

Hybrid Reasoning in BACK*

Bernhard Nebel, Kai von Luck
IBM Deutschland GmbH
Bereich Wissenschaft - LILOG
Schloßstraße 70
D-7000 Stuttgart 1

Abstract

Hybrid knowledge representation formalisms consist of two or more different subformalisms for representing different kinds of knowledge or knowledge in different kinds of representation formats. For a semantically well-founded hybrid formalism not only a precise semantics for each of the participating subformalisms has to be given but a semantics for the interrelationship between these subformalisms as well. A hybrid representation system therefore has to be implemented as a *hybrid reasoning system* taken into account these semantic models. The BACK system as an instance of this class of systems will be described with respect to the underlying semantic model and the demands for a reasoning component as one part of the realization of the formalism. The consequences and limits for the implementation of the BACK system are discussed.

1 Introduction

Aaron Sloman pointed out very clearly in [22] the need for different knowledge representation formalisms for the adequate representation of a realistic portion of the world. This position, contrasting the view that one uniform formalism is sufficient (e.g., [8]), is nowadays widely accepted. The most visible consequence of this is the emerging number of systems—so-called *hybrid systems*—supporting more than one knowledge representation formalism. Problems with this hybrid approach to knowledge representation and reasoning are that often system “seem to appeal to ‘the more the merrier,’ without any clear idea of how merrier is better” [4]. In particular, it is often not obvious how the different knowledge representation formalisms act together to form a coherent whole.

First of all, let’s clarify some terms. If we talk about a *knowledge representation formalism*, we mean a formal language with given *syntax* and *semantics*.¹ This language can be analyzed for its expressiveness and naturalness with respect to its intended purpose, and for decidability and inherent complexity of inference algorithms with respect to the formal semantics.

A *knowledge representation system* is a “materialization” of a formalism supporting the *interpretation* of well-formed expressions of a knowledge representation language by means of *inference algorithms* realizing the semantics to a certain extent. These inference algorithms should be *sound*—they should not lead to wrong propositions—and in the ideal case they should be *complete*—be able to deduce any true proposition. As it turns out, however, in the real world the latter goal can only be achieved if either the formalism is very simple or if we allow for arbitrary long computations. Therefore, often the solution is to provide only the obvious, easy to compute inferences and ignore the difficult ones.

Of course, a knowledge representation system is more than just a mechanized reasoner. Another task for such a system is the maintenance of represented knowledge, i.e., it has to account for *additions* to and *updates* of the represented knowledge. In fact, this is the main difference between a knowledge representation system and a static deductive calculus (or a programming language). Furthermore, any knowledge representation system claiming to be usable has to provide a friendly interface to the human user.² In the following, however, we will ignore the latter and focus on the former two points.

In the ideal case, a hybrid knowledge representation system is an implementation of a hybrid knowledge representation formalism, consisting of two or more different subformalisms. However, the mere combination of formalisms does not necessarily result in a hybrid formalism. A kind of “glue” is needed in order to constitute a hybrid formalism consisting of

*This work was done by the authors at the Technical University Berlin and partially supported by the EEC. It was part of the ESPRIT project 311, which involves the following participants: Nixdorf, Olivetti, Bull, Technische Universität Berlin, Università di Bologna, Universität Hildesheim and Università di Torino.

¹In the sense of, e.g., [14] or [7].

²And usually, a large fraction of code in a knowledge representation system is devoted to the user interface.

2. a *common semantics* for the overall formalism, explaining in a semantic sound manner the relationship between expressions of different subformalisms.

The representational theory should explain why there are different subformalisms, what their benefits are, and how they relate to each other. An answer should at least refer to adequacy criteria such as (cf. [13]):

1. *epistemological adequacy*, i.e., that the subformalisms are necessary to represent epistemological different kind of knowledge (e.g., analytic and contingent knowledge), or
2. *heuristic adequacy*, i.e., that the different subformalisms permit representation of the same knowledge in different ways for reasons of efficiency.

A necessary precondition for gluing things together is that their shapes fit, a fact which might be violated in designing a hybrid formalism, at least in the case where the subformalisms are intended to represent epistemological different kinds of knowledge. For example, if one subformalism permitted definition of terms by using time relationships, but none of the other subformalisms referred to time at all, the subformalisms would be in some sense *unbalanced*. This, however, can be easily uncovered by inspecting the common semantics.

