Preference-based Inconsistency Assessment in Multi-Context Systems*

Thomas Eiter, Michael Fink, and Antonius Weinzierl

Institute of Information Systems
Vienna University of Technology
Favoritenstraße 9-11, A-1040 Vienna, Austria
{eiter,fink,weinzierl}@kr.tuwien.ac.at

Abstract. Resolving inconsistency in knowledge-integration systems is a major issue, especially when interlinking heterogeneous, autonomous sources. The latter can be done using a multi-context system, also in presence of non-monotonicity. Recent work considered diagnosis and explanation of inconsistency in such systems in terms of faulty information exchange. To discriminate between different solutions, we consider inconsistency assessment using preference. We present means to a) filter undesired diagnoses b) select the most preferred ones given an arbitrary preference order and c) use CP-nets for efficient selection. Furthermore, we show how to incorporate the assessment into a Multi-Context System by a transformational approach. In a range of settings, the complexity does not increase compared to the basic case and key properties like decentralized information exchange and information hiding are preserved.

Key words: Inconsistency Management, Multi-Context Systems, Hybrid Reasoning Systems, Nonmonotonic Reasoning, Preferences

1 Introduction

Inconsistencies in heterogeneous, nonmonotonic knowledge-integration systems often do not have a single cause, but emerge from interaction, i.e., by the exchange of knowledge between knowledge bases. The nonmonotonic Multi-Context System (MCS) framework of [4], which extends seminal works by [9, 6], is a logic-based approach to flexibly model the information exchange between heterogeneous (nonmonotonic) knowledge bases, which exist a priori and incorporate external knowledge via so-called bridge rules. Recently, formal notions for explaining inconsistency in such MCSs in terms of faulty bridge rules have been developed [8], serving the purpose of inconsistency analysis with the eventual aim of resolving inconsistency. However, multiple possibilities for this call for a further assessment, taking application specific criteria into account.

To the best of our knowledge, no general method has been proposed to assess inconsistencies in MCSs which is flexible enough to adapt to application specific

* Supported by the Vienna Science and Technology Fund (WWTF), grant ICT08-020.

criteria. Although, for instance, [2] provides methods based on local trust and provenance to determine preferred models for a MCS avoiding inconsistency, the proposal requests to choose one out of four predefined evaluation algorithms. Our work instead aims at general techniques for assessing inconsistency in MCSs that can be 'instantiated' to encode application-specific properties for preferred consistency restorations.

For example, consider a health-care decision-support system that interlinks knowledge sources about patient histories, lab test results, a disease ontology, and a decision support system for patient treatment. Here, an inconsistency might easily arise if some scenario of contradicting information has not been anticipated. E.g., the ontology classifies symptoms as atypical pneumonia, which requires strong antibiotics, but a patient is allergic to it; the treatment system may then raise an inconsistency. One possibility to resolve it is to ignore the imported disease information. While technically fine, this solution might be unacceptable, as a constraint "No illness of a patient may be ignored" should be fulfilled.

To account for such selection criteria on consistency restorations, we take a preference-based approach. Two basic elements of preference-based selection can be found in the literature: filters, which discard unpreferred solutions that fail some preference condition, and qualitative comparison relations establishing preference orders to single out the most appealing solutions. Our main contributions enabling these for inconsistency assessment in MCSs are summarized as follows.

- We formalize both preference approaches above in the setting of MCSs. For preference orders, we further investigate the application of conditional preference networks (CP-nets), which exhibit appealing features of locality and privacy. CP-nets [3] 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 further show how to realize the preference approaches inside the MCS framework by using meta-reasoning on consistency restorations. For a given MCS and a filter, preference order, or CP-net, a rewriting yields a transformed system such that consistency restorations of the latter directly correspond to preferred consistency restorations of the original system (wrt. the given filter, preference order, or CP-net).
- For preference notions that are not inherently centralized, the realization allows that preferred solutions are found in a decentralized, localized manner, maintaining privacy and information hiding. Thus we preserve key properties of MCSs also for inconsistency assessment.

Our results not only refine existing methods for inconsistency handling in MCSs without complexity increase, but also show the versatility of the basic framework to couch advanced reasoning tasks, including self-reflective assessment.

2 Preliminaries

This section introduces MCS and diagnoses in general; it is largely based on [8].

A heterogeneous nonmonotonic MCS [4] consists of *contexts*, which comprise knowledge bases in underlying *logics*, and *bridge rules* to control the information flow between contexts.

A logic $L = (\mathbf{KB}_L, \mathbf{BS}_L, \mathbf{ACC}_L)$ consists, in an abstract view, of

- a set $\mathbf{KB}_L$ of knowledge bases of L, each being a set (of "formulas"),
- a set $\mathbf{BS}_L$ of possible belief sets, whose elements are "beliefs", and
- a "semantics" function $\mathbf{ACC}_L : \mathbf{KB}_L \to 2^{\mathbf{BS}_L}$ which assigns each knowledge base a set of acceptable belief sets.

This concept of a *logic* captures many monotonic and nonmonotonic logics, e.g., classical logic, description logics, modal logics, default logics, circumscription, and logic programs under the answer set semantics.

A *bridge rule* can add information to a context, depending on the belief sets which are accepted at other contexts. Let $L = (L_1, \ldots, L_n)$ be a sequence of logics. An L_k-bridge rule r over L is of the form

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

where $1 \leq c_i \leq n$, p_i is an element of some belief set of L_{c_i}, k refers to the context receiving information s. We denote by $hd\,(r)$ the formula s in the head of r.

