

Representing and Using Nonfunctional Requirements: A Process-Oriented Approach

John Mylopoulos, Lawrence Chung, and Brian Nixon

Abstract—The paper proposes a comprehensive framework for representing and using nonfunctional requirements during the development process. The framework consists of five basic components which provide for the representation of nonfunctional requirements in terms of interrelated goals. Such goals can be refined through refinement methods and can be evaluated in order to determine the degree to which a set of nonfunctional requirements is supported by a particular design. Evidence for the power of the framework is provided through the study of accuracy and performance requirements for information systems.

I. INTRODUCTION

THE complexity of an information system is determined partly by its functionality—i.e., what the system does—and partly by global requirements on its development or operational cost, performance, reliability, maintainability, portability, robustness, and the like. These *nonfunctional requirements*¹ play a crucial role during system development, serving as selection criteria for choosing among myriads of decisions. Errors of omission or commission in laying down and taking properly into account such requirements are generally acknowledged to be among the most expensive and difficult to correct once the information system has been completed. Surprisingly, nonfunctional requirements have received little attention by researchers and are definitely less well understood than other, less critical factors in software development. As far as software engineering practice is concerned, they are generally stated only informally during requirements analysis, are often contradictory, difficult to enforce during software development and to validate for the user once the final system has been built. The only glimmer of technical light in an otherwise bleak landscape originates in technical work on software quality metrics that allow the quantification of the degree to which a software system meets nonfunctional requirements [3], [5], [26].

There is not a formal definition or a complete list of nonfunctional requirements. In a report published by the Rome Air Development Center (RADC) [7], nonfunctional requirements (“software quality attributes” in their terminology) are classified into consumer-oriented (or *software quality factors*) and technically-oriented attributes (or *software quality criteria*). The former refers to nonfunctional requirements

observable by the consumer, such as efficiency, correctness, and interoperability. The latter addresses system-oriented requirements such as anomaly management, completeness, and functional scope. Table I shows the RADC consumer-oriented attributes. The nonfunctional requirements listed in the table apply to all software systems. However, additional requirements may apply for special classes of software. For instance, precision would be an important nonfunctional requirement for a numerical analysis software package, while accuracy (of maintained information) might feature prominently during the development of an information system.

Two basic approaches characterize the formal treatment of nonfunctional requirements and we shall refer to them as *product-oriented* and *process-oriented*. The first attempts to develop formal definitions of nonfunctional requirements so that a software system can be evaluated as to the degree to which it meets its requirements. For example, measuring software visibility may include, among other things, measuring the amount of branching in a software system. This might be achieved globally with a criterion such as: “There shall be no more than X branches per 1,000 lines of code” or locally with a criterion such as “There shall be no more than Y% of system modules that violate the above criterion.”

The product-oriented approach has received almost exclusive attention in the literature and is nicely overviewed in [26]. Earlier work by Boehm *et al.* [5] considered quality characteristics of software, noting that designer-awareness alone improved the quality of the final product. Also supporting a quantitative approach to software quality, Basili and Musa [3] advocate models and metrics of the software engineering process from a management perspective. It is interesting that Hauser *et al.* [21] provide a methodology for reflecting customer attributes in different phases of automobile design.

An alternative approach, explored in this paper, is to develop techniques for justifying design decisions during the software development process. Instead of evaluating the final product, the emphasis here is on trying to rationalize the development process in terms of nonfunctional requirements. Design decisions may affect positively or negatively particular nonfunctional requirements. Design decisions may affect positively or negatively particular nonfunctional requirements. These positive and negative dependencies can serve as basis for arguing that a software system indeed meets a certain nonfunctional requirement or explaining why it does not.

Orthogonally, treatments of nonfunctional requirements can be classified into *quantitative* and *qualitative* ones. Most of the

Manuscript received October 1, 1991; revised February 5, 1992. Recommended by A. Borgida and M. Jarke.

The authors are with the Department of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 1A4.

IEEE Log Number 9200171.

¹Also referred to as *constraints* [41], *goals* [31], and *quality attributes* [26] in the literature.