There is a more subtle sense of how two formalisms in such a hybrid system can be unbalanced, which has to do with the fact that most knowledge representation systems are necessarily *incomplete* in their reasoning in order to provide answers in reasonable time. Because of this incompleteness there could be situations where one subformalism allows to express something which obviously should have some impact on another subformalism according to the common semantics, but the system does not realize this because it's reasoning is incomplete in this aspect. These "black holes" may be there because the incompleteness has principal reasons or because the subformalism is not heuristically adequate for this aspect. In any case, the subformalisms of the system appear to be unbalanced. Although the term *balancedness* is a little bit vague, it can be captured by the following *principle of balancedness in hybrid representation systems*:

If a representation construct in a subformalism of a hybrid formalism suggests that its usage has some impact on knowledge represented with another subformalism (according to the common semantics), then this should be realized by the reasoning component of the underlying system.

An example for a system with unbalanced subformalisms is KL-TWO [23]: While it is possible to define concepts with a very rich language, only a fraction of it is used for stating contingent propositions. In particular, the number restrictions used in the NIKL subformalism has only a very limited impact on the PENNI subformalism because the latter is not heuristically adequate to deal with cardinalities.

In the last few years, much effort has been devoted to the development of hybrid systems. Besides systems favoring multiple representations for the sake of naturalness and efficiency of the represented knowledge (for instance the CAKE system [20]), systems combining formalisms for the representation of knowledge according to the distinction Frege made between meaning (*Sinn*) and reference (*Bedeutung*) were developed. In particular, the connection of KL-ONE [3] derivatives as formalisms for representing terminological knowledge (*TBoxes*) with formalisms for representing assertions about the actual state of the world (*ABoxes*) has been investigated³ (e.g., KRYPTON [4], KL-TWO [23], MESON [5], BACK [11] and KANDOR [17]). The main points in this research were the design of an appropriate ABox, sometimes requiring severe restrictions of the expressivity in the TBox (e.g., in KRYPTON), and developing means for connecting the reasoning of TBox and ABox.

In the following we will present one particular solution to these problems pursued in BACK⁴. The following design criteria have been taken into account in developing BACK:

1. The subformalism of the BACK system should be balanced.
2. The BACK formalism should permit tractable inference algorithms covering almost all possible inferences⁵.
3. The ABox formalism of BACK should be able to represent incomplete knowledge in a limited manner (cf. [10]).
4. The BACK system should allow for extending the knowledge base incrementally (we do not consider retractions here, but cf. [16]).

³A rough similarity of this distinction is known in the database area with the distinction of database schemata and database contents, but these approaches take the schema definition as a source of integrity constraints and not as a formalism for the intensional definitions of terms.

⁴The **B**erlin **A**dvanced **C**omputational **K**nowledge Representation System.

⁵Unfortunately, complete and tractable inference algorithms are possible only for very simple TBoxes (cf. [2] and [15]).

The rest of the paper is divided in four parts. The next section gives a brief introduction to the KL-ONE-alike TBox employed in BACK. Then the BACK ABox is discussed, and it is shown how we achieved the goals concerning the formalism stated above. Finally, we investigate what kind of inferences the combination of the ABox and the TBox permit, and how these *hybrid inferences* can be realized algorithmically. This leads to a new view on the *realization* inference as discussed in [12], which can be characterized as a “symbolic constraint propagation” process.

2 Representation and Reasoning in the BACK TBox

KL-ONE as described in [3] is perfectly well suited for the introduction of a terminology. It allows defining *concepts* by stating *super-concepts* and the restrictions on relationships to other concepts, which are called *roles*. Concepts and roles correspond roughly to *generic frames* and *slots* in the frame terminology. Unlike frames, however, concepts are understood as purely *intensional*. Furthermore, concepts may be introduced as *primitive* or *defined*. In the former case, the concept description specifies only the *necessary* conditions, while in the latter case the concept description is *necessary* and *sufficient*.

The small formal language given in Figure 1, which is actually a subset of the BACK TBox formalism, will be used as our TBox formalism in this paper.