A *multi-context system (MCS)* is a collection $M = (C_1, \ldots, C_n)$ of contexts $C_i = (L_i, kb_i, br_i)$, $1 \leq i \leq n$, where $L_i = (\mathbf{KB}_i, \mathbf{BS}_i, \mathbf{ACC}_i)$ is a logic, $kb_i \in \mathbf{KB}_i$ a knowledge base, and br_i is a set of L_i-bridge rules over $(L_1, \ldots, L_n)$. In addition, for each $H \subseteq \{hd\,(r) \mid r \in br_i\}$ we have $kb_i \cup H \in \mathbf{KB}_i$, i.e., bridge rule heads are compatible with knowledge bases. By $br_M = \bigcup_{i=1}^{n} br_i$ we denote the set of bridge rules of M.

A *belief state* of an MCS $M = (C_1, \ldots, C_n)$ is a sequence $S = (S_1, \ldots, S_n)$ such that $S_i \in \mathbf{BS}_i$. A bridge rule (1) is *applicable* in a belief state S iff for $1 \leq i \leq j$: $p_i \in S_{c_i}$ and for $j < l \leq m$: $p_l \notin S_{c_l}$.

Example 1. Consider two scientists, Prof. K and Dr. J, planning to write a paper. We formalize their reasoning in an MCS M using two contexts each employing answer set semantics. Dr. J will write most of the paper and Prof. K only participates if either he finds time or if Dr. J thinks the paper needs improvement (bridge rule r_1). Dr. J knows that the participation of Prof. K results in a good paper (r_2 and kb_J) and he will name Prof. K as author if she participates (r_3). The knowledge bases of the contexts are:

$$kb_K = \{has\_time. \quad contribute \leftarrow improve. \quad contribute \leftarrow has\_time.\}$$
$$kb_J = \{good \leftarrow coauthored.\}$$

The bridge rules are $r_1 = (K{:}improve) \leftarrow \mathbf{not}\ (J{:}good).$,
$r_2 = (J{:}coauthored) \leftarrow (K{:}contribute).$, and $r_3 = (J{:}name\_K) \leftarrow (K{:}contribute).$

Equilibrium semantics selects certain belief states of an MCS $M = (C_1, \ldots, C_n)$ as acceptable. Intuitively, an equilibrium is a belief state $S = (S_1, \ldots, S_n)$ where each context C_i respects all bridge rules applicable in S and accepts S_i. Formally, S is an equilibrium of M, iff for $1 \leq i \leq n$,

$$S_i \in \mathbf{ACC}_i(kb_i \cup \{hd\,(r) \mid r \in br_i \text{ applicable in } S\}).$$

Example 2 (Ex. 1 ctd.). The MCS has just one equilibrium $S = (\{has\_time, contribute\}, \{coauthored, name\_K, good\})$ where both scientists author a good paper. Bridge rules r_2 and r_3 are applicable in S.

Inconsistency in an MCS is the lack of an equilibrium.

Example 3 (Ex. 1 ctd.). Assume Prof. K has no time, so she only contributes, if Dr. J considers the paper to be not good, i.e, $kb_K = \{contribute \leftarrow improve. contribute \leftarrow has\_time.\}$. Then there is a loop with an odd number of negations via bridge rules r_1 and r_2. This makes the MCS inconsistent.

For any MCS M and set R of bridge rules (fitting M), we denote by $M[R]$ the MCS obtained from M by replacing br_M with R (e.g., $M[br_M] = M$ and $M[\emptyset]$ is M with no bridge rules); by $M \models \bot$ we denote that M has no equilibrium (is inconsistent). For any set of bridge rules A, $heads(A) = \{\alpha \leftarrow \top \mid \alpha \leftarrow \beta \in A\}$ are the rules in A in unconditional form.

Diagnoses. As well-known, in nonmonotonic reasoning, adding knowledge can both cause and prevent inconsistency; the same is true for removing knowledge. The consistency-based explanation of inconsistency, therefore considers pairs (D_1, D_2) of sets of bridge rules, such that if the rules in D_1 are deactivated, and the rules in D_2 are added in unconditional form, the MCS becomes consistent (i.e., it admits an equilibrium). Adding rules unconditionally makes sense due to non-monotonicity; the idea is related to that of consistency restoring rules [1].

Formally, a *diagnosis* of an MCS M is a pair $D = (D_1, D_2)$, $D_1, D_2 \subseteq br_M$, s.t. $M[br_M \setminus D_1 \cup heads(D_2)] \not\models \bot$; by $D^{\pm}(M)$ we denote the set of all diagnoses. To obtain a more relevant set of diagnoses, pointwise subset-minimal diagnoses are preferred; we denote by $D^{\pm}_m(M)$ the set of all such diagnoses of an MCS M.

In our example $D^{\pm}_m(M) = \{(\{r_1\}, \emptyset), (\{r_2\}, \emptyset)\,(\emptyset, \{r_2\}), (\emptyset, \{r_1\})\}$; the first two diagnoses break the cycle by removing a rule, the last two "stabilize" it.

3 Filtering and Comparing Diagnoses

In this section we introduce ways to assess consistency restorations of inconsistent MCSs. First, we consider selection-based preference and provide a method to check whether diagnoses adhere to user-defined criteria. This allows to filter out undesired diagnoses. Then we turn to comparison-based preference, addressing the general problem of using arbitrary preference relations on diagnoses, before we focus on CP-nets, representing (semi-)local preference relations.

