Labelled Logics of Defeasible Goals*

Leendert W.N. van der Torre
Max Planck Institute for Computer Science
Im Stadtwald, D-66123 Saarbrücken, Germany
torre@mpi-sb.mpg.de

Abstract

In this paper we study conflicts between goals. In line with negative results obtained in the logic of preference, we argue that most proof rules of the logic of goals – such as strengthening of the antecedent (monotony), transitivity and the conjunction rule – only hold in a restricted sense. We study restricted applicability in Gabbay's labelled deductive systems, and we show how to resolve conflicts in labelled logics of conditional goals.

1 Introduction

In the usual approaches to planning in AI, a planning agent is provided with a description of some state of affairs, a *goal state*, and charged with the task of discovering (or performing) some sequence of actions to achieve that goal. Recently several *logics* for goals (and the closely related desires) have been proposed [DW91, DSW91, Pea93, Bou94, TP94b, TP94a, Lan96, HB97], and it has been observed that conflicts between goals deserve special attention. Consider the following two examples given by Bacchus and Grove [BG96].

Surgery. A person may prefer not having surgery over having surgery, but this preference might be reversed in the circumstances where surgery improves one's long term health. We say that the more specific preference overrides the more general one, which can be formalized with specificity-based principles developed in logics of defeasible reasoning. We can defeasibly infer that the person prefers no surgery only as long as it is not known that surgery improves his or her long term health.

Marriage. A girl named Sue prefers to be married to John over not, she prefers to be married to Fred over not, and at the same time she reasonably prefers to be married to neither over being married to both. If all we know is that she prefers to be married to John and that she prefers to be married to Fred, then it is reasonable to defeasibly infer that she likes to be married to both. However, if we additionally know that she prefers to be married to neither of them, then this inference is defeated.

*Appeared in Proceedings of the Fourth Dutch-German Workshop on Nonmonotonic Reasoning Techniques and Their Applications (DGNMR'99).

The examples illustrate that intuitively most proof rules of the logic of *defeasible* goals – i.e. goals which can be overridden by other goals – only hold in a restricted sense. This is in line with negative results obtained in the logic of preference, where counterexamples have been given to nearly all proposed rules [Mul79]. In semantic terms the problem of the logic of preference is that there is no straight forward way to lift a preference relation between worlds, for example represented by a utility function, to a preference relation between sets of worlds or propositions. Pearl [Pea93] also observed counterexamples to most proof rules when utilities and probabilities are combined in a so-called decision-theoretic logic of preference.

This restricted applicability of rules explains why the proposed logics of defeasible goals are either too strong – the logic supports the proof rules – or very weak – it does not support them. In this paper we study restricted applicability in Gabbay's labelled deductive systems [Gab96] and we show how to incorporate a conflict resolution method in the labelled logic of conditional goals [vdT98a, vdT98b]. The resulting logic has the following desirable properties.

The identity (i.e. α is a goal if α) does not hold;

Strengthening of the antecedent (i.e. monotony) is restricted to incorporate, amongst others, the specificity principle;

Conflict averse strategies do not hold;[1]

The conjunction and transitivity rules are defined in terms of strengthening of the antecedent and a new generalized cut rule, and therefore only hold in a restricted sense;

Weakening of the consequent (i.e. consequential closure) holds to support reasoning about sub-goals.

Restricted applicability is of course well-known from logics of defeasible reasoning, where monotony and transitivity only hold in a restricted sense and where restricted alternatives of monotony have been developed, e.g. cautious and rational monotony. However, it is important to note that the logic of defaults is quite different from the logic of defeasible goals. For example, according to the generally accepted Kraus-Lehmann-Magidor paradigm [KLM90] the identity (called supra-classicallity) and the conjunction rule both hold in the logic of defaults. This difference supports Mullen [Mul79]'s claim that there is no generally applicable preference logic, and that (as he concluded) a logic of preference can only be tested in the process of testing the more general theory in which it is embedded.

This paper is organized as follows. In Section 2 we introduce a labelled logic of defeasible goals. In Section 3 we discuss and illustrate the desirable properties and in Section 4 we discuss how goal independence can be used to make the logic more adventurous.