$\begin{aligned} \langle terminology \rangle & ::= \{ \langle term-introduction \rangle \mid \langle restriction \rangle \}^* \\ \langle term-introduction \rangle & ::= \langle concept-introduction \rangle \mid \\ & \quad \langle role-introduction \rangle \\ \langle concept-introduction \rangle & ::= \langle atomic-concept \rangle = \langle concept \rangle \mid \\ & \quad \langle atomic-concept \rangle \leq \langle concept \rangle \\ \langle role-introduction \rangle & ::= \langle atomic-role \rangle = \langle role \rangle \mid \\ & \quad \langle atomic-role \rangle \leq \langle role \rangle \\ \langle concept \rangle & ::= \langle atomic-concept \rangle \mid \\ & \quad (\mathbf{and} \langle concept \rangle^*) \mid \\ & \quad (\mathbf{all} \langle role \rangle \langle concept \rangle) \mid \\ & \quad (\mathbf{atleast} \langle number \rangle \langle role \rangle) \mid \\ & \quad (\mathbf{atmost} \langle number \rangle \langle role \rangle) \\ \langle role \rangle & ::= \langle atomic-role \rangle \\ \langle restriction \rangle & ::= (\mathbf{disjoint} \langle concept \rangle \langle concept \rangle) \end{aligned}$
--

Figure 1: BNF Definition of \mathcal{TF}

Using this language, we may, for instance, introduce the (slightly artificial) concept *Modern-team*.⁶ Assuming that *Human*, *Woman*, and *Team* are already introduced as concepts and that *member* is a role with *leader* as a *subrole*, a definition of *Modern-team* might look as follows:

Modern-team = (**and** **Team**
 (**atmost** 4 **member**)
 (**all** **member** **Human**)
 (**atmost** 1 **leader**)
 (**all** **leader** **Woman**))

This example shows some of the important *concept-forming operators* which may be used in order to create new concepts, namely *specialization* of a concept, expressed by the occurrence of a concept name inside an **and** expression; *number restriction* of a role, expressed by an **atmost** or **atleast** expression; and *value restriction* of a role, expressed by an **all** expression. In KL-ONE, concepts and roles are usually graphically depicted (cf. [3]), which may be sometimes easier to comprehend. However, it does not make any difference whether a linear notation or a graphical network is used, provided that a formal semantics can be given.

⁶In the following, concept names will be capitalized.

one concept is more special or more general than another. For instance,

Small-team = (**and** Team (**atmost** 5 member))

is a more general concept than Modern-team as defined above, even though this relationship has not been *explicitly* stated. Any *object* in the world which could be described as a Modern-team is a Small-team as well.

In order to capture this relationship formally, let us define the *semantics* of our small TBox language. We will follow here the ideas spelled out in [4], specifying the semantics model-theoretically. That is, any concept and role is associated with its *extension*—with the class of objects and relationships denoted. The *subsumption* relationships should then be the *necessary* set inclusion of the extensions.

Definition 1 (Extension Function) *Let \mathcal{T} be a terminology using \mathcal{TF} syntax.. Let \mathcal{D} , the domain, be any set. Let \mathcal{E} be any function from concepts and roles to subsets of \mathcal{D} and $\mathcal{D} \times \mathcal{D}$, respectively. Then \mathcal{E} is called an extension function wrt to \mathcal{T} iff:*

$$\begin{aligned} \mathcal{E}[c_1] \cap \mathcal{E}[c_2] &= \emptyset \text{ if '(\b{disjoint } c_1 c_2)' is in } \mathcal{T} \\ \mathcal{E}[t] &= \mathcal{E}[def] \text{ if ' } t = def \text{ ' is in } \mathcal{T} \\ \mathcal{E}[t] &\subseteq \mathcal{E}[def] \text{ if ' } t \leq def \text{ ' is in } \mathcal{T} \\ \mathcal{E}[(\b{and } c_1 \dots c_n)] &= \bigcap_{i=1}^n \mathcal{E}[c_i] \\ \mathcal{E}[(\b{all } r c)] &= \{x \in \mathcal{D} \mid \forall y : \langle x, y \rangle \in \mathcal{E}[r] \Rightarrow y \in \mathcal{E}[c]\} \\ \mathcal{E}[(\b{atleast } n r)] &= \{x \in \mathcal{D} \mid \|\{y \in \mathcal{D} \mid \langle x, y \rangle \in \mathcal{E}[r]\}\| \geq n\} \\ \mathcal{E}[(\b{atmost } n r)] &= \{x \in \mathcal{D} \mid \|\{y \in \mathcal{D} \mid \langle x, y \rangle \in \mathcal{E}[r]\}\| \leq n\} \end{aligned}$$

Recalling our argument that subsumption between concepts should be mirrored as a *necessary* set inclusion between extensions of concepts and roles, we can now precisely say what is meant by that.

