\}$, respectively.

Furthermore, $C_b(r)$ denotes the set of contexts referenced in r 's body, i.e., $C_b(r) = \{c_i \mid (c_i : p_i) \in \text{body}^\pm(r)\}$. Note that different from the work of Brewka and Eiter (2007), the head of r contains not only the knowledge-base formula s but also the context identifier k . This choice merely is syntactic sugar and allows easier identification of the context where r belongs to. For later technical use, we denote by $\text{cf}(r)$ the *condition-free* bridge rule resulting from r by removing all elements in its body, i.e., $\text{cf}(r)$ is $(k : s) \leftarrow \cdot$ and for any set of bridge rules R , we let $\text{cf}(R) = \bigcup_{r \in R} \text{cf}(r)$.

We emphasize that bridge rules only deal with elements of knowledge bases and elements of belief sets, both of which are considered to be atomic expressions from the perspective of MCS. Incorporating variables into bridge rules is possible but requires restrictions on context logics or additional machinery for variable substitution (for details, see Fink, Ghionna, & Weinzierl, 2011; Barilaro, Fink, Ricca, & Terracina, 2013; Schüller & Weinzierl, 2011).

Since bridge rules are the only way to exchange information between contexts and bridge rules only refer to beliefs, the contents of a context (i.e., its knowledge-base and semantics) is completely hidden from other contexts. In addition to that, there is no central point of information exchange, hence the MCS framework is somewhat decentralized.

With bridge rules to connect contexts at hand, Multi-Context Systems are defined as follows.

Definition 3 (cf. Brewka & Eiter, 2007). A Multi-Context System is a collection $M = (C_1, \dots, C_n)$ of contexts $C_i = (L_i, kb_i, br_i)$, $1 \leq i \leq n$, where (i) $L_i = (\mathbf{KB}_i, \mathbf{BS}_i, \mathbf{ACC}_i)$ is an abstract logic, (ii) $kb_i \in \mathbf{KB}_i$ is a knowledge base, and (iii) br_i is a set of L^i -bridge rules over $L = (L_1, \dots, L_n)$. Furthermore, for each $H \subseteq \{\varphi(r) \mid r \in br_i\}$ it holds that $kb_i \cup H \in \mathbf{KB}_i$ (i.e., knowledge bases are closed under adding bridge rule heads).

In the sequel, $br(M) = \bigcup_{i=1}^n br_i$ denotes the set of all bridge rules of M ; $C(M) = \{1, \dots, n\}$ denotes the set of all context identifiers of M ; and $br_i(M)$ denotes the set of bridge rules of context i of M , i.e., $br_i(M) = \{r \in br(M) \mid C_h(r) = i\}$.

Example 3. The MCS described in Example 1 can now be formalized. Let $M = (C_1, C_2, C_3)$ be an MCS with three contexts: a patient knowledge-base C_1 , a logic program C_2 suggesting proper medication, and a logic program C_3 handling the billing. Context C_1 uses the abstract logic L_Σ^{pl} , while both C_2 and C_3 use L_Σ^{asp} . We restrict our example to a single patient with the knowledge bases kb_1 , kb_2 , and kb_3 as given in Figure 1 on page 354 for the contexts C_1 , C_2 , and C_3 , respectively. Intuitively, the knowledge base kb_1 of context C_1 states that the patient has severe hyperglycemia, that she is allergic to animal insulin, and that her health insurance is with company B. Context C_2 's knowledge base kb_2 suggests to apply either human or animal insulin if the patient has hyperglycemia and requires that the applied insulin does not cause an allergic reaction. Context C_3 does the billing and encodes that insurance B only pays animal insulin.

The MCS M contains five rather simple bridge rules shown in Figure 1 (bridge rules are presented in the format “name: head \leftarrow body.”). Their task is to carry information from one context to another. Bridge rule r_1 , for example, carries information about hyperglycemia of the patient from the patient knowledge-base C_1 to the medication recommender system C_2 . Bridge rule r_2 is the sole non-monotonic one and it turns the absence of an allergy to animal insulin (in C_1) into the allowance to administer this kind of insulin (in C_2). A graphical depiction of M is shown in Picture 1 inside Figure 1. The latter also shows the minimal diagnoses of M (cf. below for further details).

The semantics of an MCS $M = (C_1, \dots, C_n)$ is defined in terms of special *belief states*, which are sequences $S = (S_1, \dots, S_n)$ of belief sets $S_i \in \mathbf{BS}_i$, $1 \leq i \leq n$; intuitively, each S_i must be a locally accepted belief set where the bridge rules of context S_i are respected.

To formalize this, we call a bridge rule r of form (2) *applicable* in a belief state S , denoted by $S \vdash r$, if (i) for each $(j:p) \in body^+(r)$ it holds that $p \in S_j$, and (ii) for each $(j:p) \in body^-(r)$ it holds that $p \notin S_j$. For a set R of bridge rules and a belief state S , $app(R, S)$ denotes the set of bridge rules of R that are applicable in S , i.e., $app(R, S) = \{r \in R \mid S \vdash r\}$.


We can now define the desired belief states of an MCS as follows.

Definition 4 (cf. Brewka & Eiter, 2007). A belief state $S = (S_1, \dots, S_n)$ of M is an equilibrium if for every belief set S_i , $1 \leq i \leq n$, it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i, S)\})$. The set of all equilibria of an MCS M is denoted by $EQ(M)$.

To create bridge rules that are always resp. never applicable, we also allow $r = (k:s) \leftarrow \top$, resp. $r' = (k:s) \leftarrow \perp$, where $S \vdash r$ resp. $S \not\vdash r'$ for every belief state S . Here \top denotes the empty body and \perp a body containing $(\ell:p)$, **not** $(\ell:p)$ where p is any belief of any context C_ℓ . For simplicity, we assume that the bridge rules r and r' have no body literals.

Figure 1: The Hospital MCS $M = (C_1, C_2, C_3)$ with knowledge bases kb_i and bridge rules r_j .

$$\begin{aligned}
 kb_1 &= \{hyperglycemia, allergic_animal_insulin, insurance_B\} \\
 kb_2 &= \{give_human_insulin \vee give_animal_insulin \leftarrow hyperglycemia. \\
 &\quad \perp \leftarrow give_animal_insulin, not\ allow_animal_insulin\} \\
 kb_3 &= \{bill \leftarrow bill_animal_insulin. \\
 &\quad bill_more \leftarrow bill_human_insulin. \\
 &\quad \perp \leftarrow insurance_B, bill_more.\} \\
 r_1: &\quad (2 : hyperglycemia) \leftarrow (1 : hyperglycemia). \\
 r_2: &\quad (2 : allow_animal_insulin) \leftarrow \mathbf{not} (1 : allergic_animal_insulin). \\
 r_3: &\quad (3 : bill_animal_insulin) \leftarrow (2 : give_animal_insulin). \\
 r_4: &\quad (3 : bill_human_insulin) \leftarrow (2 : give_human_insulin). \\
 r_5: &\quad (3 : insurance_B) \leftarrow (1 : insurance_B).
 \end{aligned}$$


 Picture 1: The MCS M visualized.

The set of minimal diagnoses of M is:

$$D_m^\pm(M) = \{(\{r_1\}, \emptyset), (\{r_4\}, \emptyset), (\{r_5\}, \emptyset), (\emptyset, \{r_2\})\}.$$

Application of these diagnoses intuitively results in:

- $(\{r_1\}, \emptyset)$ — illness of the patient is ignored.
- $(\{r_4\}, \emptyset)$ — medication is not billed.
- $(\{r_5\}, \emptyset)$ — insurance company receives bill it will not pay.
- $(\emptyset, \{r_2\})$ — patient is given medication she is allergic to.

No diagnosis is clearly the best, it depends on one's preferences.

Given an MCS $M = (C_1, \dots, C_n)$ over abstract logics $L = (L_1, \dots, L_n)$, a set R of bridge rules is *compatible* with M , if a partitioning R_1, \dots, R_n of $R = \bigcup_{k=1}^n R_k$ exists where every $r \in R_k$ is an L^k -bridge rule over L . For such R , we write $M[R]$ for the MCS that results by replacing its bridge rules with R . E.g., $M[br(M)] = M$ and $M[\emptyset]$ is M with no bridge rules.

We say that M is *inconsistent*, denoted $M \models \perp$, if M has no equilibrium, i.e., $\text{EQ}(M) = \emptyset$. The converse, that M is consistent, is denoted by $M \not\models \perp$, i.e., $\text{EQ}(M) \neq \emptyset$.

For a consistency-based explanation of inconsistency pairs (D_1, D_2) of sets of bridge rules are considered, such that if we deactivate the rules in D_1 and add the rules in D_2 in condition-free form, the MCS becomes consistent (i.e., admits an equilibrium).

Definition 5. *Given an MCS M , a diagnosis of M is a pair (D_1, D_2) , $D_1, D_2 \subseteq br(M)$, such that $M[br(M) \setminus D_1 \cup cf(D_2)] \not\models \perp$. We denote by $D^\pm(M)$ the set of all diagnoses.*

An alternative reading of Def. 5 is that a diagnosis indicates which bridge rules are assumed to require modification in order to obtain a consistent MCS, i.e., a diagnosis constitutes a way to *repair* an MCS if its bridge rules are modified according to the diagnosis. Adding rules condition-free is the most severe form of modification of a rule's body, but as shown by Eiter et al. (2014), this notion also allows one to capture more fine-grained forms of modification.

We call any pair $D = (D_1, D_2) \in 2^{br(M)} \times 2^{br(M)}$ a *candidate diagnosis* (regardless of whether $D \in D^\pm(M)$ holds). We denote the MCS resulting from the application of a candidate diagnosis $(D_1, D_2) \subseteq (br(M), br(M))$ by $M[D_1, D_2]$, which equals the MCS $M[br(M) \setminus D_1 \cup cf(D_2)]$.

Following Occam's razor, one may consider those diagnoses to be the preferred ones that require the least modifications. This motivates the notion of minimal diagnosis.

Definition 6. *Given an MCS M , a diagnosis $D \in D^\pm(M)$ is (pointwise) subset minimal, if no $D' \subset D$ is in $D^\pm(M)$; by $D_m^\pm(M)$ we denote all such D , i.e., $D_m^\pm(M) = \{D \in D^\pm(M) \mid \forall D' \in D^\pm(M) : D' \subseteq D \Rightarrow D \subseteq D'\}$.*

Here, given pairs $A = (A_1, A_2)$ and $B = (B_1, B_2)$ of sets, the pointwise subset relation $A \subseteq B$ holds iff $A_1 \subseteq B_1$ and $A_2 \subseteq B_2$; moreover, $A \subset B$ holds iff $A \subseteq B \wedge A \neq B$, where $A = B$ holds iff $A_1 = B_1 \wedge A_2 = B_2$.

Example 4. *Reconsider the MCS M of Example 3. Since the patient has hyperglycemia and is allergic to animal insulin, the belief set containing `give_human_insulin` is the only one acceptable at C_2 , i.e., the human insulin must be given. Since the insurance company does not cover human insulin, the billing context C_3 admits no acceptable belief set and the MCS M therefore is inconsistent. As shown at the bottom of Figure 1, the minimal diagnoses of M are $D_m^\pm(M) = \{D^{(1)}, D^{(2)}, D^{(3)}, D^{(4)}\}$ with $D^{(1)} = (\{r_1\}, \emptyset)$, $D^{(2)} = (\{r_4\}, \emptyset)$, $D^{(3)} = (\{r_5\}, \emptyset)$, and $D^{(4)} = (\emptyset, \{r_2\})$.*

Applying the diagnosis $D^{(i)}$ for $1 \leq i \leq 4$, i.e., considering for $D^{(i)} = (D_1^{(i)}, D_2^{(i)})$ the MCS $M[br(M) \setminus D_1^{(i)} \cup cf(D_2^{(i)})]$, yields that the illness of the patient is ignored ($D^{(1)}$), that the medication is not billed ($D^{(2)}$), that the insurance receives a bill it will not pay ($D^{(3)}$), and that the patient is given a medication she is allergic to ($D^{(4)}$).

3. Preferences

Clearly, in general not all diagnoses of an MCS are equally appealing, as applying the selected repair might have serious consequences, e.g., in the MCS M of Example 4 if the illness of the patient is

ignored. It is not easy to identify the best diagnosis in $D_m^\pm(M)$: if the health of the patient is most important, then those diagnoses only causing a wrong billing are preferred; on the other hand, if costs matter, one might consider any diagnosis leading to a wrong billing as unacceptable.

In the literature two basic ways occur frequently: one is to separately consider each outcome (i.e., diagnosis) and discard it whenever it fails some preference condition; the other is to compare outcomes with each other and decide which is the most appealing. We call the former a filter, since it filters unwanted diagnoses, and the other a preference. Both notions of diagnosis can be defined in general by relying on some notion of plausibility (see e.g., for abduction Bylander et al., 1991).

As for preference among diagnoses, two immediate questions are (i) how to model preference formally and (ii) how to obtain or elicit a concrete preference description for a particular MCS. Regarding (i), we aim in this work to be rather general and resort to a prototypical model in terms of a preference relation, formalized as binary relation \preceq on the set of diagnoses that is reflexive and transitive; this abstract notion allows us to capture a large number of formalisms that have been developed for specifying preference. Regarding (ii), there is no simple or straightforward answer, since the construction of a particular preference order may very much depend on the particular MCS and its application domain. It may appear far easier instead to obtain a preference order on individual bridge rules; since diagnoses amount to special sets of bridge rules, it is possible to construct a preference order between diagnoses from such preferences, by using techniques that lift preference on elements to sets of elements, as in the work of Brewka, Truszczyński, and Woltran (2010), or to apply iterative improvement techniques; we refer here to the work by Brafman, Domshlak, Shimony, and Silver (2006), and Brewka et al. (2010), where also respective preference representation languages are considered. Guaranteeing that such preference is desired and preference elicitation tailored specifically for MCS, however, requires a study of its own and is beyond the scope of this article.

We thus assume that preferences (or filters) over diagnoses are already given and concentrate on the computational challenges of imposing them on an MCS. To illustrate some particular formalism for preference specification where elicitation respectively preference generation has been thoroughly studied, we consider here the widely known CP-nets (cf. Boutilier et al., 2004; Allen, 2016) in which preference is specified by statements like “if bridge rules r_1 and r_2 are removed, then I prefer bridge rule r_3 to be condition-free”.

Since preferences allow to compare diagnoses, but they do not allow the exclusion of diagnoses from being considered, preferences alone are not sufficient. If one wants to ensure that certain diagnoses are excluded from being considered acceptable, the need for a way to filter out certain diagnoses arises. For specifying a filter, we again use the most general approach, which is a Boolean function on diagnoses.

In this section we introduce the definitions of filters and preference orders in general, as well as some specific preference formalisms. The following sections then show how they can be realized in MCS in such a way that any formalism used to define the preference order or filter can be incorporated thanks to using the abstract logic of an MCS context. Furthermore, our approach preserves core properties of MCS like information hiding and allows for a decentralized evaluation.

3.1 Filters on Diagnoses

Filters allow the MCS designer to apply sanity checks on diagnoses; they act as hard constraints: diagnoses that fail to satisfy the conditions are filtered out and discarded for consistency restoration.

3.1.1 PROTECTING BRIDGE RULES

In a first attempt, we may consider *protecting* some bridge rules from modification, i.e., we disallow a diagnosis to contain them. The adapted notion of diagnosis is as follows.

Definition 7. Let M be an MCS with protected rules $br_P \subseteq br(M)$. A diagnosis excluding protected rules br_P is a diagnosis $(D_1, D_2) \in D^\pm(M)$, where $D_1, D_2 \subseteq br(M) \setminus br_P$. We denote the set of all such diagnoses by $D^\pm(M, br_P)$. The set of all minimal such diagnoses is $D_m^\pm(M, br_P) = \{D \in D^\pm(M, br_P) \mid \nexists D' \in D^\pm(M, br_P) : D' \subset D\}$.

Example 5. Consider the hospital MCS M of Example 3 again. One might decide that bridge rules for health-related information-flow are protected, i.e., $br_P = \{r_1, r_2\}$.

The set of minimal protected diagnoses then is:

$$D_m^\pm(M, br_P) = \{(\{r_4\}, \emptyset), (\{r_5\}, \emptyset)\}$$

In the following we also write diagnosis with protected bridge rules meaning a diagnosis excluding protected rules. The following property is easy to verify.

Proposition 1. Let M be an inconsistent MCS with protected rules br_P . Then $D^\pm(M, br_P) \subseteq D^\pm(M)$ and $D_m^\pm(M, br_P) \subseteq D_m^\pm(M)$, i.e., every diagnosis is a diagnosis excluding protected rules and every minimal diagnosis excluding protected rules is a minimal diagnosis.

Proof. Let $D \in D^\pm(M, br_P)$, then by definition $D \in D^\pm(M)$. Given $D = (D_1, D_2)$ with $D \in D_m^\pm(M, br_P)$, assume towards contradiction that there exists $(D'_1, D'_2) \in D_m^\pm(M)$ such that $(D'_1, D'_2) \subset (D_1, D_2)$. Observe that $D'_1, D'_2 \subseteq br(M) \setminus br_P$, hence $(D'_1, D'_2) \in D^\pm(M, br_P)$. This contradicts that $D \in D_m^\pm(M, br_P)$, thus it follows that $D \in D_m^\pm(M)$. \square

In Section 6 it is shown that deciding whether $D \in D^\pm(M, br_P)$ and $D \in D^\pm(M)$ have the same complexity, i.e., protected bridge rules do not increase the complexity.

3.1.2 FILTERS IN GENERAL

We now introduce filters in general. A candidate diagnosis (D_1, D_2) is considered whether it fails some conditions; if so, it is filtered out and not considered for consistency restoration; thus a filter can be seen as hard constraints on diagnoses.

Example 6. Consider two scientists, Prof. K and Dr. J, who plan to write a paper. We formalize their reasoning in an MCS M with two contexts C_1 and C_2 that employ L_Σ^{asp} for answer set semantics. Dr. J will write most of the paper and Prof. K will engage if she finds time or if Dr. J thinks the paper needs improvement (r_1). Dr. J knows that involving Prof. K results in a good paper (r_2 and kb_1) and he will list her as an author if she participates (r_3). The knowledge bases of the contexts are:

$$\begin{aligned} kb_1 &= \{contribute \leftarrow improve.; \quad contribute \leftarrow has_time.\} \\ kb_2 &= \{good \leftarrow coauthored.\} \end{aligned}$$

The bridge rules of M are:

$$\begin{aligned} r_1 : & & (1 : improve) & \leftarrow \mathbf{not} (2 : good). \\ r_2 : & & (2 : coauthored) & \leftarrow (1 : contribute). \\ r_3 : & & (2 : name_K) & \leftarrow (1 : contribute). \end{aligned}$$




Figure 2: Contexts and bridge rules of the MCS $M = (C_1, C_2)$ from Example 6.

Figure 2 depicts the contexts and bridge rules of M . It appears that M is inconsistent, intuitively because the cycle through bridge rules r_1 and r_2 has an odd number of negations.

The set of minimal diagnoses of M is: $D_m^\pm(M) = \{(\{r_1\}, \emptyset), (\{r_2\}, \emptyset), (\emptyset, \{r_2\}), (\emptyset, \{r_1\})\}$. The first two diagnoses break the cycle by removing a rule, the last two “stabilize” it.

We aim for a general notion of a filter, therefore we define a filter to be a Boolean function on candidate diagnoses.

Definition 8. Let M be an MCS with bridge rules $br(M)$. A diagnosis filter for M is a function $f: 2^{br(M)} \times 2^{br(M)} \rightarrow \{0, 1\}$ and the set of filtered diagnoses is $D_f^\pm(M) = \{(D_1, D_2) \in D^\pm(M) \mid f(D_1, D_2) = 1\}$. By $D_{m,f}^\pm(M)$ we denote the set of all subset-minimal such diagnoses.

Given a candidate diagnosis $D = (D_1, D_2) \in 2^{br(M)} \times 2^{br(M)}$, we also write $f(D)$ to denote $f(D_1, D_2)$. Writing the set $D_{m,f}^\pm(M)$ explicitly, we obtain:

$$D_{m,f}^\pm(M) = \{D \in D^\pm(M) \mid f(D) = 1 \wedge \nexists D' \in D^\pm(M) : (D' \subset D \wedge f(D') = 1)\} \quad (3)$$

Example 7. Consider the MCS of Example 6 and the diagnoses $D = (\{r_2\}, \emptyset)$ and $D' = (\emptyset, \{r_2\})$, where the contribution of Prof. K is either enforced or forbidden. For both cases, the authorship information conveyed by r_3 is wrong. Using a filter, we can declare diagnoses undesired if they modify r_2 without modifying r_3 accordingly as follows:

$$f(D_1, D_2) = \begin{cases} 0 & \text{if } r_3 \in D_1, r_2 \notin D_1 \text{ or } r_3 \notin D_1, r_2 \in D_1; \\ 0 & \text{if } r_3 \in D_2, r_2 \notin D_2 \text{ or } r_3 \notin D_2, r_2 \in D_2; \\ 1 & \text{otherwise.} \end{cases}$$

In particular it holds that $f(D) = 0 = f(D')$.

Note that filters generalize diagnoses with protected bridge rules. Indeed, let M be an MCS with protected bridge rules br_P . Then we construct a filter f^{br_P} in the following way:

$$f^{br_P}(D_1, D_2) = \begin{cases} 0 & \text{if } \exists r \in br_P : r \in (D_1 \cup D_2); \\ 1 & \text{otherwise.} \end{cases}$$

It is easy to see that $D \in D^\pm(M, br_P)$ holds iff $f^{br_P}(D) = 1$. From the definition of f^{br_P} one can also see that diagnoses with protected bridge rules are some kind of modular filter, where each bridge rule of a diagnosis D can be checked independently of the other bridge rules.

It also holds that every filtered diagnosis is an ordinary diagnosis, but minimal filtered diagnoses are not necessarily minimal diagnoses. Thus an analog to Proposition 1 does not hold, as shown by the following example.

Example 8. Reconsider the MCS M and the filter f of Example 7. The set of minimal filtered diagnoses is as follows: $D_{m,f}^{\pm}(M) = \{(\{r_1\}, \emptyset), (\emptyset, \{r_1\}), (\{r_2, r_3\}, \emptyset), (\emptyset, \{r_2, r_3\})\}$. While $(\{r_2, r_3\}, \emptyset)$ is not in $D_m^{\pm}(M)$, it is a subset-minimal diagnosis respecting the condition expressed by the filter f . Intuitively, the latter diagnoses modify the authorship information in a consistent way and are minimal in the sense that no unnecessary modification is applied.

One could argue whether minimal filtered diagnoses should select from the set of regular minimal diagnoses only those which pass the filter, i.e., select the set $\{D \in D_m^{\pm}(M) \mid f(D) = 1\}$. This looks appealing, but no minimal diagnosis may pass the filter while (non-minimal) diagnoses do. The resulting set of filtered minimal diagnoses then is empty while there are useful diagnoses that satisfy the filter and do not incur unnecessary modifications other than to satisfy the filter condition and to make the MCS consistent. Therefore $D_{m,f}^{\pm}$ consists of the latter diagnoses, and thus seems to be more appropriate.

3.2 Preferences on Diagnoses

To compare diagnoses and select the most appealing one(s), we use preferences. In the spirit of MCS we also want this approach to be open to any kind of formalism for specifying preference. In general, preference is just a binary order relation on diagnoses. To avoid counter-intuitive results like A being preferred over B and B being preferred over C , but A not being preferred over C , we require that preferences are transitive. Since virtually every other preference formalism yields an order relation, we first introduce the general formalization and later show how the specific formalism of CP-nets fits into our approach.

Definition 9. A preference order over diagnoses for an MCS M is a transitive and reflexive binary relation \preceq on $2^{br(M)} \times 2^{br(M)}$; for $D, D' \in 2^{br(M)} \times 2^{br(M)}$ we say that D is preferred to D' if $D \preceq D'$.

Given a preference order \preceq , we denote by \prec the irreflexive and anti-symmetric version of \preceq , i.e., $D \prec D'$ holds iff $D \preceq D'$ and $D' \not\preceq D$ hold. Using a preference order \preceq , we now define what constitutes a most preferred diagnosis. The intuition is that such a diagnosis incurs a minimal set of modifications and no other diagnosis exists that is strictly more preferred. We first introduce \preceq -preferred diagnoses, which are those diagnoses such that no other diagnosis is strictly more preferred. The most preferred diagnoses then are the subset-minimal ones from the set of \preceq -preferred diagnoses.

Definition 10. Let M be an inconsistent MCS and let $D \in D^{\pm}(M)$. Then D is called \preceq -preferred if for all $D' \in 2^{br(M)} \times 2^{br(M)}$ with $D' \prec D$ it holds that $D' \notin D^{\pm}(M)$. Furthermore, D is minimal \preceq -preferred if D is subset-minimal among all \preceq -preferred diagnoses. The set of all \preceq -preferred diagnoses is denoted by $D_{\preceq}^{\pm}(M)$ and the set of all minimal \preceq -preferred by $D_{m,\preceq}^{\pm}(M)$.

Observe that we do not require that \preceq is acyclic and by relying on \prec we consider all diagnoses in a cycle to be equally preferred.

Example 9. Consider the hospital MCS M of Example 3 again, where bridge rules r_1 and r_2 transport information regarding the patient's health and bridge rules r_3, r_4 , and r_5 cover the information flow for billing. If we consider it most important that the information flow regarding health information is changed as little as possible, a preference order \preceq as follows might be used:

$$(D_1, D_2) \preceq (D'_1, D'_2) \text{ iff } \{r_1, r_2\} \cap (D_1 \cup D_2) \subseteq (D'_1 \cup D'_2) \cap \{r_1, r_2\}$$

We observe that following this definition, the following preferences (and several more) hold:

$$\begin{array}{lll}
 (\{r_4, r_5\}, \emptyset) \preceq (\{r_1\}, \emptyset) & (\{r_4\}, \emptyset) \preceq (\{r_1\}, \emptyset) & (\{r_5\}, \emptyset) \preceq (\{r_1\}, \emptyset) \\
 (\{r_4, r_5\}, \emptyset) \preceq (\emptyset, \{r_2\}) & (\{r_4\}, \emptyset) \preceq (\emptyset, \{r_2\}) & (\{r_5\}, \emptyset) \preceq (\emptyset, \{r_2\}) \\
 (\{r_4\}, \emptyset) \preceq (\{r_5\}, \emptyset) & (\{r_5\}, \emptyset) \preceq (\{r_4\}, \emptyset) &
 \end{array}$$

Note that \preceq indeed yields cyclic preferences among those candidate diagnoses that are incomparable; in particular $(\{r_4\}, \emptyset)$ and $(\{r_5\}, \emptyset)$. We have that

$$D_{\preceq}^{\pm}(M) = \{(D_1, D_2) \mid D_1, D_2 \subseteq \{r_3, r_4, r_5\} \text{ and } r_4 \in D_1 \setminus D_2 \text{ or } r_5 \in D_1 \setminus D_2\}$$

Note that $(\{r_5\}, \emptyset)$, $(\{r_4\}, \emptyset)$, and $(\{r_4, r_5\}, \emptyset)$ are all in $D_{\preceq}^{\pm}(M)$. Selecting the subset-minimal diagnoses from $D_{\preceq}^{\pm}(M)$ we obtain $D_{m, \preceq}^{\pm}(M) = \{(\{r_5\}, \emptyset), (\{r_4\}, \emptyset)\}$. This agrees with our intuition that a minimal set of modifications should be applied and we favor to modify bridge rules for billing information rather than modifying health-related bridge rules.

For use in the following sections, we also state the sets $D_{\preceq}^{\pm}(M)$ and $D_{m, \preceq}^{\pm}(M)$ explicitly.

$$\begin{aligned}
 D_{\preceq}^{\pm}(M) &= \{D \in D^{\pm}(M) \mid \forall D' \in D^{\pm}(M) : \neg(D' \preceq D \wedge D' \neq D \wedge D \not\preceq D')\} & (4) \\
 D_{m, \preceq}^{\pm}(M) &= \{D \in D_{\preceq}^{\pm}(M) \mid \forall D' \in D_{\preceq}^{\pm}(M) : D' \subseteq D \Rightarrow D' = D\}
 \end{aligned}$$

In Section 5 we show how preferences can be realized in general.

3.2.1 SAMPLE INSTANTIATIONS OF PREFERENCE ORDERS: CP-NETS

We now briefly demonstrate how our notion of preference can capture some practical preference formalisms. Conditional preference networks (CP-nets) (Boutilier et al., 2004) are a widespread preference formalism with many appealing features. CP-nets capture a natural class of preference statements like “If my new car is from Japan, I prefer hybrid over diesel engine, assuming all else is equal”. We briefly present the essential concepts of CP-nets in compact form to show their flavor. A CP-net has a set of outcome variables where each variable ranges over some domain. In our example, we have the variables “origin country” and “engine type” with origin country including “Japan” and engine type including “diesel” and “hybrid”. A distinguishing feature of CP-nets is the dependency of preferences, e.g., the above preference on the engine type only upholds if the outcome of the origin country is “Japan”. This dependency is expressed in CP-nets as a directed graph $N = (V, E)$ on *outcome variables* V . Each outcome variable $v \in V$ is associated with a set $dom(v)$ of possible outcomes. A *total outcome* o assigns each $v \in V$ a value from $dom(v)$, denoted by o_v ; for a subset $V' \subseteq V$ of the variables, $o_{V'}$ denotes the restriction of o to V' (note that $o = o_V$). Furthermore, each variable v has an associated conditional-preference table $cpt(v)$ which contains a total preference order over $dom(v)$ for every possible outcome of the parent variables $Pa(v)$ on which v depends, i.e., $Pa(v) = \{v' \mid (v', v) \in E\}$. A preference order \preceq over total outcomes *satisfies* N , if for all total outcomes o, o' and all variables $v \in V$ it holds that $o_{Pa(v)} = o'_{Pa(v)}$ implies that $o \preceq o'$ holds iff o_v is preferred to o'_v in $cpt(v)$. Informally, \preceq satisfies N if it agrees with all conditional-preference tables of N ; for formal details and further background see the work of Boutilier et al. (2004) and Allen (2016).

Note that dependencies in CP-nets are natural to humans as CP-nets have successfully been used for preference elicitation, e.g. by Domshlak et al. (2001). CP-nets also allow one to compare total

outcomes, so we may ask whether an outcome o is preferred to an outcome o' under all preferences that satisfy N . If this is the case, then o is said to *dominate* o' , written as $N \models o \succsim o'$ (*entailment*).

CP-nets may be used to specify preference among diagnoses of an MCS M as follows: each bridge rule $r \in br(M)$ is assigned two outcome variables V_1^r and V_2^r , where the domain of V_1^r is $\{inD1, not.inD1\}$ and the domain of V_2^r is $\{inD2, not.inD2\}$. We call any CP-net $N = (V, E)$ with $V = \{V_1^r, V_2^r \mid r \in br(M)\}$ and domains as before *fully compatible* to M , or just a *fully compatible* CP-net. Every total outcome of such a CP-net corresponds one-to-one to a candidate diagnosis of M .

Definition 11. *Given an MCS M and a CP-net N that is fully compatible to M , we say a diagnosis $D \in D^\pm(M)$ is N -preferred iff there exists no $D' \in D^\pm(M)$ such that $N \models D' \succsim D$ holds and $N \models D \succsim D'$ does not hold. Let $D^N(M)$ denote the set of all N -preferred diagnoses of M . Then the set $D_{ird}^\pm(M, N)$ of irredundant N -preferred diagnoses, consists of the subset-minimal diagnoses of $D^N(M)$. Formally, $D_{ird}^\pm(M, N) = \{D \in D^N(M) \mid \forall D' \in D^N(M) : D' \subseteq D \Rightarrow D = D'\}$.*

Observe that given a CP-net N that is compatible to the MCS M , we can readily define a preference order \preceq^N that is equivalent to N as follows: for all $D, D' \in 2^{br(M)} \times 2^{br(M)}$ it holds that $D \preceq^N D' \Leftrightarrow N \models D \succsim D'$. Since the entailment of the CP-net is transitive and reflexive, \preceq^N is transitive and reflexive, therefore \preceq^N is a preference relation in the sense of Definition 9. Hence, we can use \preceq^N and the notion of most preferred diagnosis to select the irredundant N -preferred diagnoses, formally:

Proposition 2. *Given a CP-net N compatible to an MCS M , let $D \preceq^N D'$ hold iff $N \models D \succsim D'$ holds. Then $D^N(M) = D_{\preceq^N}^\pm(M)$ and $D_{m, \preceq^N}^\pm(M) = D_{ird}^\pm(M, N)$.*

Deciding whether a global outcome o is preferred over o' by a given CP-net N , i.e., deciding $N \models o' \succsim o$, is no easy task in general. In the work of Goldsmith et al. (2008) it is shown that this task is **PSPACE**-complete. Restricting the CP-net, however, decreases the computational complexity, e.g., the same decision problem is **NP**-complete for binary-valued directed-path singly connected CP-nets and even in quadratic time for binary-valued tree-structured CP-nets as shown by Boutilier et al. (2004). Notice that fully-compatible CP-nets are binary-valued.

4. Meta-Reasoning for Diagnosis

To realize filters and preference orders inside an MCS, some MCS context must be able to reason on diagnoses of the MCS. We achieve this by a rewriting technique, transforming an MCS M into an extended MCS M' , where certain new context(s) can do meta-reasoning on diagnoses of the original MCS M . The underlying idea here is that a diagnosis D applied to M' has the same effects as if D would be applied to M , but in M' there are additional contexts that observe the behavior of the bridge rules in M to reason about the observed diagnosis D . A significant advantage of this approach is that the observation contexts may use any abstract logic for reasoning about the observed diagnoses. Thus our approach can capture a wide range of formalisms to specify preferences by filters or preference orders, and it allows the creator of an MCS to use whichever formalism she or he sees to fit best. Furthermore, the rewriting is not intrusive, since it only requires that each rule is duplicated and one additional positive literal added in it.

The transformation given in this section realizes filters in general by using diagnoses with protected bridge rules. Preferences also require some additional notions of diagnoses that allow to

prioritize some bridge rules. This prioritization in principle establishes a lexicographic order on candidate diagnoses. We present in fact two possible ways to realize general preferences. The first adds exponentially many bridge rules, while the second adds only linearly many bridge rules but comes at the cost of duplicating the contexts of the original MCS.

In the following section, we present a uniform encoding for meta-reasoning on diagnoses, which serves as the basis for realizing filters and preferences as well as for proving correctness.

4.1 Injecting Diagnoses

We can encode the modifications of a diagnosis directly in an MCS such that observations are perfect, which means that the original system is not just observed but actively modified. Conceptually, given an MCS $M = (C_1, \dots, C_n)$ all its bridge rules are rewritten and protected such that a diagnosis is applied only to the bridge rules of an additional context C_{n+1} . This context C_{n+1} then is able to definitely observe the modifications and to exhibit this observation to all other contexts via its acceptable belief set. The bridge rules of the original system are rewritten to consider the belief set of C_{n+1} . So they either behave like being removed or like being made condition-free, depending on what C_{n+1} believes. Furthermore, the definition of the observation context C_{n+1} is generic and only specifies some necessary properties. This provides the user with the possibility to instantiate it with a logic and a knowledge base of her choice, which is in line with the spirit of the MCS framework.

To encode (observe) diagnoses, the context C_{n+1} needs bridge rules to which a diagnosis can be applied and which can be observed reliably. To this end, for every $r \in br(M)$ we have the following two bridge rules to encode/observe whether r is removed or made unconditional (i.e., condition-free).

$$d1(r) : (n+1 : not_removed_r) \leftarrow \top. \quad (5)$$

$$d2(r) : (n+1 : uncond_r) \leftarrow \perp. \quad (6)$$

For a set $R \subseteq br(M)$, let $d1(R) = \{d1(r) \mid r \in R\}$ and $d2(R) = \{d2(r) \mid r \in R\}$.

Since the meta-reasoning encoding is used as uniform foundation for filters and preferences, we introduce a property θ that describes the additional behavior of the context C_{n+1} . This allows us to later specify the required behavior for both filters and preferences. The preference encoding requires further bridge rules for mapping preferences to bridge rules; this set of additional bridge rules is called \mathcal{K}_p , so we obtain an MCS $M^{mr(\theta, \mathcal{K}_p)}$ as the meta-reasoning encoding of M as follows.