[1] In the surgery example above it seems reasonable to infer that the agent prefers not to be in a situation where surgery improves his long term health, but when rewards are higher than penalties it is beneficial to be in a conflict.

2 Labelled logic of defeasible goals (LLDG)

The langauge of the logic consists of conditional goals $G(\alpha|\beta)_L$, read as 'α is preferred if β in the proof context L.'

Definition 1 (Language) *Let $\mathcal{L}$ be a propositional base logic. The language of* LLDG *consists of the labelled dyadic goals $G(\alpha|\beta)_L$, with α and β sentences of $\mathcal{L}$, and L a set of sets of sentences of $\mathcal{L}$. We write $\models$ for entailment in $\mathcal{L}$.*

Each formula $G(\alpha\,|\,\beta)_L$ occurring as a premise has a label that consists of its own (propositionally consistent) fulfillment.

Definition 2 (Premise) *A formula $G(\alpha|\beta)_{\{\{\alpha\wedge\beta\}\}}$, where $\alpha \wedge \beta$ is consistent in $\mathcal{L}$, is called a premise of* LLDG.

The distinctive property of labelled deductive systems with respect to other logical systems is that we have to define how proof rules operate on the labels, and how the proof rules use the information in the labels. The label of a derived goal is the union (OR) or the product (SA, CTA) of strengthenings and the labels of the premises used in this proof rule, where the product is defined by

$$\{S_1, \ldots, S_n\} \times \{T_1, \ldots, T_m\} = \{S_1 \cup T_1, \ldots, S_1 \cup T_m, \ldots, S_n \cup T_m\}$$

For non-defeasible goals a consistency check ensures that $G(\alpha|\beta)_L$ may only be derived if each $F \in L$ is consistent: it must always be possible to fulfill a derived goal together with each of the goals it is derived from, though – to support reasoning by cases – not necessarily all of them at the same time.

In the logic of defeasible goals developed here, a goal can be overridden by other goals. The inference relation is relative to a set of goals (the goal base, intuitively the set of premises) and we may only apply a rule if the label of the derived goal does not make the fulfillment of a goal of the goal base impossible. If the goal base is fixed, then the inference relation (which we represent by $\vdash^B_{\text{LLDG}}$ for the goal base B) is monotonic. However, the intended interpretation of this goal base is that it reflects the set of premises, and it thus grows when the set of premises grows. Under this intended interpretation, the inference relation (which we represent by $\mid\!\sim_{\text{LLDG}}$) is non-monotonic. It is well-known from logics of defeasible reasoning that there are many different ways to define non-monotonic closures, ranging from safe to adventurous. When applying a safe or cautious rule, such as System Z [Pea90], the set of derived conclusions is relatively small. The proposed construction here is cautious in the sense that it blocks the inheritance of goals to sub-ideal subclasses. This is discussed in detail in Section 4, where we also show how to define more adventurous extensions.

Definition 3 (LLDG) *Let B be a set of* LLDG*-formulas called the goal base. The labelled logic of defeasible goals* LLDG *consists of the inference rules below, extended with the following condition R.*

R: *$G(\alpha|\beta)_L$ may only be derived if for each $F \in L$ (of the form $F_1 \cup \{\beta\}$, $F_1 \cup F_2$ or F_1) and for each goal $G(\alpha'|\beta') \in B$, $F \cup \{\alpha' \wedge \beta'\}$ is only inconsistent if $F_1 \cup \{\alpha' \wedge \beta'\}$ (and $F_2 \cup \{\alpha' \wedge \beta'\}$) is already inconsistent.*

