

Formal Reasoning Techniques for Goal Models^{*}

Paolo Giorgini¹, John Mylopoulos², Eleonora Nicchiarelli¹, and
Roberto Sebastiani¹

¹ Department of Information and Communication Technology
University of Trento - Italy
{pgiorgini,eleonora,rseba}@dit.unitn.it

² Computer Science Department - University of Toronto - Canada
jm@cs.toronto.edu

Abstract. Over the past decade, goal models have been used in Computer Science in order to represent software requirements, business objectives and design qualities. Such models extend traditional AI planning techniques for representing goals by allowing for partially defined and possibly inconsistent goals. This paper presents a formal framework for reasoning with such goal models. In particular, the paper proposes a qualitative and a numerical axiomatization for goal modeling primitives and introduces label propagation algorithms that are shown to be sound and complete with respect to their respective axiomatizations. In addition, the paper reports on experimental results on the propagation algorithms applied to a goal model for a US car manufacturer.

1 Introduction

The concept of goal has been used in many areas of Computer Science for quite some time. In AI, goals have been used in planning to describe desirable states of the world since the 60s [11]. In AI planning problems, an agent is given a description of the environment and a set of admissible actions, and searches for a *plan*—i.e., a sequence of actions and intermediate sub-goals—which allows for achieving a final goal from a given initial state. Very efficient planning algorithms and tools, including those based on graph-based planning [3] and on SAT conversions [9] have been conceived (see, e.g., [17] for an overview). More recently, goals have been used in Software Engineering to model (early) requirements [6, 1, 2, 8] and non-functional requirements [10] for a software system. For instance, an early requirement for a library information system might be “*Every book request will eventually be fulfilled*”, while “*The new system will be highly reliable*” is an example of a non-functional requirement. Goals are also useful for knowledge management systems which focus on strategic knowledge, e.g., “*Increase profits*”, or “*Become the largest auto manufacturer in North America*” [7].

^{*} This is an extended and updated version of the paper: Giorgini, P., Mylopoulos, J., Nicchiarelli, E., and Sebastiani, R., “Reasoning with Goal Models”. In S. Spaccapietra, S. T. March, and Y. Kambayashi (Eds.), *Conceptual Modeling - ER 2002*, proceedings of the 21st International Conference on Conceptual Modeling, LNCS 2503 Springer, 2002.

Traditional goal analysis consists of decomposing goals into subgoals through an AND- or OR-decomposition. If goal G is AND-decomposed (respectively, OR-decomposed) into subgoals G_1, G_2, \dots, G_n , then if all (at least one) of the subgoals are satisfied so is goal G . Given a goal model consisting of goals and AND/OR relationships among them, and a set of initial labels for some nodes of the graph (S for Satisfied, D for Denied) there is a simple label propagation algorithm which can generate labels for all other nodes of the graph [12].

Unfortunately, this simple framework for modeling and analyzing goals won't work for many domains where goals are not formalizable and the relationships among them can't be captured by semantically well-defined relations such as AND/OR ones. For example, goals such as "*Highly reliable system*" have no formally defined predicate which prescribes their meaning, though one may want to define necessary conditions for such a goal to be satisfied. Moreover, such a goal may be related to other goals, such as "*Thoroughly debugged system*", "*Thoroughly tested system*" in the sense that the latter obviously contribute to the satisfaction of the former, but this contribution is partial and qualitative. In other words, if the latter goals are satisfied, they certainly contribute to the satisfaction of the former goal, but don't guarantee its satisfaction. This framework will also not work in situations where there are contradictory contributions to a goal. For instance, we may want to allow for multiple decompositions of a goal G into sets of subgoals, where some decompositions suggest satisfaction of G while others suggest denial.

We are interested in a modeling framework for goals which includes AND/OR relationships among goals, but also allows for other, more qualitative goal relationships and can accommodate contradictory situations [4, 16]. We accomplish this by introducing goal relationships labelled "+" and "-" which model respectively a situation where a goal contributes positively or negatively towards the satisfaction of another goal.

A major problem that arises from this extension is giving a precise semantics to the new goal relationships. This is accomplished in two different ways in this paper. In section 3 we offer a qualitative formalization and a label propagation algorithm which is shown to be sound and complete with respect to the formalization. In section 4, we offer a quantitative semantics for the new relationships which is based on a probabilistic model, and a label propagation algorithm that is also shown to be sound and complete. Section 5 presents experimental results on our label propagation algorithms using medium- to large-size goal models, while section 6 summarizes the contributions of the paper and sketches directions for further research.

2 An Example

Suppose we are modeling the strategic objectives of a US car manufacturer, such as Ford or GM. Examples of such objectives are increase return on investment or increase customer loyalty. Objectives can be represented as goals, and can be analyzed using goal relationships such as AND, OR, "+" and "-". In addition,

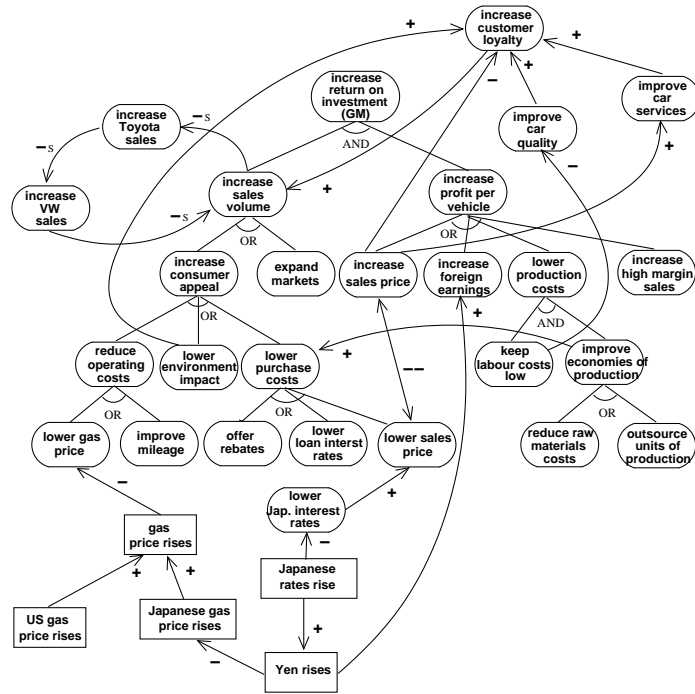


Fig. 1. A partial goal model for GM.

we will use “++” (respectively “--”) as a binary goal relationship such that if ++(G,G’) (--(G,G’)) then satisfaction of G implies satisfaction (denial) of G’.

For instance, increase return on investment may be AND-decomposed into increase sales volume and increase profit per vehicle. Likewise, increase sales volume might be OR-decomposed into increase consumer appeal and expand markets. This decomposition and refinement of goals can continue until we have goals that are tangible (i.e., someone can satisfy them through an appropriate course of action) and/or are observable (i.e., they can be confirmed satisfied/denied by simply observing the application domain).

For vaguely stated goals, such as increase customer loyalty we may want to simply model other relevant goals, such as improve car quality, improve car service and relate them through “+” and “-” relationships, as shown in Figure 1. These goals may influence positively or negatively some of the goals that have already been introduced during the analysis of the goal increase return on investment. Such lateral goal relationships may introduce cycles in our goal models.

Examples of observable goals are Yen rises, gas prices rise etc. When such a goal is satisfied, we will call it an *event* (the kind of event you may read about in a news story) and represent it in our graphical notation as a rectangle (see lower portion of Figure 1).