3.1 Filtering Diagnoses

Filters allow a designer of an MCSs to apply sanity checks on diagnoses, thus they can be seen as hard constraints on diagnoses: diagnoses that fail to satisfy the conditions are filtered out and not considered for consistency restoration.

Definition 1. *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 $D^{\pm}_{f}(M) = \{D \in D^{\pm}(M) \mid f(D) = 1\}$. By $D^{\pm}_{f,m}(M)$ we denote the set of all subset-minimal such diagnoses.*

Example 4 (Ex. 3 ctd.). Consider 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, in particular $f(D) = f(D') = 0$.

As it is a key strength of MCS to integrate different knowledge bases in a decentralized manner, users of MCS will want to specify their constraints on diagnoses in a logic of their choice, decentralized, and under the provision that they do not have to disclose information considered private. In Section 4 we realize filters within the MCS formalism, such that these properties are retained.

3.2 Comparing Diagnoses

To compare minimal diagnoses, we first consider an arbitrary preference order to select most preferred diagnoses, and further on focus on CP-nets. A *preference order* over diagnoses for an MCS M is a transitive binary relation $\preceq$ on $2^{br_M} \times 2^{br_M}$; we say that D is preferred to D' iff $D \preceq D'$.

Definition 2. *Let M be an inconsistent MCS. A diagnosis $D \in D^{\pm}(M)$ of M is called pre-most preferred iff for all $D' \in 2^{br_M} \times 2^{br_M}$ with $D' \preceq D \wedge D \not\preceq D'$ it holds that $D' \notin D^{\pm}(M)$. A diagnosis $D \in D^{\pm}(M)$ is called* most preferred, *iff D is subset-minimal among all pre-most preferred diagnoses.*

Given that MCSs are decentralized systems, users may want to express preferences on diagnoses solely based on a local set of bridge rules, assuming all other things equal. Such preferences can be formalized using CP-nets, which are an extension of *ceteris paribus* orders ("all else being equal"). They represent local preference and have successfully been used for preference elicitation (e.g. [7]).

Example 5. Assume an MCS where several corporations make contracts using bridge rules. Contract details, such as when a contract will start, how long it is valid, who owns what to whom, etc, are encoded with bridge rules. For instance, C_1 is leasing a car from C_2 with the following properties encoded as bridge rules $r_1 = (C_1 : \mathit{pay}(\mathit{car}, 500)) \leftarrow (C_2 : \mathit{price}(\mathit{car}, 500))$ and $r_2 = (C_1 : \mathit{due}(\mathit{car}, \mathit{monthly})) \leftarrow (C_2 : \mathit{due}(\mathit{car}, \mathit{monthly}))$. If r_2 is removed to restore consistency, r_1 becomes meaningless and possibly confuses further reasoning. Removing both rules is then preferred to removing only r_2.

A *CP-net* is a directed graph (V, E) where V is a finite set of variables (attributes) and $E \subseteq V \times V$ is the conditional dependency between variables. For $v \in V$ we denote the set of parents of v by $\mathit{pa}(v) = \{v' \in V \mid (v', v) \in E\}$.

Furthermore, the set of outcomes of a variable v is denoted by $dom(v)$. Preferences on the outcomes of a variable are specified in terms of total preorders, which allow indifference. A relation $\precsim$ is a *total preorder*, iff it is transitive, reflexive, and for any two elements o, o' of $\precsim$ it holds that $o \precsim o' \vee o' \precsim o$.

Each vertex v in a CP-net (V, E) is associated with a *conditional preference table* (CPT) p_v that maps each combination of outcomes of parents of v, i.e., $o \in dom(p_1) \times \ldots \times dom(p_n)$, to a total preorder $\precsim_v(o) \subseteq dom(v) \times dom(v)$ over the outcomes of v.

We associate a CP-net (V, E) with an MCS M, if every variable $v \in V$ is assigned a set of bridge rules $rules(v) \subseteq br_M$, such that the assignment of rules is disjoint, i.e., $\forall v, v' \in V : rules(v) \cap rules(v') = \emptyset$. Moreover, $dom(v)$ for every v is given by $dom(v) = \{unchanged_r, removed_r, unconditional_r \mid r \in rules(v)\}$. In the following we confine here to CP-nets that are acyclic, i.e, the directed graph (V, E) contains no cycles, and whose preference graph over outcomes is acyclic.

Example 6 (Ex. 5 ctd.). Recall r_1 and r_2 encoding properties of a leasing contract. If r_2 is removed, r_1 is preferred to be removed, too. Consider an associated CP-net $N = (\{v_1, v_2\}, \{(v_2, v_1)\})$, i.e. $pa(v_1) = \{v_2\}$ and $pa(v_2) = \emptyset$, where $rules(v_1) = \{r_1\}$, $rules(v_2) = \{r_2\}$. Assuming that adding rules unconditionally is always considered to be the worst option, v_1's conditional preference table is:

$$p_{v_1}(unchanged_{r_2}) = unchanged_{r_1} \prec_{v_1} removed_{r_1} \prec_{v_1} unconditional_{r_1} \quad (2)$$

$$p_{v_1}(removed_{r_2}) = removed_{r_1} \prec_{v_1} unchanged_{r_1} \prec_{v_1} unconditional_{r_1} \quad (3)$$

$$p_{v_1}(unconditional_{r_2}) = unchanged_{r_1} \prec_{v_1} removed_{r_1} \prec_{v_1} unconditional_{r_1} \quad (4)$$

For v_2 the table p_{v_2} is $unchanged_{r_2} \prec_{v_2} removed_{r_2} \prec_{v_2} unconditional_{r_2}$.