Definition 12. *Let $M = (C_1, \dots, C_n)$ be an MCS, let \mathcal{K}_p be a set of bridge rules such that the following holds for all $r \in \mathcal{K}_p$: $\varphi(r) \notin \{not_removed_{r'}, uncond_{r'} \mid r' \in br(M)\}$, $body(r) = \{\perp\}$, and $C_h(r) = n+1$. Furthermore, let θ be a ternary property over $2^{br(M)} \times 2^{br(M)} \times 2^{\mathcal{K}_p}$. Then, the MCS $M^{mr(\theta, \mathcal{K}_p)} = (C'_1, \dots, C'_n, C_{n+1})$ is a meta-reasoning encoding if the following holds:*

- (i) *for every $C_i = (L_i, kb_i, br_i)$ with $1 \leq i \leq n$ it holds that $C'_i = (L_i, kb_i, br'_i)$ where br'_i contains for every $r \in br_i$ of form (2) the following two bridge rules:*

$$(i : s) \leftarrow (c_1 : p_1), \dots, (c_j : p_j), \mathbf{not} (c_{j+1} : p_{j+1}), \dots, \mathbf{not} (c_m : p_m), \\ \mathbf{not} (n+1 : removed_r). \quad (7)$$

$$(i : s) \leftarrow (n+1 : uncond_r). \quad (8)$$

- (ii) $C_{n+1} = (L_{n+1}, kb_{n+1}, br_{n+1})$ is an arbitrary context such that:

- (a) *the logic $L_{n+1} = (\mathbf{KB}_{n+1}, \mathbf{BS}_{n+1}, \mathbf{ACC}_{n+1})$ is an arbitrary logic where*

- \mathbf{KB}_{n+1} fulfills that $A \subseteq \{\text{uncond}_r, \text{not_removed}_r \mid r \in \text{br}(M)\} \cup \{\varphi(r) \mid r \in \mathcal{K}_p\}$
implies that there exists a knowledge-base $kb \in \mathbf{KB}_{n+1}$ with $A \subseteq kb$.
 - \mathbf{BS}_{n+1} fulfills that $B \subseteq \{\text{removed}_r, \text{uncond}_r \mid r \in \text{br}(M)\}$
implies that there exists a belief set $bs \in \mathbf{BS}_{n+1}$ with $B \subseteq bs$.
 - \mathbf{ACC}_{n+1} is such that some $kb \in \mathbf{KB}_{n+1}$ fulfills the property (*) in item (d) below.
- (b) the set of bridge rules is $\text{br}_{n+1} = d1(\text{br}(M)) \cup d2(\text{br}(M)) \cup \mathcal{K}_p$ and the only rules with head formulas not_removed_r and uncond_r are of form (5) and (6).
- (c) the knowledge base $kb_{n+1} \in \mathbf{KB}_{n+1}$ is arbitrary but fulfills the following property:

(*) for every $H \subseteq \{\varphi(r) \mid r \in \text{br}_{n+1}\}$ it holds that

$S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \cup H)$ iff $\theta(R_1, R_2, R_3)$ is true

where $R_1 = \{r \in \text{br}(M) \mid \text{not_removed}_r \notin H\}$,

$R_2 = \{r \in \text{br}(M) \mid \text{uncond}_r \in H\}$,

$R_3 = \{r \in \mathcal{K}_p \mid \varphi(r) \in H\}$, and

$S_{n+1} = \{\text{removed}_r \mid r \in R_1\} \cup \{\text{uncond}_r \mid r \in R_2\}$.

The protected bridge rules br_P of $M^{mr(\theta, \mathcal{K}_p)}$ are all rules of form (7) and (8).

The bridge rules of form (7) and (8) for a bridge rule $r \in \text{br}(M)$ either behave like removed or like made unconditional, depending on what C_{n+1} believes. The bridge rule (7) behaves like r being removed by a diagnosis if C_{n+1} believes it is removed (i.e., it does not fire even if the original body of r is satisfied). The bridge rule (8) behaves like r being made condition-free by a diagnosis if C_{n+1} believes it is condition-free (i.e., it fires regardless of whether the original body of r is satisfied).

The property (*) guarantees that (i) the beliefs of C_{n+1} coincide with the diagnosis encoded by bridge bridge rules of form (5) and (6), and (ii) that there is an acceptable belief set whenever θ holds. This definition of $M^{mr(\theta, \mathcal{K}_p)}$ is thus generic and the Propositions 3-5 below are very general. The advantage of this approach is that we have a common foundation for encoding filters and preferences, such that several propositions hold for both encodings. Furthermore, as shown later in detail, the property θ to realize a filter f amounts to simply stating that $\theta(D_1, D_2, \emptyset)$ holds iff $f(D_1, D_2) = 1$, while θ and \mathcal{K}_p suffice to capture preferences. The context C_{n+1} is similar to an interface in programming, which defines certain properties but is open to an arbitrary implementation. A concrete realization of C_{n+1} using the logic L_Σ^{asp} for Answer-Set Programming is illustrated in Example 19 in Appendix B. Specifically, there the rules (22)–(34) “implement the interface”.

Example 10. Recall the MCS $M = (C_1, C_2)$ of Example 6. Let $\mathcal{K}_p = \emptyset$ and $\theta(D_1, D_2, \emptyset)$ always hold. Then the meta-reasoning encoding $M^{mr(\theta, \mathcal{K}_p)} = (C'_1, C'_2, C_3)$ is such that the contexts C_1, C_2 , equal modulo bridge rules the contexts C'_1, C'_2 , respectively. Recall that the bridge rules of M are:

$$\begin{array}{ll}
 r_1 : & (1 : \text{improve}) \leftarrow \mathbf{not} (2 : \text{good}). \\
 r_2 : & (2 : \text{coauthored}) \leftarrow (1 : \text{contribute}). \\
 r_3 : & (2 : \text{name_K}) \leftarrow (1 : \text{contribute}).
 \end{array}$$

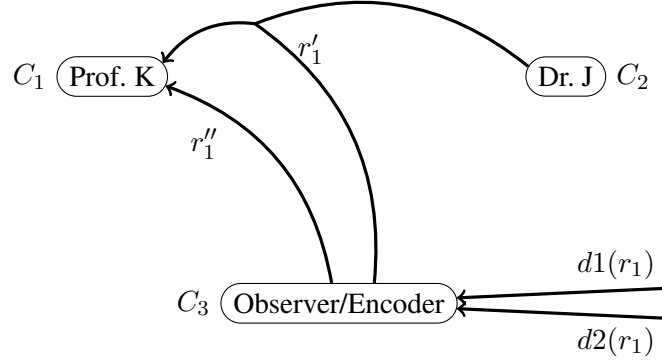


Figure 3: Contexts of the meta-reasoning encoding $M^{mr(\theta, \mathcal{K}_p)} = (C_1, C_2, C_3)$ from Example 10. Only bridge rules $r'_1, r''_1, d1(r_1), d2(r_1)$ of $M^{mr(\theta, \mathcal{K}_p)}$ that stem from bridge rule $r_1 \in br(M)$ are shown.

The bridge rules of $M^{mr(\theta, \mathcal{K}_p)}$ are then as follows:

$$\begin{aligned}
 r'_1 &: & (1: improve) &\leftarrow \mathbf{not} (2: good), \mathbf{not} (3: removed_{r_1}). \\
 r''_1 &: & (1: improve) &\leftarrow (3: uncond_{r_1}). \\
 r'_2 &: & (2: coauthored) &\leftarrow (1: contribute), \mathbf{not} (3: removed_{r_2}). \\
 r''_2 &: & (2: coauthored) &\leftarrow (3: uncond_{r_2}). \\
 r'_3 &: & (2: name_K) &\leftarrow (1: contribute), \mathbf{not} (3: removed_{r_3}). \\
 r''_3 &: & (2: name_K) &\leftarrow (3: uncond_{r_3}). \\
 d1(r_i) &: & (3: not_removed_{r_i}) &\leftarrow \top. & i \in \{1, 2, 3\} \\
 d2(r_i) &: & (3: uncond_{r_i}) &\leftarrow \perp. & i \in \{1, 2, 3\}
 \end{aligned}$$

Notice that only the last six bridge rules of $M^{mr(\theta, \mathcal{K}_p)}$ are not protected, i.e., the first six bridge rules are guaranteed not to be modified in a diagnosis with protected bridge rules. Figure 3 depicts the contexts and, for better visibility, only those bridge rules of $M^{mr(\theta, \mathcal{K}_p)}$ that stem from $r_1 \in br(M)$ are shown.

Note that the observation context of a meta-reasoning encoding only knows that some bridge rules exist but it has no knowledge of the actual information exchange or the contents of any of the other contexts. Since the meta-reasoning encoding is the basis for all later encodings, it provides us with a mechanism to determine preferred solutions for both preference orders and filters that maintains privacy and information hiding of the contexts. For filters that are not inherently centralized, the realization even allows for finding preferred solutions in a decentralized, localized manner (cf. Section 7.1 on decomposing the central observation contexts). Thus we mostly preserve key properties of MCS also for inconsistency assessment and selection of preferred diagnoses.

Def. 12 forbids the observation context C_{n+1} to exhibit any belief not of the form $removed_r$ or $uncond_r$ for $r \in br(M)$. In practice however, auxiliary beliefs are useful when realizing a preference or filter. Since the applicability of bridge rules does not depend on beliefs that do not occur in any bridge rule, this restriction can be lifted to allow for auxiliary beliefs in C_{n+1} .

In the remainder of this section, we show some properties of $M^{mr(\theta, \mathcal{K}_p)}$, which are the basis for proving the correctness of the subsequent realizations of filters and preferences. First, there is a one-to-one correspondence between diagnoses of M and diagnoses of $M^{mr(\theta, \mathcal{K}_p)}$.

Proposition 3. *Let M be an MCS and $M^{mr(\theta, \mathcal{K}_p)}$ be a meta-reasoning encoding with protected bridge rules br_P , and let $D_1, D_2 \subseteq br(M)$, $\mathcal{K} \subseteq \mathcal{K}_p$. Then,*

- (1) *let $S = (S_1, \dots, S_n)$ be a belief state of M and let $S' = (S_1, \dots, S_n, S_{n+1})$ where $S_{n+1} = \{\text{removed}_r \mid r \in D_1\} \cup \{\text{uncond}_r \mid r \in D_2\}$. Then, $S \in \text{EQ}(M[D_1, D_2])$ and $\theta(D_1, D_2, \mathcal{K})$ holds if and only if $S' \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}])$ holds.*
- (2) *$(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ holds if and only if $(D_1, D_2) \in D^\pm(M)$ and $\theta(D_1, D_2, \mathcal{K})$ hold.*

From this, the following correspondence between minimal θ -satisfying diagnoses of M and minimal diagnoses of $M^{mr(\theta, \mathcal{K}_p)}$ holds.

Proposition 4. *Let $M^{mr(\theta, \mathcal{K}_p)}$ be a meta-reasoning encoding of an MCS M . Then the set of minimal θ -satisfying diagnoses with protected bridge rules br_P of $M^{mr(\theta, \mathcal{K}_p)}$ is*

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{ (d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exist no } (D'_1, D'_2) \in D^\pm(M), \mathcal{K}' \subseteq \mathcal{K}_p \text{ such that} \\ (D'_1, D'_2 \cup \mathcal{K}') \subset (D_1, D_2 \cup \mathcal{K}) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \text{ holds} \}.$$

This result can be strengthened given that θ obeys some property. We say that θ is *functional* (or a *function*), if for every $D_1, D_2 \subseteq br(M)$ there exists at most one $\mathcal{K} \subseteq \mathcal{K}_p$ such that $\theta(D_1, D_2, \mathcal{K})$ holds. We say that θ is *functional increasing* if $\mathcal{K} \subseteq \mathcal{K}'$ holds whenever θ is functional, $\theta(D_1, D_2, \mathcal{K})$, $\theta(D'_1, D'_2, \mathcal{K}')$, and $(D_1, D_2) \subseteq (D'_1, D'_2)$, where $D_1, D_2, D'_1, D'_2 \subseteq br(M)$, $\mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p$.

Proposition 5. *Let $M^{mr(\theta, \mathcal{K}_p)}$ be a meta-reasoning encoding of an MCS M such that θ is functional increasing. Then, the set of minimal θ -satisfying diagnoses with protected bridge rules br_P of $M^{mr(\theta, \mathcal{K}_p)}$ is*

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{ (d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2) \subset (D_1, D_2) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \text{ holds for some } \mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p \}.$$

Given these relationships between diagnoses of M and $M^{mr(\theta, \mathcal{K}_p)}$ with respect to property θ , we show in the next section several ways how $M^{mr(\theta, \mathcal{K}_p)}$ can be used to realize preferences.

5. Preference Realization

In the previous section we introduced a transformation that enables meta-reasoning on diagnoses. In this section, we first present how filters can be realized and then proceed to preferences, where we first introduce a plain encoding using exponentially many bridge rules to realize total preference orders and then introduce an encoding that allows to realize arbitrary preference orders at the expense of cloning the contexts of the original MCS.

5.1 Filter Encoding

We use the meta-reasoning encoding to realize filters, by simply requiring that the observation context becomes inconsistent if the observed diagnosis does not pass the filter, i.e., we use $M^{mr(\theta, \mathcal{K}_p)}$ where $\mathcal{K} = \emptyset$ and θ is such that $\theta(D_1, D_2, \mathcal{K})$ holds if and only if $f(D_1, D_2) = 1$. Since no further bridge rules are needed to realize filtered diagnoses, it is sufficient to pick $\mathcal{K}_p = \emptyset$.

Definition 13. *Let M be an MCS and let f be a filter. Let $\mathcal{K}_p = \emptyset$ and let $\theta(D_1, D_2, \emptyset)$ hold iff $f(D_1, D_2) = 1$. Then $M^{mr(\theta, \mathcal{K}_p)}$ is the filter-encoding of M wrt. f , which we also denote by M^f .*

Example 11. *Reconsider the MCS $M = (C_1, C_2)$ of Example 7 where two scientists write a paper and diagnoses are to be filtered by a filter f if the authorship information is modified by a diagnosis in an incoherent way. The filter f (see Example 7) is defined as follows:*

$$f(D_1, D_2) = \begin{cases} 0 & \text{if } r_3 \in D_1, r_2 \notin D_1 \text{ or } r_3 \notin D_1, r_2 \in D_1 \\ 0 & \text{if } r_3 \in D_2, r_2 \notin D_2 \text{ or } r_3 \notin D_2, r_2 \in D_2 \\ 1 & \text{otherwise} \end{cases}$$

The resulting filter encoding M^f is the MCS $M^{mr(\theta, \mathcal{K}_p)} = (C'_1, C'_2, C_3)$, which has the same contexts and bridge rules as the MCS of Example 10. It only differs in the contents of the observation/encoding context C_3 which now realizes the filter f . We use ASP again for the logic of $C_3 = (L_{\Sigma}^{asp}, kb_3, br_3)$.

Recall that the knowledge-base formulas added by bridge rules to C_3 are either of the form $uncond_r$ or $not_removed_r$ and this information has to be exposed accordingly in the accepted belief set. Also remember that the definition of the meta-reasoning encoding requires that every accepted belief set only consists of beliefs in $\{removed_r, uncond_r \mid r \in br(M)\}$, but since no other bridge rule of $M^{mr(\theta, \mathcal{K}_p)}$ uses any other belief, we may allow further beliefs in the accepted belief set, i.e., our ASP program may use additional atoms.

The knowledge base kb_3 of C_3 then is:

$$kb_3 = \left\{ \begin{array}{ll} removed_{r_1} \leftarrow not\ not_removed_{r_1}. & \perp \leftarrow removed_{r_3}, not\ removed_{r_2}. \\ removed_{r_2} \leftarrow not\ not_removed_{r_2}. & \perp \leftarrow not\ removed_{r_3}, removed_{r_2}. \\ removed_{r_3} \leftarrow not\ not_removed_{r_3}. & \perp \leftarrow uncond_{r_3}, not\ uncond_{r_2}. \\ & \perp \leftarrow not\ uncond_{r_3}, uncond_{r_2}. \end{array} \right\}$$

The first three rules of kb_3 ensure that the removal information is correct while nothing is needed to ensure that the information about condition-free bridge rules is exposed (if bridge rule r_i is made unconditional, then the fact $uncond_{r_i}$ is added to kb_3 by the bridge rule $d2(r_i) \in br_3(M^{mr(\theta, \mathcal{K}_p)})$ being applicable and hence $uncond_{r_i}$ is also present in the answer set and thus in the belief set of C_3).

The four constraints of kb_3 finally encode the filter condition and they ensure that the context has no acceptable belief set if the corresponding diagnoses are applied.

Observe that the definition of θ uses the notion of f , which is an abstraction / generalization of some desired actual behavior. Instead of f , it is possible to use the desired actual behavior directly to realize the context C_{n+1} of $M^{mr(\theta, \mathcal{K}_p)}$, i.e., for a concrete use case where some logic is used to describe which diagnoses should be filtered out, it is not really necessary to first abstract the concrete

case to a filter f , build θ accordingly and then derive a concrete instantiation of C_{n+1} . Rather it is sufficient to take the definition of the meta-reasoning encoding and interpret it as the definition of the interfacing between the logic that does the filtering and the rest of the MCS framework. The reason why we introduced filters in general lies in the fact that this allows us to prove that all such filterings can be realized correctly. The following theorem now shows that diagnoses with protected bridge rules of M^f indeed correspond one-to-one to filtered diagnoses of M .

Theorem 1. *Let M be an MCS, let f be a filter and let M^f be the corresponding filter-encoding. Then, $D_{m,f}^\pm(M) = \{(D_1, D_2) \mid (d1(D_1), d2(D_2)) \in D_m^\pm(M^f, br_P)\}$.*

To obtain all minimal-filtered diagnoses of an MCS M wrt. the filter f , it is therefore sufficient to compute all subset-minimal diagnoses (with protected bridge rules) of the MCS $M^f = M^{mr(\theta, \mathcal{K}_p)}$. Note that this encoding does not come with increased computational cost, since M and M^f have the same number of bridge rules that possibly occur in a diagnosis with protected bridge rules. Consider M^f and the respective bridge rules, i.e., the set $br(M^f) \setminus br_P = d1(br(M)) \cup d2(br(M))$: since $body(r) = \{\top\}$ for $r \in d1(br(M))$ and $body(r) = \{\perp\}$ for $r \in d2(br(M))$, it holds for every $(R_1, R_2) \in D_m^\pm(M^f, br_P)$ that $r \in R_1$ implies $r \in d1(br(M))$ and $r \in R_2$ implies $r \in d2(br(M))$ (this follows from Lemma 8 in Appendix A.2). Hence, there are $2^{|d1(br(M))|} \times 2^{|d2(br(M))|}$ possibly relevant diagnoses for M^f while there are $2^{|br(M)|} \times 2^{|br(M)|}$ possible diagnoses for M ; since $|d1(br(M))| = |d2(br(M))| = |br(M)|$, the candidate space, i.e., the number of candidate diagnoses, for deciding whether a minimal-filtered diagnosis exists for M has the same size as the candidate space for deciding whether a minimal diagnosis with protected bridge rules exists for M^f .

5.2 Plain-Preference Encoding

We now show how to use the meta-reasoning encoding $M^{mr(\theta, \mathcal{K}_p)}$ for realizing preference orders. The set \mathcal{K}_p plays a crucial role, since it is used to map a given preference order on diagnoses to the \subseteq relation on \mathcal{K}_p . This allows us to select minimal \preceq -preferred diagnoses by considering \subseteq -minimal diagnoses of $M^{mr(\theta, \mathcal{K}_p)}$. Since the \subseteq -minimality on \mathcal{K}_p should take precedence over the remaining modified bridge rules of $M^{mr(\theta, \mathcal{K}_p)}$, we introduce a lexicographic order on bridge rules in which the latter are after those of \mathcal{K}_p . As we show in Section 6, the complexity of identifying a diagnosis with respect to prioritized bridge rules \mathcal{K}_p is not higher than identifying a minimal diagnosis.

In the following, we present a plain and simple encoding, which comes at the cost of \mathcal{K}_p being exponentially larger than $br(M)$, i.e., $M^{mr(\theta, \mathcal{K}_p)}$ contains exponentially many more bridge rules than M . We also prove that the approach is correct for total preference orders. Before presenting the plain encoding, first we introduce the notion of a prioritized-minimal diagnosis, and second we show how a total order can be mapped to the \subseteq relation.

In the following, we write $(D_1, D_2) \subseteq_{br_H} (D'_1, D'_2)$ as shorthand for $(D_1 \cap br_H, D_2 \cap br_H) \subseteq (D'_1 \cap br_H, D'_2 \cap br_H)$, i.e., we denote by \subseteq_{br_H} the restriction of \subseteq to the set br_H ; furthermore, we write $=_{br_H}$ for an analogous restriction on $=$.

To realize a total preference order, the following definition is sufficient where we select from the set of minimal diagnoses with protected bridge rules those that are minimal with respect to the prioritized bridge rules. The bridge rules that are marked as prioritized take precedence for minimality. A prioritized-minimal diagnosis is subset-minimal with respect to prioritized bridge rules (regardless of minimality of the remaining bridge rules).

Definition 14. Let M be an MCS with bridge rules $br(M)$, protected rules $br_P \subseteq br(M)$, and prioritized rules $br_H \subseteq br(M)$. The set of prioritized-minimal diagnoses is $D^\pm(M, br_P, br_H) = \{D \in D_m^\pm(M, br_P) \mid \forall D' \in D_m^\pm(M, br_P) : D' \subseteq_{br_H} D \Rightarrow D' =_{br_H} D\}$.

We now show how an arbitrary order relation over a pair of sets may be mapped to the \subseteq -relation on an exponentially larger set, i.e., we map \preceq on the diagnoses of an MCS M , to another exponentially larger set.

Definition 15. Let \preceq be a preference relation on $2^{br(M)} \times 2^{br(M)}$ and let $g : 2^{br(M)} \times 2^{br(M)} \rightarrow \mathcal{K}_p$ be a bijective mapping where \mathcal{K}_p is arbitrary. Then, the subset-mapping $map_{\preceq}^g : 2^{br(M)} \times 2^{br(M)} \rightarrow 2^{\mathcal{K}_p}$ is defined as follows. For every $(D_1, D_2) \in 2^{br(M)} \times 2^{br(M)}$:

$$map_{\preceq}^g(D_1, D_2) = \{K \in \mathcal{K}_p \mid K = g(D'_1, D'_2) \text{ for some } (D'_1, D'_2) \preceq (D_1, D_2)\}.$$

Observe that $map_{\preceq}^g(D_1, D_2)$ collects $g(D'_1, D'_2)$ of all (D'_1, D'_2) “below” (D_1, D_2) .

The following lemma shows that the subset-mapping correctly maps a preference relation on diagnoses to the subset-relation on an exponentially larger set. This allows to decide whether a diagnosis is more preferred than another solely based on subset relationship.

Lemma 1. Let \preceq be a preference on candidate diagnoses of an MCS M , and let g be a bijective mapping $g : 2^{br(M)} \times 2^{br(M)} \rightarrow \mathcal{K}_p$ for any set \mathcal{K}_p . Then, for any $(D_1, D_2), (D'_1, D'_2) \in 2^{br(M)} \times 2^{br(M)}$ it holds that $(D_1, D_2) \preceq (D'_1, D'_2)$ iff $map_{\preceq}^g(D_1, D_2) \subseteq map_{\preceq}^g(D'_1, D'_2)$.

We now use map_{\preceq}^g to map the preference of a total order \preceq to the set \mathcal{K}_p which occurs in the meta-reasoning transformation $M^{mr(\theta, \mathcal{K}_p)}$. To that end, we choose $\theta(D_1, D_2, \mathcal{K})$ such that it holds iff $map_{\preceq}^g(D_1, D_2) = \mathcal{K}$, for $\mathcal{K} \subseteq \mathcal{K}_p$. By that, every diagnosis of $M^{mr(\theta, \mathcal{K}_p)}$ with protected bridge rules $(d1(D_1), d2(D_2) \cup \mathcal{K})$ contains the preference \preceq encoded in \mathcal{K} . Selecting a diagnosis of $M^{mr(\theta, \mathcal{K}_p)}$ where \mathcal{K} is minimal then selects a preferred diagnosis according to \preceq .

Definition 16. Let M be an MCS and let \preceq be a preference order over diagnoses of M . Furthermore, let

$$\mathcal{K}_p = \{(n+1 : diag_{D_1, D_2}) \leftarrow \perp. \mid D_1, D_2 \subseteq br(M)\} \quad (9)$$

and let $g : 2^{br(M)} \times 2^{br(M)} \rightarrow \mathcal{K}_p$ be a bijective function such that $g(D_1, D_2) = (n+1 : diag_{D_1, D_2}) \leftarrow \perp.$ for all $D_1, D_2 \subseteq br(M)$. Let $\theta(D_1, D_2, \mathcal{K})$ hold iff $map_{\preceq}^g(D_1, D_2) = \mathcal{K}$. Then the MCS $M^{mr(\theta, \mathcal{K}_p)}$ is called the plain encoding of M wrt. \preceq , which we also denote by $M^{pl\preceq}$; all bridge rules of \mathcal{K}_p are prioritized, i.e., $br_H = \mathcal{K}_p$.

Note that since map_{\preceq}^g is a function, also θ is equivalent to a function $2^{br(M)} \times 2^{br(M)} \rightarrow \mathcal{K}_p$.

Example 12. We consider the hospital MCS M of Example 3 again using a preference order on diagnoses similar to the one of Example 9, i.e., we prefer changing bridge rules regarding health, r_1, r_2 , as little as possible. To make the preference total, we use cardinality-minimality, i.e.,

$$(D_1, D_2) \preceq (D'_1, D'_2) \text{ iff } |\{r_1, r_2\} \cap (D_1 \cup D_2)| \leq |(D'_1 \cup D'_2) \cap \{r_1, r_2\}|.$$

The resulting MCS $M^{mr(\theta, \mathcal{K}_p)}$ is outlined in Figure 4, where only bridge rules stemming from r_5 of $br(M)$ and some of the bridge rules of the observation context C_4 are indicated. Note that $br_4(M^{mr(\theta, \mathcal{K}_p)})$ contains for every possible diagnosis of M a distinguished bridge rule. For $C_4 = (L_{\Sigma}^{asp}, kb_4, br_4)$, we use ASP again to show a possible realization; kb_4 consists of the rules:




Figure 4: Contexts and some bridge rules of the plain encoding $M^{pl_{\preceq}} = (C_1, C_2, C_3, C_4)$ of the hospital MCS wrt. \preceq from Example 12. For illustration purposes, only bridge rules stemming from r_5 and some from \mathcal{K}_p are shown; dashed lines indicate bridge rules r_1, \dots, r_4 from M whose corresponding bridge rules in $M^{pl_{\preceq}}$ are not shown.

$$\begin{aligned}
 & removed_r \leftarrow not\ not_removed_r. & r \in br(M) \\
 & \perp \leftarrow cur_diag_{D_1, D_2}, not\ diag_{D_1, D_2}. & D_1, D_2 \subseteq br(M), \\
 & cur_diag_{D'_1, D'_2} \leftarrow cur_diag_{D_1, D_2}. & (D'_1, D'_2) \preceq (D_1, D_2), \\
 & cur_diag_{D_1, D_2} \leftarrow removed_{r_1, \dots, r_k}, removed_{r'_1, \dots, r'_m}, uncond_{r'_1, \dots, r'_m}. & D_1, D_2 \subseteq br(M), D_1 = \{r_1, \dots, r_k\}, D_2 = \{r'_1, \dots, r'_m\}.
 \end{aligned}$$

Intuitively, the rules of the first line ensure that diagnosis observation is exposed correctly in an accepted belief set of C_4 ; the constraints following ensure the presence of condition-free bridge rules. Rules of the third line guarantee that all bridge rules corresponding to more-preferred diagnoses also need to be condition-free; under ASP semantics, these rules effect $map_{\preceq}^g(D_1, D_2)$. Finally, the rules of the last line recognize one of the exponentially many candidate diagnoses.

The next theorem shows the relation between minimal \preceq -preferred diagnoses of M wrt. a total preference \preceq and prioritized-minimal diagnoses of $M^{pl_{\preceq}}$. Observe that map_{\preceq}^g is injective since \preceq is reflexive, thus $map_{\preceq}^g(D_1, D_2)$ contains $g(D_1, D_2)$, which by g being a bijection is different for every candidate diagnosis (D_1, D_2) . Therefore, map_{\preceq}^g is bijective on its range and it allows to establish a one-to-one relation between minimal \preceq -preferred diagnoses of M and prioritized-minimal ones of $M^{pl_{\preceq}}$. Intuitively, this shows that for a total preference order, the set of prioritized-minimal diagnoses of the plain encoding of M wrt. \preceq can be used to select the minimal \preceq -preferred diagnoses of M .

Theorem 2. For every MCS M and total preference \preceq on its diagnoses, $D^\pm(M^{pl\preceq}, br_P, br_H) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D_{m,\preceq}^\pm(M), map_{\preceq}^g(D_1, D_2) = \mathcal{K}\}$.

To select minimal \preceq -preferred diagnoses based on an arbitrary preference order, another encoding can be utilized, which we describe next.

5.3 Clone-Preference Encoding

The clone encoding requires that the original MCS M is cloned, but comes at the advantage of only requiring linearly many bridge rules, specifically it holds that $|\mathcal{K}_p| = 4|br(M)| + 1$. First an MCS $2M$ is built that consists of two independent copies of M , and then the meta-reasoning encoding is applied on $2M$, i.e., the resulting MCS is $(2M)^{mr(\theta, \mathcal{K}_p)}$. We show that the minimal \preceq -preferred diagnoses can be selected from $(2M)^{mr(\theta, \mathcal{K}_p)}$ using a slightly more involved diagnosis notion with prioritized bridge rules. The complexity of selecting such diagnoses increases, but it is still worst-case optimal as shown later.

The basic idea of the clone encoding is that the original MCS is duplicated such that the observation context sees two diagnoses of the original MCS at the same time and is able to compare them. Intuitively, if we combine two MCS M and M' into a single one M'' , then every diagnosis of the combined MCS M'' is the combination of a diagnosis of M with a diagnosis of M' . Establishing this technically requires some care, since one needs to account for the fact that contexts are identified by their position: Hence, M'' cannot simply contain the bridge rules of M and M' . We thus introduce context shifting and build an operator \otimes to combine two MCS. We then show some general properties of the operator, and finally give the clone encoding, which adds a certain observation context to the combination $M \otimes M'$ of the MCS M whose minimal \preceq -preferred diagnoses we are interested in.

For shifting contexts, we use a permutation $I : \mathbb{N} \rightarrow \mathbb{N}$, i.e., I is a bijective mapping. Given a bridge rule r of form (2), then $I(r)$ is the bridge rule

$$(I(k) : s) \leftarrow (I(c_1) : p_1), \dots, (I(c_j) : p_j), \mathbf{not} (I(c_{j+1}) : p_{j+1}), \dots, \mathbf{not} (I(c_m) : p_m);$$

furthermore, for a set R of bridge rules we have $I(R) = \{I(r) \mid r \in R\}$ and for a context $C_i = (L_i, kb_i, br_i)$ we have $I(C_i) = (L_i, kb_i, I(br_i))$. Given an MCS $M = (C_1, \dots, C_n)$, a permutation I is *compatible* with M if $I(x) \leq n$ holds for all $x \leq n$, i.e., I is a permutation on $C(M)$; the “shuffled” version of M wrt. a compatible I then is $I(M) = (I(C_{I^{-1}(1)}), \dots, I(C_{I^{-1}(n)}))$. Given a belief state $S = (S_1, \dots, S_n)$ we have $I(S) = (S_{I^{-1}(1)}, \dots, S_{I^{-1}(n)})$.

To combine two existing MCS $M = (C_1, \dots, C_n)$ and $M' = (C'_1, \dots, C'_m)$ into a new one, we use the following \otimes operator:

$$M \otimes M' = (C_1, \dots, C_n, I(C'_1), \dots, I(C'_m)) \quad \text{where } I(x) = \begin{cases} n+x & \text{for } 1 \leq x \leq m, \\ x-m & \text{for } m+1 \leq x \leq n+m, \\ x & \text{otherwise.} \end{cases}$$

In the following, we call I the *permutation wrt. $M \otimes M'$* . Note that by construction the permutation I wrt. $M \otimes M'$ is compatible with $M \otimes M'$. Recall that $M[R_1, R_2] = M[br(M) \setminus R_1 \cup cf(R_2)]$. Regarding modifications and candidate diagnoses, we then observe that $M[A_1, A_2] \otimes M'[B_1, B_2] = (M \otimes M')[A_1 \cup I(B_1), A_2 \cup I(B_2)]$ where I is the mapping wrt. $M \otimes M'$.

The following lemma shows that shifting has no influence on acceptability.

Lemma 2. *Given an MCS $M = (C_1, \dots, C_n)$ and a compatible permutation I , it holds that $S \in \text{EQ}(M)$ iff $I(S) \in \text{EQ}(I(M))$.*

It immediately follows that $S \in \text{EQ}(M[D_1, D_2])$ iff $I(S) \in \text{EQ}(I(M[D_1, D_2]))$ for any candidate diagnosis (D_1, D_2) .

The main observation on the \otimes operator is that $M \otimes M'$ admits exactly those diagnoses which are a combination of a diagnosis of M and a diagnosis of M' .

Proposition 6. *Given two MCS M and M' , then $D^\pm(M \otimes M') = \{(A_1 \cup I(B_1), A_2 \cup I(B_2)) \mid (A_1, A_2) \in D^\pm(M), (B_1, B_2) \in D^\pm(M')\}$ where I is the permutation wrt. $M \otimes M'$.*

Next we define for an MCS $M = (C_1, \dots, C_n)$ the MCS $2M = (C_1, \dots, C_{2n})$ as $2M = M \otimes M$, i.e., $2M$ consists of two clones of M . For easier reference, we write $2.r$ to denote the clone of the bridge rule r , i.e., $2.r = I(r)$ where I is the permutation wrt. $M \otimes M$. Note that $2.br(M)$ is the set of bridge rules of M shifted by n , i.e., $2.br(M)$ is the set of bridge rules of the second clone of M .

The next lemma, which follows from Proposition 6, shows that diagnoses of $2M$ correspond to diagnoses of M in such a way that every diagnosis of $2M$ is composed of two diagnoses of M .

Lemma 3. *Let M be an MCS. Then $(D_1, D_2) \in D^\pm(2M)$ holds iff there exist $(D'_1, D'_2) \in D^\pm(M)$ and $(D''_1, D''_2) \in D^\pm(M)$ such that $D_1 = D'_1 \cup 2.D''_1$ and $D_2 = D'_2 \cup 2.D''_2$.*

The underlying idea of the encoding is that a specific prioritized bridge rule t_{max} indicates whether the diagnosis applied to the second clone is preferred over the diagnosis applied to the first clone. Additionally, the diagnosis of the first clone is exhibited via prioritized bridge rules, while the diagnosis of the second clone is only exhibited via non-prioritized bridge rules.

If the diagnosis applied to the second clone is more preferred than the one applied to the first, then t_{max} needs not become condition-free. Thus, if for a given diagnosis of the first clone, there exists some more preferred diagnosis of the second clone, then there exists a diagnosis where t_{max} is not included. A diagnosis D such that no more preferred diagnosis D' exists is maximal wrt. the inclusion of t_{max} , because there exists no more preferred diagnosis D' of M that could occur at the second clone. Selecting a diagnosis that modifies a minimal set of prioritized bridge rules and that contains t_{max} thus selects a \preceq -preferred diagnosis. We define t_{max} as follows:

$$t_{max} : \quad (2n+1 : ismax) \leftarrow \perp.$$

To represent the diagnosis of the first clone, we use the following prioritized bridge rules. For a bridge rule $r \in br(M)$ let $in_1(r)$, $\bar{in}_1(r)$, $in_2(r)$, and $\bar{in}_2(r)$ denote the following bridge rules:

$$\begin{array}{ll} in_1(r) : & (2n+1 : in_1(r)) \leftarrow \perp. & in_2(r) : & (2n+1 : in_2(r)) \leftarrow \perp. \\ \bar{in}_1(r) : & (2n+1 : \bar{in}_1(r)) \leftarrow \perp. & \bar{in}_2(r) : & (2n+1 : \bar{in}_2(r)) \leftarrow \perp. \end{array}$$

We identify a candidate diagnosis $(D_1, D_2) \in 2^{br(M)} \times 2^{br(M)}$ using these bridge rules by the set $\mathcal{K}(D_1, D_2) = \{in_1(r) \mid r \in D_1\} \cup \{\bar{in}_1(r) \mid r \notin D_1\} \cup \{in_2(r) \mid r \in D_2\} \cup \{\bar{in}_2(r) \mid r \notin D_2\}$. The clone encoding then formally is as follows.

Definition 17. Let $M = (C_1, \dots, C_n)$ be an MCS and \preceq a preference order. The clone encoding of M wrt. \preceq is the MCS $2M^{mr(\theta, \mathcal{K}_p)}$ where $2M = (C_1, \dots, C_{2n}) = M \otimes M$,

$$\mathcal{K}_p = \bigcup_{r \in br(M)} \{(2n+1 : q) \leftarrow \perp., \mid q \in \{in_1(r), \overline{in}_1(r), in_2(r), \overline{in}_2(r)\}\} \cup \{t_{max}\}$$

and for any $R_1, R_2 \subseteq br(2M)$, and $R_3 \subseteq \mathcal{K}_p$, $\theta(R_1, R_2, R_3)$ holds iff $R_1 = D_1 \cup 2.D'_1$, $R_2 = D_2 \cup 2.D'_2$ and either

- $(D_1, D_2) = (D'_1, D'_2)$ and $R_3 = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$ or
- $(D'_1, D'_2) \prec (D_1, D_2)$ and $R_3 = \mathcal{K}(D_1, D_2)$.

The protected bridge rules br_P are all bridge rules except those of context C_{2n+1} ; the prioritized bridge rules are $br_H = \mathcal{K}_p$. We denote the clone encoding of M wrt. \preceq by $M^\preceq = 2M^{mr(\theta, \mathcal{K}_p)}$.

Note that the second case above with $(D'_1, D'_2) \prec (D_1, D_2)$ implies that $(D_1, D_2), (D'_1, D'_2)$ are two diagnoses of M , because the MCS $2M$ only admits a diagnosis if $(D_1, D_2) \in D^\pm(M)$ and $(D'_1, D'_2) \in D^\pm(M)$ both hold (cf. Lemma 3). Also observe that $M^\preceq = (2M)^{mr(\theta, \mathcal{K}_p)} = (M \otimes M)^{mr(\theta, \mathcal{K}_p)}$ is linear in the size of M , as for every bridge rule in M there exist $2 \cdot 4 + 4$ bridge rules in M^\preceq , (the factor 2 is from $M \otimes M$, the factor 4 is from the meta-reasoning encoding itself and the +4 is due to \mathcal{K}_p). In total $|br(M^\preceq)| = 12 \cdot |br(M)| + 1$, where the +1 is due to t_{max} .