Figure 1 shows a partial and fictitious goal model for GM focusing on the goal increase return of investment. In order to increase return of investment, GM has to satisfy both goals increase sales and increase profit per vehicle. In turn, increase

sales volume is OR-decomposed into increase consumer appeal and expand markets, while the goal increase profit per vehicle is OR-decomposed into increase sales price, lower production costs, increase foreign earnings, and increase high margin sales. Additional decompositions are shown in the figure. For instance, the goal increase consumer appeal can be satisfied by satisfying lower environment impact, trying to lower purchase costs, or reducing the vehicle operating costs (reduce operating costs).

The graph shows also lateral relationships among goals. For example, the goal increase customer loyalty has positive (+) contributions from goals lower environment impact, improve car quality and improve car services, while it has a negative (−) contribution from increase sales price. The root goal increase return on investment (GM) is also related to goals concerning others auto manufacturer, such as Toyota and VW. In particular, if GM increases sales, then Toyota loses a share of the North American market; if Toyota increases sales (increase Toyota sales), it does so at the expense of VW; finally, if VW increases sales (increase VW sales), it does so at the expense of GM. These lateral contributions introduce a cycle in the GM goal model.

So far, we have assumed that every goal relationship treats S and D in a dual fashion. For instance, if we have $+(G, G')$, then if G is satisfied, G' is partially satisfied, and (dually) if G is denied G' is partially denied. Note however, that sometimes a goal relationship only applies for S (or D). In particular, the − contribution from increase sales (GM) to increase sales (Toyota) only applies when increase sales (GM) is satisfied (i.e., if GM hasn't increased sales, this doesn't mean that Toyota has.) To capture this kind of relationship, we introduce $-_S$, $-_D$, $+_S$, $+_D$ (see also Figure 1). Details about the semantics of these relationships are given in the next section.

3 Qualitative Reasoning with Goal Models

Formally, a *goal graph* is a pair $\langle \mathcal{G}, \mathcal{R} \rangle$ where \mathcal{G} is a set of goals and \mathcal{R} is a set of goal relations over \mathcal{G} . If $(G_1, \dots, G_n) \xrightarrow{r} G$ is a goal relation in \mathcal{R} , we call $G_1 \dots G_n$ *source goals* and G the *target goal* of r . To simplify the discussion, we consider only binary *OR* and *AND* goal relations. This is not restrictive, as all the operators we consider in this section and in Section 4 —i.e., \wedge , \vee , *min*, *max*, \otimes , \oplus — are associative and can be thus trivially used as n-ary operators.

3.1 Axiomatization of goal relationships

Let G_1, G_2, \dots denote goal labels. We introduce four distinct predicates over goals, $FS(G)$, $FD(G)$ and $PS(G)$, $PD(G)$, meaning respectively that there is (at least) *full* evidence that goal G is satisfied and that G is denied, and that there is at least *partial* evidence that G is satisfied and that G is denied. We also use the proposition \top to represent the (trivially true) statement that there is at least null evidence that the goal G is satisfied (or denied). Notice that the predicates state that there is *at least* a given level of evidence, because in a goal graph

Goal	Invariant Axioms	
$G :$	$FS(G) \rightarrow PS(G)$	(1)
	$FD(G) \rightarrow PD(G)$	(2)
Goal relation	Relation Axioms	
$(G_2, G_3) \xrightarrow{and} G_1 :$	$(FS(G_2) \wedge FS(G_3)) \rightarrow FS(G_1)$	(3)
	$(PS(G_2) \wedge PS(G_3)) \rightarrow PS(G_1)$	(4)
	$FD(G_2) \rightarrow FD(G_1), \quad FD(G_3) \rightarrow FD(G_1)$	(5)
	$PD(G_2) \rightarrow PD(G_1), \quad PD(G_3) \rightarrow PD(G_1)$	(6)
$(G_2, G_3) \xrightarrow{or} G_1 :$	$FS(G_2) \rightarrow FS(G_1), \quad FS(G_3) \rightarrow FS(G_1)$	(7)
	$PS(G_2) \rightarrow PS(G_1), \quad PS(G_3) \rightarrow PS(G_1)$	(8)
	$(FD(G_2) \wedge FD(G_3)) \rightarrow FD(G_1)$	(9)
	$(PD(G_2) \wedge PD(G_3)) \rightarrow PD(G_1)$	(10)
$G_2 \xrightarrow{+S} G_1 :$	$PS(G_2) \rightarrow PS(G_1)$	(11)
$G_2 \xrightarrow{-S} G_1 :$	$PS(G_2) \rightarrow PD(G_1)$	(12)
$G_2 \xrightarrow{++S} G_1 :$	$FS(G_2) \rightarrow FS(G_1),$	(13)
	$PS(G_2) \rightarrow PS(G_1)$	(14)
$G_2 \xrightarrow{--S} G_1 :$	$FS(G_2) \rightarrow FD(G_1),$	(15)
	$PS(G_2) \rightarrow PD(G_1)$	(16)
$G_2 \xrightarrow{+D} G_1 :$	$PD(G_2) \rightarrow PD(G_1)$	(17)
$G_2 \xrightarrow{-D} G_1 :$	$PD(G_2) \rightarrow PS(G_1)$	(18)
$G_2 \xrightarrow{++D} G_1 :$	$FD(G_2) \rightarrow FD(G_1),$	(19)
	$PD(G_2) \rightarrow PD(G_1)$	(20)
$G_2 \xrightarrow{-+D} G_1 :$	$FD(G_2) \rightarrow FS(G_1),$	(21)
	$PD(G_2) \rightarrow PS(G_1)$	(22)

Fig. 2. Ground axioms for the invariants and the propagation rules in the qualitative reasoning framework.

there may be multiple sources of evidence for the satisfaction/denial of a goal. We introduce a total order $FS(G) \geq PS(G) \geq \top$ and $FD(G) \geq PD(G) \geq \top$, with the intended meaning that $x \geq y$ iff $x \rightarrow y$.

We want to allow the deduction of *positive* ground assertions of type $FS(G)$, $FD(G)$, $PS(G)$ and $PD(G)$ over the goal constants of a goal graph. We refer to externally provided assertions as *initial conditions*. To formalize the propagation of satisfiability and deniability evidence through a goal graph $\langle \mathcal{G}, \mathcal{R} \rangle$, we introduce the axioms described in Figure 2. (By “dual” we mean that we invert satisfiability with deniability.)

As indicated in Section 2, the propagation rules for goal satisfaction through $++$, $--$, $+$, $-$ relationships may or may not be symmetric w.r.t. those for denial. For example, the relationship $G_2 \xrightarrow{+} G_1$ might propagate only the satisfaction of G_2 , only the denial, or both.

Thus, for every relation type $r \in \{++, --, +, -\}$, it makes sense to have three possible labels: " r_S ", " r_D ", and " r ", meaning respectively that satisfaction is propagated, that denial is propagated, and that both satisfaction and denial are propagated. (We call the first two cases *asymmetric*, the latter *symmetric*.) For example, $G_2 \stackrel{-S}{\mapsto} G_1$ means that if G_2 is satisfied, then there is some evidence that G_1 is denied, but if G_2 is denied, then nothing is said about the satisfaction of G_1 ; $G_2 \stackrel{-D}{\mapsto} G_1$ means that if G_2 is denied, then there is some evidence that G_1 is satisfied, but if G_2 is satisfied, then nothing is said about the denial of G_1 ; $G_2 \stackrel{-}{\mapsto} G_1$ means that, if G_2 is satisfied [denied], then there is some evidence that G_1 is denied [satisfied]. In other words, a symmetric relation $G_2 \stackrel{r}{\mapsto} G_1$ is a shorthand for the combination of the two corresponding asymmetric relationships $G_2 \stackrel{r_S}{\mapsto} G_1$ and $G_2 \stackrel{r_D}{\mapsto} G_1$.