A CP-net N induces a preference graph G_N over outcomes, where each global outcome is a node in the preference graph. An arc from outcome o_i to o_j indicates that a preference for o_j over o_i can be determined directly from one conditional preference table of the CP-net (cf. [3]). The transitive closure G_N^+ of a preference graph induces a partial order on global outcomes. Furthermore, for a CP-net associated with an MCS, every global outcome represents a potential diagnosis.

Proposition 1. *Let M be an inconsistent MCS, and let N be a CP-net associated with M. Then G_N^+ induces a preference order $\prec$ over diagnoses of M.*

By $D^{\pm}_{opt}(M, N)$ we denote the subset-minimal among the most preferred diagnoses according to G_N^+, i.e., for which no other diagnosis is more preferred.

The semantics of CP-nets may also be defined in terms of *flips*: Let $|V| = m$, and let $\boldsymbol{a} = (a_1, \ldots, a_i, \ldots, a_m)$ and $\boldsymbol{b} = (a_1, \ldots, a_{i-1}, b_i, a_{i+1}, \ldots, a_m)$ be two global outcomes with $a_j, b_j \in dom(v_j)$, that differ only in the outcome of one variable. The flipping of v_i from a_i to b_i is *improving*, iff in the CPT of v_i outcome b_i is preferred over a_i, given all other parent variables set as in $\boldsymbol{a}$ and $\boldsymbol{b}$. The converse notion of an improving flip is called a *worsening* flip. A global outcome is *optimal*, if no improving flips are possible. Notably, for CP-nets an optimal outcome is reachable from any outcome by a finite sequence of improving flips.

In terms of flips, the most preferred diagnoses $D^{\pm}_{opt}(M, N)$ of an MCS are: $D^{\pm}_{opt}(M, N) = \min_{\subseteq}\{D \in D^{\pm}(M) \mid \forall D' \in D^{\pm}(M) : \mathit{iflips}(D, D') = \emptyset\}$, where $\mathit{iflips}(D, D')$ denotes the set of sequences of improving flips from D to D'.

4 MCS-Realization

We now present ways to realize filters, preference orders, and CP-nets. All realizations use a rewriting technique transforming an MCS M into an extended MCS M', where certain new contexts can do meta-reasoning on diagnoses of the original M. This is achieved in a way, such that a diagnosis of M' directly corresponds to a diagnosis of M, and subset-minimal diagnoses of M' coincide with the preferred diagnoses of M.

Meta-reasoning as described below allows certain contexts to observe whether a bridge rule of M is part of a diagnosis. For this, the context observes the body and head beliefs of a bridge rule. For a diagnosis (D_1, D_2) and a bridge rule r, if the body of r is satisfied, but its head is not believed, then $r \in D_1$; if the body is not satisfied, but the head is believed, then $r \in D_2$. The observation of body and head beliefs is accomplished by additional bridge rules in M', that are not subject to diagnosis. We thus adapt the notion of diagnosis such that certain bridge rules, tagged as protected, are never part of it.

Definition 3. *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 minimal such diagnoses by $D^{\pm}_{m}(M, br_P)$.*

A direct consequence is the following:

Proposition 2. *Let M be an inconsistent MCS with protected rules br_P. Then $D^{\pm}_{(m)}(M, br_P) \subseteq D^{\pm}_{(m)}(M)$, i.e., every (minimal) diagnosis excluding protected rules is a (minimal) diagnosis.*

Furthermore one can show that the duality between diagnoses and inconsistency explanations (cf. [8]) also holds for diagnoses and inconsistency explanations excluding protected rules, and that computing such diagnoses has the same complexity as computing ordinary diagnoses.

Meta-Reasoning Transformation: Using additional protected bridge rules in order to observe a bridge rule r with head $(k : s)$, we aim at monitoring the import of belief s into context k. Accessing k directly however, will in general not serve this purpose, since s could be in an accepted belief set of k also without import. In order to observe r properly, we therefore introduce a relay context for k, which can then be accessed by an observer.

Given an MCS M and a set of bridge rules br_o to be observed, an observation context ob for br_o is a context with bridge rules $br_{ob} = br^{ob}_b \cup br^{ob}_h$ with $br^{ob}_b = \{r^{ob}_b \mid r \in br_o\}$ and $br^{ob}_h = \{r^{ob}_h \mid r \in br_o\}$, where r^{ob}_b and r^{ob}_h are of the form

$$(ob : \mathit{body}_r) \leftarrow (c_1 : p_1), \ldots, (c_j : p_j), \mathbf{not}\ (c_{j+1} : p_{j+1}), \ldots, \mathbf{not}\ (c_m : p_m). \quad (5)$$

$$(ob : \mathit{head}_r) \leftarrow (\mathit{relay}_k : s). \quad (6)$$

for a bridge rule (1), respectively. Here, $relay_k$ is the relay context for context k (cf. below). Context ob is conservative iff $\mathbf{ACC}_{ob}(S) \neq \emptyset$ for every $S \subseteq \{hd\,(br_{ob})\}$.