Example 13. Reconsider the MCS M from Example 3 shown in Figure 1. Applying the clone encoding on M wrt. a preference order \preceq results in the MCS $M^\preceq = (C_1, C_2, C_3, C_4, C_5, C_6, C_7)$ depicted in Figure 5. It is based on two clones of M , where the first comprises the contexts C_1, C_2, C_3 and the second the contexts C_4, C_5, C_6 . The context C_7 finally is the observation/encoding context. For each bridge rule $r \in br(M)$ and each clone of M , there are two bridge rules simulating that r is removed respectively made condition-free in the clone. Hence the bridge rules in M^\preceq for the contexts C_1, \dots, C_6 are:

$$\begin{aligned} r'_1 &: && (2 : hyperglycemia) \leftarrow (1 : hyperglycemia), \mathbf{not} (7 : removed_{r_1}). \\ r''_1 &: && (2 : hyperglycemia) \leftarrow (7 : uncond_{r_1}). \\ r'_2 &: && (2 : allow_animal_insulin) \leftarrow \mathbf{not} (1 : allergic_animal_insulin), \mathbf{not} (7 : removed_{r_2}). \\ r''_2 &: && (2 : allow_animal_insulin) \leftarrow (7 : uncond_{r_2}). \\ r'_3 &: && (3 : bill_animal_insulin) \leftarrow (2 : give_animal_insulin), \mathbf{not} (7 : removed_{r_3}). \\ r''_3 &: && (3 : bill_animal_insulin) \leftarrow (7 : uncond_{r_3}). \\ r'_4 &: && (3 : bill_human_insulin) \leftarrow (2 : give_human_insulin), \mathbf{not} (7 : removed_{r_4}). \\ r''_4 &: && (3 : bill_human_insulin) \leftarrow (7 : uncond_{r_4}). \\ r'_5 &: && (3 : insurance_B) \leftarrow (1 : insurance_B), \mathbf{not} (7 : removed_{r_5}). \\ r''_5 &: && (3 : insurance_B) \leftarrow (7 : uncond_{r_5}). \\ I(r_1)' &: && (5 : hyperglycemia) \leftarrow (4 : hyperglycemia), \mathbf{not} (7 : removed_{r_1}). \\ I(r_1)'' &: && (5 : hyperglycemia) \leftarrow (7 : uncond_{r_1}). \\ I(r_2)' &: && (5 : allow_animal_insulin) \leftarrow \mathbf{not} (4 : allergic_animal_insulin), \mathbf{not} (7 : removed_{r_2}). \\ I(r_2)'' &: && (5 : allow_animal_insulin) \leftarrow (7 : uncond_{r_2}). \end{aligned}$$

$$\begin{aligned}
 I(r_3)' &: (6: \text{bill_animal_insulin}) \leftarrow (5: \text{give_animal_insulin}), \mathbf{not} (7: \text{removed}_{r_3}). \\
 I(r_3)'' &: (6: \text{bill_animal_insulin}) \leftarrow (7: \text{uncond}_{r_3}). \\
 I(r_4)' &: (6: \text{bill_human_insulin}) \leftarrow (5: \text{give_human_insulin}), \mathbf{not} (7: \text{removed}_{r_4}). \\
 I(r_4)'' &: (6: \text{bill_human_insulin}) \leftarrow (7: \text{uncond}_{r_4}). \\
 I(r_5)' &: (6: \text{insurance_B}) \leftarrow (4: \text{insurance_B}), \mathbf{not} (7: \text{removed}_{r_5}). \\
 I(r_5)'' &: (6: \text{insurance_B}) \leftarrow (7: \text{uncond}_{r_5}).
 \end{aligned}$$

For each bridge rule r above there are two (non-protected) bridge rules at the observation context C_7 indicating whether r is regarded as removed respectively condition-free. Overall, the set of bridge rules of C_7 is:

$$\begin{aligned}
 \{ (7: \text{ismax}) \leftarrow \perp. \} \cup \{ & (7: \text{not_removed}_{r_1}) \leftarrow \top. & (7: \text{uncond}_{r_1}) \leftarrow \perp. \\
 & (7: \text{not_removed}_{r_2}) \leftarrow \top. & (7: \text{uncond}_{r_2}) \leftarrow \perp. \\
 & \dots & \\
 & (7: \text{not_removed}_{I(r_4)}) \leftarrow \top. & (7: \text{uncond}_{I(r_4)}) \leftarrow \perp. \\
 & (7: \text{not_removed}_{I(r_5)}) \leftarrow \top. & (7: \text{uncond}_{I(r_5)}) \leftarrow \perp. \\
 & \dots & \\
 & (7: \text{in}_1(r_1)) \leftarrow \perp. & (7: \overline{\text{in}}_1(r_1)) \leftarrow \perp. \\
 & (7: \text{in}_2(r_1)) \leftarrow \perp. & (7: \overline{\text{in}}_2(r_1)) \leftarrow \perp. \\
 & \dots & \\
 & (7: \text{in}_1(r_5)) \leftarrow \perp. & (7: \overline{\text{in}}_1(r_5)) \leftarrow \perp. \\
 & (7: \text{in}_2(r_5)) \leftarrow \perp. & (7: \overline{\text{in}}_2(r_5)) \leftarrow \perp. & \}
 \end{aligned}$$

The bridge rules of \mathcal{K}_p are now all rules among them with head formula $\text{in}_i(r)$ or $\overline{\text{in}}_i(r)$ for $i \in \{1, 2\}$ and $r \in \text{br}(M)$, plus the bridge rule t_{max} . These bridge rules are the prioritized ones i.e., $\text{br}_H = \mathcal{K}_p$. A detailed description for a concrete preference order \preceq is given in Appendix B (Example 19).

For selecting minimal \preceq -preferred diagnoses based on an arbitrary preference order, Definition 14 is strengthened in two steps: first, if two diagnoses are equal considering their prioritized bridge rules, then subset-minimality on the remaining bridge rules is taken into account. Second, since we only want to select diagnoses where no more preferred ones exist, we consider only prioritized-minimal diagnoses that contain the bridge rule t_{max} .

For the first step, let M be an MCS with bridge rules $\text{br}(M)$, protected rules br_P , and prioritized rules $\text{br}_H \subseteq \text{br}(M)$. The set of subset-minimal prioritized-minimal diagnoses then is:

$$\begin{aligned}
 D_m^\pm(M, \text{br}_P, \text{br}_H) = \{ & D \in D_m^\pm(M, \text{br}_P) \mid \text{Min}_{\text{br}_H, \text{br}_P}(M, D) \wedge \forall D' \in D_m^\pm(M, \text{br}_P) : \\
 & \text{Min}_{\text{br}_H, \text{br}_P}(M, D') \Rightarrow (D' \subseteq_{\text{br}(M) \setminus \text{br}_H} D \Rightarrow D =_{\text{br}(M) \setminus \text{br}_H} D') \} \quad (10)
 \end{aligned}$$

where $\text{Min}_{\text{br}_H, \text{br}_P}(M, X)$ denotes that X is minimal among all protected diagnoses with respect to br_H , i.e., $\text{Min}_{\text{br}_H, \text{br}_P}(M, X) = \forall D \in D_m^\pm(M, \text{br}_P) : D \subseteq_{\text{br}_H} X \Rightarrow X =_{\text{br}_H} D$. The first condition ensures that a diagnosis D is prioritized-minimal and for all other prioritized-minimal diagnoses D' it holds that D is minimal wrt. non-prioritized bridge rules.

For the second step, we just add to $D_m^\pm(M, \text{br}_P, \text{br}_H)$ the condition that D and D' make t_{max} condition-free. Formally:




Figure 5: The MCS $M^{\preceq} = (C_1, C_2, \dots, C_7)$ of Example 13. Some bridge rules of the observation context C_7 are shown and the bridge rules stemming from r_5 ; dashed and gray lines indicate the other bridge rules of $M \otimes M$ whose resulting bridge rules in M^{\preceq} are omitted. The prioritized bridge rules of M^{\preceq} are t_{max} and all bridge rules $in_i(r_j)$ and $\overline{in}_i(r_j)$.

Definition 18. Given an MCS M with protected bridge rules br_P and prioritized bridge rules br_H , the set of subset-minimal prioritized-minimal (mpm) diagnoses wrt. t_{max} is

$$D_{m,t_{max}}^{\pm}(M, br_P, br_H) = \{D \in D_m^{\pm}(M, br_P) \mid \text{Min}_{br_H, br_P}(M, D) \wedge t_{max} \in D \\ \wedge \forall D' \in D_m^{\pm}(M, br_P) : (\text{Min}_{br_H, br_P}(M, D') \wedge t_{max} \in D') \\ \Rightarrow (D' \subseteq_{br(M) \setminus br_H} D \Rightarrow D =_{br(M) \setminus br_H} D')\}$$

where $t_{max} \in D$ stands for $D = (D_1, D_2) \wedge t_{max} \in D_2$ and $\text{Min}_{br_H, br_P}(M, X)$ is as above.

Intuitively, D is an mpm-diagnosis, if it respects protected bridge rules and contains t_{max} , if it is preferred, i.e., it is minimal wrt. prioritized bridge rules br_H among all other diagnoses of the MCS M , and if for all other preferred diagnoses that contain t_{max} it holds that D is subset-minimal wrt. regular bridge rules.

Example 14. Consider again the clone encoding M^{\preceq} of Example 13 with a preference order \preceq on the diagnoses of the MCS M that prefers changing bridge rules regarding health, i.e., r_1, r_2 , as little

as possible. Formally, let $(D_1, D_2) \preceq (D'_1, D'_2)$ hold iff $\{r_1, r_2\} \cap (D_1 \cup D_2) \subseteq (D'_1 \cup D'_2) \cap \{r_1, r_2\}$. We illustrate the notion of mpm-diagnosis by examining three candidate diagnoses of M^\preceq .

(1) Consider a diagnosis on M^\preceq corresponding to the diagnosis $(D_1, D_2) = (\{r_4\}, \emptyset)$ on M , viz. the diagnosis $D^p = (\{d1(r_4), d1(I(r_4))\}, \mathcal{K}(\{r_4\}, \emptyset) \cup \{t_{max}\})$. Then D^p is an mpm-diagnosis of M^\preceq , because (i) it contains no protected bridge rules; (ii) θ holds according to the conditions of the clone encoding, specifically for $R_1 = D_1 \cup 2.D_1, R_2 = D_2 \cup 2.D_2$ and $R_3 = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$, hence $D^p \in D_m^\pm(M^\preceq, br_P)$; (iii) D^p is minimal among diagnoses respecting prioritized bridge rules; and (iv) no other such diagnosis is a subset of D^p on non-prioritized bridge rules.

(2) The diagnosis $D^o = (\{d1(r_4), d1(r_5), d1(I(r_4)), d1(I(r_5))\}, \mathcal{K}(\{r_4, r_5\}, \emptyset) \cup \{t_{max}\})$ contains no protected bridge rules and θ holds as defined by the clone encoding. Furthermore, D^o is minimal among the diagnoses that respect prioritized bridge rules, since the encoding ensures that $\mathcal{K}(\{r_4\}, \emptyset)$ and $\mathcal{K}(\{r_4, r_5\}, \emptyset)$ are incomparable, because $\overline{in}1(r_5) \in \mathcal{K}(\{r_4\}, \emptyset) \setminus \mathcal{K}(\{r_4, r_5\}, \emptyset)$ while $in1(r_5) \in \mathcal{K}(\{r_4\}, \emptyset) \setminus \mathcal{K}(\{r_4, r_5\}, \emptyset)$. However, D^o is not an mpm-diagnosis, because D^p is minimal with respect to prioritized rules and is smaller on the non-prioritized rules than D^o , i.e., $\{d1(r_4), d1(I(r_4))\} \subset \{d1(r_4), d1(r_5), d1(I(r_4)), d1(I(r_5))\}$.

(3) Consider a diagnosis D^n stemming from the diagnosis $(\{r_1\}, \emptyset)$, which is not preferred according to \preceq , i.e., it holds that $(\{r_4\}, \emptyset) \preceq (\{r_1\}, \emptyset)$. Let $D^n = (\{d1(r_1), d1(I(r_1))\}, \mathcal{K}(\{r_1\}, \emptyset) \cup \{t_{max}\})$. Then, D^n is not an mpm-diagnosis, because it is not minimal among diagnoses respecting prioritized bridge rules. Consider the diagnosis where the diagnosis $(\{r_4\}, \emptyset)$ is applied to the second clone of M , i.e., $D^s = (\{d1(r_1), d1(I(r_4))\}, \mathcal{K}(\{r_1\}, \emptyset))$. Observe that θ holds for D^s according to Definition 17, because $(\{r_4\}, \emptyset) \prec (\{r_1\}, \emptyset)$ and $R_3 = \mathcal{K}(\{r_1\}, \emptyset)$. It also holds that $D^s \subset_{br_H} D^n$, because $\mathcal{K}(\{r_1\}, \emptyset) \subset \mathcal{K}(\{r_1\}, \emptyset) \cup \{t_{max}\}$ and therefore $Min_{br_H, br_P}(M^\preceq, D^n)$ does not hold. Consequently, D^n is not an mpm-diagnosis.

As we show in the next section, the notion of mpm-diagnosis is computationally harder than the notion of prioritized-minimal diagnosis. Nevertheless, the problem itself (i.e., identifying a minimal \preceq -preferred diagnosis) is shown to be as hard as this notion, which means the notion of mpm-diagnosis is worst-case optimal.

Note that $D, D' \in D^\pm(M, br_P)$ implies that $D \subseteq_{br(M) \setminus br_H} D'$ holds iff $D \subseteq_{br(M) \setminus br_H \setminus br_P} D'$ holds, because $D = (D_1, D_2) \in D^\pm(M, br_P)$ implies that $D_1 \cap br_P = \emptyset = D_2 \cap br_P$. The same also holds for $=_{br(M) \setminus br_H}$ and $=_{br(M) \setminus br_H \setminus br_P}$.

As it appears, $D^\pm(M^\preceq, br_P, br_H)$ suffices to obtain those diagnoses of M that are \preceq -preferred. In the following, we write $t(D_1, D_2)$ as a shorthand for the corresponding candidate diagnosis in the MCS M^\preceq , i.e., $t(D_1, D_2) = (d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2) \cup \mathcal{K}(D_1, D_2) \cup \{t_{max}\})$.

Theorem 3. *Let M be an MCS and let \preceq be a preference order on the diagnoses of M . Then $D \in D^\pm(M)$ is \preceq -preferred iff $t(D) \in D^\pm(M^\preceq, br_P, br_H)$ holds.*

Note that $t(D) \in D^\pm(M^\preceq, br_P, br_H)$ implies that $t_{max} \in t(D)$; but there also are diagnoses $T \in D^\pm(M^\preceq, br_P, br_H)$ such that $t_{max} \notin T$. Nevertheless, it follows directly from the definition of M^\preceq that for any $T \in D^\pm(M^\preceq, br_P, br_H)$ with $t_{max} \in T$ there exist $D_1, D_2 \subseteq br(M)$ such that $T = t(D_1, D_2)$. Hence, diagnoses of $D^\pm(M^\preceq, br_P, br_H)$ that contain t_{max} correspond one-to-one to \preceq -preferred diagnoses of M .

Example 15 (ctd.). *Consider the diagnosis $D = (\{r_4, r_5\}, \emptyset)$. It is \preceq -preferred, as it does not modify any of the bridge rules in $\{r_1, r_2\}$. The corresponding diagnosis $t(D)$ is the diagnosis D^o*

from Example 14, i.e. $t(D) = (\{d1(r_4), d1(r_5), d1(I(r_4)), d1(I(r_5))\}, \mathcal{K}(\{r_4, r_5\}, \emptyset) \cup \{t_{max}\})$. It holds that $D^o \in D^\pm(M^\preceq, br_P, br_H)$, as stated by Theorem 3.

The next theorem shows that the clone encoding M^\preceq and the notion of mpm-diagnosis $D_{m, t_{max}}^\pm$ allow to select all minimal \preceq -preferred diagnoses of M . This theorem therefore establishes that the clone encoding is sound and complete.

Theorem 4. *Let M be an MCS and let \preceq be a preference order on diagnoses of M . Then $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$ holds iff $t(D_1, D_2) \in D_{m, t_{max}}^\pm(M^\preceq, br_P, br_H)$ holds.*

Recall that given a CP-net N that is compatible with an MCS M , the minimal \preceq -preferred diagnoses according to \preceq^N and the irredundant N -preferred diagnoses coincide, i.e., $D_{ird}^\pm(M, N) = D_{m, \preceq^N}^\pm(M)$ (cf. Proposition 2). One thus can realize the selection of “optimal” diagnoses according to a CP-net using the clone encoding M^{\preceq^N} and the methods provided in this section. Also note that M^{\preceq^N} has size only linearly larger than M .

Since the approaches only specify some of the behavior of the observation context, the concrete choice of a logic to realize the observation remains to the user. This is especially useful for preference formalisms like CP-nets where algorithms may be chosen according to the computational complexity of the employed CP-net.

6. Computational Complexity

To select preferred and most preferred diagnoses, the previous section introduced several advanced notions of diagnosis. In this section we investigate the computational complexity of these notions. As it turns out, considering protected bridge rules as well as prioritized bridge rules does not increase the computational complexity of identifying a diagnosis.¹ Identifying subset-minimal diagnoses among those with protected and prioritized bridge rules, however, incurs additional cost. Since selecting most preferred diagnoses is hard for the same complexity class in the basic case, the additional cost are expected and our approach is thus worst-case optimal. We begin by recalling the necessary notions of complexity analysis in MCS.

6.1 Complexity Classes and Context Complexity

Recall that **P**, **EXPTIME**, and **PSPACE** are the classes of problems that can be decided using a deterministic Turing machine in polynomial time, exponential time, and polynomial space, respectively. Furthermore **NP** (resp., **coNP**) is the class of problems that can be decided on a non-deterministic Turing machine in polynomial time, where one (resp., all) computation paths accept. The polynomial hierarchy is built as follows: $\Sigma_0^P = \Pi_0^P = \mathbf{P}$, and for all $i \geq 1$, $\Sigma_i^P = \mathbf{NP}^{\Sigma_{i-1}^P}$ is **NP** with a Σ_{i-1}^P oracle and Π_i^P is **co- Σ_i^P** .

Given a complexity class C , $\mathbf{D}(C)$ denotes the “difference class” of C , i.e., $\mathbf{D}(C) = \{L_1 \times L_2 \mid L_1 \in C, L_2 \in \mathbf{co-}C\}$ is the complexity class of decision problems that are the “conjunction” of a problem L_1 in C and a problem L_2 in $\mathbf{co-}C$. We use the notation that $\mathbf{D}(\mathbf{NP}) = \mathbf{D}_1^P$ and

1. In line with and for comparability to the work of Eiter et al. (2014), we concentrate on recognizing diagnoses and omit deciding (advanced) diagnosis existence. Briefly, the latter problem is for context complexity C in \mathbf{NP}^C for polynomial-time filters f (in particular, for protected bridge rules), which collapses to C if C is closed under conjunction and projection; thus for all considered notions, the existence problem is in this case C -complete.

$\mathbf{D}(\Sigma_1^{\mathbf{P}}) = \mathbf{D}_1^{\mathbf{P}}$. A prototypical problem that is complete for $\mathbf{D}_1^{\mathbf{P}}$ is deciding, given a pair (F_1, F_2) of propositional Boolean formulas, whether F_1 is satisfiable and F_2 is unsatisfiable.

Since MCS are composed of contexts where each context is a KR formalism, the complexity of deciding whether an MCS is consistent clearly depends on the complexity of the KR formalisms employed in its contexts. This intuition is captured by the notion of *context complexity*, which measures deciding whether a set of beliefs is acceptable under a given knowledge-base of a context and a given set of formulas added via bridge rules.

Let $OUT_i = \{p \mid (i:p) \in body(r) \text{ for some } r \in br(M)\}$ denote the set of beliefs of context C_i which occur in the body of some bridge rule of the MCS. Context complexity is defined wrt. output-projected beliefs, i.e., belief sets projected to output beliefs (for details see Eiter et al., 2014), formally:

Definition 19 (cf. Eiter et al., 2014). *Given a context $C_i = (kb_i, br_i, L_i)$ and a pair (H, T_i) , with $H \subseteq \{\varphi(r) \mid r \in br_i\}$ and $T_i \subseteq OUT_i$, the context complexity $\mathcal{CC}(C_i)$ of C_i is the computational complexity of deciding whether there exists an $S_i \in \mathbf{ACC}_i(kb_i \cup H)$ such that $S_i \cap OUT_i = T_i$.*

We view here (as in Eiter et al., 2014) the computational complexity of a problem Π technically as the set of all problems Π' that are polynomial-time many-one reducible to Π , i.e., as the problems that are not “harder” than Π ; in particular SAT has the computational complexity \mathbf{NP} .

Furthermore, the logics L_i of all contexts are considered to be given implicitly and thus the instance size of a given MCS M is $|M| = |kb_M| + |br(M)|$ where $|kb_M|$ denotes the size of the knowledge bases in M and $|br(M)|$ denotes the size of its set of bridge rules.

Given an MCS M , we say M has *upper context complexity* C , denoted $\mathcal{CC}(M) \leq C$, if $\mathcal{CC}(C_i) \subseteq C$ for every context C_i of M ; we say M has *lower context complexity* C , denoted $\mathcal{CC}(M) \geq C$, if $C \subseteq \mathcal{CC}(C_i)$ for some context C_i of M . We say that M has *context complexity* C , denoted $\mathcal{CC}(M) = C$, iff $\mathcal{CC}(M) \leq C$ and $\mathcal{CC}(M) \geq C$. That is, if $\mathcal{CC}(M) = C$ all contexts in M have complexity at most C , and some context in M has C -complete complexity with respect to polynomial-time many-one reductions, which requires that the class C has complete problems.

Restricting disjunctive ASP to the ground case admits $\Sigma_2^{\mathbf{P}}$ -complete acceptability checking (see Dantsin, Eiter, Gottlob, & Voronkov, 2001; Gottlob, 1992), hence the context complexity of a context using L_{Σ}^{asp} is $\Sigma_2^{\mathbf{P}}$ -complete given that all kb -elements are ground; in the (function-free) non-ground case, the context complexity is $\mathbf{NEXPTIME}^{\mathbf{NP}}$ (Eiter, Gottlob, & Mannila, 1997). Acceptability checking of a context using L_{Σ}^{pl} amounts to entailment checking for all literals present in the belief set and non-entailment checking for all literals absent in the belief set, i.e., it amounts to an UNSAT and an independent SAT check, hence the context complexity is $\mathbf{D}^{\mathbf{P}}$.

Example 16. *The MCS $M = (C_1, C_2, C_3)$ of Example 3 is such that $\mathcal{CC}(C_1) = \mathbf{NP}$ and $\mathcal{CC}(C_2) = \mathcal{CC}(C_3) = \Sigma_2^{\mathbf{P}}$. As $\mathbf{NP} \subseteq \Sigma_2^{\mathbf{P}}$, it holds that $\mathcal{CC}(M) \leq \Sigma_2^{\mathbf{P}}$, and as C_2 is $\Sigma_2^{\mathbf{P}}$ -complete, we obtain $\mathcal{CC}(M) \geq \Sigma_2^{\mathbf{P}}$; hence $\mathcal{CC}(M) = \Sigma_2^{\mathbf{P}}$.*

The problem of deciding whether for a given MCS M and a pair (D_1, D_2) of bridge rules, it holds that (D_1, D_2) is a minimal diagnosis, i.e., deciding whether $(D_1, D_2) \in D_m^{\pm}(M)$, is denoted by $\mathbf{MCS}_{\mathbf{D}_m}$. As shown in [Prop. 11 by Eiter et al., 2014] if $\mathcal{CC}(M) = \mathbf{P}$, then $\mathbf{MCS}_{\mathbf{D}_m}$ is $\mathbf{D}_1^{\mathbf{P}}$ -complete; if $\mathcal{CC}(M) = C$ and C is a class with complete problems and closed under conjunction and projection, then the problem of $\mathbf{MCS}_{\mathbf{D}_m}$ is $\mathbf{D}(C)$ -complete. Intuitively, a class C is *closed under conjunction*, if all its decision problems are such that checking multiple instances of the problem at the same time is a problem in C . For example, checking whether a propositional formula F is

Context complexity $\mathcal{CC}(M)$	Deciding $(D_1, D_2) \stackrel{?}{\in}$			
	$D_m^\pm(M)$	$D_m^\pm(M, br_P)$	$D^\pm(M, br_P, br_H)$	$D_{m, t_{max}}^\pm(M, br_P, br_H)$
	MCS D_m	MCS DP_m	MCS DPH	MCS $DPH_{m, t_{max}}$
P	D_1^P -complete	D_1^P -complete	D_1^P -complete	Π_2^P -complete
NP	D_1^P -complete	D_1^P -complete	D_1^P -complete	Π_2^P -complete
$\Sigma_i^P, i \geq 1$	D_i^P -complete	D_i^P -complete	D_i^P -complete	Π_{i+1}^P -complete
PSPACE	PSPACE -complete			
EXPTIME	EXPTIME -complete			
Shown by	Eiter et al. (2014)	Theorem 5	Theorem 6	Theorems 7 + 8

Table 1: Complexity results of deciding whether a candidate diagnosis is subset-minimal, additionally protected, prioritized-minimal, or an mpm-diagnosis. Problem MCS D_{MPREF} has the same complexity as MCS $DPH_{m, t_{max}}$ if deciding $D \prec D'$ is in $\mathcal{CC}(M)$.

satisfiable is in **NP**; given two independent formulas F and G , checking whether both are satisfiable also is in **NP** since it amounts to checking whether $F \wedge G$ is satisfiable. A class C is *closed under projection*, if intuitively for every problem in C , the decision problem on projected instances (similar as for output-projected equilibria) is contained in C . For example, given a formula F in propositional logic over variables $var(F)$, finding an assignment V_A over (projected) variables $A \subsetneq var(F)$ such that (i) there exists an assignment $V_{\bar{A}}$ to the variables $\bar{A} = var(F) \setminus A$ and (ii) $V_A \cup V_{\bar{A}} \models F$, is as hard as finding an (overall) assignment V over $var(F)$ such that $V \models F$. For further details we refer to the work of Eiter et al. (2014). Specifically, for $\mathcal{CC}(M) = \Sigma_i^P$ it holds that MCS D_m is in D_i^P . Furthermore, since D_i^P is closed under conjunction and projection, it holds that MCS D_m is D_i^P -complete if at least one context in M is complete for Σ_i^P .

6.2 Overview of Results

We now investigate the complexity of our enhanced notions of diagnosis. More specifically, we study the complexity of the following decision problems, given an MCS M , a candidate diagnosis $D \in 2^{br(M)} \times 2^{br(M)}$, and depending on the problem additionally given protected bridge rules $br_P \subseteq br(M)$, prioritized bridge rules $br_H \subseteq br(M)$, and $t_{max} \in br(M)$:

- MCS DP_m : deciding whether D is a subset-minimal diagnosis with protected bridge rules, i.e., deciding whether $D \in D_m^\pm(M, br_P)$ holds.
- MCS DPH : deciding whether D is a prioritized-minimal diagnosis, i.e., deciding whether $D \in D^\pm(M, br_P, br_H)$ holds.
- MCS $DPH_{m, t_{max}}$: deciding whether D is an mpm-diagnosis (a subset-minimal prioritized-minimal diagnosis wrt. t_{max}), i.e., deciding whether $D \in D_{m, t_{max}}^\pm(M, br_P, br_H)$ holds.
- MCS D_{MPREF} : given an arbitrary preference order \preceq deciding whether $D \in D_{m, \preceq}^\pm(M)$ holds.

We show that MCSDP_m is not harder than MCSD_m , i.e., deciding whether a candidate diagnosis D is a subset-minimal diagnosis with protected bridge rules is not harder than deciding whether D is a subset-minimal diagnosis (Thm. 5). We also demonstrate that the same is true for prioritized-minimal diagnoses, i.e., MCSDPH is as hard as MCSD_m (Thm. 6). This notion of diagnosis can be applied to the plain encoding $M^{pl\preceq}$ for total preference orders to select minimal \preceq -preferred diagnoses according to a total preference order \preceq . The drawback of this approach, however, is the exponential number of bridge rules in $M^{pl\preceq}$.

Since the clone encoding M^\preceq incurs no exponential blow-up of bridge rules, it is reasonable to expect that the computational complexity of $\text{MCSDPH}_{m,t_{max}}$ is higher than the one of MCSD_m . Indeed, for context complexity $\mathcal{CC}(M)$ in Σ_1^P we prove that $\text{MCSDPH}_{m,t_{max}}$ is in Π_{i+1}^P while MCSD_m is in D_1^P (Thm. 7). Specifically, for $\mathcal{CC}(M)$ in NP the complexity of $\text{MCSDPH}_{m,t_{max}}$ is Π_2^P while MCSD_m is in D_1^P .

Since deciding $t(D) \in D_{m,t_{max}}^\pm(M^\preceq, br_P, br_H)$ only serves to decide $D \in D_{m,\preceq}^\pm(M)$, we also investigate the lower bound for the latter problem, i.e., $\text{MCSD}_{\text{MPREF}}$. We prove that it is Π_2^P -hard (Thm. 8) if $\mathcal{CC}(M)$ is in P ; hence we obtain that the clone encoding using M^\preceq and $D_{m,t_{max}}^\pm(M^\preceq, br_P, br_H)$ is in fact worst-case optimal. Furthermore, we also show that $\text{MCSD}_{\text{MPREF}}$ is hard for Π_{i+1}^P if $\mathcal{CC}(M)$ is hard for Σ_i^P .

Table 1 summarizes the results for the introduced notions of diagnosis and for context complexity being in one of several complexity classes. Note that the results for **PSPACE** and **EXPTIME** in the last column follow from the fact that $\text{coNP}^{\text{PSPACE}} = \text{PSPACE}$ and $\text{coNP}^{\text{EXPTIME}} = \text{EXPTIME}$ for membership while hardness can be shown using a trivial MCS where the acceptability function of some context is hard for **PSPACE** resp. **EXPTIME**. Our results are derived using several reductions and a genuine algorithm, which are presented in the remainder of this section; proofs can be found in the appendix.

6.3 Derivation of Results

For the problem of recognizing minimal diagnoses with protected bridge rules we have the following result.


Theorem 5. *MCSDP_m is equivalent to MCSD_m under polynomial-time reductions.*

Indeed, MCSDP_m is polynomially reducible to MCSD_m , by simply checking first whether the candidate diagnosis contains protected bridge rules and then solve MCSD_m to check whether it is subset-minimal. A formal reduction is provided in the proof of Theorem 5. Conversely, every instance of MCSD_m is an instance of MCSDP_m with $br_P = \emptyset$, and thus MCSD_m trivially reduces to MCSDP_m in polynomial time.

Next we consider the problem MCSDPH . We will show that this problem has the same complexity as MCSDP_m . To this end we first present a polynomial-time many-one reduction DPH2DP_m from MCSDPH to MCSDP_m . We remark that a direct membership proof would be simpler, but the reduction is of interest in its own right.

6.3.1 UNDERLYING IDEA OF DPH2DP_m

Given an MCS M with protected bridge rules br_P and prioritized bridge rules br_H , we simulate the modifications of regular bridge rules inside the resulting MCS. The set R_{reg} of regular (non-prioritized, non-protected) bridge rules is $R_{reg} = br(M) \setminus br_H \setminus br_P$ and their modifications can



$$\begin{array}{l}
 r_1 : \\
 r_2 :
 \end{array}
 \quad
 \begin{array}{l}
 (2 : b) \leftarrow (1 : a). \\
 (3 : d) \leftarrow (2 : c).
 \end{array}$$

$$\begin{array}{l}
 kb_{n+1} = \{ \text{removed}_{r_1} \leftarrow \text{not not_removed}_{r_1}. \quad \text{removed}_{r_2} \leftarrow \text{not not_removed}_{r_2}. \} \\
 kb_{n+2} = \{ \text{not_removed}_{r_1} \leftarrow \text{not removed}_{r_1}. \quad \text{removed}_{r_1} \leftarrow \text{not not_removed}_{r_1}. \} \\
 \quad \quad \quad \{ \text{uncond}_{r_1} \leftarrow \text{not not_uncond}_{r_1}. \quad \text{not_uncond}_{r_1} \leftarrow \text{not uncond}_{r_1}. \}
 \end{array}$$

Figure 6: The reduction from MCSDPH to MCSDP_m exemplified on an MCS $M = (C_1, C_2, C_3)$ with two bridge rules $br(M) = \{r_1, r_2\}$. Shown are the MCS $M' \otimes M$ employed in the reduction DPH2DP_m (components are indicated in gray), bridge rules of M , and the knowledge bases of C_{n+1} and C_{n+2} .

be simulated by using a meta-reasoning transformation $M^{mr(\theta, \mathcal{K}_p)} = (C_1, \dots, C_{n+1})$, where the bridge rules of C_{n+1} correspond to modifications of bridge rules in R_{reg} . They take their values from an additional context C_{n+2} that generates all possible modifications, i.e., every possible modification corresponds to an acceptable belief set of C_{n+2} . We protect in the resulting MCS $M' = (C_1, \dots, C_{n+2})$ all bridge rules except those that correspond to modifications of bridge rules in br_H , i.e., every diagnosis of M' corresponds to one (or more) diagnoses of M , but the diagnoses of M' only contain bridge rules corresponding to subsets of br_H . Consequently, any minimal diagnosis of M' is \subseteq_{br_H} -minimal wrt. M . To ensure that the diagnosis indeed is \subseteq -minimal, we further add a copy of M , i.e., the resulting MCS is $M' \otimes M$ where M' ensures minimality wrt. \subseteq_{br_H} and M ensures minimality wrt. \subseteq . An illustration of the resulting MCS is given in Figure 6 for a concrete MCS that is considered later in Example 17 below in detail.

6.3.2 FORMAL DETAILS OF DPH2DP_m

Given an MCS M and a set $R_{reg} \subseteq br(M)$, let $\mathcal{K}_p = \emptyset$ and let θ be such that for all $D_1, D_2 \subseteq br(M)$ the property $\theta(D_1, D_2, \emptyset)$ holds. We craft an MCS based on the meta-reasoning MCS $M^{mr(\theta, \mathcal{K}_p)} = (C_1, \dots, C_n, C_{n+1})$ to obtain an MCS where the modification of all bridge rules in R_{reg} is hidden in the set of possible belief states. To this end, we introduce another context C_{n+2} without bridge rules whose acceptable belief sets encode all respective modifications of bridge rules of R_{reg} . Formally, $C_{n+2} = (L_{\Sigma}^{asp}, kb_{n+2}, \emptyset)$ where

$$kb_{n+2} = \left\{ \begin{array}{ll} not_removed_r \leftarrow not_removed_r. & removed_r \leftarrow not_not_removed_r. \\ uncond_r \leftarrow not_not_uncond_r. & not_uncond_r \leftarrow not_uncond_r. \end{array} \middle| r \in R_{reg} \right\}.$$

Observe that for every $D_1, D_2 \subseteq R_{reg}$, there is a belief set S_{n+2} with

$$S_{n+2} \cap (\{not_removed_r, uncond_r \mid r \in R_{reg}\}) = \{not_removed_r \mid r \in R_{reg} \setminus D_1\} \\ \cup \{uncond_r \mid r \in D_2\}.$$

In addition to that, since C_{n+2} has no bridge rules, it follows that $S_{n+2} \in \mathbf{ACC}_{n+2}(kb_{n+2} \cup app(br_{n+2}, S'))$ holds for all belief states $S' = (S'_1, \dots, S'_{n+2})$ where $S'_{n+2} = S_{n+2}$.

Recall that all bridge rules of C_{n+1} are either of the form $(n+1 : not_removed_r) \leftarrow \top$. or $(n+1 : uncond_r) \leftarrow \perp$. where $r \in br(M)$. Let $C_{n+1} = (L, kb_{n+1}, br_{n+1})$; then $C'_{n+1} = (L, kb_{n+1}, br'_{n+1})$ where

$$br'_{n+1} = \{(n+1 : not_removed_r) \leftarrow (n+2 : not_removed_r). \mid r \in br(M), r \in R_{reg}\} \quad (11)$$

$$\cup \{(n+1 : uncond_r) \leftarrow (n+2 : uncond_r). \mid r \in br(M), r \in R_{reg}\} \quad (12)$$

$$\cup \{(n+1 : not_removed_r) \leftarrow \top. \mid r \in br(M), r \notin R_{reg}\} \quad (13)$$

$$\cup \{(n+1 : uncond_r) \leftarrow \perp. \mid r \in br(M), r \notin R_{reg}\}. \quad (14)$$

Intuitively, C'_{n+1} equals C_{n+1} but the bridge rules occurring in R_{reg} refer to C_{n+2} . Similar to the meta-reasoning encoding, we denote by $d1(r)$ and $d2(r)$ the corresponding bridge rule of the form in (13) and in (14), respectively. We extend these notions to sets of bridge rules and let $di(R) = \{di(r) \mid r \in R\}$ for any $R \subseteq br(M)$ and $i = 1, 2$. For example, $d1(br(M) \setminus R_1)$ denotes all bridge rules of line (13).

Finally, we call $M' = (C_1, \dots, C_n, C'_{n+1}, C_{n+2})$ the *meta-guessing MCS for M and R_{reg}* . The effect of the redirection to C_{n+2} is that the acceptable belief sets of C_{n+2} guess all possible modifications. The rest of M' behaves like an ordinary meta-reasoning encoding, where protected bridge rules of M' are $br_{P'} = br_{M'} \setminus (d1(br(M) \setminus R_{reg}) \cup d2(br(M) \setminus R_{reg}))$, i.e., all bridge rules are protected except those in C_{n+1} that do not correspond to bridge rules in R_{reg} .