Definition 2 (Subsumption) *Let \mathcal{T} be a terminology. Then we say the term t_1 is subsumed by the term t_2 in \mathcal{T} iff for any domain \mathcal{D} and any extension function \mathcal{E} of \mathcal{T} it holds that $\mathcal{E}[t_1] \subseteq \mathcal{E}[t_2]$.*

Inference algorithms computing this relationship are described, for instance, in [21]. As it turns out, a complete inference algorithm, even for the small language described, is intractable [15]. However, as mentioned above, we will be satisfied with a *sound* algorithm as long as all obvious relationships are uncovered—a claim very similar to the one made by Allen about his time reasoning system [1].

Summing up, KL-ONE-alike formalisms are very well in the area of defining concepts and reasoning about relationships between these concepts, but they are weak in stating any contingent truth about the world, for example that “KIM is a Woman, and that she is a member of some Modern-team, but not the leader of that Team.” Another example would be that “there is at least one Modern-team.” In order to compensate for this, an ABox has to be employed, a topic we will discuss next.

3 The Design of the BACK ABox

The formalisms for representing terminological knowledge in the hybrid systems mentioned differ only in what concept- and role-forming operators they provide—sometimes motivated by arguments concerning the computational complexity. The situation with ABoxes, however, is quite different. Here we meet a variety of approaches, e.g.:

KRYPTON	Full first-order predicate logic.
KL-TWO	Variable-free predicate logic with equality.
KANDOR	Object-centered, frame-like schema.
MESON	Object-centered, frame-like schema.

Using predicate logic (first-order or restricted) as the ABox has several advantages. One is that it gives a clear account to a *common semantics* for the entire formalism⁸. Another benefit of predicate logic is its *plasticity* as Hayes called it [6]—the possibility to give *partial descriptions* and to extend a knowledge base *incrementally*. On the other hand, there are a lot of disadvantages. One obvious problem with unrestricted first-order predicate logic are the computational costs, a fact which exclude such systems from being used as a knowledge representation system in practical AI systems. Besides that, there are also problems with the expressiveness of predicate logic. For instance, there is no easy way to state that a given description

⁷Any other meaningful relationship between concepts or property of a concept can be reduced to subsumption.

⁸In these approaches concepts are viewed as 1-place predicates and roles as 2-place predicates.

for the fact that *number restrictions* on roles are omitted from the KRYPTON TBox and that they are almost mere comments in the KL-TWO TBox.

For frame-like ABoxes as employed in KANDOR or MESON the situation is opposite. The *unique name assumption* and *closed world assumption* (cf. [19]) are taken for granted, very similar to conventional data bases. This forces the user of such a system to give a *complete* description of the state of affairs and prohibits incremental “monotonic” additions. However, this kind of representation—called “vivid” by Levesque in [9]—allows for elegant and simple inference algorithms. In the case of hybrid KL-ONE systems it implies that number restrictions and value restrictions of the TBox can be fully utilized in the ABox. A short note about formal semantics might be in order here. While frame-like schemata are often viewed as some ad hoc data structure, it is nevertheless possible to specify a precise formal semantics for at least a subset (cf. [7]).

In BACK we tried to combine the benefits of both approaches. An object-centered language was chosen to describe objects of the domain. However, instead of insisting that all role fillers for a given object have to be specified, the information can be *incomplete* and may be refined later. For this purpose the following operators are provided:

1. stating the cardinality of a role-filler set, e.g., “MARY has at least 2 children;”
2. disjunctive information on role fillers, as e.g., “MARY is married to JOHN or TOM;”
3. stating the exhaustiveness of the provided information selectively (cf. [6]), e.g., “MARY has as a friend TOM, and these are all friends MARY has.”