(1) and (2) state that full satisfiability and deniability imply partial satisfiability and deniability respectively. For an AND relation, (3) and (4) show that the full and partial satisfiability of the target node require respectively the full and partial satisfiability of all the source nodes; for a " $+_S$ " relation, (11) show that only the partial satisfiability (but not the full satisfiability) propagates through a " $+_S$ " relation. Combining (1) with (3), and (1) with (11), we have, respectively,

$$(G_2, G_3) \stackrel{and}{\mapsto} G_1 : (FS(G_2) \wedge PS(G_3)) \rightarrow PS(G_1) \quad (23)$$

$$G_2 \stackrel{+_S}{\mapsto} G_1 : FS(G_2) \rightarrow PS(G_1). \quad (24)$$

Thus, an AND relation propagates the minimum satisfiability value (and the maximum deniability one), while a " $+_S$ " relation propagates at most a partial satisfiability value. To this extent, notice that a " $+_S$ " relation can be seen as an AND relation with an unknown partially satisfiable goal.

From now on, we implicitly assume that axioms (1) and (2) are always applied whenever possible. Thus, we say that $PS(G_1)$ is deduced from $FS(G_2)$ and $FS(G_3)$ by applying (3) —meaning “applying (3) and then (1)” — or that $PS(G_1)$ is deduced from $FS(G_2)$ and $PS(G_3)$ by applying (4) —meaning “applying (1) and then (4)”.

We say that an atomic proposition of the form $FS(G)$, $FD(G)$, $PS(G)$ and $PD(G)$ *holds* if either it is an initial condition or it can be deduced via modus ponens from the initial conditions and the ground axioms of Figure 2. We assume conventionally that \top always holds. Notice that all the formulas in our framework are propositional Horn clauses, so that deciding if a ground assertion holds not only is decidable, but also it can be decided in polynomial time.

We say that there is a *weak conflict* if either $PS(G)$ and $PD(G)$, $FS(G)$ and $PD(G)$, $PS(G)$ and $FD(G)$ hold for some goal G . We say that there is a *strong conflict* if $FS(G)$ and $FD(G)$ hold for some G .

3.2 The label propagation algorithm

Based on the logic framework of Section 3.1, we have developed an algorithm for propagating through a goal graph $\langle \mathcal{G}, \mathcal{R} \rangle$ labels representing evidence for

	$(G_2, G_3) \xrightarrow{and} G_1$	$G_2 \xrightarrow{+S} G_1$	$G_2 \xrightarrow{-S} G_1$	$G_2 \xrightarrow{++S} G_1$	$G_2 \xrightarrow{--S} G_1$
$Sat(G_1)$	$\min \left\{ \begin{array}{l} Sat(G_2), \\ Sat(G_3) \end{array} \right\}$	$\min \left\{ \begin{array}{l} Sat(G_2), \\ P \end{array} \right\}$	N	$Sat(G_2)$	N
$Den(G_1)$	$\max \left\{ \begin{array}{l} Den(G_2), \\ Den(G_3) \end{array} \right\}$	N	$\min \left\{ \begin{array}{l} Sat(G_2), \\ P \end{array} \right\}$	N	$Sat(G_2)$
	$(G_2, G_3) \xrightarrow{or} G_1$	$G_2 \xrightarrow{+D} G_1$	$G_2 \xrightarrow{-D} G_1$	$G_2 \xrightarrow{++D} G_1$	$G_2 \xrightarrow{--D} G_1$
$Sat(G_1)$	$\max \left\{ \begin{array}{l} Sat(G_2), \\ Sat(G_3) \end{array} \right\}$	N	$\min \left\{ \begin{array}{l} Den(G_2), \\ P \end{array} \right\}$	N	$Den(G_2)$
$Den(G_1)$	$\min \left\{ \begin{array}{l} Den(G_2), \\ Den(G_3) \end{array} \right\}$	$\min \left\{ \begin{array}{l} Den(G_2), \\ P \end{array} \right\}$	N	$Den(G_2)$	N

Table 1. Propagation rules in the qualitative framework.

the satisfiability and deniability of goals. To each node $G \in \mathcal{G}$ we associate two variables $Sat(G), Den(G)$ ranging in $\{F, P, N\}$ (full, partial, none) such that $F > P > N$, representing the current evidence of satisfiability and deniability of goal G . For example, $Sat(G_i) \geq P$ states that there is at least partial evidence that G_i is satisfiable. Starting from assigning an initial set of input values for $Sat(G_i), Den(G_i)$ to (a subset of) the goals in \mathcal{G} , we propagate the values through the goal relations in \mathcal{R} according to the propagation rules of Table 1.

The schema of the algorithm is described in Figure 3. *Initial*, *Current* and *Old* are arrays of $|\mathcal{G}|$ pairs $\langle Sat(G_i), Den(G_i) \rangle$, one for each $G_i \in \mathcal{G}$, representing respectively the initial, current and previous labeling states of the graph. We call the pair $\langle Sat(G_i), Den(G_i) \rangle$ a *label* for G_i . Notationally, if W is an array of labels $\langle Sat(G_i), Den(G_i) \rangle$, by $W[i].sat$ and $W[i].den$ we denote the first and second field of the i th label of W .

The array *Current* is first initialized to the initial values *Initial* given as input by the user. At each step, for every goal G_i , $\langle Sat(G_i), Den(G_i) \rangle$ is updated by propagating the values of the previous step. This is done until a fixpoint is reached, in the sense that no further updating is possible ($Current == Old$).

The updating of $\langle Sat(G_i), Den(G_i) \rangle$ works as follows. For each relation R_j incoming in G_i , the satisfiability and deniability values sat_{ij} and den_{ij} derived from the old values of the source goals are computed by applying the rules of Table 1. The result is compared with the old value, and the maximum is returned as new value for G_i .

3.3 Termination and complexity

Theorem 1. *Label_Graph($\langle \mathcal{G}, \mathcal{R} \rangle, Initial$) terminates after at most $6|\mathcal{G}|+1$ loops.*

Proof. First, from lines 6 and 14 in Figure 3 we have that, for every goal G_i ,

$$\begin{aligned} Current[i].sat &= \max(\dots, Old[i].sat), \\ Current[i].den &= \max(\dots, Old[i].den) \end{aligned}$$

```

1  label_array Label_Graph(graph ⟨G, R⟩, label_array Initial)
2    Current=Initial;
3    do
4      Old=Current;
5      for each  $G_i \in \mathcal{G}$  do
6        Current[i] = Update_Label(i, ⟨G, R⟩, Old);
7      until not (Current==Old);
8    return Current;
9
10 label Update_Label(int i, graph ⟨G, R⟩, label_array Old)
11   for each  $R_j \in \mathcal{R}$  s.t. target( $R_j$ ) ==  $G_i$  do
12     satij = Apply_Rules_Sat( $G_i, R_j, Old$ );
13     denij = Apply_Rules_Den( $G_i, R_j, Old$ );
14   return ⟨ max(maxj(satij), Old[i].sat), max(maxj(denij), Old[i].den) ⟩

```

Fig. 3. Schema of the label propagation algorithm.

so that their values are monotonically non-decreasing. In order not to terminate, at least one value per step should monotonically increase. Each of the $2|\mathcal{G}|$ variables $Sat(G_i)$ and $Den(G_i)$ admits 3 possible values and at each non-final loop at least one value increases. Thus the procedure must terminate after at most $6|\mathcal{G}| + 1$ loops. \square

Notice that the upper bound is very pessimistic, as many value updates are done in parallel. In Section 5, we report on experiments we have conducted which suggest that the algorithm generally converges after a few loops.