Now the reduction DPH2DP_m from MCSDPH to MCSDP_m is as follows:

$$(M, (D_1, D_2), br_P, br_H) \mapsto (M' \otimes M, (D'_1, D'_2), br_{P''})$$

where M' is the meta-guessing MCS wrt. $R_{reg} = br(M) \setminus br_P \setminus br_H$ and $br_{P''} = br_{P'} \cup I(br_P)$ where I is the mapping wrt. $M' \otimes M$ and $br_{P'}$ is the set of protected bridge rules of the meta-guessing MCS M' ; furthermore $D'_1 = I(D_1) \cup d1(D_1 \cap br_H)$ and $D'_2 = I(D_2) \cup d2(D_2 \cap br_H)$, i.e., (D'_1, D'_2) contains a candidate diagnosis of M and a candidate diagnosis over br_H with modifications to the remaining bridge rules of M being simulated by M' .

Observe that the size of $(M' \otimes M, (D'_1, D'_2), br_{P''})$ is polynomial in the size of $(M, (D_1, D_2), br_P, br_H)$, because $M' \otimes M$ only has four times as many bridge rules as M and all other sets are subsets of these bridge rules. Furthermore, $(M' \otimes M, (D'_1, D'_2), br_{P''})$ can be computed in polynomial time in the size of $(M, (D_1, D_2), br_P, br_H)$; more precisely, even in linear time.

Example 17. We illustrate the reduction $DPH2DP_m$ and the MCS resulting from on a simple MCS $M = (C_1, C_2, C_3)$ with three contexts and two bridge rules $br(M) = \{r_1, r_2\}$ as follows:

$$\begin{array}{ll} r_1 : & (2 : b) \leftarrow (1 : a). \\ r_2 : & (3 : d) \leftarrow (2 : c). \end{array}$$

No bridge rule is protected, i.e., $br_P = \emptyset$, and r_2 is prioritized, i.e., $br_H = \{r_2\}$, thus $R_{reg} = \{r_1\}$. Figure 6 illustrates M (on the right) and the meta-guessing MCS M' for M and R_{reg} (on the left); their combination $M' \otimes M$ (the overall Figure 6), is the MCS constructed in the reduction $DPH2DP_m$. It also shows a possible realization of the contexts C_{n+1} and C_{n+2} using ASP. Since C_{n+1} stems from the meta-reasoning encoding where θ holds for all potential diagnoses, the two rules in kb_{n+1} are sufficient to exhibit the observed modifications of r_1 and r_2 as beliefs. Intuitively, the rules in kb_{n+2} guess all potential modifications of r_1 and exhibit them to C_{n+1} as beliefs.

The following lemma shows that $DPH2DP_m$ indeed is a correct reduction from $MCSDPH$ to $MCSDP_m$.

Lemma 4. $DPH2DP_m$ is a polynomial-time reduction from $MCSDPH$ to $MCSDP_m$.

On the other hand, one can easily reduce $MCSDP_m$ to $MCSDPH$. We thus obtain that $MCSDPH$ indeed has the same complexity as deciding $D \in D_m^\pm(M, br_P)$ and hence whether $D \in D_m^\pm(M)$ holds.

Theorem 6. $MCSDPH$ is equivalent to $MCSDP_m$ under polynomial-time reductions.

6.3.3 FURTHER COMPLEXITY RESULTS

A stepping stone for analyzing $MCSDPH_{m,t_{max}}$ is the decision problem $MCSDPH_{t_{max}}$, which we consider next. $MCSDPH_{t_{max}}$ is defined as follows: given an MCS M , a candidate diagnosis $D \in 2^{br(M)} \times 2^{br(M)}$ with $D = (D_1, D_2)$, protected bridge rules $br_P \subseteq br(M)$, prioritized bridge rules $br_H \subseteq br(M)$, and $t_{max} \in br(M)$; decide whether (i) $t_{max} \in D_2$ and (ii) for all $T \in D_m^\pm(M, br_P)$ it holds that $T \subseteq_{br_H} D \Rightarrow T =_{br_H} D$. Notice that $MCSDPH_{t_{max}}$ basically amounts to checking the presence of t_{max} in a candidate diagnosis of $MCSDPH$. As the following lemma shows, former is not harder than the latter.

Lemma 5. $MCSDPH_{t_{max}}$ is polynomial-time reducible to $MCSDPH$ and thus in the complexity class C , if $MCSDPH$ is in C and C is closed under polynomial reductions.

Note that all classes in Section 6.1 above are closed under polynomial-time reductions.

We use an $MCSDPH_{t_{max}}$ -oracle in Alg. 1 to obtain membership results of $MCSDPH_{m,t_{max}}$.

Theorem 7. If $MCSDPH$ is in C , then $MCSDPH_{m,t_{max}}$ is in coNP^C .

Algorithm 1: Deciding whether $(D_1, D_2) \notin D_{m,t_{max}}^\pm(M, br_P, br_H)$ holds.

Input : MCS M , (D_1, D_2) , br_P , and br_H with $D_1, D_2 \subseteq br(M)$, $br_P, br_H \subseteq br(M)$.

Output: YES if $(D_1, D_2) \notin D_{m,t_{max}}^\pm(M, br_P, br_H)$

- 1 **if** $\text{oracle}_{\text{MCSDPH}_{t_{max}}}((D_1, D_2), M, br_P, br_H) = \text{NO}$ **then output YES**
 - 2 **guess** $T_1, T_2 \subseteq br(M)$
 - 3 **if** $\text{oracle}_{\text{MCSDPH}_{t_{max}}}((T_1, T_2), M, br_P, br_H) = \text{YES} \wedge (T_1, T_2) \neq (D_1, D_2) \wedge (T_1, T_2) \subseteq_{br(M) \setminus br_H} (D_1, D_2) \wedge (T_1, T_2) \not\subseteq_{br(M) \setminus br_H} (D_1, D_2)$ **then output YES**
-

Proof. Algorithm 1 decides whether $(D_1, D_2) \notin D_{m,t_{max}}^\pm(M, br_P, br_H)$ holds using an oracle for $\text{MCSDPH}_{t_{max}}$. Intuitively, (D_1, D_2) is not an mpm-diagnosis if it either is no subset-minimal prioritized-minimal containing t_{max} , which is checked in the first line using the oracle, or if there exists a subset-minimal prioritized-minimal diagnosis $(T_1, T_2) \subset (D_1, D_2)$ that also contains t_{max} . In the second line such a (T_1, T_2) is guessed and in the third line it is verified that the guessed candidate indeed has the above properties. Checking whether $(D_1, D_2) \in D_{m,t_{max}}^\pm(M, br_P, br_H)$ holds is possible by Algorithm 1 and negating its output.

By assumption MCSDPH is in \mathbf{C} , thus by Lemma 5 it holds that $\text{MCSDPH}_{t_{max}}$ is in \mathbf{C} , i.e., the complexity of the oracle in Algorithm 1 is in \mathbf{C} . Since Algorithm 1 uses a polynomial-size guess for (T_1, T_2) its complexity clearly is $\mathbf{NP}^{\mathbf{C}}$. Consequently, deciding whether $(D_1, D_2) \in D_{m,t_{max}}^\pm(M, br_P, br_H)$ holds is in $\mathbf{coNP}^{\mathbf{C}}$. \square

The previous decision problems arise from our approach to realize the selection of preferred and filtered diagnoses of an MCS. To give a full picture, we also investigate the complexity of the basic problem, i.e., of $\text{MCSD}_{\text{MPREF}}$.

As the following theorem shows, $\text{MCSD}_{\text{MPREF}}$ itself is $\Pi_2^{\mathbf{P}}$ -hard even if both the context complexity and deciding whether $D \preceq D'$ holds are tractable. This result also shows that our approach of realizing the selection of minimal \preceq -preferred diagnoses is worst-case optimal.

Theorem 8. *If $\mathcal{CC}(M)$ is hard for $\Sigma_i^{\mathbf{P}}$ ($\Pi_i^{\mathbf{P}}$) then $\text{MCSD}_{\text{MPREF}}$ is hard for $\Pi_{i+1}^{\mathbf{P}}$ ($\Pi_{i+2}^{\mathbf{P}}$) with $i \geq 0$. Moreover, $\text{MCSD}_{\text{MPREF}}$ is $\Pi_2^{\mathbf{P}}$ -hard even if both $\mathcal{CC}(M)$ and deciding $D' \preceq D''$ are in \mathbf{P} .*

For establishing completeness of $\text{MCSD}_{\text{MPREF}}$, we use the clone encoding of the previous section as a polynomial-time reduction to $\text{MCSDPH}_{m,t_{max}}$.

Corollary 1. *Let M be an MCS with $\mathcal{CC}(M) = \Sigma_i^{\mathbf{P}}$, $i \geq 0$ (resp., $\mathcal{CC}(M) = \mathbf{PSPACE}$, $\mathcal{CC}(M) = \mathbf{EXPTIME}$), and a preference order \preceq such that deciding $D \prec D'$ is in $\Sigma_i^{\mathbf{P}}$ (resp., \mathbf{PSPACE} , $\mathbf{EXPTIME}$). Then $\text{MCSD}_{\text{MPREF}}$ is complete for $\Pi_{i+1}^{\mathbf{P}}$ (resp., \mathbf{PSPACE} , $\mathbf{EXPTIME}$). In particular, $\text{MCSD}_{\text{MPREF}}$ is $\Pi_{i+1}^{\mathbf{P}}$ -complete if deciding $D \preceq D'$ is in \mathbf{P} and $\mathcal{CC}(M) = \Sigma_i^{\mathbf{P}}$, $i \geq 0$.*

Examples of preference orders as hard as \mathbf{PSPACE} are CP-nets in general while restricted variants are in \mathbf{NP} or even \mathbf{P} (cf. Section 3.2.1).

We can also use the clone encoding to show the completeness of $\text{MCSDPH}_{m,t_{max}}$.

Corollary 2. *$\text{MCSDPH}_{m,t_{max}}$ is $\Pi_{i+1}^{\mathbf{P}}$ -complete if $\mathcal{CC}(M) = \Sigma_i^{\mathbf{P}}$, $i \geq 1$, and $\Pi_2^{\mathbf{P}}$ -complete if $\mathcal{CC}(M) = \mathbf{P}$ or $\mathcal{CC}(M) = \mathbf{NP}$.*

The hardness result of Π_{i+2}^P for $\text{MCS}_{\text{D}_{\text{MPREF}}}$ with $\mathcal{CC}(M) = \Pi_i^P$ might seem to contradict Corollary 2, which shows, using the clone encoding, that $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is in Π_{i+1}^P for $\mathcal{CC}(M) = \Sigma_i^P$. However this is no contradiction since the basic problem of recognizing minimal diagnoses, i.e., MCS_{D_m} , is not known to be in Σ_i^P for $\mathcal{CC}(M) = \Pi_i^P$. In the work of Eiter et al. (2014) it is shown that MCS_{D_m} is in $\mathbf{D}(\mathbf{C})$ if \mathbf{C} is closed under conjunction and projection, which presumably is not the case for Π_i^P , $i \geq 0$ (while it is for Σ_i^P). Hence for $\mathcal{CC}(M) = \Pi_i^P$, MCS_{D_m} is not in $\mathbf{D}(\Pi_i^P)$, thus $\text{MCS}_{\text{DPH}_{t_{\max}}}$ is presumably not in Π_{i+1}^P . On the other hand, Π_i^P is in Σ_{i+1}^P , consequently MCS_{D_m} is in $\mathbf{D}(\Sigma_{i+1}^P)$ and $\text{MCS}_{\text{DPH}_{t_{\max}}}$ in Π_{i+2}^P .

7. Discussion and Related Work

In this section we first discuss options for decomposing the central observation context of the meta-reasoning transformation and then consider related work.

7.1 Decomposing the Central Observation Context

A key strength of MCS is the capability of integrating different knowledge bases in a decentralized manner. Accordingly, scenarios for MCS where a centralized specification of preferences on diagnoses may be unwanted, e.g., if different companies agree to share data, their preferences might expose some information they are actually not willing to share. The approaches presented here use a central observation context that knows all bridge rules, or more specifically, that knows the labels of all bridge rules, and for each of them whether and how it is modified. The observer, however, does neither know the structure, i.e., the contents of the head and body, of the bridge rules, nor the actual status of the information exchange, and it cannot see any beliefs of any context. Hence, our approach supports almost full information hiding and only decentralization is lost.

Criteria for decomposing a context have been investigated by Weinzierl (2014). The results there can be applied to the meta-reasoning transformation that we described in Section 5.1 in order to decompose the observation context of the filter encoding M^f . If the underlying filter can be broken up, the central observation context thus may be replaced with several contexts, each covering only a partition of the bridge rules in $br(M)$. If there is a partition $br(M) = A \cup B$ (where A, B are disjoint and nonempty) such that a given filter f satisfies that for all $D_1, D_2 \subseteq br(M)$ it holds that $f(D_1, D_2) = 1$ iff $f(D_1 \cap A, D_2 \cap A) = 1$ and $f(D_1 \cap B, D_2 \cap B) = 1$, then the observation context of M^f is decomposable. Informally, f is such that the modifications of bridge rules in A can be checked independently from those in B and vice versa.

Notice that for any “reasonable” logical formalism which realizes f , checking whether $f(D_1 \cap A, D_2 \cap A) = 1$ resp. $f(D_1 \cap B, D_2 \cap B) = 1$ can be realized by two (independent) knowledge bases; the latter are the decomposition of the observation context. Depending on f , this decomposition may be repeated several times, where each time one context is decomposed into two independent contexts until the observation of diagnoses is fully decentralized.

Example 18. Consider the MCS $M^f = (C_1, C_2, C_3)$ of Example 11 realizing the filter f on the MCS M whose bridge rules are $br(M) = \{r_1, r_2, r_3\}$. Recall that f is defined by:

$$f(D_1, D_2) = \begin{cases} 0 & \text{if } r_3 \in D_1, r_2 \notin D_1 \text{ or } r_3 \notin D_1, r_2 \in D_1, \\ 0 & \text{if } r_3 \in D_2, r_2 \notin D_2 \text{ or } r_3 \notin D_2, r_2 \in D_2, \\ 1 & \text{otherwise.} \end{cases}$$

Obviously, $br(M)$ can be partitioned into $A = \{r_2, r_3\}$ and $B = \{r_1\}$, because for all $D_1, D_2 \subseteq br(M)$ holds that $f(D_1 \cap B, D_2 \cap B) = 1$ and $f(D_1 \cap A, D_2 \cap A) = f(D_1, D_2)$.

The resulting decomposition of bridge rules in C_3 is: $br_3^A = \{d1(r_2), d2(r_2), d1(r_3), d2(r_3)\}$ and $br_3^B = \{d1(r_1), d2(r_1)\}$. Since the knowledge base kb_3 of M^f uses ASP, we can easily get the knowledge bases kb_3^A and kb_3^B by partitioning kb_3 :

$$kb_3^A = \left\{ \begin{array}{ll} removed_{r_2} \leftarrow not\ not_removed_{r_2} & removed_{r_3} \leftarrow not\ not_removed_{r_3} \\ \perp \leftarrow removed_{r_3}, not\ removed_{r_2} & \perp \leftarrow not\ removed_{r_3}, removed_{r_2} \\ \perp \leftarrow uncond_{r_3}, not\ uncond_{r_2} & \perp \leftarrow not\ uncond_{r_3}, uncond_{r_2} \end{array} \right\}$$

$$kb_3^B = \{removed_{r_1} \leftarrow not\ not_removed_{r_1}\}$$

The resulting decomposed MCS is $M' = (C'_1, C'_2, C_3^A, C_3^B)$, where all bridge rules from C_3 either belong to C_3^A or C_3^B and all beliefs of C_3 that are referred to in other bridge rules of M^f either refer to C_3^A or C_3^B in M' . The diagnoses of M' correspond one-to-one to those of M^f . As diagnoses with protected bridge rules are directly based on ordinary diagnoses, these results thus extend to diagnoses with protected bridge rules. The MCS M' can be used to obtain minimal filtered diagnoses of M , where the filter itself is realized in a decentralized way.

In principle, decomposition may also be applied to the clone encoding M^{\preceq} , but the bridge rule t_{max} disallows a simple decomposition. Nonetheless, it seems possible to achieve decomposition using additional protected bridge rules for information exchange between the decomposed contexts; a formal result, however, remains to be established.

7.2 Related Work

We now briefly consider further, previous work and then discuss related work on multi-context systems and other KR formalisms.

7.2.1 PREVIOUS WORK

In the work of Weinzierl (2014) several of the notions presented here have been investigated and exemplified in more depth; among others, formal results on the decomposition method introduced in Section 7.1 above. While in this article fully-compatible CP-nets are shown as sample formalism to specify preferences over diagnoses, there are other possibilities of using CP-nets to compare diagnoses, e.g. by assigning each bridge rule r only one variable V^r and a domain of $\{unmodified, removed, condition-free\}$. This kind of CP-nets, however, cannot represent a candidate diagnosis (D_1, D_2) with $r \in D_1 \cap D_2$, i.e., where a bridge rule is both removed and condition-free, while fully compatible CP-nets can represent all candidate diagnoses.

There is another sample instantiation of preference orders that is based on units of modified bridge rules. The idea is that bridge rules are grouped according to the information they carry; e.g. in Example 4 there are two units: health-related information (r_1 and r_2) and billing-related information (r_3, r_4 , and r_5). A preference relation over diagnoses is established by considering how many such units are modified (hence potentially broken); preferred are those diagnoses modifying the least number of units.

Two transformations for meta-reasoning on diagnoses in an MCS M were developed, where the first merely adds bridge rules and contexts to observe the information exchange between contexts of M . The disadvantage of this transformation is that there are concrete MCS for which the observation

cannot identify each diagnosis correctly. The second transformation, which is the one presented in this article, is more general and allows for the correct identification of diagnoses for all MCS; however, this comes at the price of rewriting the bridge rules.

Below we discuss two closely related approaches that rely on preference to ensure consistency of MCS and a third approach that integrates preferences into MCS contexts directly. We also sketch how our approach can be applied to further extensions of the MCS framework and we relate our approach to preference-based inconsistency management in other KR formalisms. An extended discussion is available in the work of Weinzierl (2014).

7.2.2 PREFERENTIAL MCS

In the work of Mu, Wang, and Wen (2015) an approach at preference-based inconsistency management in MCS is introduced: *Preferential Multi-Context Systems (PMCS)* are similar to ordinary MCS where an additional preference order \leq_s restricts the information flow. The relation \leq_s is a total preorder on a partitioning of the contexts of M , i.e., \leq_s compares sets of contexts and all contexts in the same set are treated as equally preferred. The information flow then is restricted from more preferred to less-or-equally preferred contexts, i.e., a PMCS is stratified. Note that this total preorder differs from our notion of a total preference, since we consider preference over candidate diagnoses, not over sets of contexts.

Based on the ordering, one may ask for a *maximal consistent section*, which is the maximal initial segment of the ordering of preferred contexts that still admit an equilibrium. Furthermore, the notion of a *c-diagnosis* is introduced, which is a diagnosis that does not modify bridge rules of the maximal consistent section. Note that Mu et al. (2015) only consider diagnoses that remove bridge rules, i.e., diagnoses of the form (D_1, \emptyset) .

In the same work, it is noted that a filter f on diagnoses may be used to select c-diagnoses, by simply filtering out all diagnoses that modify bridge rules of the maximal consistent section. This however, requires to know the maximal consistent section in advance. Intuitively, c-diagnoses can be fully captured by preference orders as follows. We recall the notation of an *i-cut* for PMCS first: given a PMCS M with total preorder \leq_s on sets of contexts of M , the *i-cut* of M , denoted by $M(i)$ contains all contexts that are in the *i*-th and lower stratum according to \leq_s . For example, $M(1)$ contains the most preferred contexts, $M(2)$ contains contexts of $M(1)$ and all that are less preferred than the ones in $M(1)$ but more preferred than any other contexts, and so on. Notice that $M(2) \supseteq M(1)$ holds, i.e., $M(i)$ contains all contexts of $M(j)$ for $j \leq i$. Now a preference order \preceq is defined on candidate diagnoses as follows: $(D_1, D_2) \preceq (D'_1, D'_2)$ iff $D_2 = \emptyset$ and for every $1 \leq i \leq m$ such that $D_1 \cap \{r \in br_\ell \mid C_\ell \in M(i)\} \neq \emptyset$, it holds that $D'_1 \cap \{r \in br_\ell \mid C_\ell \in M(i)\} \neq \emptyset$. The intuition is that \preceq prefers (D_1, D_2) over (D'_1, D'_2) if every *i-cut* $M(i)$ that is modified by the former is also modified by the latter. This effectively guarantees that the most preferred diagnoses according to \preceq only modify bridge rules from less preferred contexts. In fact, no most preferred diagnosis modifies any bridge rule of the maximal consistent section, because such a diagnosis is always preferred. Thus, the set of most preferred diagnoses according to \preceq should coincide with the set of c-diagnoses. Clarifying this and a more extensive comparison to PMCS remains for future work.

Another approach for handling inconsistency in MCS based on a preference order was presented by Caire, Bikakis, and Traon (2013). The notion of *conviviality* stemming from multi-agent systems is used to model and measure information dependencies in MCS: intuitively, conviviality measures

how much contexts exchange information with each other. Every ordinary diagnosis D of an inconsistent MCS M is then associated with the conviviality of the resulting MCS where D is applied, i.e., the conviviality of $M[D]$. A diagnosis is regarded as being optimal, if its associated conviviality is maximal. Suppose that $Conv(M)$ is the conviviality of an MCS M , we can capture the approach by selecting \preceq -preferred diagnoses where \preceq is a preference on diagnoses as follows: $D \preceq D'$ iff $Conv(M[D]) \geq Conv(M[D'])$. Subsequently, $D_{\preceq}^{\pm}(M)$ is the set of optimal diagnoses according to Caire et al. (2013).

7.2.3 MULTI-CONTEXT SYSTEMS WITH PREFERENCES

Le, Son, and Pontelli (2015) recently extended the abstract logics of MCS with ranking information to *ranked logics* and defined an equilibrium semantics for the ensuing *Multi-Context Systems with Preferences (MCSP)*. Formally, a ranked logic is a tuple $L = (\mathbf{KB}_L, \mathbf{BS}_L, \mathbf{ACC}_L, <_L)$ such that $(\mathbf{KB}_L, \mathbf{BS}_L, \mathbf{ACC}_L)$ is an abstract logic for MCS and $<_L$ is a partial order over pairs (kb, bs) of a knowledge base kb and an acceptable belief set for kb , i.e., $bs \in \mathbf{ACC}_L(kb)$. The semantics of an MCSP $M = (C_1, \dots, C_n)$ is given in terms of particular equilibria of the induced ordinary MCS $M' = (C'_1, \dots, C'_n)$, in which the ranking information is dropped. An equilibrium S of M' is *weakly preferred* (resp. *strongly preferred*) over an equilibrium S' of M' , if S is preferred over S' in some (resp. every) ranking $<_i$ of the contexts C_i , i.e., $(kb_i, S_i) <_{L_i} (kb_i, S'_i)$ holds where $C_i = (L_i, kb_i, br_i)$. Furthermore, S is *most preferred*, if there is no other equilibrium S' of M' that is weakly (resp. strongly) preferred over S ; the semantics of M is given by its most preferred equilibria.

Compared to our work, there is a noticeable difference: in MCSP preference information is directly integrated into the semantics to select among equilibria (provided some exist), while in our approach preference is used to select diagnoses for systems that do not admit an equilibrium, i.e., systems that are inconsistent. This is similar with distributed abductive diagnosis (DAD), which is briefly sketched by Le et al. (2015) as an application for MCSP; they discuss how one can use an MCSP in order to obtain an abductive explanation for a set of observations on an MCS from a given set of hypotheses. While DAD strives to explain beliefs in an equilibrium, the diagnoses that we consider address the lack on an equilibrium. Le et al. (2015) do not consider whether their approach could be used to diagnose an inconsistent MCS. As their hypotheses are elements of knowledge bases respectively belief sets but not observations of inconsistency of the system as such, this is not apparent. It remains to be clarified whether (under suitable restrictions) and how this is possible. As regards preference, in principle, one could use $<_{L_i}$ over equilibria to select those diagnoses that admit the best ranked equilibria. More formally, we can construct a preference order \preceq such that $D \preceq D'$ holds if some (resp. all) $S \in \text{EQ}(M[D])$ are weakly preferred (resp. strongly preferred) over some (resp. all) $S' \in \text{EQ}(M[D'])$.

7.2.4 DEFEASIBLE MCS

In ordinary MCS, all bridge rules that are applicable in a belief state add their head formulas to the respective contexts. Different from that, *Defeasible MCS (dMCS)* have bridge rules which only add their head formulas if no inconsistency arises, i.e., bridge rules are defeasible. By that, defeasible MCS are an important contribution to inconsistency management in MCS since these MCS are inherently consistent. They have been investigated in the works of Bikakis and Antoniou (2008, 2010), and Bikakis, Antoniou, and Hassapis (2009, 2011), which address inconsistency in a homogeneous

MCS setting. The semantics of dMCS is given by resorting to an argumentation-inspired approach (Bikakis & Antoniou, 2010). Each context is a local theory composed of strict and defeasible rules, where the conclusion of an applicable strict rule is always considered while for defeasible rules their conclusion is only considered if there is no contrary evidence. Bridge rules, or mapping rules, are (local) defeasible rules whose body literals refer to other contexts. The decision which rules to ignore is based for every context on a *strict total order* of all contexts.

The set of (mapping) rules that are ignored thus corresponds to a unique deletion-only diagnosis whose declarative description is more involved compared to our notion. Since local information is important for identifying the defeasible rules that are ignored, an encoding within our framework is possible but requires an involved MCS where contexts expose private information. One notable advantage of defeasible MCS is that for acyclic systems, only a polynomial number of computation steps is required for answering queries that are a single literal. The second component of diagnoses, i.e., rules that are forced to be applicable, however, have no counterpart in the inconsistency management approach for dMCS. Furthermore, the strict total order over contexts forces the user to make (perhaps unwanted) decisions at design time; alternative orders would require a redesign and separate evaluation. Bikakis and Antoniou (2011) also relaxed the strict total order condition to allow for partial preference orders over contexts. Four variants, based on argumentation, are introduced to resolve inconsistency and it is shown that partial preference orders come at the cost of more information exchange being rejected. Our approach avoids this and allows to respect various kinds of orderings and preferences; it is not committed to a particular formalism and in principle any formalism that can be couched into a context of an MCS can be employed.

7.2.5 FURTHER MCS EXTENSIONS

In recent years, some significant extensions of the Multi-Context Systems framework itself have been proposed. We give a short overview of these extensions and sketch how the notions of diagnosis and preference can be adapted.

Managed Multi-Context Systems (mMCS) are an extension of MCS where each context is equipped with a management component that is called a context manager (Brewka, Eiter, Fink, & Weinzierl, 2011b). This manager allows applicable bridge rules to not only add information, but to apply arbitrary operations on the knowledge base. In mMCS the heads of bridge rules are operational statements of the form $o(s)$, where s is a knowledge-base element as in ordinary MCS and o is the name of an operation to apply, e.g., $revise(\neg p)$ indicates that the knowledge base is to be revised with the formula $\neg p$. Many kinds of operations can be captured by mMCS, e.g. updating logic programs, belief revision, or database view updates. Most notably, if all context managers ensure locally that some acceptable belief set exists for the context, then inconsistency in an mMCS may only arise from cyclic information flow. Notably, mMCS can be translated faithfully to MCS, hence the diagnosis notion of MCS and the techniques for selecting most preferred diagnoses also extend to mMCS.

Reactive Multi-Context Systems (rMCS), which were introduced and investigated by Brewka (2013), Ellmauthaler (2013), and Brewka, Ellmauthaler, and Pührer (2014a, 2014b), and *evolving Multi-Context Systems* (eMCS), which were defined and studied by Goncalves, Knorr, and Leite (2014b, 2014a) are similar formalisms. Both are an extension of mMCS to allow change over time; although rMCS and eMCS have been developed independently and their formalizations differ somehow, they are, in essence, quite similar. Both adopt a discrete time ontology where at each

step a set of observations is taken into account. Observations then either influence sensor atoms in bridge rules (rMCS) or the knowledge bases of designated sensor contexts with fixed acceptability functions (eMCS). A semantics is defined that pairs at each time step the observations with an equilibrium for that step, taking into account the equilibrium of the previous step, i.e., semantics is an incremental sequence of equilibria. For eMCS this sequence is called an *evolving equilibrium* while for rMCS it is called a *run*.

We note that one may “unroll” the time steps of a given rMCS or eMCS M such that for an observation sequence of k time steps, the unrolled MCS M^u contains k copies of the ordinary contexts and bridge rules of M . Sensor atoms and sensor contexts are at each time step according to the step in the observation sequence and additional bridge rules for inertia and incremental change of beliefs and knowledge bases carry information from copy i to copy $i + 1$ in the unrolled system M^u . Then, M^u is an ordinary MCS where the notions of diagnosis, filter, and preference can be applied, hence these notions also extend to rMCS and eMCS.

Since any ordinary bridge rule $r \in br(M)$ is duplicated k times in M^u , a diagnosis can independently modify the copy of r at time i from the copy of r at time j , for $i \neq j$. Such independent modifications may be unwanted and one may consider only a diagnosis to be valid that modifies all copies uniformly, or alternatively consider only diagnoses that keeps bridge rules unmodified until some time point ℓ and for all time points $i \geq \ell$ the same modification is applied. Notice that one can easily craft a filter f for such unrolled M^u that ensures either of the above conditions. Vice versa, it is also possible to define notions of diagnosis with those properties directly for rMCS and eMCS, each yielding another notion of diagnosis. More work on this is required, but outside the scope of this article.

Tasharrofi and Ternovska (2014) introduced the notion of supported equilibrium semantics, which requires a notion of *support* throughout contexts. In principle, this notion of support enables a new notion of diagnosis that also considers modifications of knowledge bases to restore global consistency. Preferences and filter may then be defined on top of such a diagnosis notion.

An event-based approach to the semantics of MCS is given by Ellmauthaler and Pührer (2014), where so-called *asynchronous Multi-Context Systems* (aMCS) are introduced. However, the semantics of aMCS is highly operational, which makes it rather difficult to see how the declarative notion of diagnosis could be reasonably extended to this setting.

Jin, Wang, and Wen (2012) developed a framework for possibilistic reasoning in MCS termed *poss-MCS* where each context is a possibilistic logic program and information exchange is realized using possibilistic bridge rules. Intuitively, every (bridge) rule of a possibilistic context is an ordinary (bridge) rule that has an associated degree of necessity $\alpha \in [0, 1]$. Based on those necessity degrees, possibilistic (local) models and possibilistic equilibria can be defined. Although this potentially hints at using possibilistic equilibria for inconsistency resolution, it is shown that every possibilistic equilibrium is an ordinary equilibrium. Consequently, the notion of possibilistic equilibrium cannot remedy the absence of an ordinary equilibrium.

7.2.6 OTHER KR FORMALISMS

Clearly, the use of preferences to resolve inconsistency has been suggested and elaborated for rule-based systems and knowledge-exchange systems in numerous works before. We briefly mention here two, but note that they are only remotely related to MCS.

Balduccini and Gelfond (2003) extended the ASP-based language A-Prolog with *consistency-restoring* (CR) rules. Such rules are normally not applicable, but if the head restores consistency of an otherwise inconsistent ASP program, then a rule may become applicable. The semantics of CR-Prolog is given via a translation to abductive logic programs (cf. Kakas, Kowalski, & Toni, 1992) and takes a ranking over the CR rules into account. We note that consistency-restoring rules are similar in behavior to making a bridge rule condition-free. Hence, the diagnoses of an MCS under a specifically crafted preference order are capable of capturing the semantics of CR rules in certain cases.

Peer-to-peer data integration systems, e.g. by Calvanese et al. (2004), allow for dynamically changing the data integration scenario in which peers can enter or leave the system anytime. An automatic approach for reasoning with inconsistent knowledge in a peer-to-peer system was presented by Binas and McIlraith (2008) where knowledge from other peers is ranked according a preference order. A semantics is given in terms of extensions of a Dung-style abstract argumentation framework (Dung, 1995), which designates formulas that are “distributed entailed”. In principle, preference orders over diagnoses of an MCS can be used to simulate the ranking of formulas that occur in the head of bridge rules, yet this approach is limited to contexts where a notion of peer support can be defined and successfully incorporated into the preference order.

8. Conclusion

In this work, we have addressed the problem of identifying and selecting those repairs of an inconsistent Multi-Context System (MCS) which are most preferred. In general, there can be by far too many possible repairs (also called diagnoses) such that they can all be manually inspected and a “best one” selected by hand. Supporting a preference mechanism to select appealing diagnoses is therefore vital for inconsistency management in MCS. To identify and select among all diagnoses of an MCS the most preferred ones, we have considered filters, which allow to discard diagnoses that do not fulfill certain criteria, and preference orders, which allow to compare diagnoses. As MCS are a flexible framework for interlinking information from heterogeneous formalisms in different application contexts, in this spirit the user should have a choice for the formalism to specify both types of preferences.

To achieve this, we followed an internalization approach: if the required conditions or preferences can be expressed via a context of an MCS, then they can be employed for the selection of preferred diagnoses, where in principle any (abstract) context logic may be used. To this end, several techniques for meta-reasoning about diagnoses in MCS have been developed which transform a given MCS M and a filter (resp., preference order) into an MCS M' such that the diagnoses of M' correspond one-to-one to the filtered (resp., most-preferred) diagnoses of M . We first presented filters and preference orders on diagnoses in their most general form, which allows to capture well-known formalisms for preferences specification like CP-nets (Boutilier et al., 2004). We then presented an approach at meta-reasoning where a direct encoding of bridge rule modifications allows for perfect observation. The approach requires some enhanced notions of diagnosis, namely diagnoses where some bridge rules are protected and diagnoses where some bridge rules are considered to be of higher priority than the rest. An analysis of the computational complexity of these notions revealed that (subset-)minimal diagnoses with protected bridge rules have the same complexity as (subset-)minimal diagnoses; prioritized-minimal diagnoses have the same complexity, but are not sufficiently strong to characterize the most-preferred diagnoses in general. The respective

MCS M and . . .	Transformation	Size	Diagnosis notion	Complexity
filter f	M^f (Def. 13)	linear	$D_m^\pm(M^f, br_P)$	\mathbf{D}_i^P
total pref. order \preceq	$M^{pl\preceq}$ (Def. 16)	exponential	$D_m^\pm(M^{pl\preceq}, br_P, br_H)$	\mathbf{D}_i^P
preference order \preceq	M^\preceq (Def. 17)	linear	$D_{m,t_{max}}^\pm(M^\preceq, br_P, br_H)$	$\mathbf{\Pi}_{i+1}^P$

Table 2: Overview of the meta-reasoning transformations to select filtered and most-preferred diagnoses. Size is in terms of $|br(M)|$, and complexity wrt. context complexity of M if $\mathcal{CC}(M) = \Sigma_i^P$, $i \geq 1$, and deciding $f(D) \stackrel{?}{=} 1$ resp. $D \prec D'$ is in $\mathcal{CC}(M)$.

notion are mpm-diagnoses, which have higher complexity than subset-minimal diagnoses. On the other hand, identifying most-preferred diagnoses is as hard as identifying mpm-diagnoses; hence our meta-reasoning approach is worst-case optimal from a complexity point of view. Table 2 gives an overview of the developed meta-reasoning techniques and their respective overall complexities.

8.0.1 OUTLOOK

Regarding future work, some issues are still open. First, we currently allow arbitrary preferences on diagnoses, but these preferences cannot take the behavior of the repaired MCS into account. For example, a diagnosis of the hospital MCS might be less preferred if vital information is “lost” due to the diagnosis, e.g., a patient having some illness is known to the context with patient data, but the medication context suggests no treatment for the patient. In this case, the required preference cannot be defined on sets of bridge rules alone, but one needs to take the resulting equilibrium into account. In principle, the meta-reasoning transformation presented here can be extended to consider also beliefs that result from the witnessing equilibrium, e.g. by adding protected bridge rules from all contexts of the original MCS to the observer context. Since a diagnosis possibly admits multiple equilibria, a correct encoding is neither obvious nor is it independent of the formalization of said preferences.

Another issue concerns an implementation of the presented transformations as well as an implementation of the advanced notions for diagnosis selection. Due to our complexity results, one may in principle exploit the implementation of subset-minimal diagnoses in the MCS-IE tool by Bögl et al. (2010) and Eiter et al. (2014). The latter is based on encoding the diagnoses of an MCS to answer sets of a HEX-program as in the work of Eiter et al. (2010). HEX-programs extend ASP with an API-style interface to external information; this interface can be exploited to model the acceptability function of a range of context logics. As the encoding represents the removal resp. modification of bridge rules by designated atoms (similar as in Lemma 4), it can be adapted to leave protected bridge rules untouched; consequently, the diagnoses with protected bridge rules, and by the system’s filtering the subset-minimal such diagnoses can be computed. Combined with our polynomial reductions for all but mpm-diagnoses to subset-minimal diagnoses, we can get an implementation of all advanced notions of diagnosis, except for mpm-diagnoses; moreover, even a distributed evaluation method (Salcher, 2016), based on the results by Dao-Tran, Eiter, Fink, and Krennwallner (2015) can be imagined. Practical restrictions and particular instantiations of the context logics will allow for additional optimizations, which however we did not consider in this foundational analysis of the problems. Indeed, such restrictions may also lower the complexity. For example, total prefer-

ence orders potentially warrant this, as on the one hand, for such orders the computationally easier notion of prioritized-minimal diagnoses is sufficient to select the most-preferred diagnoses, and on the other hand, the hardness results for general preference orders rely on a non-total preference order. Further work is needed to refine the picture in this regard.