TABLE I
THE RADCS SOFTWARE QUALITY CONSUMER-ORIENTED ATTRIBUTES [7]

Acquisition	User Concern	Quality Attribute
Performance—How well does it function?	How well does it utilize a resource? How secure is it? What confidence can be placed in what it does? How well will it perform under adverse conditions? How easy is it to use it?	Efficiency Integrity Reliability Survivability Usability
Design—How valid is the design?	How well does it conform to requirements? How easy is it to repair? How easy is it to verify its performance?	Correctness Maintainability Verifiability
Adaptation—How adaptable is it?	How easy is it to expand or upgrade its capability or performance? How easy is it to change? How easy is it to interfere with another system? How easy is it to transport? How easy is it to convert for use with another application?	Expandability Flexibility Interoperability Portability Reusability

product-oriented approaches alluded to earlier are quantitative in the sense that they study quantitative metrics for measuring the degree to which a software system satisfies a nonfunctional requirement. The process-oriented treatment proposed here, on the other hand, is definitely qualitative, adopting ideas from qualitative reasoning [1]. It should be acknowledged that a process-oriented treatment of nonfunctional requirements need not be qualitative. Indeed, one could imagine quantitative measures for, say, software visibility that can be used as the system is being developed to offer advance warning that nonfunctional requirements are not being met. Qualitative techniques were chosen here primarily because it was felt that the problem of quantitatively measuring an incomplete software system is even harder than that of measuring the final product.

Of course, neither product-oriented quantitative metrics nor process-oriented qualitative measures have a monopoly on properly treating nonfunctional requirements. They are best seen as complementary, both contributing to an evolving comprehensive framework for dealing with nonfunctional requirements.

Two sources of ideas were particularly influential on our work. The first involves recent work on decision support systems, such as that described in [19], [28] and [29]. Lee's work, for example, adopts an earlier model for representing design rationale [38] and extends in by making explicit the goals presupposed by arguments. The work reported here can be seen as an attempt to adopt this model to the representation and use of nonfunctional requirements. The second source of ideas is the DAIDA environment for information system development [23] which has provided us with a comprehensive software development framework covering both notations for requirements modeling, design, implementation and decision support, as well as a starting point on how the treatment of nonfunctional requirements might be integrated into that framework. Users of the DAIDA environment are offered three languages through which they can elaborate requirements, design, and implementation specifications. In developing a design specification, the user consults and is constrained by corresponding requirements specifications. Likewise, the generation of an implementation is guided by a corresponding design specification. *Dependency links* represent design decision and relate implementation objects to their design




Fig. 1. Employee and report hierarchy.

counterparts and design objects to their requirements counterparts. The framework proposed in this paper focuses on these dependency links and how they can be justified in terms of nonfunctional requirements. An early description of the framework and an account of how it relates to DAIDA can be found in [12].

The example used throughout this paper is an *expense management system* for a hypothetical research project, similar to the one used in [6]. According to the example, project members from organizations based in different countries register for and attend various meetings. They then submit their expense summaries to an expense management system, which maintains all such information and generates expense reports for each member, meeting, and project. As shown in Fig. 1, there are several kinds of employees, including secretaries, engineers, and researchers, who are in turn classified into computer researchers, math researchers, and so on.

Establishment of the framework is achieved in two steps. First, the framework is presented in Section II. The presentation includes motivation, the framework's structure and short suggestive examples. This framework is then elaborated and illustrated in the following two sections by examining its application respectively to accuracy and performance requirements for information systems. The final section summarizes the contributions of this research and presents a number of open questions and directions for further research.

II. REPRESENTING NON-FUNCTIONAL REQUIREMENTS: A PROCESS-ORIENTED FRAMEWORK

Formally, the proposed framework consists of five major components:² a set of *goals* for representing nonfunctional requirements, design decisions, and arguments in support of or against other goals; a set of *link types* for relating goals or goal relationships (hereafter links) to other goals; a set of generic *methods* for refining goals into other goals; a collection

²An earlier version of portions of this and the next section have appeared in [13].

of *correlation rules* for inferring potential interactions among goals; and finally, a *labelling procedure* which determines the degree to which any given nonfunctional requirement is being addressed by a set of design decisions. The examples throughout this section concentrate on accuracy and to a lesser extent operating cost requirements for information systems.

During the design process, goals are organized into a goal graph structure, very much in the spirit of AND/OR trees used in problem solving [34]. Unlike traditional problem solving and planning frameworks, however, goals representing nonfunctional requirements can rarely be said to be “accomplished” or “satisfied” in a clearcut sense. Instead, different design decisions contribute positively or negatively towards a particular goal. Accordingly, for the rest of the discussion we will speak of goal *satisficing* [42]³ to suggest that generated software is expected to satisfy within acceptable limits, rather than absolutely, nonfunctional requirements.