3.4 Soundness and completeness

We call a *value statement* an expression of the form $(v \geq c)$, $v \in \{Sat(G_i), Den(G_i)\}$ for some goal G_i and $c \in \{F, P, N\}$, with the intuitive meaning “there is at least evidence c for v ”. Thus from now on we rewrite the assertion $FS(G)$, $PS(G)$, \top as $(Sat(G) \geq F)$, $(Sat(G) \geq P)$, $(Sat(G) \geq N)$ and $FD(G)$, $PD(G)$, \top as $(Den(G) \geq F)$, $(Den(G) \geq P)$, $(Sat(G) \geq N)$ respectively.

We say that $(v_1 \geq c_1)$ is *deduced* from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by a relation axiom meaning that the corresponding assertions are deduced. For instance, $(Sat(G_1) \geq P)$ is deduced from $(Sat(G_2) \geq F)$ by axiom (11) as $PS(G_1)$ is deduced from $FS(G_2)$ by axiom (11).

We say that $(v_1 \geq c_1)$ *derives* from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by means of one propagation rule $x = f(y, z)$ [$x = f(y)$] of Table 1 if $c_1 = f(c_2, c_3)$ [$c_1 = f(c_2)$]. For instance, $(Sat(G_1) \geq P)$ derives from $(Sat(G_2) \geq P)$ and $(Sat(G_3) \geq F)$ by means of the first propagation rule. Notice that this is possible because the

operators *min* and *max* are *monotonic*, that is, e.g., $\max(v_1, v_2) \geq \max(v'_1, v'_2)$ iff $v_1 \geq v'_1$ and $v_2 \geq v'_2$.

To this extent, the rules in Table 1 are a straightforward translation of the axioms (1)-(22), as stated by the following result.

Lemma 1. $(v_1 \geq c_1)$ derives from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by means of the propagation rules of Table 1 if and only if $(v_1 \geq c_1)$ is deduced from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] with the application of one of the relation axioms (3)-(22).

Proof. For short, we consider only the AND and $+_S$ cases for $Sat(G_1)$, as the other cases are either analogous or trivial and can be verified by the reader.

$(G_2, G_3) \xrightarrow{and} G_1$: If either $(Sat(G_2) \geq N)$ or $(Sat(G_3) \geq N)$, then $(Sat(G_1) \geq N)$ is derived. This matches one-to-one the fact that from \top nothing else is deduced;

otherwise, if either $(Sat(G_2) \geq P)$ or $(Sat(G_3) \geq P)$, then $(Sat(G_1) \geq P)$ is derived. This matches one-to-one the fact that from $(Sat(G_2) \geq P)$ and $(Sat(G_3) \geq P)$ axiom (4) is applied, so that $(Sat(G_1) \geq P)$ is deduced; finally, if both $(Sat(G_2) \geq F)$ and $(Sat(G_3) \geq F)$, then $(Sat(G_1) \geq F)$. This matches one-to-one the fact that from $(Sat(G_2) \geq F)$ and $(Sat(G_3) \geq F)$, axiom (3) is applied, so that $(Sat(G_1) \geq F)$ is deduced.

$G_2 \xrightarrow{+_S} G_1$: If $(Sat(G_2) \geq N)$, then $(Sat(G_1) \geq N)$. Again, this matches one-to-one the fact that from \top nothing else is deduced;

otherwise, if either $(Sat(G_2) \geq P)$ or $(Sat(G_2) \geq F)$, then we have that $Sat(G_1) \geq \min(Sat(G_2), P) \geq P$. This matches one-to-one the fact that from either $(Sat(G_2) \geq P)$ or $(Sat(G_2) \geq F)$ axiom (11) is applied and $(Sat(G_1) \geq P)$ is deduced. \square

Given an array of labels W , we say that $(Sat(G_i) \geq c)$ [resp. $(Den(G_i) \geq c)$] is true in W if and only if $W[i].sat \geq c$ [resp. $W[i].den \geq c$]. This allows us to state the correctness and completeness theorem for *Label_Graph()*.

Theorem 2. Let *Final* be the array returned by *Label_Graph*((\mathcal{G}, \mathcal{R}), *Initial*). $(v \geq c)$ is true in *Final* if and only if $(v \geq c)$ can be deduced from *Initial* by applying the relation axioms (3)-(22).

Proof. First we define inductively the notion of “deduced from *Initial* in k steps”: (i) an assertion in *Initial* can be deduced from *Initial* in 0 steps; (ii) an assertion can be deduced from *Initial* in up to $k + 1$ steps if either it can be deduced from *Initial* in up to k steps or it can be deduced by applying a relation axiom to some assertions, all of which can be deduced from *Initial* in up to k steps.

Let $Current_k$ be the value of *Current* after k loops. We show that $(v \geq c)$ is true in $Current_k$ if and only if $(v \geq c)$ can be deduced from *Initial* in up to k steps. The thesis follows from the fact that $Final = Current_k$ for some k .

We reason by induction on k . The base case $k = 0$ is obvious as $Current_0 = Initial$. By inductive hypothesis, we assume the thesis for k and we prove it for $k + 1$.

- If.** Consider $(v \geq c)$ such that $(v \geq c)$ is deduced from *Initial* in up to $k + 1$ steps. Thus, $(v \geq c)$ is obtained by applying some relation axiom *AX* to some assertion(s) $(v_1 \geq c_1)$ [and $(v_2 \geq c_2)$], which can be both deduced from *Initial* in up to k steps. Thus, by inductive hypothesis, $(v_1 \geq c_1)$ [and $(v_2 \geq c_2)$] occur in *Current_k*. Then, by Lemma 1, $(v \geq c)$ derives from $(v_1 \geq c_1)$ [and $(v_2 \geq c_2)$] by means of the propagation rules of Table 1. Thus, if v is *Sat*(G_i) [resp. *Den*(G_i)], then c is one of the values *sat_{ij}* [resp. *den_{ij}*] of lines 14, 15 in Figure 3, so that *Current_{k+1}*[i].*sat* $\geq c$ [*Current_{k+1}*[i].*den* $\geq c$]. Thus $(v \geq c)$ is true in *Current_{k+1}*.
- Only if.** Consider a statement $(v \geq c)$ true in *Current_{k+1}*. If $(v \geq c)$ is true also in *Current_k*, then by inductive hypothesis $(v \geq c)$ can be deduced from *Initial* in up to k steps, and hence in $k + 1$ steps. Otherwise, let R_j be the rule and let $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] the statements(s) true in *Current_k* from which v has been derived. By inductive hypothesis $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] can be deduced from *Initial* in up to k steps. By Lemma 1 $(v \geq c)$ can be deduced from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by the application of one relation axiom, and thus can be deduced from *Initial* in up to $k + 1$ steps. \square

Thus, from Theorem 2, the values returned by *Label_Graph*($\langle \mathcal{G}, \mathcal{R} \rangle, \text{Initial}$) are the maximum evidence values which can be deduced from *Initial*.

4 Quantitative Reasoning with Goal Models

The qualitative approach of Section 3 allows for setting and propagating partial evidence about the satisfiability and deniability of goals and the discovery of conflicts.