Acknowledgements

We are grateful for the meticulous feedback we received from several anonymous reviewers and we thank them for their time and their suggestions that helped to improve this article significantly. This work has been partially supported by the Austrian Science Fund (FWF) project P27730 and the Academy of Finland, project 251170. Most of the contributions of the second author were done at TU Wien where initial investigations were also funded by the Vienna Science and Technology Fund (WWTF) under grant ICT08-020.

Appendix A. Proofs

A.1 Proofs of Section 3

Proof of Proposition 2. We first show that $D^N(M) = D_{\leq N}^\pm(M)$. We write down $D^N(M)$ in set-notation and obtain:

$$\begin{aligned} D^N(M) &= \{D \in D^\pm(M) \mid \nexists D' \in D^\pm(M) : N \models D' \lesssim D \wedge \neg(N \models D \lesssim D')\} \\ &= \{D \in D^\pm(M) \mid \forall D' \in D^\pm(M) : \neg N \models D' \lesssim D \vee N \models D \lesssim D'\} \end{aligned}$$

Regarding $D_{\leq N}^\pm(M)$ we have that:

$$\begin{aligned} D_{\leq N}^\pm(M) &= \{D \in D^\pm(M) \mid \forall D' \in D^\pm(M) : \neg(D' \lesssim^N D \wedge D \not\lesssim^N D' \wedge D' \neq D)\} \\ &= \{D \in D^\pm(M) \mid \forall D' \in D^\pm(M) : \neg(N \models D' \lesssim D \wedge \neg N \models D \lesssim D' \wedge D' \neq D)\} \\ &= \{D \in D^\pm(M) \mid \forall D' \in D^\pm(M) : \neg N \models D' \lesssim D \vee N \models D \lesssim D' \vee D' = D\} \end{aligned}$$

It remains to show that given any $D, D' \in D^\pm(M)$, the following two formulas are equivalent:

$$\neg N \models D' \lesssim D \vee N \models D \lesssim D' \tag{15}$$

$$\neg N \models D' \lesssim D \vee N \models D \lesssim D' \vee D' = D \tag{16}$$

Clearly, (15) implies (16), it thus remains to show that (16) implies (15). The latter clearly holds if $\neg N \models D' \lesssim D$ holds or $N \models D \lesssim D'$ holds. Therefore, it only remains to show that in the case where both do not hold, (15) is implied by (16): from $N \models D' \lesssim D$ and $\neg N \models D \lesssim D'$ follows $D' = D$, hence by $N \models D' \lesssim D$ it then follows that $N \models D \lesssim D'$, i.e., (15) is satisfied in this case. Consequently, (16) implies (15) and thus, both conditions are equivalent. Therefore, it holds that $D^N(M) = D_{\leq N}^\pm(M)$.

It then follows trivially from the definitions of $D_{ird}^\pm(M, N)$ and $D_{m, \leq N}^\pm(M)$ that they are the same, because $D_{ird}^\pm(M, N)$ is the set of \subseteq -minimal diagnoses of $D_{\leq N}^\pm(M)$ while $D_{m, \leq N}^\pm(M)$ is the set of \subseteq -minimal diagnoses of $D^N(M)$. \square

A.2 Proofs of Section 4

The following lemma shows that the applicable bridge rules of M under a diagnosis (D_1, D_2) add exactly those knowledge-base elements that are also added under the corresponding diagnosis $(d1(D_1), d2(D_2) \cup \mathcal{K})$ of $M^{mr(\theta, \mathcal{K}_p)}$, where $\mathcal{K} \subseteq \mathcal{K}_p$ is arbitrary.

Lemma 6. *Let M be an MCS and $M^{mr(\theta, \mathcal{K}_p)}$ be a meta-reasoning encoding wrt. θ and \mathcal{K}_p . Furthermore, let $D_1, D_2 \subseteq br(M)$, let $\mathcal{K} \subseteq \mathcal{K}_p$, let $S = (S_1, \dots, S_n)$ be a belief state of M , and let $S' = (S_1, \dots, S_n, S_{n+1})$ be a belief state of $M^{mr(\theta, \mathcal{K}_p)}$ where $S_{n+1} = \{uncond_r \mid r \in D_2\} \cup \{removed_r \mid r \in D_1\}$. Then, for all $1 \leq i \leq n$, $\{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\} = \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}$.*

Proof. Let $D_1, D_2 \subseteq br(M)$, let $\mathcal{K} \subseteq \mathcal{K}_p$, let $S = (S_1, \dots, S_n)$ be a belief state of M , and let $S' = (S_1, \dots, S_n, S_{n+1})$ be a belief state of $M^{mr(\theta, \mathcal{K}_p)}$ where $S_{n+1} = \{uncond_r \mid r \in D_2\} \cup \{removed_r \mid r \in D_1\}$. Furthermore, let i be arbitrary such that $1 \leq i \leq n$ holds. We show that $\{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\} = \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$ holds.

“ \supseteq ”: Let $s \in \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$. Then $s = \varphi(r)$ for some bridge rule r such that either $r \in br(M) \setminus D_1$ and $S \vdash_{\rightarrow} r$, or $r = cf(r_2)$ where $r_2 \in D_2$. In the former case, consider the bridge rule r_1 of form (7) wrt. r . By construction, $body(r_1) = body(r) \cup \{\mathbf{not} (n+1 : removed_r)\}$ and $\varphi(r_1) = \varphi(r)$. Since $r \notin D_1$, $removed_r \notin S_{n+1}$, and since S and S' agree on S_i for $i \in \{1, \dots, n\}$, i.e., $S =_{\{1, \dots, n\}} S'$, it follows that $S' \vdash_{\rightarrow} r_1$. Therefore $\varphi(r_1) = \varphi(r) = s \in \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}$. In the latter case, where $r = cf(r_2)$ and $r_2 \in D_2$ hold, observe that $r_2 \in D_2$ implies that $uncond_{r_2} \in S_{n+1}$. Consider the bridge rule r'_2 of form (8) wrt. r_2 and observe that $\varphi(r'_2) = \varphi(r_2) = s$ while $body(r'_2) = \{(n+1 : uncond_{r_2})\}$. Since $uncond_{r_2} \in S_{n+1}$, it holds that $S' \vdash_{\rightarrow} r'_2$, hence $s \in \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}$. Thus it follows that $\{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\} \subseteq \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}$.

“ \subseteq ”: Let $s \in \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}$. Then there exists some $r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')$ such that $s = \varphi(r)$. Note that r either is of form (7) or of form (8). In the former case, it holds that $S' \vdash_{\rightarrow} r$ and $removed_{r_1} \notin S_{n+1}$ where $r_1 \in br_i(M)$ and r is the bridge rule of form (7) wrt. r_1 . Since S and S' agree on all belief sets from S_1 to S_n , i.e., $S =_{\{1, \dots, n\}} S'$, and $body(r) = body(r_1) \cup \{\mathbf{not} (n+1 : removed_r)\}$, it holds that $S \vdash_{\rightarrow} r$. Since $removed_{r_1} \notin S_{n+1}$ it furthermore holds that $r_1 \notin D_1$. This implies that $r_1 \in br_i(M[D_1, D_2])$ and consequently it holds that $r_1 \in app(br_i(M[D_1, D_2]), S)$, thus $s = \varphi(r) = \varphi(r_1) \in \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$. If r is of form (8), $body(r) = \{(n+1 : uncond_{r_2})\}$ where $r_2 \in br_i(M)$ and r is the bridge rule of form (8) wrt. r_2 . Since $r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')$ and $r \notin d2(D_2) \cup \mathcal{K}$, it follows that $S' \vdash_{\rightarrow} r$, hence $uncond_{r_2} \in S_{n+1}$ and thus $r_2 \in D_2$. Therefore, it holds that $cf(r_2) \in app(br_i(M[D_1, D_2]), S)$ and consequently $\varphi(r_2) = \varphi(r) = s \in \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$. In both cases it holds that $\{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\} \subseteq \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$. \square

The next lemma shows that every protected diagnosis of a meta-reasoning MCS is exhibited in the belief set of the observation context of every witnessing equilibrium of said diagnosis.

Lemma 7. *Let $M = (C_1, \dots, C_n)$ be an MCS and $M^{mr(\theta, \mathcal{K}_p)} = (C_1, \dots, C_n, C_{n+1})$ be a meta-reasoning encoding. Given that $D_1, D_2 \subseteq br(M)$, $\mathcal{K} \subseteq \mathcal{K}_p$, and $S = (S_1, \dots, S_n, S_{n+1})$ is a belief state of $M^{mr(\theta, \mathcal{K}_p)}$,*

$$S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \cup \{\varphi(r) \mid r \in app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S)\})$$

holds iff $S_{n+1} = \{uncond_r \mid r \in D_2\} \cup \{removed_r \mid r \in D_1\}$ and $\theta(D_1, D_2, \mathcal{K})$ holds.

Proof. By definition of \mathbf{ACC}_{n+1} (cf. Definition 12)

$$S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \cup \{\varphi(r) \mid r \in app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S)\})$$

holds iff $S_{n+1} = \{removed_r \mid r \in R_1\} \cup \{uncond_r \mid r \in R_2\}$ and $\theta(R_1, R_2, R_3)$ is true, where

$$\begin{aligned} R_1 &= \{r \in br(M) \mid not_removed_r \notin H\}, \\ R_2 &= \{r \in br(M) \mid uncond_r \in H\}, \\ R_3 &= \{r \in \mathcal{K}_p \mid \varphi(r) \in H\}, \text{ and} \\ H &= \{\varphi(r) \mid r \in app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S)\}. \end{aligned}$$

To prove this lemma, it therefore suffices to show that $R_1 = D_1$, $R_2 = D_2$, and $R_3 = \mathcal{K}$.

Consider the set B of bridge rules of context C_{n+1} in the MCS resulting from the application of the diagnosis:

$$\begin{aligned} B &= br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]) \\ &= \left(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}) \setminus d1(D_1) \right) \cup cf(d2(D_2) \cup \mathcal{K}) \\ &= \left((d1(br(M)) \cup d2(br(M)) \cup \mathcal{K}_p) \setminus d1(D_1) \right) \cup cf(d2(D_2) \cup \mathcal{K}). \end{aligned}$$

Observe that every bridge rule $r \in B$ is such that either $body(r) = \{\perp\}$ or $body(r) = \{\top\}$. Hence, for any belief state S the set of applicable bridge rules, call it B_{app} , is exactly the set of rules whose body is \top . Formally,

$$B_{app} = \{r \in B \mid body(r) = \{\top\}\} = app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S).$$

Recall that $r \in d1(br(M)) \cup d1(D_1) \cup cf(d2(D_2) \cup \mathcal{K})$ implies that $body(r) = \{\top\}$, while $r \in d2(br(M)) \cup \mathcal{K}_p$ implies that $body(r) = \{\perp\}$. Therefore,

$$B_{app} = d1(br(M)) \setminus d1(D_1) \cup cf(d2(D_2) \cup \mathcal{K})$$

and consequently it holds for the set H of heads of applicable bridge rules that

$$\begin{aligned} H &= \{\varphi(r) \mid r \in app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S)\} \\ &= \{\varphi(r) \mid r \in B_{app}\} \\ &= \{\varphi(r) \mid r \in (d1(br(M)) \setminus d1(D_1) \cup cf(d2(D_2) \cup \mathcal{K}))\} \\ &= \{not_removed_r \mid r \in br(M) \setminus D_1\} \cup \{uncond_r \mid r \in D_2\} \cup \{\varphi(r) \mid r \in \mathcal{K}\}. \end{aligned}$$

Note that the heads of br_{n+1} are *unique*, i.e., for all $r \neq r', r, r' \in br_{n+1}$ holds $\varphi(r) \neq \varphi(r')$. This is because the bridge rules r of \mathcal{K}_p all have unique head formulas $\varphi(r)$ and the remaining

bridge rules of br_{n+1} also have unique heads. Since the heads of br_{n+1} are unique, it holds for any $r_{\mathcal{K}} \in \mathcal{K}_p$ and $r \in br(M)$ that $uncond_r \neq \varphi(r_{\mathcal{K}}) \neq not_removed_r$ and it also holds for any $\mathcal{K}' \subseteq \mathcal{K}_p$ that the heads of \mathcal{K}' are unique. Consequently, it holds that

$$\begin{aligned} R_1 &= \{r \in br(M) \mid not_removed_r \notin H\} = \{r \in br(M) \mid r \in D_1\} = D_1 \\ R_2 &= \{r \in br(M) \mid uncond_r \in H\} = \{r \in br(M) \mid r \in D_2\} = D_2 \\ R_3 &= \{r \in \mathcal{K}_p \mid \varphi(r) \in H\} = \{r \in \mathcal{K}_p \mid r \in \mathcal{K}\} = \mathcal{K}. \end{aligned}$$

Since it only remained to show that $R_1 = D_1$, $R_2 = D_2$, and $R_3 = \mathcal{K}$, the lemma is therefore proven. \square

Proof of Proposition 3. (1) Since $S_{n+1} = \{uncond_r \mid r \in D_2\} \cup \{removed_r \mid r \in D_1\}$ and $S' = (S_1, \dots, S_n, S_{n+1})$, all pre-conditions of Lemma 7 and Lemma 6 are satisfied; hence we conclude the following.

By Lemma 7, $\theta(D_1, D_2, \mathcal{K})$ holds iff

$$\begin{aligned} S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \\ \cup \{\varphi(r) \mid r \in app(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}). \end{aligned} \quad (17)$$

By Lemma 6, for all $1 \leq i \leq n$ holds

$$\begin{aligned} \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\} \\ = \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}. \end{aligned}$$

which implies that for all $1 \leq i \leq n$ it holds that

$$\begin{aligned} \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}) \\ = \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}). \end{aligned}$$

This in turn implies that for all $1 \leq i \leq n$, it holds that

$$\begin{aligned} S_i \in \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}) \\ \text{iff } S_i \in \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}). \end{aligned} \quad (18)$$

From (18) and (17) we therefore obtain that: $\theta(D_1, D_2, \mathcal{K})$ holds and for all $1 \leq i \leq n$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M[D_1, D_2]), S)\}$ if and only if for all $1 \leq j \leq n+1$ it holds that

$$S_i \in \mathbf{ACC}_i(kb_i \cup \{\varphi(r) \mid r \in app(br_i(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}]), S')\}).$$

This is equivalent to: $\theta(D_1, D_2, \mathcal{K})$ and $S \in \text{EQ}(M[D_1, D_2])$ hold iff it holds that $S' \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}])$.

- (2) This is a direct consequence of (1) and the fact that a diagnosis implies the existence of a witnessing equilibrium and vice versa, i.e., $(D_1, D_2) \in D^\pm(M)$ iff there exists a belief state $S \in \text{EQ}(M[D_1, D_2])$, for any M, D_1, D_2 , and S . Thus

- $(D_1, D_2) \in D^\pm(M)$ and $\theta(D_1, D_2, \mathcal{K})$ hold
 iff $\theta(D_1, D_2, \mathcal{K})$ and $(S_1, \dots, S_n) \in \text{EQ}(M[D_1, D_2])$ hold
 iff (by (1)) $(S_1, \dots, S_n, S_{n+1}) \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[d1(D_1), d2(D_2) \cup \mathcal{K}])$ holds
 iff $(d1(D_1), d2(D_1) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)})$ holds.

It remains to show that $(d1(D_1), d2(D_1) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)})$ iff $(d1(D_1), d2(D_1) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$. This follows from $(d1(D_1) \cup d2(D_2) \cup \mathcal{K}) \cap br_P = \emptyset$ (see Definition 12) and Proposition 1, which shows that $D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) \subseteq D^\pm(M^{mr(\theta, \mathcal{K}_p)})$, i.e., every diagnosis with protected bridge rules also is a diagnosis. \square

The following lemma shows that the bridge rules of context C_{n+1} in the MCS $M^{mr(\theta, \mathcal{K}_p)}$ are such that for a minimal diagnosis $(D_1, D_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$, a bridge rule r with $body(r) = \{\top\}$ is only contained in D_1 (or not modified at all), and a bridge rule r with $body(r) = \{\perp\}$ is only contained in D_2 (or not modified at all).

Lemma 8. *Let $M^{mr(\theta, \mathcal{K}_p)}$ be a meta-reasoning encoding with protected bridge rules br_P , and let $(D_1, D_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$. Then, for every $r \in br(M^{mr(\theta, \mathcal{K}_p)}) \setminus br_P$ holds that:*

- (i) $body(r) = \{\top\}$ implies $r \notin D_2$ and
- (ii) $body(r) = \{\perp\}$ implies $r \notin D_1$.

Proof. Since $(D_1, D_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$, there exists a witnessing equilibrium $S \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[D_1, D_2])$ of (D_1, D_2) . Since (D_1, D_2) is a diagnosis with protected bridge rules, it holds that $(D_1 \cup D_2) \cap br_P = \emptyset$, which by construction of $M^{mr(\theta, \mathcal{K}_p)}$ implies that $r \in br_{n+1}$.

For a proof by contradiction, we now show the following:

- (i) if $body(r) = \{\top\}$ and $r \in D_2$ then $(D_1 \setminus \{r\}, D_2 \setminus \{r\}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$;
- (ii) if $body(r) = \{\perp\}$ and $r \in D_1$ then $(D_1 \setminus \{r\}, D_2) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$.

To show that the respective smaller diagnosis admits a witnessing equilibrium it suffices in the following to consider only applicable bridge rules of C_{n+1} , because it is the only context of $M^{mr(\theta, \mathcal{K}_p)}$ with bridge rules that are not protected.

- (i) Case $body(r) = \{\top\}$ and $r \in D_2$. Then

$$\varphi(r) \in \{\varphi(r) \mid r \in \text{app}(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[D_1, D_2]), S)\}$$

since $cf(r) \in \text{app}(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[D_1, D_2]), S)$. Now consider $(D_1 \setminus \{r\}, D_2 \setminus \{r\}) \subseteq (D_1, D_2)$ and observe that $r \in \text{app}(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[D_1 \setminus \{r\}, D_2 \setminus \{r\}]), S)$ since r is a bridge rule of the modified system and $body(r) = \{\top\}$. Thus, $S \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[D_1 \setminus \{r\}, D_2 \setminus \{r\}])$ and $(D_1 \setminus \{r\}, D_2 \setminus \{r\}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$. Note that this reasoning applies regardless of whether $r \in D_1$ holds.

- (ii) Case $body(r) = \{\perp\}$ and $r \in D_1$. Then

$$\text{app}(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[D_1 \setminus \{r\}, D_2]), S) = \text{app}(br_{n+1}(M^{mr(\theta, \mathcal{K}_p)}[D_1, D_2]), S)$$

since r either is not applicable (left-hand side), or it is not a bridge rule of the modified MCS (right-hand side). Consequently, $S \in \text{EQ}(M^{mr(\theta, \mathcal{K}_p)}[D_1 \setminus \{r\}, D_2])$ and therefore $(D_1 \setminus \{r\}, D_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$.

Each of these statements contradicts that $(D_1, D_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$, hence the statement of the lemma follows. \square

The following lemma shows that there are no diagnoses in $D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ other than those which correspond to diagnoses of M .

Lemma 9. *Let M be an MCS and $M^{mr(\theta, \mathcal{K}_p)}$ be some meta-reasoning encoding for M . For every $(R_1, R_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ there exist $D_1, D_2 \subseteq br(M)$ and $\mathcal{K} \subseteq \mathcal{K}_p$ such that $R_1 = d1(D_1)$ and $R_2 = d2(D_2) \cup \mathcal{K}$.*

Proof. Recall that br_P contains all bridge rules of form (7) and (8), hence the only bridge rules not in br_P are those of br_{n+1} , because $br_{M^{mr(\theta, \mathcal{K}_p)}} = br_P \cup br_{n+1}$. Since $br_{n+1} = d1(br(M)) \cup d2(br(M)) \cup \mathcal{K}_p$, it follows directly that for every $(R_1, R_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ there exist $D_1, D'_1, D_2, D'_2 \subseteq br(M)$ and $\mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p$ such that $R_1 = d1(D_1) \cup d2(D'_1) \cup \mathcal{K}'$ and $R_2 = d1(D'_2) \cup d2(D_2) \cup \mathcal{K}$. Observe that for all $r \in d2(D'_1) \cup \mathcal{K}'$ it holds that $body(r) = \{\perp\}$, hence by Lemma 8 it follows that $d2(D'_1) \cup \mathcal{K}' = \emptyset$. Furthermore, it holds for all $r \in d1(D'_2)$ that $body(r) = \{\top\}$, hence by Lemma 8 it follows that $d1(D'_2) = \emptyset$. Together, this means that $D'_1 = D'_2 = \mathcal{K}' = \emptyset$ and therefore it holds for every $(R_1, R_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ that there exist $D_1, D_2 \subseteq br(M)$ and $\mathcal{K} \subseteq \mathcal{K}_p$ such that $R_1 = d1(D_1)$ and $R_2 = d2(D_2) \cup \mathcal{K}$. \square

Proof of Propostion 4. By definition of minimal diagnosis, it holds that

$$\begin{aligned} D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \\ \{(R_1, R_2) \mid (R_1, R_2) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) \\ \text{and there exists no } (R'_1, R'_2) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) \\ \text{such that } (R'_1, R'_2) \subset (R_1, R_2)\} \end{aligned}$$

By Lemma 9, it holds for every $(R_1, R_2) \in D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ that there exist $D_1, D_2 \subseteq br(M)$ and $\mathcal{K} \subseteq \mathcal{K}_p$ such that $R_1 = d1(D_1)$ and $R_2 = d2(D_2) \cup \mathcal{K}$, hence we obtain that

$$\begin{aligned} D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \\ \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (d1(D_1), d2(D_2) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) \\ \text{and there exists no } (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) \\ \text{such that } (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subset (d1(D_1), d2(D_2) \cup \mathcal{K}) \\ \text{holds for some } \mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p\} \end{aligned}$$

By Proposition 3 we know that $(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P)$ holds iff $(D_1, D_2) \in D^\pm(M)$ and $\theta(D_1, D_2, \mathcal{K})$ hold. Therefore we obtain

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subset (d1(D_1), d2(D_2) \cup \mathcal{K}) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \\ \text{holds for some } \mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p\}.$$

Since $d1$ and $d2$ are bijective, $(d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subset (d1(D_1), d2(D_2) \cup \mathcal{K})$ holds iff $(D'_1, D'_2 \cup \mathcal{K}') \subset (D_1, D_2 \cup \mathcal{K})$ holds.

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2 \cup \mathcal{K}') \subset (D_1, D_2 \cup \mathcal{K}) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \\ \text{holds for some } \mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p\}.$$

□

Proof of Proposition 5. From Proposition 4 we know that

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M), \theta(D_1, D_2, \mathcal{K}) \text{ holds}, \\ [\nexists (D'_1, D'_2) \in D^\pm(M), \mathcal{K}' \subseteq \mathcal{K}_p : \\ (D'_1, D'_2 \cup \mathcal{K}') \subset (D_1, D_2 \cup \mathcal{K}) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \text{ holds}]\}.$$

Because θ is functional increasing, it holds that $(D'_1, D'_2 \cup \mathcal{K}') \subset (D_1, D_2 \cup \mathcal{K})$ holds iff $(D'_1, D'_2) \subset (D_1, D_2)$. We therefore obtain that:

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2) \subset (D_1, D_2) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \text{ holds for some } \mathcal{K}, \mathcal{K}' \subseteq \mathcal{K}_p\}.$$

□

A.3 Proofs of Section 5

A.3.1 PROOFS OF SECTION 5.1

Proof of Theorem 1. Recall that $M^f = M^{mr(\theta, \mathcal{K}_p)}$ where θ is defined such that $\theta(D_1, D_2, \emptyset)$ holds iff it holds that $f(D_1, D_2) = 1$, hence θ is functional increasing. By Lemma 5 it therefore holds that

$$D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2) \subset (D_1, D_2) \text{ and } \theta(D'_1, D'_2, \mathcal{K}') \text{ holds for some } \mathcal{K}', \mathcal{K}' \subseteq \mathcal{K}_p\}$$

which in case of M^f becomes

$$D_m^\pm(M^f, br_P) = \{(d1(D_1), d2(D_2)) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \emptyset) \text{ holds} \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2) \subset (D_1, D_2) \text{ and } \theta(D'_1, D'_2, \emptyset) \text{ holds}\}.$$

By definition of M^f it furthermore holds that $\theta(D_1, D_2, \emptyset)$ holds iff $f(D_1, D_2) = 1$, hence we obtain that

$$D_m^\pm(M^f, br_P) = \{(d1(D_1), d2(D_2)) \mid (D_1, D_2) \in D^\pm(M) \text{ and } f(D_1, D_2) = 1 \\ \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\ (D'_1, D'_2) \subset (D_1, D_2) \text{ and } f(D'_1, D'_2) = 1\}$$

Applying the definition of minimal-filtered diagnoses, we thus obtain that

$$D_m^\pm(M^f, br_P) = \{(d1(D_1), d2(D_2)) \mid (D_1, D_2) \in D_{m,f}^\pm(M)\}.$$

Note that this statement is equivalent to

$$D_{m,f}^\pm(M) = \{(D_1, D_2) \mid (d1(D_1), d2(D_2)) \in D_m^\pm(M^f, br_P)\}.$$

□

A.3.2 PROOFS OF SECTION 5.2

Proof of Lemma 1. “ \Rightarrow ”: Suppose that $(D_1, D_2) \preceq (D'_1, D'_2)$. We have to show that for every $K \in \text{map}_{\preceq}^g(D_1, D_2)$ it holds that $K \in \text{map}_{\preceq}^g(D'_1, D'_2)$. Let $K \in \text{map}_{\preceq}^g(D_1, D_2)$ hold. Then it follows by definition that $K = g(D''_1, D''_2)$ for some $(D''_1, D''_2) \in 2^{br(M)} \times 2^{br(M)}$. In the case that $(D''_1, D''_2) = (D_1, D_2)$ it follows from the reflexivity of \preceq that $(D''_1, D''_2) \preceq (D'_1, D'_2)$ and thus by definition of $\text{map}_{\preceq}^g(D'_1, D'_2)$ it holds that $K \in \text{map}_{\preceq}^g(D'_1, D'_2)$. In the case that $(D''_1, D''_2) \neq (D_1, D_2)$ it follows by the definition of $\text{map}_{\preceq}^g(D_1, D_2)$ that $(D''_1, D''_2) \preceq (D_1, D_2)$. Since $(D_1, D_2) \preceq (D'_1, D'_2)$ and \preceq is transitive, it follows that $(D''_1, D''_2) \preceq (D'_1, D'_2)$ and consequently, it holds that $K \in \text{map}_{\preceq}^g(D'_1, D'_2)$. Thus for any $K \in \text{map}_{\preceq}^g(D_1, D_2)$ it holds that $K \in \text{map}_{\preceq}^g(D'_1, D'_2)$, i.e., $\text{map}_{\preceq}^g(D_1, D_2) \subseteq \text{map}_{\preceq}^g(D'_1, D'_2)$.

“ \Leftarrow ”: Suppose that $\text{map}_{\preceq}^g(D_1, D_2) \subseteq \text{map}_{\preceq}^g(D'_1, D'_2)$. We have to show that $(D_1, D_2) \preceq (D'_1, D'_2)$. By reflexivity of \preceq it holds that $g(D_1, D_2) \in \text{map}_{\preceq}^g(D_1, D_2)$, thus $g(D_1, D_2) \in \text{map}_{\preceq}^g(D'_1, D'_2)$ and consequently $g(D_1, D_2) \in \text{map}_{\preceq}^g(D'_1, D'_2)$. By definition of $\text{map}_{\preceq}^g(D'_1, D'_2)$ it then follows that $(D_1, D_2) \preceq (D'_1, D'_2)$. □

The following lemma shows that the set $D_m^\pm(M^{pl\preceq}, br_P)$ of minimal diagnoses with protected bridge rules of $M^{pl\preceq}$ corresponds to those diagnoses of M which are at the same time, preferred according to \preceq and \subseteq -minimal. These diagnoses not yet correspond to minimal \preceq -preferred diagnoses since preference among \subseteq -incomparable diagnoses is not captured by $D_m^\pm(M^{pl\preceq}, br_P)$.

Lemma 10. *Given an MCS M and a preference \preceq on its diagnoses, it holds that*

$$D_m^\pm(M^{pl\preceq}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \wedge \\ \mathcal{K} = \text{map}_{\preceq}^g(D_1, D_2) \wedge \forall (D'_1, D'_2) \in D^\pm(M) : \\ ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \Rightarrow (D_1, D_2) = (D'_1, D'_2)\}.$$

Proof. By Proposition 4 it holds that:

$$\begin{aligned}
 D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) &= \\
 &\{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \text{ and } \theta(D_1, D_2, \mathcal{K}) \text{ holds} \\
 &\quad \text{and there exists no } (D'_1, D'_2) \in D^\pm(M) \text{ such that} \\
 &\quad (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subset (d1(D_1), d2(D_2) \cup \mathcal{K}) \text{ and} \\
 &\quad \theta(D'_1, D'_2, \mathcal{K}') \text{ holds for some } \mathcal{K}' \subseteq \mathcal{K}_p\} \\
 = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \wedge \theta(D_1, D_2, \mathcal{K}) \wedge \forall (D'_1, D'_2) \in D^\pm(M): \\
 &\quad (\exists \mathcal{K}' : \theta(D'_1, D'_2, \mathcal{K}') \wedge \\
 &\quad (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subseteq (d1(D_1), d2(D_2) \cup \mathcal{K})) \\
 &\quad \Rightarrow (d1(D'_1), d2(D'_2) \cup \mathcal{K}') = (d1(D_1), d2(D_2) \cup \mathcal{K})\}
 \end{aligned}$$

Next we substitute θ by its definition, i.e., $\theta(D_1, D_2, \mathcal{K})$ iff $map_{\leq}^g(D_1, D_2) = \mathcal{K}$.

$$\begin{aligned}
 D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) &= \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \\
 &\quad \wedge map_{\leq}^g(D_1, D_2) = \mathcal{K} \wedge \forall (D'_1, D'_2) \in D^\pm(M): \\
 &\quad (\exists \mathcal{K}' : map_{\leq}^g(D'_1, D'_2) = \mathcal{K}' \wedge \\
 &\quad (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subseteq (d1(D_1), d2(D_2) \cup \mathcal{K})) \\
 &\quad \Rightarrow (d1(D'_1), d2(D'_2) \cup \mathcal{K}') = (d1(D_1), d2(D_2) \cup \mathcal{K})\}
 \end{aligned}$$

Since $d1$ and $d2$ both are bijective, $map_{\leq}^g(D_1, D_2) = \mathcal{K}$, and $map_{\leq}^g(D'_1, D'_2) = \mathcal{K}'$, it follows that $(d1(D'_1), d2(D'_2) \cup \mathcal{K}') = (d1(D_1), d2(D_2) \cup \mathcal{K})$ holds iff $(D'_1, D'_2) = (D_1, D_2)$. Hence,

$$\begin{aligned}
 D_m^\pm(M^{mr(\theta, \mathcal{K}_p)}, br_P) &= \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \\
 &\quad \wedge map_{\leq}^g(D_1, D_2) = \mathcal{K} \wedge \forall (D'_1, D'_2) \in D^\pm(M): \\
 &\quad (\exists \mathcal{K}' : (map_{\leq}^g(D'_1, D'_2) = \mathcal{K}' \wedge \\
 &\quad (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subseteq (d1(D_1), d2(D_2) \cup \mathcal{K})) \\
 &\quad \Rightarrow (D_1, D_2) = (D'_1, D'_2))\}
 \end{aligned}$$

Towards the next step, we need to show that the following holds true for $(D_1, D_2) \in D^\pm(M)$, $(D'_1, D'_2) \in D^\pm(M)$, and $map_{\leq}^g(D_1, D_2) = \mathcal{K}$:

$$\begin{aligned}
 (map_{\leq}^g(D'_1, D'_2) = \mathcal{K}' \wedge (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subseteq (d1(D_1), d2(D_2) \cup \mathcal{K})) \\
 \Rightarrow (D_1, D_2) = (D'_1, D'_2) \quad (19)
 \end{aligned}$$

iff

$$\begin{aligned}
 ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \\
 \Rightarrow (D_1, D_2) = (D'_1, D'_2) \quad (20)
 \end{aligned}$$

Observe that $(d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subseteq (d1(D_1), d2(D_2) \cup \mathcal{K})$ holds iff $(D'_1, D'_2) \subseteq (D_1, D_2)$ and $\mathcal{K}' \subseteq \mathcal{K}$ both hold. By Lemma 1 it further holds that $\mathcal{K}' = map_{\leq}^g(D'_1, D'_2) \subseteq map_{\leq}^g(D_1, D_2) = \mathcal{K}$ iff $(D'_1, D'_2) \preceq (D_1, D_2)$. Therefore, (19) holds iff (20) holds. After substitution, it therefore holds

that:

$$\begin{aligned}
 D_m^\pm(M^{pl\preceq}, br_P) &= \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \\
 &\quad \wedge map_{\preceq}^g(D_1, D_2) = \mathcal{K} \wedge \forall (D'_1, D'_2) \in D^\pm(M) : \\
 &\quad ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \Rightarrow (D_1, D_2) = (D'_1, D'_2)\}
 \end{aligned}$$

□

Proof of Theorem 2. In the following, let θ , \mathcal{K}_p , and map_{\preceq}^g be according to $M^{pl\preceq} = M^{mr(\theta, \mathcal{K}_p)}$.

“ \Rightarrow ”: Let $(R_1, R_2) \in D^\pm(M^{pl\preceq}, br_P, br_H)$, i.e., $(R_1, R_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ and for all $(R'_1, R'_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ holds that $(R'_1, R'_2) \subseteq_{br_H} (R_1, R_2) \Rightarrow (R'_1, R'_2) =_{br_H} (R_1, R_2)$. By Lemma 10 it holds that $(R_1, R_2) = (d1(D_1), d2(D_2) \cup \mathcal{K})$ where $\mathcal{K} = map_{\preceq}^g(D_1, D_2)$ and $(D_1, D_2) \in D^\pm(M)$. To show that $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$, we have to show that (D_1, D_2) is \preceq -preferred and subset minimal among all \preceq -preferred diagnoses. Assume that (D_1, D_2) is not \preceq -preferred. Then by (4) there exists a diagnosis $(D'_1, D'_2) \in D^\pm(M)$ such that $(D'_1, D'_2) \preceq (D_1, D_2)$ and $(D_1, D_2) \not\subseteq (D'_1, D'_2)$ hold. Let $map_{\preceq}^g(D'_1, D'_2) = \mathcal{K}'$ and $map_{\preceq}^g(D_1, D_2) = \mathcal{K}$. Since it holds that $(D'_1, D'_2) \preceq (D_1, D_2)$ it follows from Lemma 1 that $\mathcal{K}' \subseteq \mathcal{K}$. From $(D_1, D_2) \not\subseteq (D'_1, D'_2)$ it also follows by the same lemma that $\mathcal{K} \not\subseteq \mathcal{K}'$ holds and thus $\mathcal{K}' \subset \mathcal{K}$ holds. This means that $(R'_1, R'_2) = (d1(D'_1), d2(D'_2) \cup \mathcal{K}') \subset_{br_H} (d1(D_1), d2(D_2) \cup \mathcal{K}) = (R_1, R_2)$ holds.

Now suppose $(R'_1, R'_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ holds; then $(R_1, R_2) \in D^\pm(M^{pl\preceq}, br_P, br_H)$ contradicts that $(R'_1, R'_2) \subset_{br_H} (R_1, R_2)$. On the other hand, $(R'_1, R'_2) \notin D_m^\pm(M^{pl\preceq}, br_P)$ implies that some $(R''_1, R''_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ exists with $(R''_1, R''_2) \subset (R'_1, R'_2)$, i.e., there exist $D''_1, D''_2 \subseteq br(M)$ such that $(D''_1, D''_2) \preceq (D'_1, D'_2) \preceq (D_1, D_2)$ and $\mathcal{K}'' \subseteq \mathcal{K}' \subset \mathcal{K}$ both hold where $\mathcal{K}'' = map_{\preceq}^g(D''_1, D''_2)$, $R''_1 = d1(D''_1)$, and $R''_2 = d2(D''_2) \cup \mathcal{K}''$. Since $\mathcal{K}'' \subset \mathcal{K}$ it therefore holds that $(R''_1, R''_2) \subset_{br_H} (R_1, R_2)$ and together with $(R''_1, R''_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ this contradicts that $(R_1, R_2) \in D^\pm(M^{pl\preceq}, br_P, br_H)$. Since every case yields a contradiction, it therefore follows that there exists no such (D'_1, D'_2) , i.e., (D_1, D_2) indeed is a \preceq -preferred diagnosis.

It remains to show that (D_1, D_2) is subset-minimal among all \preceq -preferred diagnoses. Towards contradiction, assume there exists $(D'_1, D'_2) \in D_{\preceq}^\pm(M)$ such that $(D'_1, D'_2) \subset (D_1, D_2)$. We distinguish on how \preceq relates (D_1, D_2) and (D'_1, D'_2) .