We will not describe the entire assertional formalism here,⁹ but specify only the part of the ABox which will be used in the following examples (cf. Figure 2).

$$\begin{aligned}
 \langle \text{world-description} \rangle & ::= (\langle \text{object-description} \rangle \mid \langle \text{relation-description} \rangle)^* \\
 \langle \text{object-description} \rangle & ::= (\langle \text{atomic-concept} \rangle \langle \text{object} \rangle) \\
 \langle \text{relation-description} \rangle & ::= (\langle \text{atomic-role} \rangle \langle \text{object} \rangle \langle \text{object} \rangle) \mid \\
 & \quad (\langle \text{atomic-role} \rangle \langle \text{object} \rangle (\mathbf{atleast} \langle \text{number} \rangle)) \mid \\
 & \quad (\langle \text{atomic-role} \rangle \langle \text{object} \rangle (\mathbf{atmost} \langle \text{number} \rangle)) \mid
 \end{aligned}$$

Figure 2: BNF Definition of \mathcal{AF}

Obviously, the assertional formalism \mathcal{AF} is very weak. We just can state that an object is in the extension of some concept and that two objects are related, or that one object is in relation to a number of other objects. The semantics of this formalism can be described as in the following definition.

Definition 3 (Model of a World Description) *Let \mathcal{W} be a world description in \mathcal{AF} syntax. Let \mathcal{D} , the domain, be any set. Let N_O, N_C, N_R be sets of objects, atomic concepts, and atomic roles, respectively. Let \mathcal{I} be a function*

$$\begin{aligned}
 \mathcal{I} & : N_O \rightarrow \mathcal{D} \\
 \mathcal{I} & : N_C \rightarrow 2^{\mathcal{D}} \\
 \mathcal{I} & : N_R \rightarrow 2^{\mathcal{D} \times \mathcal{D}}
 \end{aligned}$$

Then the structure $\langle \mathcal{D}, \mathcal{I} \rangle$ is called a model of \mathcal{W} iff

$$\begin{aligned}
 \mathcal{I}[o_1] \neq \mathcal{I}[o_2] & \text{ if } o_1, o_2 \in N_O \text{ and } o_1 \neq o_2 \\
 \mathcal{I}[o] \in \mathcal{I}[c] & \text{ if } (c \ o) \text{ is in } \mathcal{W} \\
 \langle \mathcal{I}[o_1], \mathcal{I}[o_2] \rangle \in \mathcal{I}[r] & \text{ if } (r \ o_1 \ o_2) \text{ is in } \mathcal{W} \\
 \|\{ \langle \mathcal{I}[o], x \rangle \in \mathcal{I}[r] \}\| \geq n & \text{ if } (r \ o \ (\mathbf{atleast} \ n)) \text{ is in } \mathcal{W} \\
 \|\{ \langle \mathcal{I}[o], x \rangle \in \mathcal{I}[r] \}\| \leq n & \text{ if } (r \ o \ (\mathbf{atmost} \ n)) \text{ is in } \mathcal{W}
 \end{aligned}$$

⁹For a more complete treatment cf. [10].

While it appears to be very obvious what kind of inferences are granted by the separate formalisms, the interesting problems stand up if we combine the reasoning of the TBox and the ABox.

The TBox alone is only good for *terminological inferences*, including detecting implicit specialization relationships between concepts, recognizing contradictory concepts and computing the properties inherited from superconcepts. The ABox alone, depending on its flavor, can draw the inferences sanctioned by first-order logic (in the KRYPTON case), by variable-free predicate logic with equality (in KL-TWO), or it can infer whether a domain object is related to other objects, and how many there are (in the case of object-centered ABoxes), as in our case.

In order to demonstrate the interaction between ABox and TBox in BACK let us investigate a small example. Let us assume the concept definitions given in Figure 3.

Man	\leq	Human
Woman	\leq	Human
(disjoint Man Woman)		
Team	$=$	(and Set (all member Human) (atleast 2 member))
Male-team	$=$	(and Team (all member Man))
Small-team	$=$	(and Team (atmost 5 member))
leader	\leq	member
Modern-team	$=$	(and Team (atmost 4 member) (atleast 1 leader) (all leader Woman))

Figure 3: A Formal Terminology Using \mathcal{TF} Syntax

Furthermore, let us assume that object descriptions are entered into the system as displayed in Figure 4. If we interpret this partial “world description” taking the terminology given in Figure 3 into account, there are a number of relationships which can be deduced. For instance, the JUNK team is obviously a Small-team as well, because Small-team subsumes Modern-team. However, there are a number of other relationships as well. In order to get an idea what can be inferred from such world descriptions combined with a terminology, it might be helpful to have a formal definition.