The inference rules of LLDG *are replacements by logical equivalents and the following four rules.*

$$\text{SA}_R : \frac{G(\alpha \mid \beta_1)_L, R}{G(\alpha \mid \beta_1 \wedge \beta_2)_{L \times \{\beta_2\}}} \qquad \text{CTA}_R : \frac{G(\alpha \mid \beta \wedge \gamma)_{L_1}, G(\beta \mid \gamma)_{L_2}, R}{G(\alpha \wedge \beta \mid \gamma)_{L_1 \times L_2}}$$

$$\text{WC}_R : \frac{G(\alpha_1 \mid \beta)_L, R}{G(\alpha_1 \vee \alpha_2 \mid \beta)_L} \qquad \text{OR}_R : \frac{G(\alpha \mid \beta_1)_{L_1}, G(\alpha \mid \beta_2)_{L_2}, R}{G(\alpha \mid \beta_1 \vee \beta_2)_{L_1 \cup L_2}}$$

For a goal base B, we say $\{G(\alpha_i \mid \beta_i) \mid 1 \leq i \leq n\} \vdash^{B}_{\text{LLDG}} G(\alpha \mid \beta)$ if there is a labelled goal $G(\alpha|\beta)_L$ that can be derived from the set of goals $\{G(\alpha_i|\beta_i)_{\{\{\alpha_i \wedge \beta_i\}\}} \mid 1 \leq i \leq n\}$. Moreover, we say $B {\mid\!\sim}_{\text{LLDG}} G(\alpha|\beta)$ iff $B \vdash^{B}_{\text{LLDG}} G(\alpha|\beta)$.

Before we consider the relation between LLDG and LLG in Proposition 2 we first prove the following useful proposition. It says that each element of the label implies the fulfillment of the goal, i.e. the conjunction of its antecedent and consequent.

Proposition 1 *For each goal $G(\alpha \mid \beta)_L$ derived in* LLDG *we have for each $F \in L$ that $F \models \alpha \wedge \beta$.*

Proof *By induction on the structure of the proof tree. The property trivially holds for the premises, and it is easily seen that the proof rules retain the property.*

The following proposition shows that LLDG derives a subset of LLG introduced in [vdT98b].

Proposition 2 *Let* LLG *be the* LLDG *in which the consistency check R is replaced by the following check R_{LLG} on the consistency of all elements of L. Note that the inference relation of* LLG *is not relative to the goal base.*

R_{LLG}: $G(\alpha|\beta)_L$ may only be derived if each $F \in L$ is consistent.

If we have $B {\mid\!\sim}_{\text{LLDG}} \alpha$, then we have $B \vdash_{\text{LLG}} \alpha$ (but obviously not necessarily vice versa).

Proof *We prove that if R holds, then R_{LLG} also holds. So assume that R holds for a derived goal $G(\alpha \mid \beta)_L$. For each $F \in L$ there exists a goal $G(\alpha' \mid \beta') \in B$ used in the proof of $G(\alpha|\beta)_L$, and we have $F \models \alpha' \wedge \beta'$ and $F \models \alpha \wedge \beta$ (Proposition 1). We have $F \cup \{\alpha' \wedge \beta'\}$ and $F \cup \{\alpha \wedge \beta\}$ are consistent according to R, and therefore F is consistent. Consequently R_{LLG} holds.*

The following proposition shows how LLDG can be defined as LLG with additional restrictions.

Proposition 3 *Let the fulfillment conditions of a goal $G(\alpha_i|\beta_i)$ in goal base B be*

$$FC(G(\alpha_i|\beta_i) \mid B) = \{\{\alpha_i \wedge \beta_i \wedge \alpha_j \wedge \beta_j\} \mid \begin{array}{l} G(\alpha_j|\beta_j) \in B \text{ and} \\ \alpha_i \wedge \beta_i \wedge \alpha_j \wedge \beta_j \text{ is consistent}\} \end{array}$$

We have $B {\mid\!\sim}_{\text{LLDG}} G(\alpha|\beta)$ if and only if there is a labelled goal $G(\alpha|\beta)_L$ that can be derived from the set of goals $\{G(\alpha_i \mid \beta_i)_{\{\{\alpha_i \wedge \beta_i\}\} \cup FC(G(\alpha_i \mid \beta_i) \mid B)} \mid 1 \leq i \leq n\}$ in LLG*, i.e. in* LLDG *in which the consistency check R is replaced by the check R_{LLG} of Proposition 2.*

Proof *We prove that R implies R_{LLG} and vice versa. $\Rightarrow$ Assume a derivation for which R_{LLG} does not hold. There is an $F \in L$ which is inconsistent, and given Proposition 2 the inconsistency is a result of the fulfillment conditions. Consequently R does not hold. $\Leftarrow$ If R_{LLG} holds then the fulfillment conditions imply R.*

The only unusual rule of LLG and LLDG is CTA, a combination of cumulative transitivity (CUT or CT) and the conjunction rule (AND). In the following sections we will see that this unusual rule leads to several desired properties of the logic.