- case $(D_1, D_2) \preceq (D'_1, D'_2) \wedge (D'_1, D'_2) \preceq (D_1, D_2)$: since $(R_1, R_2) \in D_m^\pm(M^{pl\preceq}, br_P)$, it holds by Lemma 10 that $(D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2) \Rightarrow (D'_1, D'_2) = (D_1, D_2)$ which directly contradicts that $(D'_1, D'_2) \subset (D_1, D_2)$.
- case $(D_1, D_2) \preceq (D'_1, D'_2) \wedge (D'_1, D'_2) \not\preceq (D_1, D_2)$: in this case, (D'_1, D'_2) is not \preceq -preferred, because $(D_1, D_2) \prec (D'_1, D'_2)$. Hence, it contradicts that $(D'_1, D'_2) \in D_{\preceq}^\pm(M)$.
- case $(D_1, D_2) \not\preceq (D'_1, D'_2) \wedge (D'_1, D'_2) \preceq (D_1, D_2)$: this case is analogous to the first one, i.e., $(R_1, R_2) \in D_m^\pm(M^{pl\preceq}, br_P)$ contradicts that $(D'_1, D'_2) \preceq (D_1, D_2)$ and $(D'_1, D'_2) \subset (D_1, D_2)$ both hold.
- case $(D_1, D_2) \not\preceq (D'_1, D'_2) \wedge (D'_1, D'_2) \not\preceq (D_1, D_2)$: this case contradicts with \preceq being total.

Consequently, there exists no $(D'_1, D'_2) \in D_{\preceq}^\pm(M)$ such that $(D'_1, D'_2) \subset (D_1, D_2)$ and therefore it holds that $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$.

“ \Leftarrow ”: Let $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$. We have to show that

$$(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D^\pm(M^{pl\preceq}, br_P, br_H)$$

holds with $map_{\preceq}^g(D_1, D_2) = \mathcal{K}$. By definition, it holds that

$$D^\pm(M^{pl\preceq}, br_P, br_H) = \{D \in D_m^\pm(M^{pl\preceq}, br_P) \mid \forall D' \in D_m^\pm(M^{pl\preceq}, br_P) : \\ D' \subseteq_{br_H} D \Rightarrow D' =_{br_H} D\}.$$

While by Lemma 10 it holds that:

$$D_m^\pm(M^{pl\preceq}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \mid (D_1, D_2) \in D^\pm(M) \wedge \\ \mathcal{K} = map_{\preceq}^g(D_1, D_2) \wedge \forall (D'_1, D'_2) \in D^\pm(M) : \\ ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \Rightarrow (D_1, D_2) = (D'_1, D'_2)\}.$$

Observe that $br_H = \mathcal{K}_p$ and $(d1(br(M)) \cup d2(br(M))) \cap \mathcal{K}_p = \emptyset$, hence $(d1(D_1), d2(D_2) \cup \mathcal{K}) \subseteq_{br_H} (d1(D'_1), d2(D'_2) \cup \mathcal{K}')$ holds iff $\mathcal{K} \subseteq \mathcal{K}'$ holds.

Therefore, it also holds that:

$$D^\pm(M^{pl\preceq}, br_P) = \{(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D_m^\pm(M^{pl\preceq}, br_P) \mid \\ \forall (D'_1, D'_2) \in D^\pm(M) : \\ [\forall (D''_1, D''_2) \in D^\pm(M) : ((D''_1, D''_2) \preceq (D'_1, D'_2) \\ \wedge (D''_1, D''_2) \subseteq (D'_1, D'_2)) \Rightarrow (D'_1, D'_2) = (D''_1, D''_2)] \\ \Rightarrow (map_{\preceq}^g(D'_1, D'_2) \subseteq \mathcal{K} \Rightarrow \mathcal{K} = map_{\preceq}^g(D'_1, D'_2))\}. \quad (21)$$

First, we show that $(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D_m^\pm(M^{pl\preceq}, br_P)$, which by Lemma 10 holds iff the following holds: $(D_1, D_2) \in D^\pm(M) \wedge map_{\preceq}^g(D_1, D_2) = \mathcal{K} \wedge \forall (D'_1, D'_2) \in D^\pm(M) : ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \Rightarrow (D_1, D_2) = (D'_1, D'_2)$. Since it holds that $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$, it also holds that $(D_1, D_2) \in D^\pm(M)$, and $\mathcal{K} = map_{\preceq}^g(D_1, D_2)$ by construction.

It remains to show that $\forall (D'_1, D'_2) \in D^\pm(M) : ((D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)) \Rightarrow (D_1, D_2) = (D'_1, D'_2)$. Assume towards contradiction that there exists some $(D'_1, D'_2) \in D^\pm(M)$ such that $(D'_1, D'_2) \preceq (D_1, D_2) \wedge (D'_1, D'_2) \subseteq (D_1, D_2)$ and $(D_1, D_2) \neq (D'_1, D'_2)$, i.e., it holds for (D'_1, D'_2) that $(D'_1, D'_2) \subset (D_1, D_2) \wedge (D'_1, D'_2) \preceq (D_1, D_2)$. We distinguish whether $(D_1, D_2) \preceq (D'_1, D'_2)$ also holds: if $(D_1, D_2) \preceq (D'_1, D'_2)$ holds, (D'_1, D'_2) is \preceq -preferred since (D_1, D_2) is. Since $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$, (D_1, D_2) is subset-minimal among all \preceq -preferred diagnoses, which contradicts that $(D'_1, D'_2) \subset (D_1, D_2)$ holds. In the case that $(D_1, D_2) \not\preceq (D'_1, D'_2)$, it holds that $(D_1, D_2) \notin D_{\preceq}^\pm(M)$, since it holds that $(D'_1, D'_2) \preceq (D_1, D_2) \wedge (D_1, D_2) \neq (D'_1, D'_2) \wedge (D_1, D_2) \not\preceq (D'_1, D'_2)$. This contradicts that $(D_1, D_2) \in D_{m, \preceq}^\pm(M)$. Hence it follows that no such (D'_1, D'_2) exists. Consequently, it holds that $(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D_m^\pm(M^{pl\preceq}, br_P)$.

According to (21), it remains to show that for all $(D'_1, D'_2) \in D^\pm(M)$ it holds that

$$[\forall (D''_1, D''_2) \in D^\pm(M) : ((D''_1, D''_2) \preceq (D'_1, D'_2) \wedge (D''_1, D''_2) \subseteq (D'_1, D'_2)) \\ \Rightarrow (D'_1, D'_2) = (D''_1, D''_2)] \Rightarrow (map_{\preceq}^g(D'_1, D'_2) \subseteq \mathcal{K} \Rightarrow \mathcal{K} = map_{\preceq}^g(D'_1, D'_2)).$$

Towards contradiction, assume there exists $(D'_1, D'_2) \in D^\pm(M)$ such that $\forall (D''_1, D''_2) \in D^\pm(M) : ((D''_1, D''_2) \preceq (D'_1, D'_2) \wedge (D''_1, D''_2) \subseteq (D'_1, D'_2)) \Rightarrow (D'_1, D'_2) = (D''_1, D''_2)$ holds and also $map_{\preceq}^g(D'_1, D'_2) \subsetneq \mathcal{K}$ holds. Since $map_{\preceq}^g(D'_1, D'_2) \subsetneq \mathcal{K}$, it follows by Lemma 1 that $(D'_1, D'_2) \preceq (D_1, D_2)$ and $(D_1, D_2) \not\preceq (D'_1, D'_2)$ both hold, which implies $(D_1, D_2) \notin D_{m, \preceq}^\pm(M)$, in contradiction to the assumption. Therefore, no such (D'_1, D'_2) can exist.

This proves that $(d1(D_1), d2(D_2) \cup \mathcal{K}) \in D^\pm(M^{pl\preceq}, br_P, br_H)$, which completes the proof. \square

A.3.3 PROOFS OF SECTION 5.3

Proof of Lemma 2. Observe that I is a bijection on $\{1, \dots, n\}$ which simply renames context identifiers. Therefore, one can directly conclude that $S \in \text{EQ}(M)$ holds iff $I(S) \in \text{EQ}(I(M))$ holds. In the following, we show in full detail that this renaming indeed is correct.

Let $S = (S_1, \dots, S_n)$ and $I(S) = (S_{I^{-1}(1)}, \dots, S_{I^{-1}(n)}) = (S'_1, \dots, S'_n)$. and let $1 \leq i \leq n$. Note that $S \in \text{EQ}(M)$ holds iff for all $1 \leq i \leq n$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M), S))$; additionally $I(S) \in \text{EQ}(I(M))$ holds iff for all $1 \leq j \leq n$ holds $S_j \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(I(M)), I(S)))$. Given that I is bijective and compatible to M , there exists $j \in \{1, \dots, n\}$ for every $i \in \{1, \dots, n\}$ such that $j = I(i)$ and vice versa, i.e., for every $j \in \{1, \dots, n\}$ exists a $i \in \{1, \dots, n\}$ such that $i = I^{-1}(j)$. We now show that for any $1 \leq i, j \leq n$ such that $j = I(i)$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M), S))$ iff $S_j \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(I(M)), I(S)))$. Observe that by construction of $I(M)$ it holds that $S_i = S_j$, $\mathbf{ACC}_i = \mathbf{ACC}_j$, and $kb_i = kb_j$. Hence it suffices to show that $\text{app}(br_i(M), S) = \text{app}(br_j(I(M)), I(S))$. Note that $br_j(I(M)) = I(br_i(M))$, hence there exists a bijection from $br_j(I(M))$ to $br_i(M)$, namely I ; furthermore I also maps bijectively each $r \in br_i(M)$ and every $(c : p) \in \text{body}^\pm(r)$ to $I(r)$ and $(I(c) : p)$. Since $\varphi(r) = \varphi(I(r))$ it suffices to show that $p \in S_c$ holds iff $p \in S'_{I(c)}$ holds. This is true since $S'_{I(c)} = S_{I^{-1}(I(c))} = S_c$, thus it follows that $\text{app}(br_i(M), S) = \text{app}(br_j(I(M)), I(S))$ which in turn implies that $S \in \text{EQ}(M)$ iff $I(S) \in \text{EQ}(I(M))$.

From this we also conclude that $S \in \text{EQ}(M[D_1, D_2])$ holds iff $I(S) \in \text{EQ}(I(M[D_1, D_2]))$ holds, because $M[D_1, D_2]$ is an MCS, hence the above statement also applies to $M[D_1, D_2]$. \square

To show that the set of diagnoses of $M \otimes M'$ is the product of the set of diagnoses of M and of M' , we use the following lemma, which states that if M' has no bridge rules, the set of diagnoses of M coincides with the set of diagnoses of $M \otimes M'$.

Lemma 11. *Given an MCS $M = (C_1, \dots, C_n)$ and an MCS $M' = (C'_1, \dots, C'_m)$ with $br(M') = \emptyset$. Then for every belief state (S_1, \dots, S_n) of M exist belief sets S_{n+1}, \dots, S_{n+m} such that $(S_1, \dots, S_{n+m}) \in \text{EQ}(M \otimes M')$ holds iff $(S_1, \dots, S_n) \in \text{EQ}(M)$ holds.*

Proof. Let $M^o = M \otimes M'$.

“ \Rightarrow ”: Let $S = (S_1, \dots, S_{n+m}) \in \text{EQ}(M \otimes M')$ be such that for every $1 \leq i \leq n + m$ holds $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^o), S))$. Note that by construction of M^o it holds for every bridge rule $r \in br_i(M^o)$ with $1 \leq i \leq n$ that $(c : p) \in \text{body}^\pm(r)$ implies that $c \in \{1, \dots, n\}$ holds. Hence by $br_i(M^o) = br_i(M)$ follows that $\text{app}(br_i(M^o), S) = \text{app}(br_i(M), (S_1, \dots, S_n))$. Therefore, for all $i \in C(M)$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M), (S_1, \dots, S_n)))$, i.e., $(S_1, \dots, S_n) \in \text{EQ}(M)$.

“ \Leftarrow ”: Let $S = (S_1, \dots, S_n) \in \text{EQ}(M)$ hold. Since $br(M') = \emptyset$, it holds for all $n + 1 \leq j \leq n + m$ that $br_j(M^o) = \emptyset$. Recall that contexts are consistent without bridge rules, i.e., there exists $S_j^\emptyset \in \mathbf{ACC}_j(kb_j \cup \emptyset)$ for all $n + 1 \leq j \leq n + m$. Consider the belief state $S' =$

$(S_1, \dots, S_n, S_{n+1}^\emptyset, \dots, S_{n+m}^\emptyset)$ and observe that for all $1 \leq i \leq n$ it holds that $\text{app}(br_i(M^\emptyset), S') = \text{app}(br_i(M), S)$ since $br_i(M^\emptyset) = br_i(M)$. It therefore follows that $S' \in \text{EQ}(M^\emptyset)$ holds. \square

Since shifting has no influence on acceptability, we can turn around the above lemma to show that the set of diagnoses of $M \otimes M'$ equals the set of diagnoses of M' if $br(M) = \emptyset$.

Corollary 3. *Given an MCS $M = (C_1, \dots, C_n)$ and an MCS $M' = (C'_1, \dots, C'_{n'})$ with $br(M) = \emptyset$. Then, for every belief state $(S'_1, \dots, S'_{n'})$ of M' exist belief sets S_1, \dots, S_n such that $(S_1, \dots, S_n, S'_1, \dots, S'_{n'}) \in \text{EQ}(M \otimes M')$ holds iff $(S'_1, \dots, S'_{n'}) \in \text{EQ}(M')$ holds.*

Proof. Consider a permutation I' that exchanges the positions of contexts of M and M' in $M \otimes M'$, formally: let I be the permutation wrt. $M \otimes M'$ and recall that I is compatible with $M \otimes M'$. Let $I' = I^{-1}$ and $M^s = I'(M \otimes M')$. Note that M^s equals $M' \otimes M$, hence by Lemma 2 we obtain that $(S_1, \dots, S_n, S'_1, \dots, S'_{n'}) \in \text{EQ}(M \otimes M')$ iff $I'((S_1, \dots, S_n, S'_1, \dots, S'_{n'})) \in \text{EQ}(M^s)$ iff $(S'_1, \dots, S'_{n'}, S_1, \dots, S_n) \in \text{EQ}(M' \otimes M)$.

Since $br(M) = \emptyset$ it holds by Lemma 11 that for every belief state $(S'_1, \dots, S'_{n'})$ of M' exist belief sets $S_{n'+1}, \dots, S_{n'+n}$ such that $(S'_1, \dots, S'_{n'}, S_{n'+1}, \dots, S_{n'+n}) \in \text{EQ}(M' \otimes M)$ holds iff $(S'_1, \dots, S'_{n'}) \in \text{EQ}(M')$ holds. In summary, $(S_1, \dots, S_n, S'_{n+1}, \dots, S'_{n+n'}) \in \text{EQ}(M \otimes M')$ holds iff $(S'_{n+1}, \dots, S'_{n+n'}) \in \text{EQ}(M')$ holds. \square

The proof of Proposition 6 makes use of Lemma 4 and Lemma 5 occurring in the appendix of the work by Eiter et al. (2014) about splitting sets in MCS. For convenience, we recap them here as well as the definition of a splitting set (which is similar to the notion of a splitting set in answer-set programming).

Definition 20 (cf. Eiter et al., 2014). *A set of contexts $U \subseteq C(M)$ is a splitting set of an MCS M , if every rule $r \in br(M)$ is such that $C_h(r) \in U$ satisfies $C_b(r) \subseteq U$. More formally, U is a splitting set iff $U \supseteq \bigcup \{C_b(r) \mid r \in br(M), C_h(r) \in U\}$.*

Lemma 12 (cf. Eiter et al., 2014). *Let U be a splitting set of an MCS M and let $R_1, R_2 \subseteq br(M)$. Then, U is also a splitting set of $M[R_1 \cup cf(R_2)]$.*

Lemma 13 (cf. Eiter et al., 2014). *Let M be an MCS, let B be a set of bridge rules compatible with M , and let U be a splitting set for $M[B]$. Furthermore, let $S = (S_1, \dots, S_n)$ and $S' = (S'_1, \dots, S'_{n'})$ be belief states of M , and let $b_U \subseteq R \subseteq B$. Then, $S =_U S'$ and $i \in U$ implies $\mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M[B]), S)) = \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M[R]), S'))$.*

Proof of Proposition 6. W.l.o.g. let $M = (C_1, \dots, C_n)$, let $M' = (C'_1, \dots, C'_{n'})$, and let $M^\emptyset = M \otimes M'$. Observe that by construction, there is no bridge rule whose head belongs to M (resp. M') and whose body contains a belief from M' (resp. M). Consequently, $U = \{1, \dots, n\}$ and $U' = \{n+1, \dots, n+n'\} = C(M^\emptyset) \setminus U$ are both splitting sets of M^\emptyset . Let $S^\emptyset = (S_1^\emptyset, \dots, S_{n+n'}^\emptyset)$ be an equilibrium of $M^\emptyset[\emptyset]$, which exists by our assumption that all contexts (of M and M') are consistent without bridge rules; additionally let $B = br(M^\emptyset) \setminus D_1 \cup cf(D_2)$.

“ \Rightarrow ”: Let $(D_1, D_2) \in D^\pm(M^\emptyset)$ hold. Then there exists a belief state $S = (S_1, \dots, S_{n+n'})$ such that for every $1 \leq i \leq n+n'$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^\emptyset[D_1, D_2]), S))$.

Consider $S_U = (S_1, \dots, S_n, S_{n+1}^\emptyset, \dots, S_{n+n'}^\emptyset)$ and observe that $S_U =_U S$; hence by Lemma 13 it follows for all $i \in U$ that

$$\mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^\emptyset[B]), S)) = \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^\emptyset[R]), S_U))$$

holds for all $b_U \subseteq R \subseteq B$, specifically for $R = b_U$. Note that U, U' , and b_U meant here are relative to the MCS $M^o[B]$, where by Lemma 12 U and U' are also splitting sets of $M^o[B]$. Consequently, for all $i \in U$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^o[b_U]), S_U))$ and for all $j \in C(M^o) \setminus U$ it holds that $S_j^\emptyset \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^o[b_U]), S_U))$, because $br_j(M^o[b_U]) = \emptyset$; thus it holds that $S_U \in \text{EQ}(M^o[b_U])$. Recall that b_U is defined relative to $M^o[B]$, hence $b_U = br(M) \setminus (D_1 \cap br(M)) \cup cf(D_2 \cap br(M))$, i.e., for $A_1 = D_1 \cap br(M)$ and $A_2 = D_2 \cap br(M)$ it holds that $M^o[b_U] = M^o[br(M) \setminus A_1 \cup cf(A_2)]$ and it follows that $S_U \in \text{EQ}(M^o[br(M) \setminus A_1 \cup cf(A_2)])$, i.e., it holds that $(A_1, A_2) \in D^\pm(M^o[br(M)])$. Since $M^o[br(M)] = M \otimes M'[\emptyset]$, Lemma 11 applies, i.e., it holds that $(S_1, \dots, S_n) \in \text{EQ}(M[A_1, D_2])$ and we conclude that $(A_1, A_2) \in D^\pm(M)$.

The proof that $(B_1, B_2) \in D^\pm(M')$ for $B_1 = D_1 \cap I(br(M'))$ and $B_2 = D_2 \cap I(br(M'))$ is analogous; it is based on the belief state $S_{U'} = (S_1^\emptyset, \dots, S_n^\emptyset, S_{n+1}, \dots, S_{n+n'})$ which is a witness of $(I(B_1), I(B_2)) \in D^\pm(M^o[b_{U'}])$; applying Corollary 3 (for $(M \otimes M')[b_{U'}] = M \otimes M'[B_1, B_2]$) then yields that $(B_1, B_2) \in D^\pm(M')$.

“ \Leftarrow ”: Let $(A_1, A_2) \in D^\pm(M)$ and $(B_1, B_2) \in D^\pm(M')$ hold. Then there exists some $S^A = (S_1^A, \dots, S_n^A) \in \text{EQ}(M[A_1, A_2])$ and $S^B = (S_1^B, \dots, S_{n'}^B) \in \text{EQ}(M'[B_1, B_2])$. Consider the belief state $S = (S_1, \dots, S_{n+n'})$ such that $S_i = S_i^A$ for $1 \leq i \leq n$ and $S_{n+j} = S_j^B$ for $1 \leq j \leq n'$. Observe that S is a belief state of the MCS $M^d = M^o[A_1 \cup I(B_1), A_2 \cup I(B_2)]$. Thus it suffices to show $S \in \text{EQ}(M^d)$, because this implies that $(A_1 \cup I(B_1), A_2 \cup I(B_2)) \in D^\pm(M \otimes M')$.

We first show that for all $1 \leq i \leq n$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^d), S))$. Let $B = br(M^d)$; hence $M^d = M^d[B]$, and note that U and U' are splitting sets of $M^d[B]$ by Lemma 12. Next we consider $M^d[b_U]$ (with b_U relative to M^d) and $R = b_U$. Since $M^d[R] = M^d[b_U] = (M[A_1, A_2] \otimes M'[\emptyset])$ and $S^A \in \text{EQ}(M[A_1, A_2])$, it holds by Lemma 11 that there exist $S'_{n+1}, \dots, S'_{n+n'}$ such that $S^M = (S_1, \dots, S_n, S'_{n+1}, \dots, S'_{n+n'}) \in \text{EQ}(M[A_1, A_2] \otimes M'[\emptyset])$, i.e., for all $1 \leq i \leq n$ it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^d[R]), S^M))$.

It holds that $S^M \equiv_U S$ and $b_U \subseteq R \subseteq B$; hence by Lemma 13 it holds for all $1 \leq i \leq n$ that $\mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^d[B]), S)) = \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^d[R]), S^M))$. Consequently, it holds that $S_i \in \mathbf{ACC}_i(kb_i \cup \text{app}(br_i(M^d[B]), S))$ for all $1 \leq i \leq n$.

Second, we show that for all $n+1 \leq j \leq n'$ it holds that $S_j \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^d), S))$. Consider $M^d[b_{U'}]$ (with $b_{U'}$ relative to M^d) and $R' = b_{U'}$. Since $M^d[R'] = M^d[b_{U'}] = M[\emptyset] \otimes M'[B_1, B_2]$ and $S^B \in \text{EQ}(M[B_1, B_2])$ hold, it follows by Corollary 3 that there exist S'_1, \dots, S'_n such that $S^{M'} = (S'_1, \dots, S'_n, S_{n+1}, \dots, S_{n+n'}) \in \text{EQ}(M[\emptyset] \otimes M'[B_1, B_2])$, i.e., for all $n+1 \leq j \leq n'$ it holds that $S_j \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^d), S^{M'}))$. Since it holds that $S^{M'} \equiv_{U'} S$ and $b_{U'} \subseteq R' \subseteq B$, Lemma 13 applies and it follows that for all $n+1 \leq j \leq n+n'$ it holds that $\mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^d[B]), S)) = \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^d[R']), S^{M'}))$. Consequently, it holds that $S_j \in \mathbf{ACC}_j(kb_j \cup \text{app}(br_j(M^d[B]), S))$ with $n+1 \leq j \leq n+n'$.

In summary, it holds for every $1 \leq i \leq n+n'$ that S_i is accepted, i.e. $S \in \text{EQ}(M^d)$, hence $(A_1 \cup I(B_2), B_1 \cup I(B_2)) \in D^\pm(M \otimes M')$. \square

Proof of Lemma 3. Observe that $2M = M \otimes M$ and that $2.R = I(R)$ where I is the mapping wrt. $M \otimes M$. The statement then follows directly from Proposition 6. \square

Towards proving that $D_{m, t_{max}}^\pm$ applied on M^\preceq allows to select \subseteq -minimal, preferred diagnoses of M according to \preceq , we use the following lemmas about the set $\mathcal{K}(D_1, D_2)$. Recall that $\mathcal{K}(D_1, D_2)$ is the set of prioritized bridge rules of M^\preceq that represent the candidate diagnosis (D_1, D_2) of M ,

i.e., $\mathcal{K}(D_1, D_2)$ is as follows:

$$\begin{aligned} \mathcal{K}(D_1, D_2) = & \{in_1(r) \mid r \in D_1\} \cup \{\overline{in_1}(r) \mid r \notin D_1\} \cup \\ & \{in_2(r) \mid r \in D_2\} \cup \{\overline{in_2}(r) \mid r \notin D_2\} \end{aligned}$$

The next lemma shows that the set $\mathcal{K}(D_1, D_2)$ is unique for every $D_1, D_2 \subseteq br(M)$.

Lemma 14. *Let M^{\preceq} be a clone encoding, $D_1, D_2 \subseteq br(M)$, and $R = \mathcal{K}(D_1, D_2)$. Then, there exists no $D'_1, D'_2 \subseteq br(M)$ with $(D_1, D_2) \neq (D'_1, D'_2)$ such that $R = \mathcal{K}(D'_1, D'_2)$.*

Proof. Towards contradiction, let $(D_1, D_2) \neq (D'_1, D'_2)$ be such that $\mathcal{K}(D_1, D_2) = \mathcal{K}(D'_1, D'_2)$. By $(D_1, D_2) \neq (D'_1, D'_2)$ follows that either $D_1 \neq D'_1$ or $D_2 \neq D'_2$. Let $D_1 \neq D'_1$ and observe that $\mathcal{K}(D_1, D_2) \cap \{in_1(r) \mid r \in br(M)\} = \{in_1(r) \mid r \in D_1\} \neq \{in_1(r) \mid r \in D'_1\} = \mathcal{K}(D'_1, D'_2) \cap \{in_1(r) \mid r \in br(M)\}$. Consequently $\mathcal{K}(D_1, D_2) \neq \mathcal{K}(D'_1, D'_2)$ which contradicts the assumption. The case $D_2 \neq D'_2$ is similar. It therefore follows that for $R = \mathcal{K}(D_1, D_2)$ no $D'_1, D'_2 \subseteq br(M)$ with $(D_1, D_2) \neq (D'_1, D'_2)$ exists such that $R = \mathcal{K}(D'_1, D'_2)$. \square

The next lemma shows that two sets $\mathcal{K}(D_1, D_2)$ and $\mathcal{K}(D'_1, D'_2)$ are incomparable iff (D_1, D_2) is different from (D'_1, D'_2) .

Lemma 15. *Given $M^{mr(\theta, \mathcal{K}_P)}$ and some $D_1, D_2, D'_1, D'_2 \subseteq br(M)$, let $R = \mathcal{K}(D_1, D_2)$ and let $R' = \mathcal{K}(D'_1, D'_2)$; then $R \subseteq R'$ or $R' \subseteq R$ holds iff $(D_1, D_2) = (D'_1, D'_2)$.*

Proof. Let M be an MCS, $D_1, D_2, D'_1, D'_2 \subseteq br(M)$, $R = \mathcal{K}(D_1, D_2)$, and $R' = \mathcal{K}(D'_1, D'_2)$. Observe that by definition of \mathcal{K} it holds that $|R| = |R'|$. Hence, $R \subseteq R'$ or $R' \subseteq R$ only holds iff $R = R'$. By Lemma 14 it holds that \mathcal{K} is injective, i.e., $R = R'$ iff $(D_1, D_2) = (D'_1, D'_2)$. Consequently, $R \subseteq R'$ or $R' \subseteq R$ holds iff $(D_1, D_2) = (D'_1, D'_2)$. \square

The following lemma shows the relationship between \preceq -preferred diagnoses of M and the prioritized-minimal ones of M^{\preceq} .

Lemma 16. *Given an MCS M and a preference order \preceq , $D \in 2^{br(M)} \times 2^{br(M)}$ is \preceq -preferred iff both (1) $t(D) \in D_m^\pm(M^{\preceq}, br_P)$ and (2) for every $D' \in D_m^\pm(M^{\preceq}, br_P) : D' \subseteq_{br_H} t(D) \Rightarrow D' =_{br_H} t(D)$ hold.*

Proof. “ \Rightarrow ”: Let D be \preceq -preferred, then $D \in D^\pm(M)$ holds.

We first show that $t(D) \in D_m^\pm(M^{\preceq}, br_P)$ holds: by Proposition 4 and the definition of $M^{\preceq} = M^{mr(\theta, \mathcal{K}_P)}$ it holds that $(d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2) \cup \mathcal{K}(D_1, D_2) \cup \{t_{max}\}) \in D^\pm(M^{\preceq}, br_P)$ iff

1. $(D_1 \cup 2.D_1, D_2 \cup 2.D_2) \in D^\pm(2M)$ holds,
2. $\theta(D_1 \cup 2.D_1, D_2 \cup 2.D_2, \mathcal{K}(D_1, D_2) \cup \{t_{max}\})$ holds, and
3. there exists no $(D'_1 \cup 2.D''_1, D'_2 \cup 2.D''_2) \in D^\pm(2M)$ such that (i) $(d1(D'_1 \cup 2.D''_1), d2(D'_2 \cup 2.D''_2) \cup \mathcal{K}') \subset (d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2) \cup \mathcal{K}(D_1, D_2) \cup \{t_{max}\})$ and (ii) $\theta(D'_1 \cup 2.D''_1, D'_2 \cup 2.D''_2, \mathcal{K}')$ holds for some $\mathcal{K}' \subseteq \mathcal{K}_P$.

We show that each of those statements holds:

1. Since $D \in D^\pm(M)$ holds, it follows from Lemma 3 that $(D_1 \cup 2.D_1, D_2 \cup 2.D_2) \in D^\pm(2M)$ holds.
2. Recall that $\theta(R_1, R_2, R_3)$ for $M^\preceq = (2M)^{mr(\theta, \mathcal{K}_p)}$ is defined such that it holds if $R_1 = D_1 \cup 2.D_1$, $R_2 = D_2 \cup 2.D_2$, and $R_3 = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$, hence $\theta(D_1 \cup 2.D_1, D_2 \cup 2.D_2, \mathcal{K}(D_1, D_2) \cup \{t_{max}\})$ holds.
3. Towards contradiction, assume that there exists $(D'_1 \cup 2.D''_1, D'_2 \cup 2.D''_2) \in D^\pm(2M)$ and $\mathcal{K}' \subseteq \mathcal{K}_p$ such that it holds that $(d1(D'_1 \cup 2.D''_1), d2(D'_2 \cup 2.D''_2) \cup \mathcal{K}') \subset (d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2) \cup \mathcal{K}(D_1, D_2) \cup \{t_{max}\})$ and $\theta(D'_1 \cup 2.D''_1, D'_2 \cup 2.D''_2, \mathcal{K}')$ holds. Note that from this it follows that $\mathcal{K}' \subseteq \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$ and from the definition of θ that $\mathcal{K}' \subseteq \mathcal{K}(D'_1, D'_2) \cup \{t_{max}\}$. Hence by Lemma 15, it follows that $(D'_1, D'_2) = (D_1, D_2)$. If $(D''_1, D''_2) = (D_1, D_2)$ then it holds by definition of θ that $t_{max} \in \mathcal{K}'$, i.e., $(d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2) \cup \mathcal{K}(D_1, D_2) \cup \{t_{max}\}) = (d1(D'_1 \cup 2.D''_1), d2(D'_2 \cup 2.D''_2) \cup \mathcal{K}')$ which contradicts that the latter is a proper subset of the former. If $(D''_1, D''_2) \neq (D_1, D_2)$ holds, then by definition of θ it follows that $(D''_1, D''_2) \prec (D_1, D_2) = (D'_1, D'_2)$, which contradicts that (D_1, D_2) is \preceq -preferred. It therefore follows that no such $(D'_1 \cup 2.D''_1, D'_2 \cup 2.D''_2) \in D^\pm(2M)$ exists.

Since all three statements hold, it follows that $t(D) \in D_m^\pm(M^\preceq, br_P)$ holds.

It remains to show that $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D) \Rightarrow T =_{br_H} t(D)$ holds. Assume that $T \in D_m^\pm(M^\preceq, br_P)$ is such that $T \subseteq_{br_H} t(D)$ holds. Then by definition of θ it holds that $T = (d1(T_1 \cup 2.T'_1), d2(T_2 \cup 2.T'_2) \cup \mathcal{K}(T_1, T_2) \cup T_m)$ for some $T_1, T_2, T'_1, T'_2 \subseteq br(M)$ and $T_m \subseteq \{t_{max}\}$. Since $\mathcal{K}(T_1, T_2) \subseteq \mathcal{K}_p$, it holds by $T \subseteq_{br_H} t(D)$ that $\mathcal{K}(T_1, T_2) \subseteq \mathcal{K}(D_1, D_2)$, hence by Lemma 15 it follows that $(T_1, T_2) = (D_1, D_2)$. Since (D_1, D_2) is \preceq -preferred, i.e., there exists no $(D'_1, D'_2) \in D^\pm(M)$ such that $(D'_1, D'_2) \prec (D_1, D_2)$ holds, it follows from the definition of θ that $(T'_1, T'_2) = (D_1, D_2)$ and consequently it holds that $T_m = \{t_{max}\}$. Altogether this means that $T = t(D)$ and thus it holds that $T =_{br_H} t(D)$. It therefore holds that $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D) \Rightarrow T =_{br_H} t(D)$.

“ \Leftarrow ”: Suppose $t(D_1, D_2) \in D_m^\pm(M^\preceq, br_P)$ and $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D) \Rightarrow T =_{br_H} t(D)$ with $D = (D_1, D_2)$ hold. Since $t(D_1, D_2) \in D_m^\pm(M^\preceq, br_P)$ holds, it follows from Proposition 4 that $(d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2)) \in D^\pm(2M)$, hence by Lemma 3 it holds that $(D_1, D_2) \in D^\pm(M)$.

To show that D is \preceq -preferred, consider the set F of diagnoses that are more preferred than D , i.e., $F = \{D'' \in D^\pm(M) \mid D'' \prec D\}$. Towards contradiction, assume that F is non-empty, hence there exists some subset-minimal $D' \in F$, i.e., $D' \in F$ and for all $D'' \in F$ holds $D'' \not\subseteq D'$. Next we consider $(T'_1, T'_2) = (d1(D_1 \cup D'_1), d2(D_2 \cup D'_2) \cup \mathcal{K}(D_1, D_2))$ and observe that $\theta(D_1 \cup D'_1, D_2 \cup D'_2, \mathcal{K}(D_1, D_2))$ holds, because $D' \prec D$ holds.

Since $(D_1, D_2) \in D^\pm(M)$ and $(D'_1, D'_2) \in D^\pm(M)$ it holds that $(D_1 \cup 2.D'_1, D_2 \cup 2.D'_2) \in D^\pm(2M)$. Observe that there exists no other $D'' \subset D'$ with $D \prec D''$ and $D'' \in D^\pm(M)$. Therefore, there exists no $(D''_1, D''_2) \in D^\pm(M)$ such that $(d1(D_1 \cup 2.D''_1), d2(D_2 \cup 2.D''_2) \cup \mathcal{K}(D_1, D_2)) \subset (T'_1, T'_2)$ and $\theta(D_1 \cup 2.D''_1, D_2 \cup 2.D''_2, \mathcal{K}(D_1, D_2))$ both hold. Thus Proposition 4 applies and it follows that $(T'_1, T'_2) \in D_m^\pm(M^\preceq, br_P)$. Observe that $(T'_1, T'_2) \subseteq_{br_H} t(D)$ since $T'_2 \cap br_H = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$ and for $t(D) = (T_1, T_2)$ holds $T_2 \cap br_H = \mathcal{K}(D_1, D_2)$. This directly contradicts that $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D) \Rightarrow T =_{br_H} t(D)$ holds. Thus the set F cannot be non-empty, i.e., there exists no $D' \in D^\pm(M)$ such that $D' \prec D$ holds. Therefore, D is \preceq -preferred. \square

Proof of Theorem 3. Recall that $D^\pm(M, br_P, br_H) = \{D \in D_m^\pm(M, br_P) \mid \forall D' \in D_m^\pm(M, br_P) : D' \subseteq_{br_H} D \Rightarrow D' =_{br_H} D\}$. Hence, $t(D) \in D^\pm(M^\preceq, br_P, br_H)$ holds iff $t(D) \in D_m^\pm(M^\preceq, br_P)$ holds and for every $D' \in D_m^\pm(M^\preceq, br_P)$ it holds that $D' \subseteq_{br_H} t(D) \Rightarrow D' =_{br_H} t(D)$. By Lemma 16 this condition holds iff D is \preceq -preferred. In summary, D is \preceq -preferred iff $t(D) \in D^\pm(M^\preceq, br_P, br_H)$ holds. \square

Proof of Theorem 4. “ \Rightarrow ”: Let $D = (D_1, D_2) \in D_{m, \preceq}^\pm(M)$ hold. Then $D \in D_{\preceq}^\pm(M)$ holds, i.e., D is \preceq -preferred and $D \in D^\pm(M)$ holds. From Lemma 16 we then conclude that $t(D) \in D_m^\pm(M^\preceq, br_P)$ and that the following holds: $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D) \Rightarrow T =_{br_H} t(D)$. By construction of $t(D)$ it furthermore holds that $t_{max} \in t(D)$. Hence it remains to show that $\forall T' \in D_m^\pm(M^\preceq, br_P) : [(\forall T'' \in D_m^\pm(M^\preceq, br_P) : T'' \subseteq_{br_H} T' \Rightarrow T'' =_{br_H} T') \wedge t_{max} \in T'] \Rightarrow [T' \subseteq_{br(M^\preceq) \setminus br_H} t(D) \Rightarrow t(D) =_{br(M^\preceq) \setminus br_H} T']$.

Towards contradiction, assume that $T' \in D_m^\pm(M^\preceq, br_P)$ exists with $(\forall T'' \in D_m^\pm(M^\preceq, br_P) : T'' \subseteq_{br_H} T' \Rightarrow T'' =_{br_H} T') \wedge t_{max} \in T'$ and $T' \subset_{br(M^\preceq) \setminus br_H} t(D)$. Note that the definition of θ and $t_{max} \in T'$ together imply that there exists some $D' = (D'_1, D'_2)$ with $D'_1, D'_2 \subseteq br(M)$ such that $T' = t(D')$ holds. Further note that $T' = t(D')$ satisfies all conditions of Lemma 16, thus it holds that $D' \in D^\pm(M)$ and that D' is \preceq -preferred.