We may want to provide a more fine-grained evaluation of such partial evidence. For instance, when we have $G_2 \stackrel{+s}{\vdash} G_1$, from *PS*(G_2) we can deduce *PS*(G_1), whilst one may argue that the satisfiability of G_1 is in some way less evident than that of G_2 . For example, in the goal graph of Figure 1, the satisfaction of the goal *lower environment impact* may not necessarily imply satisfaction of *increase customer loyalty*, so it may be reasonable to assume "less evidence" for the satisfaction of the latter compared to the former. Moreover, the different relations which mean partial support – i.e., $+s$, $-s$, $+D$, $-D$ – may have different strengths. For instance, in the goal graph of Figure 1, *US gas price rises* may have a bigger impact on *gas price rises* than *japanese gas price rises* if the manufacturer's market is mainly in the US.

To cope with these facts, we need a way for representing different *numerical* values of partial evidence for satisfiability/deniability and for attributing different weights to the $+s$, $-s$, $+D$, $-D$ relations. And, of course, we also need a formal framework to reason with such quantitative information.

4.1 An extended framework

In our second attempt, we introduce a *quantitative* framework, inspired by [5]. We introduce two real constants *inf* and *sup* such that $0 \leq \text{inf} < \text{sup}$. For

	$(G_2, G_3) \xrightarrow{and} G_1$	$G_2 \xrightarrow{w+S} G_1$	$G_2 \xrightarrow{w-S} G_1$	$G_2 \xrightarrow{++S} G_1$	$G_2 \xrightarrow{--S} G_1$
$Sat(G_1)$	$Sat(G_2) \otimes Sat(G_3)$	$Sat(G_2) \otimes w$	$Sat(G_2) \otimes w$	$Sat(G_2)$	$Sat(G_2)$
$Den(G_1)$	$Den(G_2) \oplus Den(G_3)$				
	$(G_2, G_3) \xrightarrow{or} G_1$	$G_2 \xrightarrow{w+D} G_1$	$G_2 \xrightarrow{w-D} G_1$	$G_2 \xrightarrow{++D} G_1$	$G_2 \xrightarrow{--D} G_1$
$Sat(G_1)$	$Sat(G_2) \oplus Sat(G_3)$		$Den(G_2) \otimes w$		$Den(G_2)$
$Den(G_1)$	$Den(G_2) \otimes Den(G_3)$	$Den(G_2) \otimes w$		$Den(G_2)$	

Table 2. Propagation rules in the quantitative framework.

each node $G \in \mathcal{G}$ we introduce two *real* variables $Sat(G), Den(G)$ ranging in the interval $\mathcal{D} =_{def}]inf, sup]$, representing the current evidence of satisfiability and deniability of the goal G . The intended meaning is that *inf* represents no evidence, *sup* represents full evidence, and different values in $]inf, sup]$ represent different levels of partial evidence.

To handle the goal relations we introduce two operators $\otimes, \oplus : \mathcal{D} \times \mathcal{D} \mapsto \mathcal{D}$ representing respectively the evidence of satisfiability of the conjunction and that of the disjunction [deniability of the disjunction and that of the conjunction] of two goals. \otimes and \oplus are associative, commutative and monotonic, and such that $x \otimes y \leq x, y \leq x \oplus y$; there is also an implicit unary operator $inv()$, representing negation, such that $inf = inv(sup)$, $sup = inv(inf)$, $inv(x \oplus y) = inv(x) \otimes inv(y)$ and $inv(x \otimes y) = inv(x) \oplus inv(y)$.

We also attribute to each goal relation $+_S, -_S, +_D, -_D$ a weight $w \in]inf, sup]$ stating the strength by which the satisfiability/deniability of the source goal influences the satisfiability/deniability of the target goal. The propagation rules are described in Table 2. As in the qualitative approach, a symmetric relation —such as, $G_2 \xrightarrow{w+} G_1$ — is a shorthand for the combination of the two corresponding asymmetric relationships sharing the same weight w —e.g., $G_2 \xrightarrow{w+S} G_1$ and $G_2 \xrightarrow{w+D} G_1$.

There are a few possible models following the schema described above. In particular, here we adopt a *probabilistic* model, where the evidence of satisfiability $Sat(G)$ [resp. deniability $Den(G)$] of G is represented as the probability that G is satisfied (respectively denied). As usual, we adopt the simplifying hypothesis that the different sources of evidence are independent. Thus, we fix $inf = 0$, $sup = 1$, and we define $\otimes, \oplus, inv()$ as:

$$p_1 \otimes p_2 =_{def} p_1 \cdot p_2, \quad p_1 \oplus p_2 =_{def} p_1 + p_2 - p_1 \cdot p_2, \quad inv(p_1) = 1 - p_1$$

that is, respectively the probability of the conjunction and disjunction of two independent events of probability p_1 and p_2 , and that of the negation of the first event. To this extent, the propagation rules in Table 2 are those of a Bayesian network, where, e.g., in $G_2 \xrightarrow{w+S} G_1$ w has to be interpreted as the conditional probability $P[G_1 \text{ is satisfied} \mid G_2 \text{ is satisfied}]$. Notice that the qualitative

Goal relation	Axioms	
$(G_2, G_3) \xrightarrow{and} G_1 :$	$(Sat(G_2) \geq x \wedge Sat(G_3) \geq y) \rightarrow Sat(G_1) \geq (x \otimes y)$	(25)
	$(Den(G_2) \geq x \wedge Den(G_3) \geq y) \rightarrow Den(G_1) \geq (x \oplus y)$	(26)
$(G_2, G_3) \xrightarrow{or} G_1 :$	$(Sat(G_2) \geq x \wedge Sat(G_3) \geq y) \rightarrow Sat(G_1) \geq (x \oplus y)$	(27)
	$(Den(G_2) \geq x \wedge Den(G_3) \geq y) \rightarrow Den(G_1) \geq (x \otimes y)$	(28)
$G_2 \xrightarrow{w+S} G_1 :$	$Sat(G_2) \geq x \rightarrow Sat(G_1) \geq (x \otimes w)$	(29)
$G_2 \xrightarrow{w-S} G_1 :$	$Sat(G_2) \geq x \rightarrow Den(G_1) \geq (x \otimes w)$	(30)
$G_2 \xrightarrow{++S} G_1 :$	$Sat(G_2) \geq x \rightarrow Sat(G_1) \geq x$	(31)
$G_2 \xrightarrow{--S} G_1 :$	$Sat(G_2) \geq x \rightarrow Den(G_1) \geq x$	(32)
$G_2 \xrightarrow{w+D} G_1 :$	$Den(G_2) \geq x \rightarrow Den(G_1) \geq (x \otimes w)$	(33)
$G_2 \xrightarrow{w-D} G_1 :$	$Den(G_2) \geq x \rightarrow Sat(G_1) \geq (x \otimes w)$	(34)
$G_2 \xrightarrow{++D} G_1 :$	$Den(G_2) \geq x \rightarrow Den(G_1) \geq x$	(35)
$G_2 \xrightarrow{--D} G_1 :$	$Den(G_2) \geq x \rightarrow Sat(G_1) \geq x$	(36)

Fig. 4. Axioms for the propagation rules in the quantitative reasoning framework.

framework of Section 3 can be seen as another such model with $\mathcal{D} = \{F, P, N\}$, $\otimes = \min()$ and $\oplus = \max()$.³

4.2 Axiomatization

As with the qualitative case, we call a *value statement* an expression of the form $(v \geq c)$, $v \in \{Sat(G_i), Den(G_i)\}$ for some G_i and $c \in [0, 1]$, with the intuitive meaning “there is at least evidence c of v ”. We want to allow the user to state and deduce non-negated value statements of the kind $(Sat(G) \geq c)$ and $(Den(G) \geq c)$ over the goal constants of the graph. As before, we call externally provided assertions about the satisfaction/denial of goals *initial conditions*.