3 The desirable properties of LLDG

In this section we discuss and illustrate the desirable properties of LLDG we already mentioned in the introduction: the identity and conflict averse strategies do not hold, but restricted strengthening of the antecedent (to support the specificity principle), restricted conjunction, restricted transitivity, and the disjunction rule (to support reasoning by cases) do hold.

3.1 The identity

From the definitions follows immediately that $G(\alpha|\alpha)$ does not hold, and that we cannot derive $G(\alpha \wedge \beta \,|\, \beta)$ from $G(\alpha \,|\, \beta)$. This contrasts LLDG with logics of goals defined as follows in an underlying preference logic, in which $\alpha_1 \succ \alpha_2$ is read as 'α_1 is preferred to α_2.'

$$G(\alpha|\beta) =_{def} (\alpha \wedge \beta) \succ (\neg\alpha \wedge \beta)$$

Note that the fact that the identity is an undesirable property of the logic of goals also makes it difficult to define a semantics for such logics, because we cannot use results obtained in e.g. conditional, counterfactual and default logics.

3.2 Restricted strengthening of the antecedent and conflict averse strategies

Strengthening of the antecedent is restricted to formalize, amongst others, overriding of general preferences by more specific ones. This is illustrated by Bacchus and Grove's surgery example.

Example 1 (Surgery) *Consider the two goals 'A person prefers not having surgery over having surgery' $G(\neg s)$, and 'this preference is reversed in the circumstances where surgery improves one's long term health' $G(s|i)$. We have the following.*

$$G(\neg s) \mid\!\sim_{\text{LLDG}} G(\neg s|i)$$
$$G(\neg s), G(s|i) \not\mid\!\sim_{\text{LLDG}} G(\neg s|i)$$

Hence, by addition of a premise we loose conclusions. The circumstances where surgery improves one's long term health (i) represent a conflict situation, because the first goal prefers $\neg s$ whereas the second goal prefers the contradictory s. We say that the latter more specific goal overrides the former more general one, because $G(\neg s|i)$ can be derived from the first goal, but not from both goals.

The surgery example can also be used to illustrate another aspect in which the logic of goals is different from, for example, the logic of defaults and the logic of prima facie obligations. When in conflict, i.e. when i is true, it is sure that some rule will be achieved and another one will be violated, and the relevant question is what we can say about this conflict situation. In the logic of defaults it seems unlikely to be in a situation in which a default is violated, and therefore it is unlikely to be in a conflict situation. In the logic of prima facie obligations, it is forbidden to violate obligations and it is therefore forbidden to be in a conflict situation (this derivation also depends on the capabilities of the agent to evade the conflict). However, in the logic of goals it may be desirable to be in a conflict situation, namely when the rewards of the achieved goal are higher than the penalties of the violated goal. We have, as desired, the following non-derivation in LLDG.

$$G(\neg s|\top), G(s|i) \not\mid\!\sim_{\text{LLDG}} G(\neg i|\top)$$

Moreover, in several logics of goals the combination of conflicting contexts leads to the counterintuitive result to prefer one of these contexts [vdTT98]. Obviously in LLDG these counterintuitive derivations are blocked, because conflicting contexts can only be combined when reasoning by cases.

3.3 Restricted conjunction and transitivity

The following proposition shows that the conjunction and transitivity rules are defined in terms of strengthening of the antecedent, and they are therefore also restricted.