From $T' = t(D') \subset_{br(M^\preceq) \setminus br_H} t(D)$ it follows that $(d1(D'_1 \cup 2.D'_1), d2(D'_2 \cup 2.D'_2)) \subset (d1(D_1 \cup 2.D_1), d2(D_2 \cup 2.D_2))$ and since $d1, d2$, and $2.$ are bijective, it holds that $(D'_1, D'_2) \subset (D_1, D_2)$. Since D' is \preceq -preferred, this contradicts that D is subset-minimal among all \preceq -preferred diagnoses, i.e., it contradicts that $D \in D_{m, \preceq}^\pm(M)$. Therefore no such T' can exist and it holds that $t(D) \in D_{m, t_{max}}^\pm(M^\preceq, br_P, br_H)$.

“ \Leftarrow ”: Let $t(D_1, D_2) \in D_{m, t_{max}}^\pm(M^\preceq, br_P, br_H)$ hold. Since $t(D_1, D_2) \in D_m^\pm(M^\preceq, br_P)$ and $t_{max} \in t(D_1, D_2)$ hold, it follows from Lemma 16 that $D = (D_1, D_2) \in D^\pm(M)$ and that D is \preceq -preferred. It remains to show that D is subset-minimal among diagnoses in $D_{\preceq}^\pm(M)$.

Towards contradiction, assume that there exists $D' \in D_{\preceq}^\pm(M)$ with $D' \subset D$. Since D' is \preceq -preferred and $D' \in D^\pm(M)$ holds, it follows from Lemma 16 that $t(D') \in D_m^\pm(M^\preceq, br_P)$ and $\forall T \in D_m^\pm(M^\preceq, br_P) : T \subseteq_{br_H} t(D') \Rightarrow T =_{br_H} t(D')$ holds. Let $T' = t(D')$. Then it holds for T' that $(\forall T'' \in D_m^\pm(M, br_P) : T'' \subseteq_{br_H} T' \Rightarrow T'' =_{br_H} T') \wedge t_{max} \in T'$. Let $T = t(D)$. Because $d1, d2$, and $2.$ are bijective and $D' \subset D$, it follows that $[T' \subseteq_{(br(M) \setminus br_H)} T \Rightarrow T =_{br(M) \setminus br_H} T']$ does not hold. This contradicts that $t(D_1, D_2) \in D_{m, t_{max}}^\pm(M^\preceq, br_P, br_H)$ holds. Therefore no such D' exists and it holds that D is subset-minimal among $D_{\preceq}^\pm(M)$, i.e., $D \in D_{m, \preceq}^\pm(M)$ holds. \square

A.4 Proofs of Section 6

Proof of Theorem 5. In the remainder of this proof we assume \mathbf{C} to be the computational complexity of MCSD_m .

Membership: In the following we give a polynomial-time reduction \leq_m^p from MCSDP_m to MCSD_m . Given an instance of MCSDP_m , i.e., given an MCS M , a set $br_P \subseteq br(M)$, and a candidate diagnosis $D \in 2^{br(M)} \times 2^{br(M)}$, we define \leq_m^p such that

$$(M, br_P, D) \mapsto \begin{cases} (M, D) & \text{if } D_1 \cap br_P = \emptyset = D_2 \cap br_P \text{ where } D = (D_1, D_2) \\ (M_\perp, (\emptyset, \emptyset)) & \text{otherwise} \end{cases}$$

where $M_\perp = (C_\perp)$, $C_\perp = (L_{\Sigma}^{asp}, kb_\perp, br_\perp)$, $br_\perp = \{(1:a) \leftarrow \top.\}$, and $kb_\perp = \{\perp \leftarrow a.\}$ is such that $(\emptyset, \emptyset) \notin D_m^\pm(M_\perp)$. Intuitively, the reduction checks whether D contains bridge rules from br_P

and if so, maps to an instance which is not in MCSD_m . If D contains no bridge rules from br_P , then \leq_m^p simply drops br_P . Since the check whether D contains bridge rules of br_P is possible in polynomial time, \leq_m^p is a polynomial-time many-one reduction.

It remains to show that indeed (M, br_P, D) is a yes-instance of MCSDP_m iff $\leq_m^p(M, br_P, D)$ is a yes-instance of MCSD_m .

“ \Rightarrow ”: Let (M, br_P, D) be a yes-instance of MCSDP_m , i.e., $D \in D_m^\pm(M, br_P)$ holds. Then, $D = (D_1, D_2)$ is such that $D_1 \cap br_P = \emptyset = D_2 \cap br_P$, hence $\leq_m^p(M, br_P, D) = (M, D)$. By Proposition 1 it holds that $D_m^\pm(M, br_P) \subseteq D_m^\pm(M)$, hence it follows that $D \in D_m^\pm(M)$ holds, i.e., (M, D) is a yes-instance of MCSD_m .

“ \Leftarrow ”: Let $\leq_m^p(M, br_P, D)$ be a yes-instance of MCSD_m . Note that it cannot be the case that $\leq_m^p(M, br_P, D) = (M_\perp, (\emptyset, \emptyset))$, because $(\emptyset, \emptyset) \notin D_m^\pm(M_\perp)$ contradicts that $\leq_m^p(M, br_P, D)$ is a yes-instance of MCSD_m . Consequently, it holds that $\leq_m^p(M, br_P, D) = (M, D)$ and thus $D = (D_1, D_2)$ is such that $D_1 \cap br_P = \emptyset = D_2 \cap br_P$. Furthermore, $D \in D_m^\pm(M)$ holds, thus it follows that $D \in D_m^\pm(M, br_P)$ holds. Assume that $D \notin D_m^\pm(M, br_P)$ holds. Then there exists $D' \subset D$ such that $D' \in D_m^\pm(M, br_P)$ holds. By Proposition 1 then follows that $D' \in D_m^\pm(M)$, which contradicts that $D \in D_m^\pm(M)$. Therefore no such D' exists and it follows that $D \in D_m^\pm(M, br_P)$ holds.

Since \leq_m^p is a polynomial reduction from MCSDP_m to MCSD_m , it follows that the computational complexity of MCSDP_m is in \mathbf{C} , i.e., the same complexity class where MCSD_m is in.

Hardness: Let $D \in D_m^\pm(M)$ be hard for some complexity class \mathbf{C} . Observe that by definition of diagnoses with protected bridge rules, it holds that $D \in D_m^\pm(M)$ is true iff $D \in D_m^\pm(M, \emptyset)$ is true. Since deciding whether $D \in D_m^\pm(M)$ is \mathbf{C} -hard, it thus follows that deciding whether $D \in D_m^\pm(M, br_P)$ also is \mathbf{C} -hard. \square

Proof of Lemma 4. “ \Rightarrow ”: Let $(M, (D_1, D_2), br_P, br_H)$ be a yes-instance of MCSDPH , i.e., it holds that $(D_1, D_2) \in D_m^\pm(M, br_P, br_H)$. We have to show that $(D'_1, D'_2) \in D_m^\pm(M' \otimes M, br_{P'})$ holds.

From $(D_1, D_2) \in D_m^\pm(M, br_P, br_H)$ and (10) it follows that $(D_1, D_2) \in D_m^\pm(M, br_P)$ holds.

By Proposition 1 it then holds that $(D_1, D_2) \in D_m^\pm(M)$, thus there exists $S = (S_1, \dots, S_n)$ with $S \in \text{EQ}(M[D_1, D_2])$. We now show that $(d1(D_1 \cap br_H), d2(D_2 \cap br_H)) \in D_m^\pm(M', br_{P'})$ holds; to that end consider the belief state $S' = (S_1, \dots, S_n, S_{n+1}, S_{n+2})$ where

$$\begin{aligned} S_{n+1} &= \{\text{removed}_r \mid r \in r \in D_1\} \cup \{\text{uncond}_r \mid r \in D_2\} \\ S_{n+2} &= \{\text{not_removed}_r \mid r \in D_1 \setminus br_H\} \cup \{\text{uncond}_r \mid r \in D_2 \setminus br_H\}. \end{aligned}$$

By construction of C_{n+2} , it holds that

$$S_{n+2} \in \mathbf{ACC}_{n+2}(kb_{n+2} \cup \text{app}(br_{n+2}(M'[d1(D_1 \cap br_H), d2(D_2 \cap br_H)]), S')).$$

Consider the set of applicable bridge rules of C_{n+1} under S' and the candidate diagnosis $(D_1 \cap br_H, D_2 \cap br_H)$ (where $R_{reg} = (br(M) \setminus br_P \setminus br_H)$):

$$\begin{aligned}
 & \{\varphi(r) \mid r \in app(br_{n+1}(M'[d1(D_1 \cap br_H), d2(D_2 \cap br_H)]), S')\} \\
 &= \{not_removed_r \mid r \in br(M), r \notin R_{reg}, r \notin D_1 \setminus br_H\} \\
 &\quad \cup \{not_removed_r \mid r \in br(M), r \in R_{reg}, r \notin D_1 \cap br_H\} \\
 &\quad \cup \{uncond_r \mid r \in br(M), r \notin R_{reg}, r \in D_2 \cap br_H\} \\
 &\quad \cup \{uncond_r \mid r \in br(M), r \in R_{reg}, r \in D_2 \setminus br_H\} \\
 &= \{not_removed_r \mid r \in br(M), r \notin D_1\} \\
 &\quad \cup \{uncond_r \mid r \in br(M), r \in D_2\} \\
 &=: H
 \end{aligned}$$

Since $S_{n+1} = \{removed_r \mid r \in r \in D_1\} \cup \{uncond_r \mid r \in D_2\}$ and $\theta(D_1, D_2, \emptyset)$ holds, it follows from the definition of C_{n+1} (cf. Definition 12 and Lemma 7) that $S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \cup H)$ holds.

Following the reasoning in Lemma 6 it is then possible to construct a proof showing that for all $1 \leq i \leq n$ it holds that

$$app(br_i(M[D_1, D_2]), S) = app(br_i(M'[d1(D_1 \cap br_H), d2(D_2 \cap br_H)]), S').$$

Since the semantics \mathbf{ACC}_i and knowledge base kb_i of each context C_i are the same in M and M' , it then follows from $S \in \text{EQ}(M[D_1, D_2])$ that for all $1 \leq i \leq n$ holds $S_i \in \mathbf{ACC}_i(kb_i \cup app(br_i(M'[d1(D_1 \cap br_H), d2(D_2 \cap br_H)]), S'))$.

In summary, it holds that $S' \in \text{EQ}(M'[d1(D_1 \cap br_H), d2(D_2 \cap br_H)])$.

Since $(D_1, D_2) \in D^\pm(M)$ holds, we then conclude from Proposition 6 that $(I(D_1) \cup d1(D_1 \cap br_H), I(D_2) \cup d2(D_2 \cap br_H)) = (D'_1, D'_2) \in D^\pm(M' \otimes M)$ holds. Note that $D'_1 \cap br_{P''} = \emptyset = D'_2 \cap br_{P''}$, hence $(D'_1, D'_2) \in D^\pm(M' \otimes M, br_{P''})$ also holds.

It remains to show that $(D'_1, D'_2) \in D_m^\pm(M' \otimes M, br_{P''})$. Towards contradiction assume that there exists $(T_1, T_2) \in D^\pm(M' \otimes M, br_{P''})$ with $(T_1, T_2) \subset (D'_1, D'_2)$, i.e., by construction of $M' \otimes M$ it either is the case that $(T_1 \cap I(br(M)), T_2 \cap I(br(M))) \subset (D'_1 \cap I(br(M)), D'_2 \cap I(br(M)))$ holds or $(T_1 \cap br_{M'}, T_2 \cap br_{M'}) \subset (D'_1 \cap br_{M'}, D'_2 \cap br_{M'})$ holds.

In the former case, Proposition 6 implies that $(I^{-1}(T_1 \cap I(br(M))), I^{-1}(T_2 \cap I(br(M)))) \in D^\pm(M)$; furthermore, since $(D'_1 \cap I(br(M)), D'_2 \cap I(br(M))) = (D_1, D_2)$ it holds that $(I^{-1}(T_1 \cap I(br(M))), I^{-1}(T_2 \cap I(br(M)))) \subset (D_1, D_2)$. This contradicts that $(D_1, D_2) \in D^\pm(M, br_P, br_H)$.

In the latter case, i.e., $(T_1 \cap br_{M'}, T_2 \cap br_{M'}) \subset (D'_1 \cap br_{M'}, D'_2 \cap br_{M'})$, it holds that $(T_1 \cap br_{M'}, T_2 \cap br_{M'}) \subset (br_H, br_H)$ since all other bridge rules of $br_{M'}$ are contained in $br_{P''}$. Let S be a witnessing equilibrium, i.e., let $S = (S_1, \dots, S_{n+2}) \in \text{EQ}(M'[(T_1 \cap br_{M'}, T_2 \cap br_{M'})])$ hold. Consider the modifications of bridge rules in $br(M) \setminus br_P \setminus br_H$ which are represented by S , i.e., consider $T'_1 = \{r \in br(M) \setminus br_P \setminus br_H \mid not_removed_r \notin S_{n+2}\}$ and $T'_2 = \{r \in br(M) \setminus br_P \setminus br_H \mid uncond_r \in S_{n+2}\}$. It holds that $((T_1 \cap br_{M'}) \cup T'_1, (T_2 \cap br_{M'}) \cup T'_2)$ is a candidate diagnosis of M . Since $S \in \text{EQ}(M'[(T_1 \cap br_{M'}, T_2 \cap br_{M'})])$ holds and M' stems from $M^{mr(\theta, \mathcal{K}_P)}$, one can show using Lemma 6 that $((T_1 \cap br_{M'}) \cup T'_1, (T_2 \cap br_{M'}) \cup T'_2) \in D^\pm(M, br_P)$ holds. This contradicts that $(D_1, D_2) \in D^\pm(M, br_P, br_H)$, because $((T_1 \cap br_{M'}) \cup T'_1, (T_2 \cap br_{M'}) \cup T'_2) \subset_{br_H} (D_1, D_2)$.

Therefore, no such (T_1, T_2) exists and it holds that $(D'_1, D'_2) \in D_m^\pm(M' \otimes M, br_{P''})$.

“ \Leftarrow ”: We prove the converse, i.e., we assume that $(M, (D_1, D_2), br_P, br_H)$ is not a yes-instance of MCSDPH and show that $\text{DPH2DP}_m(M, (D_1, D_2), br_P, br_H) = (M' \otimes M, (D'_1, D'_2), br_{P''})$ also is not a yes-instance of MCSDP $_m$. By assumption it thus holds that $(D_1, D_2) \notin D^\pm(M, br_P, br_H)$ holds. From the def. of $D^\pm(M, br_P, br_H)$ we then obtain that either (i) $(D_1, D_2) \notin D_m^\pm(M, br_P)$ holds or (ii) that there exists $(D'_1, D'_2) \in D_m^\pm(M, br_P)$ with $(D'_1, D'_2) \subset_{br_H} (D_1, D_2)$.

In case (i) $(D_1, D_2) \notin D_m^\pm(M, br_P)$, hence by Proposition 6 it holds that $(I(D_1) \cup d1(D_1 \cap br_H), I(D_2) \cup d2(D_2 \cap br_H)) \notin D_m^\pm(M \otimes M', br_{P''})$.

In case (ii) $(D'_1, D'_2) \in D_m^\pm(M, br_P)$ with $(D'_1, D'_2) \subset_{br_H} (D_1, D_2)$. W.l.o.g. we assume that (D'_1, D'_2) is minimal wrt. \subset_{br_H} , i.e., there exists no $(D''_1, D''_2) \in D_m^\pm(M, br_P)$ with $(D''_1, D''_2) \subset_{br_H} (D'_1, D'_2)$. This means that $(M, (D'_1, D'_2), br_P, br_H)$ is a yes-instance of MCSDPH. We can further assume that $(D_1, D_2) \in D_m^\pm(M, br_P)$ from (i).

Now consider $(T_1, T_2) = (I(D_1) \cup d1(D'_1 \cap br_H), I(D_2) \cup d2(D'_2 \cap br_H))$. Since $(M, (D'_1, D'_2), br_P, br_H)$ is a yes-instance of MCSDPH, the “ \Rightarrow ” direction above can be applied to it; this yields that $(d1(D'_1 \cap br_H), d2(D'_2 \cap br_H)) \in D_m^\pm(M', br_{P'})$ holds. Applying Proposition 6 and the fact that $T_1 \cap br_{P''} = \emptyset = T_2 \cap br_{P''}$ then implies that $(T_1, T_2) \in D_m^\pm(M' \otimes M, br_{P''})$ holds. Note that $(T_1, T_2) \subset (D'_1, D'_2)$ holds which in turn implies that $(D'_1, D'_2) \notin D_m^\pm(M' \otimes M, br_{P''})$ holds. In other words, $(M' \otimes M, (D'_1, D'_2), br_{P''})$ is not a yes-instance of MCSDP $_m$.

In all cases, we showed that $\text{DPH2DP}_m(M, (D_1, D_2), br_P, br_H)$ is not a yes-instance of MCSDP $_m$, which concludes the “ \Leftarrow ” direction of the proof.

In summary, we showed that $(M, (D_1, D_2), br_P, br_H)$ is a yes-instance of MCSDPH iff $(M' \otimes M, (D'_1, D'_2), br_{P''}) = \text{DPH2DP}_m(M, (D_1, D_2), br_P, br_H)$ is a yes-instance of MCSDP $_m$, i.e., DPH2DP_m is a reduction from MCSDPH to MCSDP $_m$. Since $(M' \otimes M, (D'_1, D'_2), br_{P''})$ can be computed in time linear in the size of $(M, (D_1, D_2), br_P, br_H)$, it furthermore holds that DPH2DP_m is a polynomial-time reduction. \square

Proof of Theorem 6. Membership: By Lemma 4 it holds that DPH2DP_m is a polynomial-time reduction from MCSDPH to MCSDP $_m$, hence membership immediately follows.

Hardness: Let M' and D' be any MCS and candidate diagnosis, respectively, used for showing hardness of MCSDP $_m$ for \mathbf{C} (i.e., M' is the result of the reduction showing \mathbf{C} -hardness of MCSDP $_m$ and D' is the diagnosis resulting from the reduction of M'). Now pick $br_{P'} = br_{H'} = \emptyset$.

By definition, it holds for all M, br_P, br_H and D , that $D \in D^\pm(M, br_P, br_H)$ implies $D \in D_m^\pm(M, br_P)$ which in turn implies $D \in D_m^\pm(M)$. Therefore, $D' \in D^\pm(M', br_{P'}, br_{H'})$ implies that $D' \in D_m^\pm(M')$ holds. Furthermore, since $br_{P'} = br_{H'} = \emptyset$ it also follows from the definition of prioritized-minimal diagnosis and protected diagnosis that $D' \in D_m^\pm(M')$ implies $D' \in D^\pm(M', br_{P'}, br_{H'})$. In summary, $D' \in D_m^\pm(M')$ holds iff $D' \in D^\pm(M', br_{P'}, br_{H'})$ holds. Therefore MCSDPH also is \mathbf{C} -hard. \square

Proof of Lemma 5. For membership, we give a reduction \leq_m^p from $\text{MCSDPH}_{t_{max}}$ to MCSDPH as follows:

$$(M, D, br_P, br_H, t_{max}) \mapsto \begin{cases} (M, D, br_P, br_H) & \text{if } D = (D_1, D_2), t_{max} \in D_2 \\ (M_\perp, (\emptyset, \emptyset), br_{M_\perp}, \emptyset) & \text{otherwise} \end{cases}$$

where M_\perp is the inconsistent MCS from the proof of Theorem 5, i.e., $(M_\perp, (\emptyset, \emptyset), br_{M_\perp}, \emptyset)$ is not a yes-instance of MCSDPH since the MCS is inconsistent but all its bridge rules are protected. Clearly, \leq_m^p is a polynomial-time reduction.

“ \Rightarrow ”: Let $(M, D, br_P, br_H, t_{max})$ be a yes-instance of $\text{MCSDPH}_{t_{max}}$, this means that $D \in D^\pm(M, br_P, br_H)$ and $t_{max} \in D_2$ with $D = (D_1, D_2)$ hold. Then $D \in D^\pm(M, br_P, br_H)$ also holds, i.e., (M, D, br_P, br_H) is a yes-instance of MCSDPH .

“ \Leftarrow ”: Let $(M, D, br_P, br_H, t_{max})$ be not a yes-instance of $\text{MCSDPH}_{t_{max}}$, i.e., let it not be the case that $D \in D^\pm(M, br_P, br_H)$ and $t_{max} \in D_2$ with $D = (D_1, D_2)$ both hold. In case $t_{max} \notin D_2$ it holds that $(\emptyset, \emptyset) \notin D^\pm(M_\perp, br_{M_\perp}, \emptyset)$ since M_\perp is inconsistent but all its bridge rules are protected, i.e., \leq_m^p maps to a no-instance of MCSDPH . In case $t_{max} \in D_2$ holds, it follows that $D \in D^\pm(M, br_P, br_H)$ does not hold by the assumption. Therefore (M, D, br_P, br_H) is not a yes-instance of MCSDPH . Hence in all cases, $\leq_m^p(M, br_P, br_H, t_{max})$ is not a yes-instance of MCSDPH . \square

Proof of Theorem 8. Since QBF problems correlate to complexity classes in the polynomial hierarchy, we reduce (different) QBF problems to $\text{MCSD}_{\text{MPREF}}$ to prove the following hardness statements: if $\mathcal{CC}(M)$ is hard for Σ_i^P (Π_i^P) then $\text{MCSD}_{\text{MPREF}}$ is hard for Π_{i+1}^P (Π_{i+2}^P) with $i \geq 0$; and $\text{MCSD}_{\text{MPREF}}$ is Π_2^P -hard even if $\mathcal{CC}(M)$ and deciding whether $D' \preceq D''$ hold are both in \mathbf{P} .

QBF: A formula G is a quantified Boolean formula (QBF) if it is of the form $Q_1 \vec{X}_1 \dots Q_n \vec{X}_n : F$ where for each $1 \leq i \leq n$, $Q_i \in \{\forall, \exists\}$ is a quantifier, \vec{X}_i is a set of Boolean variables, and F is a propositional formula over the set of variables $V = \bigcup_{i \in \{1, \dots, n\}} \vec{X}_i$. We assume that the quantifiers alternate, i.e., $Q_i \neq Q_{i+1}$ for all $1 \leq i < n$. \mathbf{QBF}_k denotes all QBF with $k \geq 1$ quantifiers, $\mathbf{QBF}_{2, \forall}$ denotes all QBF with 2 quantifiers and $Q_1 = \forall$, and $\mathbf{QBF}_{k, \forall}$ denotes all QBF with k quantifiers and $Q_1 = \forall$. Given a formula G in $\mathbf{QBF}_{k, \forall}$ of the form as above, we denote the subformula $\forall \vec{X}_3 \dots Q_k \vec{X}_k : F$ by $\text{rem}_{2, \forall}(G)$. Note that for a $\mathbf{QBF}_{2, \forall}$ formula G , $\text{rem}_{2, \forall}(G) = F$. For readability and simplicity, we denote the variables \vec{X}_1 by \vec{X} and the variables \vec{X}_2 by \vec{Y} .

A valuation is an assignment of variables to $\{\top, \perp\}$, we denote a assignment to variables X by $V_X : X \rightarrow \{\top, \perp\}$. Let $\psi[x/t]$ denote the substitution of the propositional variable x by $t \in \{\top, \perp\}$ in ψ . Then the substitution by an assignment V_X over $X = \{x_1, \dots, x_k\}$ is $\psi[x_1/V_X(x_1)] \dots [x_k/V_X(x_k)]$, with shorthand notation $\psi[V_X]$. The semantics of QBF is inductively given in terms of valuations. A QBF G of the above form with n quantifiers *evaluates to true* if: G is quantifier-free ($n = 0$) and G is a true propositional formula; if $Q_1 = \forall$ ($Q_1 = \exists$) and for all (for some) valuations V_{X_1} it holds that $G'[V_{X_1}]$ evaluates to true where $G' = Q_2 \vec{X}_2 \dots Q_n \vec{X}_n : F$. Note that G being in \mathbf{QBF}_k implies that $G'[V_{X_1}]$ is in \mathbf{QBF}_{k-1} for any valuation V_{X_1} .

Reduction: We define a logic $L_\Sigma^{qbf} = (\mathbf{KB}, \mathbf{BS}, \mathbf{ACC})$ for QBFs over a set of variables Σ , which enables QBF as a query language (cf. Egly, Eiter, Tompits, & Woltran, 2000) and is based on the idea of combining a database, under the closed-world assumption, with a theory (cf. Bonatti & Eiter, 1995). Formally, $\mathbf{KB} = 2^Q$ with Q being the set of quantified Boolean formulas that can be built over Σ , i.e., each $kb \in \mathbf{KB}$ is a set of QBF; $\mathbf{BS} = \{\emptyset\}$, i.e., the only belief set is the empty set indicating evaluation to true; and $\mathbf{ACC}(kb)$ intuitively takes the conjunction CF of all formulas in kb except those that are unit, i.e., of the form (χ) with some variable χ , creates an assignment V from the unit clauses, and accepts the single belief set if and only if $CF[V]$ evaluates to true. Formally, $\mathbf{ACC}(kb) = \{\emptyset\}$ if the QBF $CF[V]$ evaluates to true with $CF = \bigwedge_{f \in \{f \in kb \mid f \text{ is not unit}\}} f$ and valuation $V_Z : Z \rightarrow \{\top, \perp\}$ such that $V(\chi) = \top$ if $(\chi) \in kb$ and $V(\chi) = \perp$ otherwise where Z is the set of free (un-quantified) variables of CF .

Note that if each $f \in kb$ is in \mathbf{QBF}_k then $CF[V]$ is in \mathbf{QBF}_k since $CF[V]$ contains no more quantifier alternations than any $f \in kb$. By that, the computational complexity of evaluating \mathbf{ACC}

is in Σ_k^P and Π_k^P if each $f \in kb$ is in $\mathbf{QBF}_{k,\exists}$ and $\mathbf{QBF}_{k,\forall}$, respectively. Also note that the construction of $CF[V]$ is possible in linear time.

We now define an MCS M^G whose single context utilizes the evaluation of $rem_{2,\forall}(G)$ given a valuation of all variables in \vec{X} and \vec{Y} . Given a \mathbf{QBF}_k -formula G with $k \geq 2$, and \vec{X} , \vec{Y} , and $rem_{2,\forall}(G)$ as above. Let br_1^X and br_1^Y be defined as follows:

$$br_1^X = \bigcup_{x \in \vec{X}} \{(1:(x)) \leftarrow \top., (1:(\bar{x})) \leftarrow \top.\} \quad br_1^Y = \bigcup_{y \in \vec{Y}} \{(1:(y)) \leftarrow \top., (1:(\bar{y})) \leftarrow \top.\}$$

Then $M^G = (C_1)$ with the context $C_1 = (L_{\Sigma}^{qbf}, kb_1, br_1)$, $br_1 = br_1^X \cup br_1^Y$, and kb_1 as follows:

$$kb_1 = \left\{ \begin{aligned} & \left(\bigwedge_{\chi \in \vec{X} \cup \vec{Y}} (\chi \leftrightarrow \neg \bar{\chi}) \wedge \bigwedge_{\chi \in \vec{X} \cup \vec{Y}} (\chi \leftrightarrow \chi') \wedge rem_{2,\forall}(G^t) \right) \\ & \vee \left(\bigwedge_{\chi \in \vec{X}} (\neg \chi \wedge \neg \bar{\chi}) \right) \\ & \vee \left(\bigwedge_{\chi \in \vec{X}} (\chi \leftrightarrow \neg \bar{\chi}) \wedge \bigwedge_{\chi \in \vec{Y}} (\chi \wedge \bar{\chi}) \right) \end{aligned} \right\}$$

where G^t is equal to G except that every variable $\chi \in \vec{X} \cup \vec{Y}$ is replaced with a new variable χ' . Intuitively, C_1 evaluates the remainder of G if a consistent valuation in terms of χ and $\bar{\chi}$ is given for $\vec{X} \cup \vec{Y}$ (first line) and it becomes true also for two other cases: (i) no value for any variable in \vec{X} is given (second line), and (ii) a consistent valuation for \vec{X} is given and all values for \vec{Y} are present (third line). Note that br_1 and kb_1 are both polynomial (even linear) in the size of G .

Notation: For a set R of bridge rules $\varphi(R)$ denotes the set of head-formulas of the bridge rules of R , i.e., $\varphi(R) = \{\varphi(r) \mid r \in R\}$. For $H \subseteq \varphi(br_1)$ we say that H is *consistent wrt. a set of variables* X if for all $\chi \in X$ it holds that $(\chi) \in H$ iff $(\bar{\chi}) \notin H$. We call H *consistent* if it is consistent wrt. $\vec{X} \cup \vec{Y}$. If H is consistent, then the corresponding valuation $V_H : \vec{X} \cup \vec{Y} \rightarrow \{\top, \perp\}$ is $V_H(\chi) = \top$ if $(\chi) \in H$ and $V_H(\chi) = \perp$ otherwise.

One can show that the semantics of C_1 is as follows for any $H \subseteq \varphi(br_1)$:

$$\mathbf{ACC}(kb_1 \cup H) = \begin{cases} \{\emptyset\} & \text{if } H \text{ is consistent and } rem_{2,\forall}(G)[V_H] \text{ evaluates to true, or} \\ & \text{if } H \cap \varphi(br_1^X) = \emptyset, \text{ or} \\ & \text{if } H \text{ is consistent wrt. } \vec{X} \text{ and } H \supseteq \varphi(br_1^Y), \\ \emptyset & \text{otherwise.} \end{cases}$$

Intuitively, the above holds for the following reasons: $kb_1 \cup H$ is such that all formulas $f \in H$ are unit clauses, the variables in each clause are distinct, and each variable inside a unit clause is free wrt. kb_1 , because the only quantifiers occur in $rem_{2,\forall}(G^t)$ which does quantify over χ' but neither quantifies χ nor $\bar{\chi}$, for any $\chi \in \vec{X} \cup \vec{Y}$.

We use a certain diagnosis of M^G , namely $D_{valid} = (br_1^X, \emptyset)$, to indicate whether G evaluates to true. In order to obtain that $D_{valid} \in D_{m,\preceq}^{\pm}(M^G)$ iff G evaluates to true, we use the following preference order \preceq :

$$\begin{aligned} (D_1, D_2) \preceq (D'_1, D'_2) \text{ holds iff } & D_2 = D'_2 = \emptyset, (D_1, D_2) \neq D_{valid} \neq (D'_1, D'_2), \\ & D_1 \cap br_1^X = D'_1 \cap br_1^X, \text{ and } D_1 \cap br_1^Y \subseteq D'_1 \cap br_1^Y \text{ all hold, or} \\ & (D_1, D_2) = (D'_1, D'_2). \end{aligned}$$

Note that \preceq is transitive and reflexive, i.e., indeed it is a preference order.

Correctness of Reduction: We now show for any QBF in $\mathbf{QBF}_{k,\forall}$ with $k \geq 2$ that $D_{\text{valid}} \in D_{m,\preceq}^{\pm}(M^G)$ holds iff G evaluates to true. In some abuse of notation, in the following we write $M[\vec{D}]$ to denote the MCS obtained from modifying M according to a candidate diagnosis D , i.e., $M[D]$ with $D = (D_1, D_2)$ here denotes $M[D_1, D_2]$. Furthermore, if $H \subseteq br(M^G)$ is consistent wrt. a set Z of variables, we denote by V_Z^H the corresponding valuation, i.e., $V_Z^H(\chi) = \top$ iff $\chi \in \varphi(H)$ and $V_Z^H(\chi) = \perp$ iff $\bar{\chi} \in \varphi(H)$ with $\chi \in Z$.

“ \Rightarrow ”: Let $D_{\text{valid}} \in D_{m,\preceq}^{\pm}(M^G)$ hold. Towards contradiction, assume that G does not evaluate to true, i.e., there exists a valuation V_X for \vec{X} such that no valuation V_Y for \vec{Y} makes $rem_{2,\forall}(G)[V_X \cup V_Y]$ evaluate to true. Let $R \subseteq br_1^X$ be such that $V_X^{\varphi(R)} = V_X$ and consider the diagnosis $D = (br_1^X \setminus R, \emptyset)$. Let $H = \{\varphi(r) \mid r \in app(br_1(M^G[D]), S_\emptyset)\}$ and observe that H is consistent wrt. \vec{X} since R is consistent. Since $(br_1^X \setminus R) \cap br_1^Y = \emptyset$, it follows that $H \cap \varphi(br_1^Y) = \varphi(br_1^Y)$ and it thus holds that $\{\emptyset\} \in \mathbf{ACC}(kb_1 \cup H)$, i.e., S_\emptyset is an equilibrium of $M^G[D]$, hence $D \in D^{\pm}(M^G)$. Further note that $D \subset D_{\text{valid}}$ holds. Since $D_{\text{valid}} \in D_{m,\preceq}^{\pm}(M^G)$ and $D \subset D_{\text{valid}}$, it follows that $D \notin D_{\preceq}^{\pm}(M^G)$ holds; i.e. there exists a diagnosis $D' \in D^{\pm}(M^G)$ such that $D' \prec D$ holds, which implies that $D' \neq D$.

Let $D' = (D'_1, D'_2)$ and $D = (D_1, D_2)$; from the definition of \preceq and the fact that $D' \neq D$ we obtain that $D'_2 = \emptyset$, $D' \neq D_{\text{valid}}$, $D'_1 \cap br_1^X = D_1 \cap br_1^X$, and $D'_1 \cap br_1^Y \subseteq D_1 \cap br_1^Y$ all hold. Let $H' = \{\varphi(r) \mid r \in app(br_1(M^G[D']), S_\emptyset)\}$ and observe that H' is consistent wrt. \vec{X} since D_1 is consistent wrt. \vec{X} and $D'_1 \cap br_1^X = D_1 \cap br_1^X$. Since $D' \neq D$ holds, it is the case that $D'_1 \cap br_1^Y \subset D_1 \cap br_1^Y$ and thus $D_1 \cap br_1^Y \neq \emptyset$, i.e., $H \cap \varphi(br_1^Y) \neq \varphi(br_1^Y)$. This contradicts with $H \cap \varphi(br_1^Y) = \varphi(br_1^Y)$ established earlier. Therefore no such D exists and consequently no valuation V_X exists such that all valuations V_Y make $rem_{2,\forall}(G)[V_X \cup V_Y]$ not evaluate to true, i.e., G evaluates to true.

“ \Leftarrow ”: Let G evaluate to true, i.e., for every valuation of \vec{X} there exists a valuation of \vec{Y} such that $rem_{2,\forall}(G)[V_X \cup V_Y]$ evaluates to true. Observe that $br_1(M^G[D_{\text{valid}}]) = br_1^Y$, hence $H = app(br_1(M^G[D_{\text{valid}}]), S_\emptyset)$ is such that $H \cap \varphi(br_1^X) = \emptyset$, thus $\mathbf{ACC}(kb_1 \cup H) = \{\emptyset\}$ and S_\emptyset is a witnessing equilibrium of $D_{\text{valid}} \in D^{\pm}(M^G)$. Furthermore, since D_{valid} is, by definition of \preceq , in no relation to any other candidate diagnosis except itself, it thus follows that $D_{\text{valid}} \in D_{\preceq}^{\pm}(M^G)$.

It remains to show that D_{valid} is subset-minimal among all diagnoses in $D_{\preceq}^{\pm}(M^G)$. Consider any $D' \subset D_{\text{valid}}$, i.e., $D' = (D'_1, \emptyset)$ where $D'_1 \subset br_1^X$. Recall that D' is not a diagnosis, if there exists no witnessing equilibrium; since S_\emptyset is the only belief state of M^G , it follows that D' is a diagnosis if and only if S_\emptyset is an equilibrium of $M^G[D']$. In the following, let $H' = app(br_1(M^G[D']), S_\emptyset)$. Since $D'_1 \subset br_1^X$ holds, it follows that $H' \supseteq \varphi(br_1^Y)$, because for any $r \in br_1^Y$ it holds that $body(r) = \{\top\}$, i.e., r is applicable in any belief state. Since $H' \supseteq \varphi(br_1^Y)$ holds, it cannot be the case that H' is consistent wrt. \vec{Y} ; thus H' is not consistent. Furthermore, by $D'_1 \subset br_1^X$ it follows that $H' \cap \varphi(br_1^X) \neq \emptyset$. By the definition of \mathbf{ACC} it then follows that D' only is a diagnosis if H' is consistent wrt. \vec{X} .

Assume that H' is consistent wrt. \vec{X} then $V_X^{H'}$ is a consistent valuation for variables in \vec{X} . Since G evaluates to true and all variables in \vec{X} are \forall -quantified, there exists a valuation V_Y for the variables of \vec{Y} such that $rem_{2,\forall}(G)[V_X^{H'} \cup V_Y]$ evaluates to true. Let $R \subset br_1^Y$ be the set of bridge rules consistent wrt. \vec{Y} such that $V_Y^{\varphi(R)} = V_Y$ and consider the candidate diagnosis $D'' = (D'_1 \cup (br_1^Y \setminus R), \emptyset)$. Let $H'' = \{\varphi(r) \mid r \in app(br_1(M^G[D'']), S_\emptyset)\}$ and observe that H''

is consistent since D'_1 and R both are consistent. Furthermore, $V_X^{H''} = V_X^{H'}$ and $V_Y^{H''} = V_Y$, thus $rem_{2,\vee}(G)[V_X^{H''} \cup V_Y^{H''}]$ evaluates to true, hence S_\emptyset is an equilibrium of $M[D'']$ and $D'' \in D^\pm(M^G)$ holds.