(Man	TOM)
(Man	DICK)
(Man	HARRY)
(Modern-team	JUNK)
(member	JUNK TOM)
(member	JUNK DICK)
(member	JUNK HARRY)
(member	JUNK KIM)
(Male-team	CHAUVIS)
(member	CHAUVIS KIM)

Figure 4: A Formal World Description

Interpreting a terminology and a world description in combination means that we have to relate the models of a world description and the extension functions of a terminology somehow. While we may view extensions of a terminology as possible structures induced by the way we have organized our vocabulary, world descriptions are partial descriptions of how the world actually is supposed to be. However, such partial descriptions should, of course, respect the relationships layed down in a terminology. Respecting a terminology means for a world description that the models we really intend to have are those which are at the same time extensions of a terminology.

Definition 4 (Models Respecting a Terminology) *Any model $\langle \mathcal{D}, \mathcal{I} \rangle$ of a world description \mathcal{W} is said to respect the terminology \mathcal{T} , written $\langle \mathcal{D}, \mathcal{I} \rangle_{\mathcal{T}}$, iff there is an extension function \mathcal{E} over \mathcal{D} such that it holds for the restriction of \mathcal{I} to atomic concepts and roles $\mathcal{I}^{N_C \cup N_R}$, and for the restriction of \mathcal{E} to introduced terms \mathcal{E}^{N_d} that: $\mathcal{I}^{N_C \cup N_R} = \mathcal{E}^{N_d}$.*

descriptions and a terminology, as spelled out in the next definition:

Definition 5 (Hybrid Entailment) *Let \mathcal{W} be a world description and \mathcal{T} be a terminology. Then we say a object or relation description d is hybridly entailed by \mathcal{W} and \mathcal{T} iff the models of \mathcal{W} respecting \mathcal{T} are the same as the models of \mathcal{W} extended by d respecting \mathcal{T} .*

Taking these definitions into account, we see that there are a couple of descriptions which are entailed. For instance, it must be the case that KIM is a Man as well, because all members of a Male-team have to be. Moreover, because at least one of the members of the JUNK team has to be a Woman, namely, the leader, there is a contraction in this world description. There is actually no model of the world description displayed in Figure 3 respecting the terminology in Figure 4.

5 Realization as Constraint Propagation

In order to find such contradictions and, more generally, to determine the concepts which most accurately describe a given object, a forward inference technique called *realization* [12] is usually employed. Realization is very similar to *classification*, an inference technique used to maintain the taxonomy of concepts in the TBox according to subsumption [21]. In fact, realization can be viewed as *abstraction*—generating a description in terms of the TBox—followed by *classification* of this description (cf. [23]).

A first approximation to the implementation of this inference could be realized as follows. After a new assertion about an individual enters the ABox (either a new concept or a new role filler) the following has to be done:

1. Propagate all role fillers of subroles to the corresponding superrole.
2. Determine the cardinality of the role filler sets for each role (this can be a range in the case of incomplete information). This cardinality information is used as the actual number restriction for the role in the abstraction process.
3. If a role-filler set is closed, i.e., all potential candidates are known, then the *generalization*¹⁰ of the descriptions of all potential role fillers can serve as the actual value restriction in the abstraction process.
4. Now, the generated number and value restrictions, the old description of the individual as well as the new one can be used to construct a concept definition which can be classified and after that serve as the new *most general specialization* (MSG) of the object under investigation.

This algorithm does take into account all information which is supplied *locally* to an individual, but ignores any non-local consequences. In the example of the last section, we would at least require that after the abstraction process the role fillers are to be specialized according to the value restriction of the Male-team. Additionally, we note that in order to detect the contradiction in the example we also have to account for the case that the specialization of an individual (KIM from Human to Man) can lead to the specialization of another individual (JUNK from Modern-team to the empty concept) the first one is a role filler of.

If we analyze this algorithm more thoroughly, we may note that this process can trigger other specializations which in turn may propagate restrictions. Because this sounds very expensive from a computational point of view, one could argue that because the system reasoning process is incomplete anyway, it is legitimate to restrict the resources allocated for the realization process and leave such situations alone. And this was indeed the first approximation to a solution we chose. However, this is not a general incompleteness of the inference algorithm, but it depends on the order of input! If, in our example, the order of input between the JUNK and the CHAUVI were reversed, the contradiction would be easily detected.