Proposition 4 *The following conjunction rule, transitivity rule and extended disjunction rule are implied by the logic* LLDG.

$$\text{AND}_R : \frac{G(\alpha_1 \mid \beta)_{L_1}, G(\alpha_2 \mid \beta)_{L_2}, R}{G(\alpha_1 \wedge \alpha_2 \mid \beta)_{L_1 \times L_2}}$$

$$\text{TRANS}_R : \frac{G(\alpha \mid \beta)_{L_1}, G(\beta \mid \gamma)_{L_2}, R}{G(\alpha \mid \gamma)_{L_1 \times L_2}}$$

$$\text{OR2}_R : \frac{G(\alpha_1 \mid \beta_1)_{L_1}, G(\alpha_2 \mid \beta_2)_{L_2}, R}{G(\alpha_1 \vee \alpha_2 \mid \beta_1 \vee \beta_2)_{L_1 \cup L_2}}$$

Proof *For the conjunction rule, we can first strengthen $G(\alpha_1|\beta)$ to $G(\alpha_1|\beta \wedge \alpha_2)$, and then apply* CTA *as follows to derive $G(\alpha_1 \wedge \alpha_2|\beta)$. The derived goal is equivalent to the goal derived by the proof rule above, because L_2 implies α_2 (Proposition 1).*

$$\frac{\dfrac{G(\alpha_1|\beta)_{L_1}}{G(\alpha_1|\beta \wedge \alpha_2)_{L_1 \times \{\alpha_2\}}}\ \text{SA} \qquad G(\alpha_2|\beta)_{L_2}}{G(\alpha_1 \wedge \alpha_2|\beta)_{L_1 \times \{\alpha_2\} \times L_2}}\ \text{CTA}$$

For the transitivity rule, we have the following derivation. Again, L_2 implies γ (Proposition 1).

$$\frac{\dfrac{\dfrac{G(\alpha|\beta)_{L_1}}{G(\alpha|\beta \wedge \gamma)_{L_1 \times \{\gamma\}}}\ \text{SA} \qquad G(\beta|\gamma)_{L_2}}{G(\alpha \wedge \beta|\gamma)_{L_1 \times \{\gamma\} \times L_2}}\ \text{CTA}}{G(\alpha|\gamma)_{L_1 \times \{\gamma\} \times L_2}}\ \text{WC}$$

For the extended disjunction rule, we have the following derivation.

$$\dfrac{\dfrac{G(\alpha_1 \,|\, \beta_1)_{L_1}}{G(\alpha_1 \vee \alpha_2 \,|\, \beta_1)_{L_1}}\,\text{WC} \qquad \dfrac{G(\alpha_2 \,|\, \beta_2)_{L_2}}{G(\alpha_1 \vee \alpha_2 \,|\, \beta_2)_{L_2}}\,\text{WC}}{G(\alpha_1 \vee \alpha_2 \,|\, \beta_1 \vee \beta_2)_{L_1 \cup L_2}}\,\text{OR}$$

The following example illustrates the derived restricted conjunction rule by Bacchus and Grove's marriage example.

Example 2 (Marriage) *We have*

$G(j), G(f) \mathrel{|\!\sim}_{\text{LLDG}} G(j \wedge f)$ *and*
$G(j), G(f), G(\neg(j \wedge f)) \mathrel{|\!\not\sim}_{\text{LLDG}} G(j \wedge f)$,

because in the following derivation the goal $G(j \,|\, f)$ cannot be derived from $G(j \,|\, \top)$ due to the existence of $G(\neg(j \wedge f))$.

$$\dfrac{\dfrac{G(j \,|\, \top)_{\{\{j\}\}}}{G(j \,|\, f)_{\{\{j,f\}\}}}\,\text{SA (dashed)} \qquad G(f \,|\, \top)_{\{\{f\}\}}}{G(j \wedge f \,|\, \top)_{\{\{j,f\}\}}}\,\text{CTA}$$

Hence, the goal to marry both is only derived non-monotonically from the first two goals, and the derivation is blocked by addition of the third goal. Notice that the derivation is not blocked by a specific goal but by a most general goal (with antecedent $\top$), and that consequently the specificity principle cannot be used to block the derivation.