Now consider whether $D'' \preceq D'$ holds: $D' = (D'_1, \emptyset)$, $D'' = (D'_1 \cup (br_1^Y \setminus R), \emptyset)$, $D'' \neq D_{valid} \neq D'$, $D'_1 \cap br_1^X = (D'_1 \cup (br_1^Y \setminus R)) \cap br_1^X$, and $D'_1 \cap br_1^Y \subseteq (D'_1 \cup (br_1^Y \setminus R)) \cap br_1^Y$ all hold. Therefore $D'' \preceq D'$ holds. Since $R \subset br_1^Y$ holds, it follows that $br_1^Y \setminus R \neq \emptyset$ and by $D'_1 \subset br_1^X$ it then follows that $(D'_1 \cup (br_1^Y \setminus R)) \cap br_1^Y \subseteq D'_1 \cap br_1^Y$ does not hold. Therefore $D' \preceq D''$ does not hold and consequently, it holds that $D'' \prec D'$ and $D' \notin D_{\preceq}^\pm(M^G)$. Since $D' \subset D_{valid}$ was chosen arbitrary, it follows that D_{valid} is subset-minimal among all diagnoses in $D_{\preceq}^\pm(M^G)$, hence $D_{valid} \in D_{m,\preceq}^\pm(M^G)$ holds.

In summary, this proves that $D_{valid} \in D_{m,\preceq}^\pm(M^G)$ holds iff G evaluates to true.

Complexity: Observe that deciding whether $D \preceq D'$ holds for the above \preceq clearly is in \mathbf{P} . Further note that M^G is polynomial in the size of G since kb_1 and br_1 are both polynomial (even linear) in the size of G . In the following we assume wlog. that all QBFs are in prenex normal form.

Let G be an arbitrary formula in $\mathbf{QBF}_{2,\vee}$, then $rem_{2,\vee}(G)$ contains no quantifiers, hence deciding whether $rem_{2,\vee}(G)$ evaluates to true under an assignment for $\vec{X} \cup \vec{Y}$ amounts to evaluating a propositional formula under a given assignment; this is possible in \mathbf{P} , hence $\mathcal{CC}(M^G) = \mathbf{P}$. Since $D_{valid} \in D_{m,\preceq}^\pm(M^G)$ iff G evaluates to true, it thus follows that $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is $\Pi_2^{\mathbf{P}}$ -hard if $\mathcal{CC}(M) = \mathbf{P}$.

Let G be an arbitrary formula in $\mathbf{QBF}_{i+2,\vee}$ for $i \geq 0$, then $rem_{2,\vee}(G)$ contains i quantifiers, hence $rem_{2,\vee}(G)$ is a formula of $\mathbf{QBF}_{i,\vee}$ and checking whether it evaluates to true is in $\Pi_i^{\mathbf{P}}$, i.e., $\mathcal{CC}(M^G)$ being hard for $\Pi_i^{\mathbf{P}}$ is sufficient for $\text{MCS}_{\text{D}_{\text{MPREF}}}$ to decide whether G evaluates to true. Thus $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is $\Pi_{i+2}^{\mathbf{P}}$ -hard for $\mathcal{CC}(M)$ being hard for $\Pi_i^{\mathbf{P}}$, $i \geq 0$.

Similarly, if G is an arbitrary formula in $\mathbf{QBF}_{i+1,\exists}$ for $i \geq 1$, then $rem_{2,\vee}(G)$ contains $i - 1$ quantifiers, i.e., $rem_{2,\vee}(G)$ is a formula in $\mathbf{QBF}_{i-1,\vee}$. Since $\mathbf{QBF}_{i,\exists}$ contains all formulas of $\mathbf{QBF}_{i-1,\vee}$ it follows that $\Sigma_i^{\mathbf{P}}$ is sufficient for checking whether $rem_{2,\vee}(G)$ evaluates to true. Thus, $\mathcal{CC}(M^G)$ being hard for $\Sigma_i^{\mathbf{P}}$ is sufficient for $\text{MCS}_{\text{D}_{\text{MPREF}}}$ to decide whether G evaluates to true. Thus $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is $\Pi_{i+1}^{\mathbf{P}}$ -hard for $\mathcal{CC}(M) = \Sigma_i^{\mathbf{P}}$, $i \geq 1$.

In summary it thus follows that if $\mathcal{CC}(M)$ is hard for $\Sigma_i^{\mathbf{P}}$ ($\Pi_i^{\mathbf{P}}$) then $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is hard for $\Pi_{i+1}^{\mathbf{P}}$ ($\Pi_{i+2}^{\mathbf{P}}$) with $i \geq 0$ and that $\text{MCS}_{\text{D}_{\text{MPREF}}}$ is $\Pi_2^{\mathbf{P}}$ -hard even if $\mathcal{CC}(M)$ and deciding whether $D' \preceq D''$ hold are both in \mathbf{P} .

For hardness in case that $\mathcal{CC}(M) = \mathbf{P}$ it is sufficient to use a stratified logic program and the logic L_Σ^{asp} as the context C_1 in the MCS M^G . Let G be in $\mathbf{QBF}_{2,\vee}$, then $rem_{2,\vee}(G)$ is a SAT formula and wlog. we assume $rem_{2,\vee}(G)$ to be in CNF. Let $F = rem_{2,\vee}(G) = \{c_1, \dots, c_m\}$ be given as a set of clauses each of the form $c_\ell = (l_{\ell_1} \vee l_{\ell_2} \vee \dots \vee l_{\ell_k})$ with $k \in \mathbb{N}$. We associate with each clause c_ℓ of this form a set of rules

$$lp(c_\ell) = \left\{ \begin{array}{ll} \text{clause_}c_\ell \leftarrow l_{\ell_1}. & \dots \quad \text{clause_}c_\ell \leftarrow l_{\ell_j}. \\ \text{clause_}c_\ell \leftarrow \overline{l_{\ell_{j+1}}}. & \dots \quad \text{clause_}c_\ell \leftarrow \overline{l_{\ell_k}}. \end{array} \right\}$$

where l_{ℓ_1} to l_{ℓ_j} are the positive literals and $l_{\ell_{j+1}}$ to l_{ℓ_k} are the negative literals of c_ℓ .

Finally, $C_1 = (L_{\Sigma}^{asp}, kb_1, br_1)$ uses the abstract logic of ASP and kb_1 is as follows:

$$\begin{aligned}
 kb_1 = \{ & \text{consistent}X \leftarrow \text{not inconsistent}X. \\
 & \text{inconsistent}X \leftarrow x, \bar{x}. & \forall x \in \vec{X} \\
 & \text{consistent}Y \leftarrow \text{not inconsistent}Y. \\
 & \text{inconsistent}Y \leftarrow y, \bar{y}. & \forall y \in \vec{Y} \\
 & \perp \leftarrow \text{not ok}. \\
 & \text{ok} \leftarrow \text{consistent}X, \text{consistent}Y, \text{true}F. \\
 & \text{ok} \leftarrow \text{not nonempty_intersect}. \\
 & \text{ok} \leftarrow \text{consistent}X, \text{not notfull}Y. \\
 & \text{nonempty_intersect} \leftarrow x. & \forall x \in \vec{X} \\
 & \text{nonempty_intersect} \leftarrow \bar{x}. & \forall x \in \vec{X} \\
 & \text{notfull}Y \leftarrow \text{not } y. & \forall y \in \vec{Y} \\
 & \text{notfull}Y \leftarrow \text{not } \bar{y}. & \forall y \in \vec{Y} \\
 & \text{true}F \leftarrow \text{clause_}c_1, \dots, \text{clause_}c_k. & \text{for } F = \{c_1, \dots, c_k\} \\
 & \} \cup \{r \in lp(c_\ell) \mid c_\ell \in F\}
 \end{aligned}$$

Observe that kb_1 is a stratifiable logic program while bridge rules only add facts, thus $\mathbf{ACC}(kb_1 \cup H)$ can be computed in polynomial time. Also note that kb_1 is linear in the size of G . \square

Proof of Corollary 1. For the membership part, observe that the size of the clone-encoding M^{\preceq} is polynomial in the size of M . As in M^{\preceq} all contexts apart from the observation context have the same knowledge bases and logics as in M , their complexities are in $\mathcal{CC}(M)$; thus it remains to show that also the complexity of the observation context is in $\mathcal{CC}(M)$. Definition 17 specifies which belief sets are acceptable for the latter in terms of the following conditions for which the property $\theta(R_1, R_2, R_3)$ holds, namely: if $R_1 = D_1 \cup 2.D'_1$, $R_2 = D_2 \cup 2.D'_2$ and either $(D_1, D_2) = (D'_1, D'_2)$ and $R_3 = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$ or $(D'_1, D'_2) \prec (D_1, D_2)$ and $R_3 = \mathcal{K}(D_1, D_2)$.

The equalities are trivially checked in polynomial time; thus if deciding $D \prec D'$ is in $\Sigma_i^{\mathbf{P}}$, then checking all conditions, i.e., whether $\theta(R_1, R_2, R_2)$ holds, is in $\Sigma_i^{\mathbf{P}}$; note in particular that if deciding $D \preceq D'$ is in \mathbf{P} , then deciding $D \not\preceq D'$ is in \mathbf{P} and consequently deciding $D \prec D'$ also is in \mathbf{P} . By construction of M^{\preceq} , it thus follows that $\mathcal{CC}(M) = \Sigma_i^{\mathbf{P}}$. The argument for \mathbf{PSPACE} and $\mathbf{EXPTIME}$ in place of $\Sigma_i^{\mathbf{P}}$ is analogous.

By Theorem 4, we have $(D_1, D_2) \in D_{m, \preceq}^{\pm}(M)$ iff $t(D_1, D_2) \in D_{m, t_{max}}^{\pm}(M^{\preceq}, br_P, br_H)$ and from Theorem 7 it follows that deciding $t(D_1, D_2) \in D_{m, t_{max}}^{\pm}(M^{\preceq}, br_P, br_H)$ is in $\Pi_{i+1}^{\mathbf{P}}$ (resp., $\mathbf{coNP}^{\mathbf{PSPACE}} = \mathbf{PSPACE}$, $\mathbf{coNP}^{\mathbf{EXPTIME}} = \mathbf{EXPTIME}$); hence deciding $(D_1, D_2) \in D_{m, \preceq}^{\pm}(M)$ is in $\Pi_{i+1}^{\mathbf{P}}$ (resp. \mathbf{PSPACE} , $\mathbf{EXPTIME}$). The hardness follows directly from Theorem 8 (resp. an MCS where the acceptability function of some context is hard for \mathbf{PSPACE} resp. $\mathbf{EXPTIME}$); hence, $\mathbf{MCS}_{\mathbf{D}_{\text{MPREF}}}$ is complete for $\Pi_{i+1}^{\mathbf{P}}$ (resp. \mathbf{PSPACE} , $\mathbf{EXPTIME}$). \square

Proof of Corollary 2. The membership follows from Theorem 7. As for the hardness part, let $\mathcal{CC}(M)$ be equal to \mathbf{P} , \mathbf{NP} , or $\Sigma_i^{\mathbf{P}}$ with $i \geq 1$ then by Theorem 8 $\mathbf{MCS}_{\mathbf{D}_{\text{MPREF}}}$ is hard for $\Pi_2^{\mathbf{P}}$, $\Pi_2^{\mathbf{P}}$, or $\Pi_{i+2}^{\mathbf{P}}$, respectively. Let \preceq be any preference order on M such that deciding whether $D \preceq D'$

holds is in \mathbf{P} . Consider the clone encoding M^{\preceq} and any candidate diagnosis (D_1, D_2) . By Theorem 4, we have $(D_1, D_2) \in D_{m, \preceq}^{\pm}(M)$ iff $t(D_1, D_2) \in D_{m, t_{max}}^{\pm}(M^{\preceq}, br_P, br_H)$. Therefore the clone encoding induces a polynomial-time reduction of $\text{MCSD}_{\text{MPREF}}$ to $\text{MCSDPH}_{m, t_{max}}$, because $t(D_1, D_2)$ and M^{\preceq} are both linear in the size of (D_1, D_2) plus M as well as $\mathcal{CC}(M^{\preceq})$ is \mathbf{P} , \mathbf{NP} , or $\Sigma_1^{\mathbf{P}}$ ($i \geq 1$) given that $\mathcal{CC}(M)$ is \mathbf{P} , \mathbf{NP} , or $\Sigma_1^{\mathbf{P}}$ ($i \geq 1$). The hardness results of Theorem 8 thus also hold for $\text{MCSDPH}_{m, t_{max}}$. \square

Appendix B. Detailed Examples

Example 19. Consider the hospital MCS (cf. Figure 1), where one can see two groups of bridge rules: health-related (r_1 and r_2) and billing-related (r_3 , r_4 , and r_5). Assume a preference order is designed under the following two assumptions: first, modifying one group of bridge rules potentially modifies all information conveyed by the group (i.e., a correct result depends on all information of a group to be conveyed without modification); and second, information of one group might influence another group, as in the hospital MCS health information influences billing.

Candidate diagnoses then either modify (and possibly break) no group, the billing group, or health and billing group together. Since billing depends on health, it is impossible to modify health without possibly breaking billing. Only the candidate diagnosis (\emptyset, \emptyset) is of the first kind. The second kind consists of all candidate diagnoses (D_1, D_2) with $D_1 \cup D_2 \subseteq \{r_3, r_4, r_5\}$ and $D_1 \cup D_2 \neq \emptyset$. All other candidate diagnoses are of the third kind.

Preferring those diagnoses which (possibly) break only the least set of groups, then prefers diagnoses of the first kind over all others, and the second over the third. Formally, we obtain a preference order \preceq_U such that $(D_1, D_2) \preceq_U (D'_1, D'_2)$ holds iff one of the following is the case: 1) $D_1 \cup D_2 = \emptyset$, or 2) $D_1 \cup D_2 \subseteq \{r_3, r_4, r_5\}$ and $D'_1 \cup D'_2 \neq \emptyset$.

The resulting MCS clone encoding wrt. \preceq_U , $M^{\preceq_U} = (C_1, C_2, C_3, C_4, C_5, C_6, C_7)$ is based on two clones of M , where the first comprises the contexts C_1, C_2, C_3 and the second the contexts C_4, C_5, C_6 . The context C_7 finally is the observation/encoding context.

We first recall the bridge rules of $2M = M \otimes M$ using the permutation I corresponding to $M \otimes M$. Accordingly $br(2M)$ is:

$$\begin{aligned}
 r_1 : & & (2 : \text{hyperglycemia}) & \leftarrow (1 : \text{hyperglycemia}). \\
 r_2 : & & (2 : \text{allow_animal_insulin}) & \leftarrow \mathbf{not} (1 : \text{allergic_animal_insulin}). \\
 r_3 : & & (3 : \text{bill_animal_insulin}) & \leftarrow (2 : \text{give_animal_insulin}). \\
 r_4 : & & (3 : \text{bill_human_insulin}) & \leftarrow (2 : \text{give_human_insulin}). \\
 r_5 : & & (3 : \text{insurance_B}) & \leftarrow (1 : \text{insurance_B}). \\
 I(r_1) : & & (5 : \text{hyperglycemia}) & \leftarrow (4 : \text{hyperglycemia}). \\
 I(r_2) : & & (5 : \text{allow_animal_insulin}) & \leftarrow \mathbf{not} (4 : \text{allergic_animal_insulin}).
 \end{aligned}$$

$$\begin{aligned}
 I(r_3) : & & (6 : \text{bill_animal_insulin}) & \leftarrow (5 : \text{give_animal_insulin}). \\
 I(r_4) : & & (6 : \text{bill_human_insulin}) & \leftarrow (5 : \text{give_human_insulin}). \\
 I(r_5) : & & (6 : \text{insurance_B}) & \leftarrow (4 : \text{insurance_B}).
 \end{aligned}$$

A graphical rendering of M^{\succeq_U} is given in Figure 5, where for readability only some of the bridge rules of M^{\succeq_U} are shown. The set of bridge rules of the observation context C_7 is as follows:

$$\begin{aligned}
 br_7(M^{\succeq_U}) = \{ & \quad (7: not_removed_{r_1}) \leftarrow \top. & \quad (7: uncond_{r_1}) \leftarrow \perp. \\
 & \quad (7: not_removed_{r_2}) \leftarrow \top. & \quad (7: uncond_{r_2}) \leftarrow \perp. \\
 & \quad \dots & \quad \dots \\
 & \quad (7: not_removed_{I(r_4)}) \leftarrow \top. & \quad (7: uncond_{I(r_4)}) \leftarrow \perp. \\
 & \quad (7: not_removed_{I(r_5)}) \leftarrow \top. & \quad (7: uncond_{I(r_5)}) \leftarrow \perp. \\
 & \quad (7: in_1(r_1)) \leftarrow \perp. & \quad (7: \overline{in}_1(r_1)) \leftarrow \perp. \\
 & \quad (7: in_2(r_1)) \leftarrow \perp. & \quad (7: \overline{in}_2(r_1)) \leftarrow \perp. \\
 & \quad \dots & \quad \dots \\
 & \quad (7: in_1(r_5)) \leftarrow \perp. & \quad (7: \overline{in}_1(r_5)) \leftarrow \perp. \\
 & \quad (7: in_2(r_5)) \leftarrow \perp. & \quad (7: \overline{in}_2(r_5)) \leftarrow \perp. & \quad \}
 \end{aligned}$$

To fully realize the property θ and the preference order \preceq_U based on real-world entities patient's health/treatment, and billing, we may use for the observation context C_7 an ASP program that consists of the following rules:

$$removed_r \leftarrow not\ not_removed_r. \quad \text{for all } r \in br(M \otimes M) \quad (22)$$

$$\perp \leftarrow removed_r, not\ in_1(r). \quad \text{for all } r \in \{r_1, \dots, r_5\} \quad (23)$$

$$\perp \leftarrow not\ removed_r, in_1(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow not\ removed_r, not\ \overline{in}_1(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow removed_r, \overline{in}_1(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow uncond_r, not\ in_2(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow uncond_r, \overline{in}_2(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow not\ uncond_r, not\ \overline{in}_2(r). \quad \text{for all } r \in \{r_1, \dots, r_5\}$$

$$\perp \leftarrow uncond_r, \overline{in}_2(r). \quad \text{for all } r \in \{r_1, \dots, r_5\} \quad (24)$$

$$mod(clone1, billing) \leftarrow removed_r. \quad \text{for all } r \in \{r_3, \dots, r_5\} \quad (25)$$

$$mod(clone1, billing) \leftarrow uncond_r. \quad \text{for all } r \in \{r_3, \dots, r_5\}$$

$$mod(clone2, billing) \leftarrow removed_r. \quad \text{for all } r \in \{I(r_3), \dots, I(r_5)\}$$

$$mod(clone2, billing) \leftarrow uncond_r. \quad \text{for all } r \in \{I(r_3), \dots, I(r_5)\}$$

$$mod(clone1, treatment) \leftarrow removed_r. \quad \text{for all } r \in \{r_1, r_2\}$$

$$mod(clone1, treatment) \leftarrow uncond_r. \quad \text{for all } r \in \{r_1, r_2\}$$

$$mod(clone2, treatment) \leftarrow removed_r. \quad \text{for all } r \in \{I(r_1), I(r_2)\}$$

$$mod(clone2, treatment) \leftarrow uncond_r. \quad \text{for all } r \in \{I(r_1), I(r_2)\} \quad (26)$$

$$\text{mod}(\text{clone1}, \text{billing}) \leftarrow \text{mod}(\text{clone1}, \text{treatment}). \quad (27)$$

$$\text{mod}(\text{clone2}, \text{billing}) \leftarrow \text{mod}(\text{clone2}, \text{treatment}). \quad (28)$$

$$\text{clones_different} \leftarrow \text{removed}_r, \text{not removed}_{r'}. \quad \text{for all } r \in \text{br}(M), r' \in I(\text{br}(M)) \quad (29)$$

$$\text{clones_different} \leftarrow \text{not removed}_r, \text{removed}_{r'}. \quad \text{for all } r \in \text{br}(M), r' \in I(\text{br}(M))$$

$$\text{clones_different} \leftarrow \text{uncond}_r, \text{not uncond}_{r'}. \quad \text{for all } r \in \text{br}(M), r' \in I(\text{br}(M))$$

$$\text{clones_different} \leftarrow \text{not uncond}_r, \text{uncond}_{r'}. \quad \text{for all } r \in \text{br}(M), r' \in I(\text{br}(M)) \quad (30)$$

$$\text{clone1_modifies_more} \leftarrow \text{mod}(\text{clone1}, U), \text{not mod}(\text{clone2}, U). \quad (31)$$

$$\text{clone2_modifies_more} \leftarrow \text{mod}(\text{clone2}, U), \text{not mod}(\text{clone1}, U).$$

$$\text{clone1_less_preferred} \leftarrow \text{clone1_modifies_more}, \text{not clone2_modifies_more}. \quad (32)$$

$$\perp \leftarrow \text{not ismax}, \text{clone1_less_preferred}, \text{clones_different}. \quad (33)$$

$$\perp \leftarrow \text{not clone1_less_preferred}, \text{clones_different}. \quad (34)$$

The intuition of the above rules is as follows: rules of form (22) expose the diagnoses of both clones; the constraints of form (23)–(24) ensure that the diagnosis of the first clone is exhibited via prioritized bridge rules; rules of form (25)–(26) deduce which units of bridge rules have been modified in the first and second clone; rules (27) and (28) take care of the dependency between the units treatment and billing; rules of form (29)–(30) infer whether the diagnosis of the first clone is different from the diagnosis of the second clone; rules (31)–(32) infer whether the modified units of the first clone is a superset of the modified units of the second clone, which means the diagnosis of the second clone is more preferred than the one of the first clone. Finally, the constraint (33) ensures that t_{max} is made condition-free if the diagnosis of the second clone is more preferred than the diagnosis of the first clone, and the constraint (34) ensures that only comparable diagnoses (or if both diagnoses are equal) yield a diagnosis of the MCS $M^{\leq v}$.

References

- Allen, T. E. (2016). *CP-Nets: From Theory to Practice*. Ph.D. thesis, University of Kentucky. <http://dx.doi.org/10.13023/ETD.2016.131>.
- Balduccini, M., & Gelfond, M. (2003). Logic programs with consistency-restoring rules. In Doherty, P., McCarthy, J., & Williams, M.-A. (Eds.), *Working Papers of the 2003 AAI Spring Symposium on Logical Formalization of Commonsense Reasoning*, pp. 9–18. AAAI Press, Menlo Park, California.
- Barilaro, R., Fink, M., Ricca, F., & Terracina, G. (2013). Towards query answering in relational multi-context systems.. In Cabalar, & Son (Cabalar & Son, 2013), pp. 168–173.
- Bikakis, A., & Antoniou, G. (2008). Distributed defeasible contextual reasoning in ambient computing. In Aarts, E. H. L., Crowley, J. L., de Ruyter, B. E. R., Gerhäuser, H., Pflaum, A., Schmidt, J., & Wichert, R. (Eds.), *Ambient Intelligence, European Conference, Aml 2008, Nuremberg, Germany, November 19-22, 2008. Proceedings*, Vol. 5355 of *Lecture Notes in Computer Science*, pp. 308–325. Springer.
- Bikakis, A., & Antoniou, G. (2010). Defeasible contextual reasoning with arguments in ambient intelligence. *IEEE Trans. Knowl. Data Eng.*, 22(11), 1492–1506.

- Bikakis, A., & Antoniou, G. (2011). Partial preferences and ambiguity resolution in contextual de-feasible logic. In Delgrande, J. P., & Faber, W. (Eds.), *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, Vol. 6645 of *Lecture Notes in Computer Science*, pp. 193–198. Springer.
- Bikakis, A., Antoniou, G., & Hassapis, P. (2009). Alternative strategies for conflict resolution in multi-context systems. In Iliadis, L. S., Maglogiannis, I., Tsoumakas, G., Vlahavas, I. P., & Bramer, M. (Eds.), *Artificial Intelligence Applications and Innovations III, Proceedings of the 5TH IFIP Conference on Artificial Intelligence Applications and Innovations (AIAI'2009), April 23-25, 2009, Thessaloniki, Greece*, Vol. 296 of *IFIP Advances in Information and Communication Technology*, pp. 31–40. Springer.
- Bikakis, A., Antoniou, G., & Hassapis, P. (2011). Strategies for contextual reasoning with conflicts in ambient intelligence. *Knowl. Inf. Syst.*, 27(1), 45–84.
- Binas, A., & McIlraith, S. A. (2008). Peer-to-peer query answering with inconsistent knowledge. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 329–339, Sydney, Australia.
- Bögl, M., Eiter, T., Fink, M., & Schüller, P. (2010). The MCS-IE system for explaining inconsistency in multi-context systems. In Janhunen, T., & Niemelä, I. (Eds.), *12th European Conference on Logics in Artificial Intelligence (JELIA 2010)*, *Lecture Notes in Artificial Intelligence*, pp. 356–359. Springer.
- Bonatti, P. A., & Eiter, T. (1995). Querying disjunctive database through nonmonotonic logics. In Gottlob, G., & Vardi, M. Y. (Eds.), *Database Theory - ICDT'95, 5th International Conference, Prague, Czech Republic, January 11-13, 1995, Proceedings*, Vol. 893 of *Lecture Notes in Computer Science*, pp. 68–81. Springer.
- Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21, 135–191.
- Brafman, R. I., Domshlak, C., Shimony, S. E., & Silver, Y. (2006). Preferences over sets. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pp. 1101–1106. AAAI Press.
- Brewka, G. (2013). Towards reactive multi-context systems.. In Cabalar, & Son (Cabalar & Son, 2013), pp. 1–10.
- Brewka, G., & Eiter, T. (2007). Equilibria in heterogeneous nonmonotonic multi-context systems.. In *AAAI Conference on Artificial Intelligence (AAAI)*, pp. 385–390.
- Brewka, G., Eiter, T., & Fink, M. (2011a). Nonmonotonic multi-context systems: A flexible approach for integrating heterogeneous knowledge sources. In Balduccini, M., & Son, T. C. (Eds.), *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*, Vol. 6565 of *Lecture Notes in Computer Science*, pp. 233–258. Springer.
- Brewka, G., Eiter, T., Fink, M., & Weinzierl, A. (2011b). Managed multi-context systems.. In Walsh (Walsh, 2011), pp. 786–791.
- Brewka, G., Ellmauthaler, S., & Pührer, J. (2014a). Multi-context systems for reactive reasoning in dynamic environments. In *21st European Conference on Artificial Intelligence (ECAI 2014)*.

- Brewka, G., Ellmauthaler, S., & Pührer, J. (2014b). Multi-context systems for reactive reasoning in dynamic environments. In *Proceedings of the International Workshop on Reactive Concepts in Knowledge Representation (ReactKnow 2014)*, pp. 23–29.
- Brewka, G., Roelofsen, F., & Serafini, L. (2007). Contextual default reasoning. In Veloso, M. M. (Ed.), *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pp. 268–273.
- Brewka, G., Truszczyński, M., & Woltran, S. (2010). Representing preferences among sets. In Fox, M., & Poole, D. (Eds.), *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*. AAAI Press.
- Bylander, T., Allemang, D., Tanner, M. C., & Josephson, J. R. (1991). The computational complexity of abduction. *Artificial Intelligence*, 49(1-3), 25–60.
- Cabalar, P., & Son, T. C. (Eds.). (2013). *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15-19, 2013. Proceedings*, Vol. 8148 of *Lecture Notes in Computer Science*. Springer.
- Caire, P., & Bikakis, A. (2011). Enhancing cooperation in distributed information systems using conviviality and multi-context systems. In Sombattheera, C., Agarwal, A., Udgata, S. K., & Lavangnananda, K. (Eds.), *Multi-disciplinary Trends in Artificial Intelligence - 5th International Workshop, MIWAI 2011, Hyderabad, India, December 7-9, 2011. Proceedings*, Vol. 7080 of *Lecture Notes in Computer Science*, pp. 14–25. Springer.
- Caire, P., Bikakis, A., & Traon, Y. L. (2013). Information dependencies in MCS: conviviality-based model and metrics. In Boella, G., Elkind, E., Savarimuthu, B. T. R., Dignum, F., & Purvis, M. K. (Eds.), *PRIMA 2013: Principles and Practice of Multi-Agent Systems - 16th International Conference, Dunedin, New Zealand, December 1-6, 2013. Proceedings*, Vol. 8291 of *Lecture Notes in Computer Science*, pp. 405–412. Springer.
- Calvanese, D., Giacomo, G. D., Lenzerini, M., & Rosati, R. (2004). Logical foundations of peer-to-peer data integration. In Beeri, C., & Deutsch, A. (Eds.), *PODS*, pp. 241–251. ACM.
- Dantsin, E., Eiter, T., Gottlob, G., & Voronkov, A. (2001). Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3), 374–425.
- Dao-Tran, M., Eiter, T., Fink, M., & Krennwallner, T. (2015). Distributed evaluation of nonmonotonic multi-context systems. *J. Artif. Intell. Res. (JAIR)*, 52, 543–600.
- Domshlak, C., Brafman, R. I., & Shimony, S. E. (2001). Preference-based configuration of web page content. In Nebel, B. (Ed.), *IJCAI*, pp. 1451–1456. Morgan Kaufmann.
- Doyle, J., Shoham, Y., & Wellman, M. P. (1991). A logic of relative desire (preliminary report). In Ras, Z. W., & Zemankova, M. (Eds.), *Methodologies for Intelligent Systems, 6th International Symposium, ISMIS '91, Charlotte, N.C., USA, October 16-19, 1991, Proceedings*, Vol. 542 of *Lecture Notes in Computer Science*, pp. 16–31. Springer.
- Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2), 321–358.
- Egly, U., Eiter, T., Tompits, H., & Woltran, S. (2000). Solving advanced reasoning tasks using quantified boolean formulas. In Kautz, H. A., & Porter, B. W. (Eds.), *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on on*

Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000, Austin, Texas, USA., pp. 417–422. AAAI Press / The MIT Press.

- Eiter, T., Fink, M., Schüller, P., & Weinzierl, A. (2010). Finding explanations of inconsistency in multi-context systems. In Lin, F., Sattler, U., & Truszczyński, M. (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9-13, 2010*, pp. 329–339. AAAI Press.
- Eiter, T., Fink, M., Schüller, P., & Weinzierl, A. (2014). Finding explanations of inconsistency in multi-context systems. *Artificial Intelligence*, 216(0), 233–274.
- Eiter, T., Fink, M., & Weinzierl, A. (2010). Preference-based inconsistency assessment in multi-context systems. In Janhunen, T., & Niemelä, I. (Eds.), *Logics in Artificial Intelligence - 12th European Conference, JELIA 2010, Helsinki, Finland, September 13-15, 2010. Proceedings*, Vol. 6341 of *Lecture Notes in Computer Science*, pp. 143–155. Springer.
- Eiter, T., Gottlob, G., & Mannila, H. (1997). Disjunctive datalog. *ACM Trans. Database Syst.*, 22(3), 364–418.
- Eiter, T., & Šimkus, M. (2015). Linking open-world knowledge bases with nonmonotonic rules. In Calimeri, F., Ianni, G., & Truszczyński, M. (Eds.), *Proceedings of the 13th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2015)*, Vol. 9345 of *LNCS*, pp. 294–308. Springer.
- Ellmauthaler, S. (2013). Generalizing multi-context systems for reactive stream reasoning applications. In Jones, A. V., & Ng, N. (Eds.), *Proceedings of the 2013 Imperial College Computing Student Workshop (ICCSW 2013)*, OpenAccess Series in Informatics (OASiCs), pp. 17–24. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Ellmauthaler, S., & Pührer, J. (2014). Asynchronous multi-context systems. In Ellmauthaler, S., & Pührer, J. (Eds.), *Proceedings of the International Workshop on Reactive Concepts in Knowledge Representation (ReactKnow 2014)*, pp. 31–37.
- Faber, W., Leone, N., & Pfeifer, G. (2004). Recursive aggregates in disjunctive logic programs: Semantics and complexity. In *European Conference on Logics in Artificial Intelligence (JELIA)*, pp. 200–212.
- Fink, M., Ghionna, L., & Weinzierl, A. (2011). Relational information exchange and aggregation in multi-context systems. In Delgrande, J. P., & Faber, W. (Eds.), *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, Vol. 6645 of *Lecture Notes in Computer Science*, pp. 120–133. Springer.
- Gelfond, M., & Lifschitz, V. (1991). Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9(3/4), 365–386.
- Ghidini, C., & Giunchiglia, F. (2001). Local models semantics, or contextual reasoning=locality+compatibility. *Artif. Intell.*, 127(2), 221–259.
- Giunchiglia, F., & Serafini, L. (1994). Multilanguage hierarchical logics or: How we can do without modal logics. *Artif. Intell.*, 65(1), 29–70.
- Goldsmith, J., Lang, J., Truszczyński, M., & Wilson, N. (2008). The computational complexity of dominance and consistency in cp-nets. *J. Artif. Intell. Res. (JAIR)*, 33, 403–432.

- Goncalves, R., Knorr, M., & Leite, J. (2014a). Evolving bridge rules in evolving multi-context systems. In Bulling, N., van der Torre, L. W. N., Villata, S., Jamroga, W., & Vasconcelos, W. (Eds.), *Computational Logic in Multi-Agent Systems - 15th International Workshop, CLIMA XV, Prague, Czech Republic, August 18-19, 2014. Proceedings*, Vol. 8624 of *Lecture Notes in Computer Science*, pp. 52–69. Springer.
- Goncalves, R., Knorr, M., & Leite, J. (2014b). Evolving multi-context systems. In Schaub, T., Friedrich, G., & O’Sullivan, B. (Eds.), *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, Vol. 263 of *Frontiers in Artificial Intelligence and Applications*, pp. 375–380. IOS Press.
- Gottlob, G. (1992). Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2, 397–425.
- Inoue, K. (2011). Logic programming for boolean networks.. In Walsh (Walsh, 2011), pp. 924–930.
- Jin, Y., Wang, K., & Wen, L. (2012). Possibilistic reasoning in multi-context systems: Preliminary report. In Anthony, P., Ishizuka, M., & Lukose, D. (Eds.), *PRICAI 2012: Trends in Artificial Intelligence - 12th Pacific Rim International Conference on Artificial Intelligence, Kuching, Malaysia, September 3-7, 2012. Proceedings*, Vol. 7458 of *Lecture Notes in Computer Science*, pp. 180–193. Springer.
- Kakas, A. C., Kowalski, R. A., & Toni, F. (1992). Abductive logic programming. *J. Log. Comput.*, 2(6), 719–770.
- Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3), 437–467.
- Kauffman, S. A. (1993). *Origins of Order: Self-Organization and Selection in Evolution*. Oxford Univ. Press. Technical monograph.
- Knorr, M., Slota, M., Leite, J., & Homola, M. (2014). What if no hybrid reasoner is available? hybrid MKNF in multi-context systems. *J. Log. Comput.*, 24(6), 1279–1311.
- Le, T., Son, T. C., & Pontelli, E. (2015). Multi-context systems with preferences. In Chen, Q., Torroni, P., Villata, S., Hsu, J. Y., & Omicini, A. (Eds.), *PRIMA 2015: Principles and Practice of Multi-Agent Systems - 18th International Conference, Bertinoro, Italy, October 26-30, 2015, Proceedings*, Vol. 9387 of *Lecture Notes in Computer Science*, pp. 449–466. Springer.
- McCarthy, J. (1993). Notes on formalizing context. In Bajcsy, R. (Ed.), *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, pp. 555–562. Morgan Kaufmann.
- Mu, K., Wang, K., & Wen, L. (2015). Preferential multi-context systems. *CoRR*, abs/1504.06700.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–132.
- Roelofsen, F., & Serafini, L. (2005). Minimal and absent information in contexts. In Kaelbling, L. P., & Saffiotti, A. (Eds.), *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, pp. 558–563. Professional Book Center.
- Salcher, F. (2016). Distributed computation of diagnoses for inconsistent multi-context systems. Master’s thesis, Vienna University of Technology. To appear.

- Schüller, P., & Weinzierl, A. (2011). Semantic reasoning with SPARQL in heterogeneous multi-context systems. In Salinesi, C., & Pastor, O. (Eds.), *Advanced Information Systems Engineering Workshops - CAiSE 2011 International Workshops, London, UK, June 20-24, 2011. Proceedings*, Vol. 83 of *Lecture Notes in Business Information Processing*, pp. 575–585. Springer.
- Serafini, L., & Homola, M. (2012). Contextualized knowledge repositories for the semantic web. *J. Web Sem.*, 12, 64–87.
- Tasharofi, S., & Ternovska, E. (2014). Generalized multi-context systems. In Baral, C., Giacomo, G. D., & Eiter, T. (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20-24, 2014*. AAAI Press.
- Velikova, M., Novák, P., Huijbrechts, B., Laarhuis, J., Hoeksma, J., & Michels, S. (2014). An integrated reconfigurable system for maritime situational awareness. In Schaub, T., Friedrich, G., & O’Sullivan, B. (Eds.), *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, Vol. 263 of *Frontiers in Artificial Intelligence and Applications*, pp. 1197–1202. IOS Press.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Walsh, T. (Ed.). (2011). *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*. IJCAI/AAAI.
- Weinzierl, A. (2014). *Inconsistency Management under Preferences for Multi-Context Systems and Extensions*. Ph.D. thesis, Vienna University of Technology. <http://media.obvsg.at/p-AC12121980-2001>.