Now, let $B_k = \{s \mid (k : s) \leftarrow \top \in heads(br_o)\}$. We say that $relay_k$ is a relay context for context C_k wrt. br_o iff $\mathbf{KB}_{relay_k} = \mathbf{BS}_{relay_k} = 2^{B_k}$, $\mathbf{ACC}_{relay_k}(S) = \{S\}$, $kb_{relay_k} = \emptyset$, and $br_{relay_k} = \{r_{relay} \mid r \in br_o\}$, where r_{relay} is of the form

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

for a bridge rule of the form (1). Furthermore, we associate with context $C_k = (L_k, kb_k, br_k)$ its relayed context $C_k^{rel} = (L_k, kb_k, (br_k \setminus br_{ob}) \cup br_k^{rel})$ wrt. br_o, where $br_k^{rel} = \{r_{rel} \mid r \in br_o\}$, and r_{rel} is for a bridge rule (1) of the form

$$(k : s) \leftarrow (relay_k : s). \quad (8)$$

Based on this, the meta-reasoning transformation of an MCS is as follows.

Definition 4. *Given an MCS $M = (C_1, \ldots, C_n)$, let $B = \{(ob_1, br_{o_1}), \ldots, (ob_m, br_{o_m})\}$ be an association of observation contexts $ob_i \notin M$ to disjoint sets of bridge rules $br_{o_i} \subseteq br_M$. The* meta-reasoning transformation *M^B of M wrt. B is the MCS $M^B = (C_1^{rel}, \ldots, C_n^{rel}, relay_1, \ldots, relay_n, ob_1, \ldots ob_m)$, where C_i^{rel} and $relay_i$ are relayed contexts and relay contexts wrt. $\bigcup_{k=1}^{m} br_{o_k}$, respectively, and $br_P^B = \bigcup_{i=1}^{n} br_i^{rel} \cup \bigcup_{k=1}^{m} br_{ob_k}$ are protected rules.*

In cases where contexts are known to not interfere with the beliefs they are importing, a simpler transformation can be obtained where observation contexts directly import from the original contexts and relay contexts are omitted.

Suppose that ob is an ASP context for the observation of a set of rules br_o of an MCS M. Then the following ASP rules allow ob to check whether r is part of a subset-minimal diagnosis:

$$r_{removed} \leftarrow body_r, not\ \ head_r. \quad (9)$$

$$r_{unconditional} \leftarrow not\ \ body_r, head_r. \quad (10)$$

$$r_{unchanged} \leftarrow not\ \ r_{removed}, not\ \ r_{unconditional}. \quad (11)$$

Note that this is correct for diagnoses $D = (D_1, D_2)$ with $D_1 \cap D_2 = \emptyset$ as otherwise for $r \in D_1 \cap D_2$, ob will not observe $r_{removed}$. We call a diagnosis $D = (D_1, D_2)$ *safe* iff $D_1 \cap D_2 = \emptyset$ holds and for any equilibrium S of $M[br_M \setminus D_1 \cup heads(D_2)]$ holds that $r \in D_1$ only if r is applicable in S and $r \in D_2$ only if r is not applicable in S. Note that all minimal diagnoses are safe. In the following theorem we consider only safe diagnoses.

Theorem 1. *Let M^B be the meta-reasoning transformation of an MCS M wrt. $B = \{(ob_1, br_{o_1}), \ldots, (ob_n, br_{o_n})\}$, let $ob_1, \ldots, ob_n$ be conservative, and let $br_o = \bigcup_{k=1}^{n} br_{o_k}$. Then,*
(i) $(D_1, D_2) \in D^{\pm}_{(m)}(M^B, br_p^B)$ implies $(D_1', D_2') \in D^{\pm}_{(m)}(M)$, where $D_i' = (D_i \cap br_M) \cup \{r \in br_M \mid r_{relay} \in D_i\}$ for $1 \leq i \leq 2$, and
(ii) $(D_1, D_2) \in D^{\pm}_{(m)}(M)$ implies $(D_1', D_2') \in D^{\pm}_{(m)}(M^B, br_p^B)$, where $D_i' = (D_i \setminus br_o) \cup \{r_{relay} \mid r \in D_i \cap br_o\}$ for $1 \leq i \leq 2$.

This theorem effectively states that the meta-reasoning transformation enables observation contexts to correctly observe the effects of diagnosis. This is the basis of the following realizations, which use non-conservative observation contexts for assessing and pruning diagnoses.

4.1 Filters

Using the above transformation, users of MCSs can analyze diagnoses inside observation contexts a way they see fit. If a diagnosis is considered inappropriate, the observer just needs to become inconsistent which prevents a corresponding diagnosis of the transformed system. This holds because the assessment uses protected bridge rules only. We next present a transformation realizing a general filter. For generality, it uses a central context m_f for analysis.