The logic LLDG is very cautious, because the counterintuitive goal cannot be derived if we replace the first goal of the goal base by $G(j \,|\, f)$. We have the following.

$G(j \,|\, f), G(f) \mathrel{|\!\sim}_{\text{LLDG}} G(j \wedge f)$
$G(j \,|\, f), G(f), G(\neg(j \wedge f)) \mathrel{|\!\not\sim}_{\text{LLDG}} G(j \wedge f)$

The fact that the conjunction rule only holds non-monotonically means that we have to be careful whether we formalize a goal for $p \wedge q$ by $G(p) \wedge G(q)$ or by $G(p \wedge q)$. Although we have:

$G(p), G(q) \mathrel{|\!\sim}_{\text{LLDG}} G(p \wedge q)$
$G(p \wedge q) \mathrel{|\!\sim}_{\text{LLDG}} G(p)$
$G(p \wedge q) \mathrel{|\!\sim}_{\text{LLDG}} G(q)$

they derive different consequences when other goals are added to the goal base, for example:

$G(p), G(q), G(r) \mathrel{|\!\sim}_{\text{LLDG}} G(r \,|\, \neg(p \wedge q))$
$G(p \wedge q), G(r) \mathrel{|\!\not\sim}_{\text{LLDG}} G(r \,|\, \neg(p \wedge q))$

and

$G(p), G(q), G(r \,|\, \neg p) \mathrel{|\!\not\sim}_{\text{LLDG}} G(r \,|\, \neg p \wedge \neg q)$

$$G(p \wedge q), G(r \mid \neg p) \mid\!\sim_{\text{LLDG}} G(r \mid \neg p \wedge \neg q)$$

Hence, both representations are incomparable, in the sense that in some circumstances the former is more adventurous, and in other circumstances the latter.

3.4 Reasoning by cases

The disjunction rule of LLDG supports reasoning by cases in a similar way as it does in most conditional logics. The following example shows how the logic reasons by cases to derive a conclusion from conflicting goals.

Example 3 (Reasoning by cases) *As desired, we have*

$$G(a \wedge c \mid b), G(a \wedge \neg c \mid \neg b) \mid\!\sim_{\text{LLDG}} G(a),$$

as a consequence of the following derivation.

$$\dfrac{\dfrac{G(a \wedge c \mid b)_{\{\{a \wedge c \wedge b\}\}}}{G(a \mid b)_{\{\{a \wedge c \wedge b\}\}}}\ \text{WC} \qquad \dfrac{G(a \wedge \neg c \mid \neg b)_{\{\{a \wedge \neg c \wedge \neg b\}\}}}{G(a \mid \neg b)_{\{\{a \wedge \neg c \wedge \neg b\}\}}}\ \text{WC}}{G(a \mid \top)_{\{\{a \wedge c \wedge b\}, \{a \wedge \neg c \wedge \neg b\}\})}}\ \text{OR}$$

The example also illustrates why the label has to consist of a set of set of formulas, and not only of a set of formulas. In the latter case, the consistency check would block the intuitive derivation.

4 Goal independence

The logic LLDG is cautious in the sense that it, just like System Z, does not have the inheritance of properties to sub-ideal subclasses. For example, we have

$$G(p), G(q) \not\mid\!\sim_{\text{LLDG}} G(p \mid \neg q),$$

because the strengthening $\neg q$ of the first premise makes it impossible to fulfill the premise $G(q)$. Depending on the interpretation of the propositions, this derivation may be desired or not. Solutions for this ambiguity problem proposed in logics of defeasible reasoning use a notion of repair [BDP98] or they add independence assumptions [DdCHP97] when this derivation is desired.

In LLDG we can add independence relations between goals of the goal base. We restrict the consistency check to a subset of all premises, the premises it depends on. As an extreme case, non-defeasible goals only depend on the premises they are derived from. In the following definition we use the characterization of LLDG in Proposition 3.