A. Goals

The space of goals includes three mutually exclusive classes, namely, *nonfunctional requirements goals (NFR goals)*, *satisficing goals*, and *argumentation goals*. In general, each goal will have an associated *sort* and zero or more *parameters* whose nature depends on the goal sort. For example, an operating cost requirement might have as parameter a desired upper bound on the annual operating costs of the system under development. Sorts may be further subdivided into subsorts, representing special cases for each goal class. For instance, the performance sort may have subsorts *Time Performance* (or simply *Time*) and *SpacePerformance* (or simply *Space*), representing respective time and space performance requirements on a particular system. *Goals*, *NFRGoals*, *SatGoals* and *ArgGoals* will refer respectively to the set of all possible goals, NFR goals, satisficing goals, and argumentation goals.

1) *Nonfunctional Requirements Goals*: The sorts for such goals range over the different categories of such requirements, including accuracy, security, development, operating or hardware costs, and performance. For our expense management system, suppose that it is expected of the system under development to maintain accurately employee data. Such a goal might be represented by:

Accuracy[**attributes(Employee)**]

where *Accuracy* is the goal sort and the parameter of **attributes(Employee)** evaluates to the set of all attributes associated with the data class **Employee**. The interpretation of this goal is that instances of the attributes of the data class **Employee**, i.e., all attributes of employees, ought to be maintained accurately in the system’s database. As another example, it may also be expected that the system under development make minimal demands on manpower. This can be treated as an operating cost requirement, and since there are several contributing factors to operating costs (manpower,

maintenance, etc.), this requirement might be represented as *OperatingCost*[**manpower**].

2) *Satisficing Goals*: These are also sorted and parameterized. In this case, however, the sorts range over different categories of design decisions that might be adopted in order to satisfy one or more nonfunctional requirements goals. The parameters associated with each sort, again, depend on the nature of the corresponding satisficing goal. For instance, one way to satisfy the accuracy goal mentioned earlier might be to validate all employee data entered into the system. This can be represented as a satisficing goal:

Validation[**attributes(Employee)**],

where *Validation* is the goal sort and **attributes(Employee)** is as before. This goal, in turn, might be refined into another satisficing goal,

Validated By[**JohnWong, attributes(Employee)**],

representing the situation that **JohnWong** will be doing the validation.

3) *Argumentation Goals (or Arguments)*: These always have the sort *Claim*, with subsorts *FormalClaim*, and *InformalClaim*, representing formally or informally stated evidence or counter-evidence for other goals or goal refinements. Consider:

Formal Claim[$\exists e$: *ValidatedBy*[**e, attributes(Employee)**]
 \wedge *EmpStatus*(**e, Sec I**)]

This argumentation goal supports the refinement from the goal of validating employee data to the one assigning **JohnWong** to the task, by claiming that class I secretaries will perform the validation. In contrast,

InformationClaim[“Rigorous examination is recommended for publications by employees.”]

is an informally-stated argumentation goal supporting the previous argumentation goal by pointing out why class I secretaries should validate employee data.

B. Link Types

As indicated earlier, design proceeds by refining one or more times each goal, the parent, into a set of other goals, the offspring. Unlike AND/OR goal trees, where the relationship between a collection of offspring and their parent can only be AND or OR, in our proposed framework there can be several different types of relationships or *link types* describing how the satisficing of the offspring (or failure thereof) relates to the satisficing of the parent goal. The need for at least some link types is evidenced in [5], which states that some quality characteristics are necessary, but not sufficient, for achieving others. Boehm *et al.* then use a four-grade scale to correlate each quality metric with quality attributes in the final product.

Links may relate a parent goal to one or several of its offspring. In fact, links may also be used to relate other links to argumentation goals, to indicate that an argument offers positive or negative support for a particular refinement of a goal. Thus, links too need to be satisfied either through a formal refinement process or through arguments provided by

³[42] actually uses the term to refer to decision methods that look for satisfactory solutions rather than optimal ones. The term is adopted here in a broadened sense since in the context of nonfunctional requirements, even the notions of a solution or optimality of a solution may be unclear.

the designer.

Let *Links* denote the set of all links and *satisfied* be a predicate which is true of satisfied goals or links and false of others. Also, let *denied* be a predicate which is true of goals and links that have been shown unsatisficeable ("unsolvable" in problem solving terminology [34]). If

$$\text{Propositions} = \text{Goals} \cup \text{Links}$$

then *satisfied* and *denied* are predicates taking a proposition as argument.