Definition 5. *Let M be an MCS and f a filter for M. A* filter-transformation *of M wrt. f is a meta-reasoning transformation M^B with $B = \{(m_f, br_M)\}$ and the logic L_f of m_f is such that for any diagnosis $D = (D_1, D_2)$:*

$$\mathbf{ACC}_{m_f}\left(kb_{m_f} \cup \{r_{body} | r \in D_1\} \cup \{r_{head} | r \in D_2\}\right) = \emptyset^1 \text{ iff } f(D) = 0.$$

Example 7 (Ex. 4 ctd.). For our scientists, we want to filter diagnoses that modify r_2 and r_3 differently, e.g., $D = (\{r_2\}, \emptyset)$. The filter-transformation yields a system with five contexts, K, $relay_K$, J, $relay_J$, and m_f. The rewritten rules for r_2 are (analogous for r_1 and r_3):

$$\begin{aligned}(relay_J : coauthored) &\leftarrow (K : contribute).\\ (J : coauthored) &\leftarrow (relay_J : coauthored).\\ (m : body_{r_2}) &\leftarrow (K : contribute).\\ (m : head_{r_2}) &\leftarrow (relay_J : coauthored).\end{aligned}$$

We use answer-set semantics for the assessment context m_f. To realize the filter function, m_f contains rules (9) - (11) for r_1 and r_2 and:

$$\begin{aligned}\bot &\leftarrow not\ same\_change.\\ same\_change &\leftarrow r_{1\,unconditional}, r_{2\,unconditional}.\\ same\_change &\leftarrow r_{1\,unchanged}, r_{2\,unchanged}.\\ same\_change &\leftarrow r_{1\,removed}, r_{2\,removed}.\end{aligned}$$

One can show that the transformation indeed realizes any given filter:

Theorem 2. *Given an inconsistent MCS M, let f be a filter on diagnoses and let M_f be a filter-transformation for f of M with protected rules br_P. Then $D \in D^{\pm}_m(M_f, br_P)$ iff $D \in D^{\pm}_{f,m}(M)$.*

If f is not given abstractly as a function, but as a family of constraints, each set of constraints is realizable using a separate assessment context, observing just the bridge rules it needs to assess. Realizing such a filter is decentralized, adheres to information hiding, and each observer's logic can be chosen as desired.

---

[1] For logics having always an acceptable belief set, inconsistency can still be created using a new bridge rule $(m_f : inc) \leftarrow (m_f : cause\_inc), \mathbf{not}\ (m_f : inc)$.

4.2 Preference Orders and CP-nets

The general idea for realizing best outcomes of a CP-net is to create an order-preserving mapping from the CP-net preference to the subset-order of diagnoses. We add a new context m_v for each variable v of the CP-net which "observes" the bridge rules with outcomes represented by v. It searches for improving flips on a path to a more preferred diagnosis. If the combined local guesses succeed, some bridge rules can be removed; ensuring that the zero-length path allows no removal, the most preferred diagnoses are those without removal. To accomplish this, we use prioritized bridge rules whose minimization has precedence.

Definition 6. *Let M be an MCS with bridge rules br_M, protected rules br_P, and prioritized rules $br_H \subseteq br_M$. The set of minimal prioritized diagnoses is*

$$D_m^{\pm}(M, br_P, br_H) = \{\ D \in D_m^{\pm}(M, br_P) \mid \forall D' \in D_m^{\pm}(M, br_P) : \\ D' \cap br_H \subseteq D \cap br_H \Rightarrow D' \cap br_H = D \cap br_H\ \}\,.$$

where $(D_1, D_2) \cap S := (D_1 \cap S, D_2 \cap S)$.

Note that given D, $D_m^{\pm}(M, br_P)$, and br_H, deciding $D \in D_m^{\pm}(M, br_P, br_H)$ is easy.

Let M be an MCS and consider an associated CP-net $N = (V, E)$. We assume that the CPT of any $v \in V$ is given by $v_n = 2^{|dom(p_v)|} \times 2 \times |dom(v)|$ "binarized" preferences (two successive outcomes); let $m = \sum_{v \in V} v_n$ be their total number. Given any total and strict preference order $<_r$ on the set $V_r \subseteq V$ of root nodes in N, we call $E' = E \cup \{(v, v') \mid v <_r v'\}$ a *root extension* of E.

Let $2M$ be the mirrored M, i.e., add a copy C_i' of each of context C_i in M with disjoint beliefs (alphabetic variants). An observer associated with $v \in V$ sees $rules(v)$ and $crules(v) = \{cr \mid r \in rules(v)\}$, where cr is the rule copy of r.

We say a set Q *encodes* o iff (a) $head_r \in Q \Leftrightarrow body_r \in Q$ for $o = unchanged_r$, (b) $head_r \notin Q$ and $body_r \in Q$ for $o = removed_r$, (c) $head_r \in Q$ and $body_r \notin Q$ for $o = unconditional_r$; that Q encodes co is analog. Eventually, let

$$\mathcal{L}_v = \{vf_{i,k},\ vg_{i,k},\ vng_{i,k},\ v'used_k,\ v'diff,\ o,\ co,\ io,\ o_k,\ co_k,\ eq,\ false\},$$

where $1 \leq i \leq v_n$, $1 \leq k \leq m$, $v' \in V$, and $o \in dom(v)$. A set $S \subseteq 2^{\mathcal{L}_v}$ is *compatible with* Q, iff