Definition 4 (Independence) *Let the extended goal base $B = (S, I)$ be a set of goals S with a reflexive and a-symmetric relation I on $S \times S$. $I(g_1, g_2)$ should be read as 'the goal g_1 is independent of the goal g_2.' The fulfillment conditions of $G(\alpha_i \mid \beta_i)$ in B are*

$$\begin{aligned} FC(G(\alpha_i \mid \beta_i) \mid B) = \{\{\alpha_i \wedge \beta_i \wedge \alpha_j \wedge \beta_j \mid\ & G(\alpha_j \mid \beta_j) \in B \textit{ and} \\ & \textit{not } I(G(\alpha_i \wedge \beta_i), G(\alpha_j, \beta_j)) \textit{ and} \\ & \alpha_i \wedge \beta_i \wedge \alpha_j \wedge \beta_j \textit{ is consistent}\} \end{aligned}$$

We have $B \mid\!\sim_{\text{LLDG}} G(\alpha \mid \beta)$ *if there is a goal* $G(\alpha \mid \beta)_L$ *that can be derived from the set of goals* $\{G(\alpha_i \mid \beta_i)_{\{\{\alpha_i \wedge \beta_i\}\} \cup FC(G(\alpha_i|\beta_i)|B)} \mid 1 \leq i \leq n\}$ *in* LLG, *i.e. in* LLDG *in which the consistency check* R *is replaced by the check* R_{LLG} *of Proposition 2.*

Goal independence is illustrated in the following example.

Example 4 *Consider the two extended goal bases* $B_1 = (\{G(p), G(q)\}, \emptyset)$ *and* $B_1 = (\{G(p), G(q)\}, \{(G(p), G(q)), (G(q), G(p))\})$*. The fulfillment conditions of* $G(p)$ *in* B_1 *and* B_2 *are respectively* $\{\{p \wedge q\}\}$ *and* $\emptyset$*. We therefore have the following.*

$$B_1 \not\mid\!\sim_{\text{LLDG}} G(p|\neg q)$$
$$B_2 \mid\!\sim_{\text{LLDG}} G(p|\neg q)$$

Hence, we have inheritance to sub-ideal subclasses ($\neg q$) *only if the two goals are independent.*

5 Related research

The first labelled logic of the type considered in this paper was proposed in [vdTT95] for obligations. In this rudimentary logic the label of a premise only contains the consequent, and the consistency check considers the label together with the antecedent. The logic does not have the disjunction rule, because there is a complicating factor with this consistency check for reasoning by cases. When reasoning by cases we intuitively want to combine cases from conflicting contexts, and sometimes we even want to combine cases of conflicting goals. For example, we want to derive $G(a|\top)$ from the conflicting goals $G(a \wedge b|c)$ and $G(a \wedge \neg b|\neg c)$ (this example is discussed in detail in Section 3.4). However, a consistency check on the consequents would obviously block the derivation. Two solutions have been proposed to cover the parallel tracks created through reasoning by cases. First, in [vdT98a] the label of premises $G(\alpha \mid \beta)$ contains the material conditional $\beta \rightarrow \alpha$. Second, in [Mak98] the label contains a separate set of formulas for each case. Hence, in the latter more flexible solution, which we adopt here too, the label contains sets of sets of boolean formulas. As far as reasoning by cases is concerned in the latter solution the labels of the premises $G(\alpha \mid \beta)$ can either contain material implications $\beta \rightarrow \alpha$, consequents α only, or fulfillments $\alpha \wedge \beta$. In this paper we choose the most conservative option for the labels of the premises.

6 Conclusions and further research

In this paper we studied restricted applicability of rules in Gabbay's labelled deductive systems. In particular, we introduced a labelled logic of defeasible goals, with several desirable properties. The logic does not have the identity or conflict averse strategies, but it has restricted strengthening of the antecedent, conjunction, transitivity and the disjunction rule.

In [vdT98b] phasing is introduced in the labelled logic of non-defeasible goals, and two phasing theorems have been proven. Phasing makes the proof theory more efficient, because only a single order of rule application has to be considered. An inspection of the proofs shows that as a corollary of Proposition 3 the two theorems also hold for LLDG developed in this paper.

There are several issues for further research.