That is certainly not the kind of behaviour we expect from a knowledge representation system. Although it is clear that we have to live with incomplete reasoners, this incompleteness should be systematic, perhaps even describable by a systematic, model-theoretic approach (cf. [18]). In conclusion, if we claim to *integrate* the reasoning of two formalisms, a minimal requirement is that the inferences are independent of the order of input.

The only solution to this problem is to employ some sort of *symbolic constraint propagation*.¹¹ After a new MSG is determined for some individual I ,

¹⁰Generalization is not a concept-forming operator and there seems to be no easy way to assign a compositional semantics to it. However, it is easy to build such a concept structurally, which is the most specialized one subsuming a given set of concepts.

¹¹A fact ignored in KANDOR and the old KL-TWO system. Meanwhile, in KL-TWO this deficiency has been removed [24].

2. additionally, all individuals, which mention I as a role filler have to be checked whether it is possible to specialize them (backward propagation); this can only happen if the corresponding role-filler set is closed;
3. finally, these steps have to be performed for all individuals which get new MSGs during this process.

For the example above, this suffices to detect the contradiction. The JUNK team will be specialized to the following concept:

(and Modern-team (all member man))

In classifying this concept, the value restriction of the `member` role is propagated to the subrole `leader` (a subrole has to adhere to the value restriction of its superrole) resulting in a value restriction of **(and Man Woman)** which apparently is contradictory.

Altogether, this process assures that incompleteness depending on the order of input cannot appear. However, it also sounds very costly from a computational point of view. It is even not very clear whether this process always terminates.

Fortunately, the process of propagating MSG is not a general constraint propagation process. A first fact about this process, we may note, is that backtracking cannot occur. That means if we encounter a contradiction there is no way to resolve it, but to reject the input which led to the contradiction. This tells us that we will not get a combinatorial explosion because of reasoning by case. There are more facts which constrain the propagation space:

- An individual which gets a new MSG by backward propagation cannot trigger new forward propagations. This is because the new MSG does not reflect the current role fillers as necessary and sufficient conditions, i.e., there cannot be any new restrictions on role fillers.
- Backward propagation leads to a new MSG, if the role-filler set is closed and all other role fillers are already more specialized than the value restriction, a situation which does not occur very often.
- Forward propagation has a significant non-local effect only if the chain of individuals the value restrictions are propagated along corresponds to a line of concepts which are specialized in parallel, which is not very likely.
- In the worst case, the number of recomputations of an MSG during forward and backward propagation is bounded by the product of the number of individuals and the number of concepts in the TBox (before the entire process starts). This, however, would result in MSGs which cannot be specialized further, i.e., the upper bound can be divided by the number of input operations.

Finally, we should discuss where incompleteness might arise in the integrated reasoning process. First, realization cannot be more complete than classification. This means that some very weird cases of contradictions cannot be recognized (cf. [15]). Furthermore, the abstraction process is a source of incompleteness as well. In our example, we could prove that KIM is necessarily a Woman after the first two inputs; however, the abstraction failed to recognize this. The reason is, that this would require some kind of reasoning by case we strictly avoided because of the computational costs.

6 Conclusion

The hybrid reasoning component of the BACK system was presented and discussed as the “materialization” of a hybrid representation formalism. In particular, the “realization” inference was discussed and it was shown that a symbolic constraint propagation technique is necessary to insure the global propagation of consequences of incremental, local additions of properties. Together with the completion and abstraction of local properties this mechanism glues together the Abox and Tbox of BACK with respect to the underlying semantics.

Acknowledgement

Many of the ideas presented here are the result of the work in the BACK project. The other members of this project, Christof Peltason, and Albrecht Schmiedel, deserve much of the credit for these ideas. Additionally, the discussions with Marc Vilain and Bob MacGregor were very helpful in shedding some light on the realization process.