- $vdiff \in S$ iff (a) $v'diff \in Q$ for some $v' \neq v$, or (b) $o' \in Q$, and $co' \in Q$, such that o corresponds to r, co' corresponds to cr, and $o \neq o'$;
- $vused_k \in S$ iff $vg_{i,k} \in Q$ for some $1 \leq i \leq v_n$ and $1 \leq k \leq m$;
- $eq \in S$ iff $eq \in Q$ or v is $<_r$-maximal in V_r and $v'diff \notin S \cup Q$ for any $v' \in V$;
- $o_1 \in S$ iff $o \in Q$;
- $o_{k+1} \in S$ for $o \in dom(v)$ iff (a) o is most preferred according to v's i-th CPT entry and $vg_{i,k} \in Q$, or (b) o is the outcome of $o_k \in S$ and $vg_{i,k} \notin Q$;
- $false \in S$ iff for some $1 \leq i \leq v_n$, $1 \leq k \leq m$: (a) $io \in Q$ and $o \in Q$, (b) $eq \in S$ and $vf_{i,k} \in Q$, (c) $vg_{i,k}, vng_{i,k} \in Q$, (d) $vng_{i,k}, vf_{i,k} \in Q$, (e) $v'used_k, vf_{i,k} \in Q$, where $v' \neq v$, (f) $v'used_k, v''used_k \in Q$, where $v' \neq v''$, (g) $vg_{i,k} \in Q$ and either v's i-th CPT entry is not applicable wrt. $\{o \mid o_k \in Q \wedge o \in \bigcup_{v' \in pa(v)} dom(v')\}$, or it is not improving wrt. $o \in dom(v)$ such that $o_k \in Q$; (h) $o \neq o'$ for o and o' from $dom(v)$, such that $o_m \in S$, and $co' \in Q$.

Definition 7. *Let $N = (V, E)$, $V = \{v_1, \ldots, v_n\}$, be a CP-net associated with an MCS M, and let E' be a root extension of E. The* CP-net transformation of M wrt. N *is the meta-reasoning transformation $2M^B$ of $2M$ wrt. $B = \{(m_{v_i}, rules(v_i) \cup crules(v_i)) \mid 1 \leq i \leq n\}$, putting in observation contexts*

- *for every $v \in V$, $1 \leq i \leq v_n$, and $1 \leq k \leq m$: protected bridge rules $(m_v : vg_{i,k}) \leftarrow not\ (m_v : vng_{i,k})$ and $(m_v : vng_{i,k}) \leftarrow not\ (m_v : vg_{i,k})$, and prioritized bridge rules $(m_v : vf_{i,k}) \leftarrow \top$ and $(m_v : io) \leftarrow \top$ for $o \in dom(v)$;*
- *for every $(v', v) \in E'$, a protected bridge rule $(m_v : eq) \leftarrow (m_{v'} : eq)$;*
- *for every $(v, v') \in E'$, and $1 \leq k \leq m$, protected bridge rules $(m_v : v'used_k) \leftarrow (m_{v'} : v'used_k)$ and $(m_v : v'diff) \leftarrow (m_{v'} : v'diff)$;*
- *for every $(v', v) \in E$, $o \in dom(v')$, and $1 \leq k \leq m$: protected bridge rules $(m_v : o_k) \leftarrow (m_{v'} : o_k)$.*

Furthermore, the logic of every m_v ($v \in V$) has $\mathbf{BS}_{m_v} = 2^{\mathcal{L}_v}$, and $\mathbf{ACC}_{m_v}(kb_{m_v} \cup Q) = S$ for any set S in $\mathbf{BS}_{m_v}$ compatible with Q and $false \notin S$.

Let br^i_H denote the set of all prioritized bridge rules of the form $(m_v : vf_{i,k}) \leftarrow \top$ in $2M^B$. Intuitively, this transformation ensures the following for any observations o and co'. If $o = o'$ (the outcomes correspond to the same diagnosis of the original system), then bridge rules br^i_H need to be removed to obtain a diagnosis of $2M^B$. If $o \neq o'$ (different diagnoses), then there are two cases: either o' is reachable from o via a sequence of improving flips (thus a diagnosis of $2M^B$ exists removing only some of br^i_H), otherwise there is no diagnosis of $2M^B$ for this observation (an inconsistency by the protected rules, which cannot be resolved by removing prioritized rules).

Example 8 (Ex. 6 ctd.). Recall the contracts example, and assume that the resulting MCS M is inconsistent having only two diagnoses $D = (\{r_1, r_2\}, \emptyset)$ and $D' = (\{r_2\}, \emptyset)$. Note that D is preferred over D' wrt. N. Consider the case where $o = D$ and $co' = D'$ are observed in $2M^B$: Flipping r_1 from *unchanged* to *removed* (using the 3rd rule of v_2's CPT) improves the outcome, hence prioritized rules $br^i_H \setminus \{m_{v_1} : v_1 f_{3,1}\}$ have to be removed to restore consistency. Since this is a subset of br^i_H, D is not most preferred (see also the following theorem).

Theorem 3. *Let M_N be the CP-net transformation of an inconsistent MCS M wrt. an associated CP-net N with protected rules br_P and prioritized rules br_H. Then, $D \in D^{\pm}_{opt}(N, M)$ iff $D' \in D^{\pm}_m(M_N, br_P, br_H)$ such that $D' \cap br^i_H = br^i_H$ and $D' \cap br_M = D$.*

The techniques of meta-reasoning and prioritized diagnoses can be used to realize arbitrary preference orders on diagnoses. This is achieved by introducing a global assessment context, i.e., an observation context for all bridge rules of the original system, and (exponentially many) new prioritized bridge rules, to which the preference order is mapped. For space reasons, however, we omit details.

5 Discussion

Computational complexity: The generalized notions of diagnoses can be realized by transformations to diagnoses of an MCS, using assessment contexts. The respective bridge rules can be set up efficiently for filters and preference orders. The number of bridge rules for filters increases by at most a factor of 4 per bridge rule while for a CP-net (V, E) it is quadratic in the size of the CPTs. Then, the complexity of preferred diagnoses does not increase over that of ordinary diagnoses, if the preference assessment in the analysis contexts has not higher complexity than regular contexts.