Decision-theoretic goals. The use of LLDG in decision-theoretic planning is our main present concern. In decision-theoretic planning usually decision-theoretic goals are used, i.e. goals based on (qualitative abstractions of) utilities and probabilities [Pea93, Bou94, BG96]. To deal with irrelevance problems usually strong independence notions are introduced, but consequently the logic is not able to reason about conflicts.

Comparison. Moreover, in further research we will contrast LLDG with other conflict-tolerant languages. We already observed that the logic of defaults satisfies the identity (called supra-classicallity) and the conjunction rule. Another distinction is that in default reasoning usually first the conflict is repaired, and thereafter classical inference is used to derive conclusions (called the coherence approach in [BG97]).

TMS. We have been informed that our labelled logic of goals has some formal similarities with truth maintenance systems. This formal relationship will be studied in further research.

References

[BDP98] S. Benferhat, D. Dubois, and H. Prade. Repairing and standard probabilistic semantics of conditional knowledge bases. Technical Report 98-07-R, IRIT, Toulouse, 1998.

[BG96] F. Bacchus and A.J. Grove. Utility independence in a qualitative decision theory. In *Proceedings of KR'96*, pages 542–552, 1996.

[BG97] S. Benferhat and L. Garcia. A coherence-based approach to default reasoning. In *Proceedings of the FAPR'97*, volume LNAI 1244, 1997.

[Bou94] C. Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the KR'94*, pages 75–86, 1994.

[DdCHP97] D. Dubois, L. Farinas del Cerro, A. Herzig, and H. Prade. Qualitative relevance and independence: a roadmap. In *Proceedings of the IJCAI'97*, pages 62–67, 1997.

[DSW91] J. Doyle, Y. Shoham, and M.P. Wellman. The logic of relative desires. In *Sixth International Symposium on Methodologies for Intelligent Systems*, Charlotte, North Carolina, 1991.

[DW91] J. Doyle and M.P. Wellman. Preferential semantics for goals. In *Proceedings of the AAAI'91*, pages 698–703, Anaheim, 1991.

[Gab96] D. Gabbay. *Labelled Deductive Systems*, volume 1. Oxford University Press, 1996.

[HB97] Z. Huang and J. Bell. Dynamic goal hierarchies. In *Intelligent Agent Systems: Theoretical and Practical Issues*, LNAI 1027, pages 88–103. Springer, 1997.

[KLM90] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[Lan96] J. Lang. Conditional desires and utilities - an alternative approach to qualitative decision theory. In *Proceedings of the ECAI'96*, pages 318–322, 1996.

[Mak98] D. Makinson. On a fundamental problem of deontic logic. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*. IOS Press, 1998.

[Mul79] J.D. Mullen. Does the logic of preference rest on a mistake? *Metaphilosophy*, 10:247–255, 1979.

[Pea90] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the TARK'90*, pages 121–135, 1990.

[Pea93] J. Pearl. From conditional oughts to qualitative decision theory. In *Proceedings of the UAI'93*, pages 12–20, 1993.

[TP94a] S.-W. Tan and J. Pearl. Qualitative decision theory. In *Proceedings of the AAAI'94*, 1994.

[TP94b] S.-W. Tan and J. Pearl. Specification and evaluation of preferences under uncertainty. In *Proceedings of the KR'94*, pages 530–539, 1994.

[vdT98a] L.W.N. van der Torre. Labeled logics of conditional goals. In *Proceedings of the ECAI'98*, pages 368–369. John Wiley & Sons, 1998.

[vdT98b] L.W.N. van der Torre. Phased labeled logics of conditional goals. In *Logics in Artifical Intelligence*, LNAI 1489, pages 92–106. Springer, 1998.

[vdTT95] L.W.N. van der Torre and Y.-H. Tan. Cancelling and overshadowing: two types of defeasibility in defeasible deontic logic. In *Proceedings of the IJCAI'95*, pages 1525–1532. Morgan Kaufman, 1995.

[vdTT98] L.W.N. van der Torre and Y.-H. Tan. Deliberate robbery, or the calculating Samaritan. In *Proceedings of the ECAI'98 Workshop on Practical Reasoning and Rationality (PRR'98)*, 1998.