To formalize the propagation of satisfiability and deniability evidence values through a goal graph, for every goal and goal relation in $\langle \mathcal{G}, \mathcal{R} \rangle$, we introduce the axioms (25)-(32) in Figure 4. Unlike those of Figure 2, the relation axioms in Figure 4 are not ground Horn clauses —thus, propositional— but rather first-order closed Horn formulas, so that they require a first-order deduction engine.

We say that a statement $(v \geq c)$ holds if either it is an initial condition or it can be deduced from the initial conditions and the axioms of Figure 4. We implicitly assume that $(Sat(G_i) \geq 0)$ and $(Den(G_i) \geq 0)$ hold for every G_i , and that the deduction engine —either human or machine— can semantically evaluate \otimes and \oplus and perform deductions deriving from the values of the evaluated terms and the semantics of \geq . For instance, we assume that, if $(G_2, G_3) \xrightarrow{and} G_1$ is in

³ Another model of interest is that based on a serial/parallel resistance model, in which $inf = 0$, $sup = +\infty$, $x \oplus y = x + y$, $x \otimes y = \frac{x \cdot y}{x + y}$ and $inv(x) = \frac{1}{x}$ [5].

\mathcal{R} , then $(Sat(G_1) \geq 0.1)$ can be deduced from $(Sat(G_2) \geq 0.5)$, $(Sat(G_3) \geq 0.4)$, as from (25) it is deduced $(Sat(G_1) \geq 0.5 \otimes 0.4)$, which is evaluated into $(Sat(G_1) \geq 0.2)$, from which it can be deduced $(Sat(G_1) \geq 0.1)$.

We say that there is a *weak conflict* if both $(Sat(G) \geq c_1)$ and $(Den(G) \geq c_2)$ hold for some goal G and constants $c_1, c_2 > 0$. We say that there is a *strong conflict* if there is a weak conflict s.t. $c_1 = c_2 = 1$.

4.3 The label propagation algorithm

Starting from the new numeric framework, we have adapted the label propagation algorithm of Figure 3 to work with numeric values. The new version of the algorithm differs from that of Section 3 in that: the elements in *Initial*, *Current* and *Old* range in $[0, 1]$; the input graph contain also weights to the $+_S$, $-_S$, $+_D$, $-_D$ goal relations; and, the propagation rules applied are those of Table 2.

4.4 Soundness and completeness

We say that $(v_1 \geq c_1)$ *derives* from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by means of one propagation rule $x = f(y, z)$ [$x = f(y)$] of Table 2 if $c_1 = f(c_2, c_3)$ [$c_1 = f(c_2)$]. For instance, $(Sat(G_1) \geq 0.4)$ derives from $(Sat(G_2) \geq 0.8)$ and $(Sat(G_3) \geq 0.5)$ by means of the first and propagation rule. Again, this is possible because the the operators \otimes and \oplus are monotonic.

Lemma 2. $(v_1 \geq c_1)$ *derives* from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] by means of the propagation rules of Table 1 if and only if $(v_1 \geq c_1)$ is deduced from $(v_2 \geq c_2)$ [and $(v_3 \geq c_3)$] with the application of one of the relation axioms (25)-(32).

Proof. Trivial observation that the axioms (25)-(32) are straightforward translation of the propagation rules in Table 1. \square

Given an array of labels W , we say that $(Sat(G_i) \geq c)$ [resp. $(Den(G_i) \geq c)$] is *true in W* if and only if $W[i].sat \geq c$ [resp. $W[i].den \geq c$]. This allows for stating the correctness and completeness theorem for *Label_Graph()*.

Theorem 3. *Let Final be the array returned by $Label_Graph(\langle \mathcal{G}, \mathcal{R} \rangle, Initial)$. $(v \geq c)$ is true in Final if and only if $(v \geq c)$ can be deduced from Initial by applying the relation axioms (25)-(32).*

Proof. Identical to that of Theorem 2, substituting the axioms (25)-(32) for the axioms (3)-(22), the rules in Table 2 for those in Table 1 and Lemma 2 for Lemma 1. \square

Again, from Theorem 3, the values returned by $Label_Graph(\langle \mathcal{G}, \mathcal{R} \rangle, Initial)$ are the maximum evidence values which can be deduced from *Initial*.

4.5 Termination

To guarantee termination, the condition ($Current == Old$) of line 7 in Figure 3 is implemented as:

$$\begin{aligned} \max_i (|Current[i].sat - Old[i].sat|) < \epsilon \text{ and} \\ \max_i (|Current[i].den - Old[i].den|) < \epsilon, \end{aligned} \quad (37)$$

ϵ being a sufficiently small real constant. (This is a standard practice to avoid numeric errors when doing computations with real numbers.) Thus, the algorithm loops until all satisfiability or deniability value variations have become negligible.

Theorem 4. *Label_Graph($\langle \mathcal{G}, \mathcal{R} \rangle$, Initial) terminates in a finite number of loops.*

Proof. For the same reason as in Theorem 1, the values $Current[i]_k.sat$ and $Current[i]_k.den$ are monotonically non-decreasing with k . As every monotonically non-decreasing upper-bounded sequence is convergent, both $Current[i]_k.sat$ and $Current[i]_k.den$ are convergent. Thus they are also Cauchy-convergent.⁴ It follows that condition (37) becomes true after a certain number of loops. \square

Remark 1. In the proof of Theorem 3, we have assumed the terminating condition ($Current == Old$), whilst in our implementation we use condition (37). Thus, whilst *Label_Graph()* stops when (37) is verified, the corresponding deductive process might go on, and thus deduce a value slightly bigger than the one returned by the algorithm. Since we can reduce ϵ at will, this “incompleteness” has no practical consequences.

5 Experimental Results

Both qualitative and quantitative algorithms have been implemented in Java and a series of tests were conducted on a Dell Inspiron 8100 laptop with a Pentium III CPU and 64 MB RAM (OS: GNU/Linux, kernel 2.4.7-10). The tests were intended to demonstrate the label propagation algorithms, also to collect some experimental results on how fast the algorithms converge.

In the following, we present the experimental results obtained applying the algorithms to the auto manufacturer goal graph introduced in Figure 1.

5.1 Experiments with the qualitative approach

The first set of experiments was carried out in order to demonstrate the qualitative label propagation algorithm. For each experiment, we assigned a set of labels to some of the goals and events of the auto manufacturer graph (Figure 1) and see their consequences for other nodes. The label propagation algorithm reached a steady state after at most five iterations.

⁴ An infinite sequence a_n is Cauchy-convergent if and only if, for every $\epsilon > 0$, there exist an integer N s.t. $|a_{n+1} - a_n| < \epsilon$ for every $n \geq N$. a_n is convergent if and only if it is Cauchy-convergent.

Table 3 reports the results of four different experiments (Experiments 1, 2, 3, and 4).

For each goal/event, the table shows the initial (Init) label assignment to the variables S and D, also their final (Fin) value after label propagation. For instance, in the first experiment, the initial assignment for the goal expand markets is S=P and D=N, while the final values for increase return on investment (GM) are S=N and D=P.