In particular, deciding whether a given pair (D_1, D_2) of bridge rules is a prioritized diagnosis (excluding protected bridge rules) is of the same complexity as recognizing ordinary diagnoses; from the results in [8], the complexity ranges, depending on the computational complexity of contexts, from **coNP** (for **P** and **NP** contexts) to $\mathbf{D_2^P}$ (for $\mathbf{\Sigma_2^P}$contexts, e.g. disjunctive ASP contexts). Detecting subset-minimal diagnosis is $\mathbf{D^P}$-complete (for **P**, **NP**, and **coNP**contexts).

From the requirement that preferred diagnoses also be minimal follows that this complexity can not be improved, even for easily evaluable preference orders.

Decentralization: A key property of MCSs is decentralized information exchange. Filters and preference orders can be realized in such a way. If a filter or preference order is composed of local tests, then the realization can be broken down to these local, decentral tests. An example of this is the realization of CP-nets, where information is only exchanged as much as necessary to realize the CP-net.

Quantitative Assessment: As system knowledge to rank diagnoses is not always available, one may consider a quantitative inconsistency measure. Following the approach of [10], which is based on the cardinalities of the minimal inconsistent sets a certain formula belongs to, a measure on minimal inconsistent sets of bridge rules may be established. A notion of such sets is given in [8], termed inconsistency explanation, which is a pair of bridge rules (E_1, E_2) where E_1 is a minimal inconsistent set of bridge rules causing inconsistency and E_2 contains rules which could resolve the inconsistency if some were applicable.

Based on this we may define an *inconsistency measure* as follows. Let M be an MCS and $r \in br_M$, and let $A_r^i(M) = \{(E_1, E_2) \in E_m^{\pm}(M) \mid r \in E_i\}$, $i = 1, 2$ where $E_m^{\pm}(M)$ is the set of subset-minimal inconsistency explanations of M:

$$m(M, r) = \Big(\sum_{(E_1,E_2)\in A_r^1(M)} \frac{1}{|E_1|}, \sum_{(E_1,E_2)\in A_r^2(M)} \frac{1}{|E_2|} \Big).$$

Thus, $m(M, r)$ measures the inconsistency of r in M by counting the relative contribution of r to the minimal inconsistent sets in M, respectively, its contribution to resolving inconsistency. A key property of a measure is monotonicity (in the form of sub-additivity), which m is lacking in general. Finding such an inconsistency measure for MCSs faces the problem of providing a monotonic measure for a non-monotonic system. It remains to explore whether restrictions on inconsistency explanations can serve this purpose.

6 Related Work and Conclusion

Although inconsistency handling and preferences are widely used in knowledge-based systems, their application to systems interlinking different knowledge bases is still rare. In [5] argumentation context systems which equip special MCS with mediators are introduced. Mediators realize two tasks: First, they guard a context by controlling its import information. Second, they restore consistency using local information. In our approach, the first task is done by import relays of the meta-reasoning transformation, and the second task is achieved by the analyzing contexts which may be local or global as needed. This allows us not only to establish global filters and preferences, but also the implementation is local and decentralized as possible.

Bikakis et al. [2] propose a certain way of inconsistency removal which is based on local trust orders. They propose several trust-based algorithms using such orders combined with provenance; their provenance-free algorithm can be realized with our approach. In our view, information hiding is an important aspect of MCSs, which however is in conflict with provenance. Thus the goals behind the works are different.

Future work includes the development and implementation of particular filter and CP-net transformations as well as further analysis of our and other possible measures of inconsistency in MCS.

References

1. Balduccini, M., Gelfond, M.: Logic programs with consistency-restoring rules. In: International Symposium on Logical Formalization of Commonsense Reasoning, AAAI 2003 Spring Symposium Series. pp. 9–18 (2003)
2. Bikakis, A., Antoniou, G., Hassapis, P.: Alternative strategies for conflict resolution in multi-context systems. In: AIAI. pp. 31–40 (2009)
3. Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H.H., Poole, D.: Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. J. Artif. Intell. Res. (JAIR) 21, 135–191 (2004)
4. Brewka, G., Eiter, T.: Equilibria in heterogeneous nonmonotonic multi-context systems. In: AAAI. pp. 385–390. AAAI Press (2007)
5. Brewka, G., Eiter, T.: Argumentation context systems: A framework for abstract group argumentation. In: LPNMR. pp. 44–57 (2009)
6. Brewka, G., Roelofsen, F., Serafini, L.: Contextual default reasoning. In: IJCAI. pp. 268–273 (2007)
7. Domshlak, C., Brafman, R.I., Shimony, S.E.: Preference-based configuration of web page content. In: Nebel, B. (ed.) IJCAI. pp. 1451–1456. Morgan Kaufmann (2001)
8. Eiter, T., Fink, M., Schüller, P., Weinzierl, A.: Finding explanations of inconsistency in nonmonotonic multi-context systems. In: KR (2010)
9. Giunchiglia, F., Serafini, L.: Multilanguage hierarchical logics or: How we can do without modal logics. Artif. Intell. 65(1), 29–70 (1994)
10. Hunter, A., Konieczny, S.: Measuring inconsistency through minimal inconsistent sets. In: Brewka, G., Lang, J. (eds.) KR. pp. 358–366. AAAI Press (2008)