- [1] James F. Allen, "Maintaining Knowledge About Temporal Intervals," *Communications of the ACM* **26**(11): 832–843, 1983.
- [2] Ronald J. Brachman and Hector J. Levesque, "The tractability of subsumption in frame-based description languages," *Proc. AAAI-84*, Austin, Tex., August 1984, pp. 34–37.
- [3] Ronald J. Brachman and James G. Schmolze, "An overview of the KL-ONE knowledge representation system," *Cognitive Science*, **9**(2): 171–216, April–June 1985.
- [4] Ronald J. Brachman, Victoria Pigman Gilbert, and Hector J. Levesque, "An Essential Hybrid Reasoning System: Knowledge and Symbol Level Accounts in KRYPTON," *Proc. 9th IJCAI*, Los Angeles, Cal., 1985, pp. 532–539.
- [5] Jürgen Edelmann and Bernd Owsnicki, "Data Models in Knowledge Representation Systems: A Case Study," in: C.-R. Rollinger, W. Horn (eds.), *GWAI-86 und 2. Österreichische Artificial-Intelligence-Tagung*, Berlin: Springer, 1986, pp. 69–74.
- [6] Patrick J. Hayes, "Some Problems and Non-Problems in Representation Theory," *Proc. AISB Summer Conference*, University of Sussex, 1974, pp. 63–79.
- [7] Patrick J. Hayes, "The Logic of Frames," in: D. Metzger (ed.), *Frame Conceptions and Text Understanding*, Berlin: deGruyter, 1979, pp. 46–61.
- [8] Robert A. Kowalski, Contribution to the SIGART Special Issue on Knowledge Representation, *SIGART Newsletter*, No 70, February 1980.
- [9] Hector J. Levesque, "Making Believers out of Computers," *Artificial Intelligence*, **30**(1): 81–108, 1986.
- [10] Kai von Luck, "Semantic Networks with Number Restricted Roles or Another Story about Clyde," in: C.-R. Rollinger, W. Horn (eds.), *GWAI-86 und 2. Österreichische Artificial-Intelligence-Tagung*, Berlin: Springer, 1986, pp. 58–68.
- [11] Kai von Luck, Bernhard Nebel, Christof Peltason, and Albrecht Schmiedel, "The Anatomy of the BACK System," KIT Report 41, Technische Universität Berlin, January 1987.
- [12] William Mark, Realization, in: J. G. Schmolze and R. J. Brachman (eds.), *Proc. 1981 KL-ONE Workshop*, BBN Technical Report 4842, Bolt, Berank and Newman, Inc., Cambridge, Mass., 1982, pp. 78–89.
- [13] John McCarthy and Patrick J. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," in: B. Meltzer, D. Michie (eds.), *Machine Intelligence 4*, Edinburgh: Edinburgh University Press, 1969, 463–502.
- [14] Drew McDermott, "Tarskian semantics, or no notation without denotation!" *Cognitive Science*, **2**(3): 277–282, July–September 1978.
- [15] Bernhard Nebel, "Computational Complexity of Terminological Reasoning in BACK," *Artificial Intelligence*, **34**(3): 371–383, April 1988.
- [16] Bernhard Nebel, *Knowledge Base Revision in Terminological Representation Systems*, Berlin: Department of Computer Science, Technische Universität Berlin, in preparation.
- [17] Peter F. Patel-Schneider, "Small can be Beautiful in Knowledge Representation," *Proc. IEEE Workshop on Principles of Knowledge-Based Systems*, Denver, Colo., 1984, pp. 11–16.
- [18] Peter F. Patel-Schneider, "A Four-Valued Semantics for Frame-Based Description Languages," *Proc. AAAI-86*, Philadelphia, Pa., August 1986, pp. 344–348.
- [19] Raymond Reiter, "Towards a Logical Reconstruction of Relational Database Theory," in M. L. Brodie, J. Mylopoulos, and J. W. Schmidt (eds.), *On Conceptual Modelling*, Berlin: Springer, 1984, pp. 191–233.
- [20] Charles Rich, The Layered Architecture of a System for Reasoning about Programs, *Proc. IJCAI-85*, Los Angeles, Cal., 1985, pp. 540–546.
- [21] James G. Schmolze, Thomas Lipkis, "Classification in the KL-ONE Knowledge Representation System," *Proc. 8th IJCAI*, Karlsruhe, 1983, pp. 330–332.

Research and Development in Expert Systems, Cambridge: Cambridge University Press, 1985.

[23] Marc B. Vilain, “The Restricted Language Architecture of a Hybrid Representation System,” *Proc. 9th IJCAI*, Los Angeles, Cal., 1985, pp. 547–551.

[24] Marc B. Vilain, pers. comm., 1987.