In the first experiment, we start with goals lower environment impact, improve mileage and lower sales price totally satisfied (S=F), goals expand markets and offer rebates partially satisfied (S=P), goals increase foreign earnings and keep labour costs low totally denied (D=F), and goals increase sales price and increase high margin sales partially denied (D=P). The propagation of such values produces a final state in which the main goal increase return on investment (GM) is partially denied.

In the second experiment we start with the same initial assignment of the first experiment, except that the event Yen rises is now totally satisfied. As reported in Table 3, this new assignment produces a different result for the root goal. Now, it is both partially satisfied and partially denied, namely we have a contradiction. Since the analysis is qualitative, it is not possible to distinguish between the two values S=P and D=P and we cannot say anything about this contradiction.

For the third experiment, we assume that goals keep labour costs low and reduce raw material costs are totally satisfied. This new assignment produce a contradiction for the root goal (totally satisfied and partially denied). Unlike from the previous experiment, we have now more evidence for the satisfaction of this goal, rather than its denial.

Finally, in the fourth experiment, though goals increase Toyota sales and increase VW sales are totally satisfied, for the root goal we still have the same contradictory result. Note that the new assignment also produces contradictory values for goals about GM, Toyota and VW sales (S=F and D=P).

5.2 Experiments with the quantitative approach

In a second set of experiments we assigned numerical weights to “+”, “-” and “-_S” lateral relationships as reported in Table 4. For instance, the goal increase sales volume contributes negatively to the goal increase Toyota sales with a weight 0.6, while the goal increase car quality contributes positively to the goal increase customer loyalty with weight 0.8. Table 5 reports the results of four different experiments.

The results of these experiments confirm those obtained with the qualitative algorithm. However, the numeric approach allows us to draw more precise conclusions about the final value of goals. This is particularly helpful for evaluating contradictions. For instance, in the second experiment (Exp 2), even though we have a contradiction for the final values of the root goal increase return on investment (GM) (i.e., S=0.8 and D=0.4), we have more evidence for its satisfaction than its denial. With the qualitative approach we had S=P and D=P, and there

Goals/Events	Exp 1		Exp 2		Exp 3		Exp 4									
	Init	Fin	Init	Fin	Init	Fin	Init	Fin								
	S	D	S	D	S	D	S	D								
increase return on investment (GM)	N	N	N	P	N	N	P	P	N	N	F	P	N	N	F	P
increase sales volume	N	N	F	P	N	N	F	P	N	N	F	P	N	N	F	P
increase profit per vehicle	N	N	N	P	N	N	P	P	N	N	F	N	N	N	F	N
increase customer appeal	N	N	F	N	N	N	F	N	N	N	F	N	N	N	F	N
expand markets	P	N	P	N	P	N	P	N	P	N	P	N	P	N	P	N
increase sales price	N	P	N	F	N	P	N	F	N	P	N	F	N	P	N	F
increase foreign earnings	N	F	N	F	N	F	P	F	N	F	P	F	N	F	P	F
lower production costs	N	N	N	F	N	N	N	F	N	N	F	N	N	N	F	N
increase high margin sales	N	P	N	P	N	P	N	P	N	P	N	P	N	P	N	P
reduce operating costs	N	N	F	N	N	N	F	N	N	N	F	N	N	N	F	N
lower environmental impact	F	N	F	N	F	N	F	N	F	N	F	N	F	N	F	N
lower purchase costs	N	N	F	N	N	N	F	N	N	N	F	N	N	N	F	N
keep labour costs low	N	F	N	F	N	F	N	F	F	N	F	N	F	N	F	N
improve economies of production	N	N	N	N	N	N	N	N	N	N	F	N	N	N	F	N
lower gas price	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
improve mileage	F	N	F	N	F	N	F	N	F	N	F	N	F	N	F	N
offer rebates	P	N	P	N	P	N	P	N	P	N	P	N	P	N	P	N
lower loan interest rates	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
lower sales price	F	N	F	N	F	N	F	N	F	N	F	N	F	N	F	N
reduce raw materials costs	N	N	N	N	N	N	N	N	F	N	F	N	F	N	F	N
outsource units of production	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
gas price rises	N	N	N	N	N	N	N	P	N	N	N	P	N	N	N	P
lower Japanese interest rates	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
US gas price rises	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Japanese gas price rises	N	N	N	N	N	N	N	P	N	N	N	P	N	N	N	P
Yen rises	N	N	N	N	F	N	F	N	F	N	F	N	F	N	F	N
Japanese rates rise	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
improve car quality	N	N	P	N	N	N	P	N	N	N	P	N	N	N	P	N
improve car services	N	N	N	P	N	N	N	P	N	N	N	P	N	N	N	P
improve customer loyalty	N	N	P	P	N	N	P	P	N	N	P	P	N	N	P	P
increase Toyota sales	N	N	N	P	N	N	N	P	N	N	N	P	F	N	F	P
increase VW sales	N	N	N	N	N	N	N	N	N	N	N	N	F	N	F	P

Table 3. Experiments with the qualitative approach

Goal/Event	Relationship	Goal/Event
increase sales volume	$\xrightarrow{0.6-S}$	increase Toyota sales
increase Toyota sales	$\xrightarrow{0.6-S}$	increase VW sales
increase VW sales	$\xrightarrow{0.6-S}$	increase sales volume
increase customer loyalty	$\xrightarrow{0.4+}$	increase sales volume
increase sales prices	$\xrightarrow{0.5-}$	increase customer loyalty
increase car quality	$\xrightarrow{0.8+}$	increase customer loyalty
improve car services	$\xrightarrow{0.7+}$	increase customer loyalty
lower environment impact	$\xrightarrow{0.4+}$	increase customer loyalty
increase sales prices	$\xrightarrow{0.3+}$	improve car services
keep labour costs low	$\xrightarrow{0.7-}$	increase car quality
improve economies of production	$\xrightarrow{0.8+}$	lower purchase costs
Yen rises	$\xrightarrow{0.8+}$	increase foreign earnings
lower Japanese interest rates	$\xrightarrow{0.4+}$	lower sales price
Japanese rates rises	$\xrightarrow{0.8-}$	lower Japanese interest rates
Japanese rates rises	$\xrightarrow{0.6+}$	Yen rises
Yen rises	$\xrightarrow{0.4-}$	Japanese gas price rises
Japanese gas price rises	$\xrightarrow{0.6+}$	gas price rises
US gas price rises	$\xrightarrow{0.6+}$	gas price rises
gas price rises	$\xrightarrow{0.8-}$	lower gas price

Table 4. Quantitative relationships for the auto manufacturer example of Figure 2

was nothing else to say about this contradiction. Analogous comments apply for the fourth experiment.

The threshold ϵ (equation 37) used in the experiments has been chosen as the smallest positive value of type float (i.e., `Float.MIN_VALUE`). With this threshold, the algorithm converged in five iterations for all four experiments.

5.3 A Final Experiment

A final set of experiments were carried out with graphs with a bigger number of cycles. The goal of the experiments was to evaluate the growth of the number of iterations with respect to the growth of the number of cycles.

We generated such graphs randomly. Starting from the auto manufacturer model of Figure 1, we generated new graphs adding randomly new contribution links between goals and thereby generating new cycles. Of course, since we were only interested on the structure of the graphs, we have not associated to them any semantics. Figure 5 shows an example of such a random graph.

The results of the experiments showed that new cycles increase significantly the number of iterations. For instance, applying the quantitative algorithm to the graph showed in Figure 5, the number of iterations for the last three experiments presented in Table 5 (Exp2, Exp2, and Exp3) increased from 6 to 33, from 6 to 80 and from 5 to 24, respectively.

Goals/Events	Exp 1				Exp 2				Exp 3				Exp 4			
	Init		Fin		Init		Fin		Init		Fin		Init		Fin	
	S	D	S	D	S	D	S	D	S	D	S	D	S	D	S	D
increase return on investment (GM)	0	0	0	.4	0	0	.8	.4	0	0	.9	.2	0	0	.9	.6
increase sales volume	0	0	1	.1	0	0	1	.1	0	0	1	.2	0	0	1	.6
increase profit per vehicle	0	0	0	.4	0	0	.8	.4	0	0	.9	0	0	0	.9	0
increase customer appeal	0	0	1	0	0	0	1	0	0	0	1	0	0	0	1	0
expand markets	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0
increase sales price	0	.5	0	.8	0	.5	0	.8	0	.5	0	.8	0	.5	0	.8
increase foreign earnings	0	.9	0	.9	0	.9	.8	.9	0	.9	.8	.9	0	.9	.8	.9
lower production costs	0	0	0	.9	0	0	0	.9	0	0	.6	0	0	0	.6	0
increase high margin sales	0	.6	0	.6	0	.6	0	.6	0	.6	0	.6	0	.6	0	.6
reduce operating costs	0	0	.8	0	0	0	.8	0	0	0	.8	0	0	0	.8	0
lower environmental impact	.9	0	.9	0	.9	0	.9	0	.9	0	.9	0	.9	0	.9	0
lower purchase costs	0	0	.9	0	0	0	.9	0	0	0	.9	0	0	0	.9	0
keep labour costs low	0	.9	0	.9	0	.9	0	.9	.9	0	.9	0	.9	0	.9	0
improve economies of production	0	0	0	0	0	0	0	0	0	0	.7	0	0	0	.7	0
improve mileage	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
lower gas price	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0
offer rebates	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0	.3	0
lower loan interest rates	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
lower sales price	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0	.8	0
reduce raw materials costs	0	0	0	0	0	0	0	0	.7	0	.7	0	.7	0	.7	0
outsource units of production	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
gas price rises	0	0	0	0	0	0	0	.2	0	0	0	.2	0	0	0	.2
lower Japanese interest rates	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
US gas price rises	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Japanese gas price rises	0	0	0	0	0	0	0	.4	0	0	0	.4	0	0	0	.4
Yen rises	0	0	0	0	1	0	1	0	1	0	1	0	1	0	1	0
Japanese rates rise	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
improve car quality	0	0	.6	0	0	0	.6	0	0	0	.6	0	0	0	.6	0
improve car services	0	0	0	.2	0	0	0	.2	0	0	0	.2	0	0	0	.2
improve customer loyalty	0	0	.5	.2	0	0	.5	.2	0	0	.5	.2	0	0	.4	.5
increase Toyota sales	0	0	0	.6	0	0	0	.6	0	0	0	.6	1	0	1	.6
increase VW sales	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	.6

Table 5. Results with the quantitative approach

Finally, a set of experiments were carried out with graphs with a bigger number of nodes (from hundred to thousand nodes). The results showed that also for bigger graphs value propagation performs in negligible time ($\leq 10^{-2}s$). This suggests us that our approach can be applied in real life applications where goals models can count more than hundred goals.

6 Conclusions

We have presented a formal framework for reasoning with goal models. Our goal models use AND/OR goal relationships, but also allow more qualitative relationships, as well as contradictory situations. A precise semantics has been given for all goal relationships which comes in a qualitative and a numerical form. Moreover, we have presented label propagation algorithms for both the qualitative and the numerical case that are shown to be sound and complete with respect to their respective axiomatization. Finally, the paper reports experimental re-




Fig. 5. A partial goal model for GM, with added cycles.

sults on the label propagation algorithms that suggest that the algorithms are scalable and can work with goal models having hundreds of nodes.

The main limitation of our framework is about the definition of contribution links and the labels assignment. This is a problem that an analyst can encounter in modeling complex situations where many goals are related one another and different perspectives have to be considered in assigning initial values to the goals. We are currently working on a methodology that can help and drive an analyst in building goal models considering different point of views and discovering relationships between goals.

Another important limitation of our approach is that we deal with conflicts but we do not resolve them. Also here, we are currently working to top-down approach that should allow us to find a set of possible assignments that can guarantee the satisfaction (or the denial) of specific goals without any conflict.

Future research directions include applying different techniques, such as Dempster Shafer theory [15], to take into consideration the sources of information in the propagation algorithms. This will allow us to consider, for instance, the reliability and/or the competence of a source of evidence and then model and analyze different scenarios [14, 13]. We also propose to apply our framework to real case studies to confirm its validity.

References

1. A. I. Anton. Goal based requirements analysis. In *Proceedings of the 2nd International Conference on Requirements Engineering ICRE'96*, pages 136–144, 1996.
2. A. I. Anton and C. Potts. The use of goals to surface requirements for evolving systems. In *Proceedings of the International Conference on Software Engineering (ICSE '98)*, pages 157–166, Kyoto, Japan, April 1998.
3. A. Blum and M. Furst. Fast planning through planning graph analysis. *Artificial Intelligence*, 90(1-2), 1997.
4. B. Boehm. Identify quality-requirements conflicts. In *Proceedings of the 2nd International Conference on Requirements Engineering ICRE'96*, Colorado spring, Colorado, 1996.
5. A. Bundy, F. Giunchiglia, R. Sebastiani, and T. Walsh. Calculating Criticalities. *Artificial Intelligence*, 88(1-2), December 1996.
6. A. Dardenne, A. van Lamsweerde, and S. Fickas. Goal-directed requirements acquisition. *Science of Computer Programming*, 20(1–2):3–50, 1993.
7. R. Jarvis, G. McArthur, J. Mylopoulos, P. Rodriguez-Gianolli, and S. Zhou. Semantic Models for Knowledge Management. In *Proc. of the Second International Conference on Web Information Systems Engineering (WISE'01)*, 2001.
8. H. Kaindl. A design process based on a model combining scenarios with goals and functions. *IEEE Transactions on Systems, Man and Cybernetic*, 30(5):537–551, September 2000.
9. H. Kautz, D. McAllester, and B. Selman. Encoding Plans in Propositional Logic. In *Proceedings International Conference on Knowledge Representation and Reasoning*. AAAI Press, 1996.
10. J. Mylopoulos, L. Chung, and B. Nixon. Representing and Using Non-Functional Requirements: A Process-Oriented Approach. *IEEE Transactions on Software Engineering*, 6(18):483–497, June 1992.

11. A. Newell and H. Simon. GPS: A Program that Simulates Human Thought. In E. Feigenbaum and J. Feldman, editors, *Computers and Thought*. McGraw Hill.
12. N. Nilsson. *Problem Solving Methods in Artificial Intelligence*. McGraw Hill, 1971.
13. C. Rolland, Grosz, and R. Kla. Experience with goal-scenario coupling. In *Proceedings of the Fourth IEEE International Symposium on Requirements Engineering*, Limerik, Ireland, 1999.
14. C. Rolland, C. Souveyet, and C. Ben Achour. Guiding goal modelling using scenarios. *IEEE Transactions on Software Engineering*, 24(12), December 1998.
15. G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press., 1976.
16. A. v. Lamsweerde, R. Darimont, and E. Letier. Managing conflicts in goal-driven requirements engineering. *IEEE Transactions on Software Engineering*, 24(11):908–926, November 1998.
17. Daniel S. Weld. Recent Advances in AI Planning. *AI Magazine*, 20(2):93–123, 1999.