Sometimes a proposition will be found to be satisficeable—thanks to one refinement—and deniable—thanks to another. For instance, the accuracy goal for employee data might be satisficeable thanks to a validation procedure adopted for all such data, but deniable because of a user interface that permits general access to this information. To deal with such conflicting cases, we need to distinguish between a proposition being satisfied or denied, on one hand, and a proposition being *potentially* satisficeable or deniable thanks to some refinement on the other. Accordingly, two more predicates, *satisficeable* and *deniable* are introduced to deal with the latter case.

The set of logical types to be used for links is presented below. For each type, axioms are provided which formalize its semantics in terms of the predicates just introduced (see Axiom 1 below).

The link type *sub* is also intended to convey the sense that G_1 contributes partially to the satisficing of G_0 . This can be expressed as follows: If *satisfied*(*sub*(G_0, G_1)) then there exist propositions G_2, \dots, G_n such that

$$\neg(\text{satisfied}(G_2) \wedge \dots \wedge \text{satisfied}(G_n)) \longrightarrow \text{satisficeable}(G_0)$$

but

$$\text{satisfied}(G_1) \wedge \text{satisfied}(G_2) \wedge \dots \wedge \text{satisfied}(G_n) \\ \wedge \text{satisfied}(\text{sub}(G_0, G_1)) \longrightarrow \text{satisficeable}(G_0)$$

In words, if G_1 is a *sub*(proposition) of G_0 then there exist propositions G_2, \dots, G_n which cannot achieve the satisficing of G_0 without the contribution of G_1 .

Two additional link types are introduced to represent negative influences of one goal on another (see Axiom 2 below).

In words, if G_1 is a *negative sub*(proposition) of G_0 then denial of G_1 leads to the satisficing of G_0 and satisficing of G_1 contributes to the denial of G_0 .

Finally, it is useful to define the *eql* (*equivalent*) link type in terms of the link types introduced here:

$$\text{eql: } \text{Propositions} \times \text{Propositions.} \\ \text{eql}(G_0, G_1) \equiv \text{sup}(G_0, G_1) \wedge \text{sup}(G_0, G_1) \\ \wedge \text{sup}(G_0, G_1) \wedge \text{sup}(G_0, G_1)$$

At times, it may be hard to determine *a priori* the logical relationship between a set of offspring and their parent goal without further expansion of the goal graph. For example, the designer may see that a certain hiring policy for technical staff is relevant, without being sure of its impact on a particular goal, say, in justifying the assignment of a class I secretary to the task of validating employee data. This situation is accommodated through three variations of an undetermined link type:

$$\text{und: } \text{Propositions} \times \text{Propositions.} \\ \text{und}(G_0, G_1) \text{ indicates the possible presence of}$$

Axiom 1

$$\text{AND: } \text{Propositions} \times 2^{\text{Propositions}}. \\ \text{satisfied}(G_1) \wedge \text{satisfied}(G_2) \wedge \dots \wedge \text{satisfied}(G_n) \wedge \text{satisfied}(\text{AND}(G_0, \{G_1, G_2, \dots, G_n\})) \longrightarrow \text{satisficeable}(G_0) \\ \text{satisfied}(\text{AND}(G_0, \{G_1, G_2, \dots, G_n\})) \wedge (\text{denied}(G_1) \vee \text{denied}(G_2) \vee \dots \vee \text{denied}(G_n)) \longrightarrow \text{deniable}(G_0)$$

$$\text{OR: } \text{Propositions} \times 2^{\text{Propositions}}. \\ \text{denied}(G_1) \wedge \text{denied}(G_2) \wedge \dots \wedge \text{denied}(G_n) \wedge \text{satisfied}(\text{OR}(G_0, \{G_1, G_2, \dots, G_n\})) \longrightarrow \text{deniable}(G_0) \\ \text{satisfied}(\text{OR}(G_0, \{G_1, G_2, \dots, G_n\})) \wedge (\text{satisfied}(G_1) \vee \text{satisfied}(G_2) \vee \dots \vee \text{satisfied}(G_n)) \longrightarrow \text{satisficeable}(G_0)$$

$$\text{sup: } \text{Propositions} \times \text{Propositions.} \\ \text{satisfied}(G_1) \wedge \text{satisfied}(\text{sup}(G_0, G_1)) \longrightarrow \text{satisficeable}(G_0)$$

$$\text{sub: } \text{Propositions} \times \text{Propositions.} \\ \text{denied}(G_1) \wedge \text{satisfied}(\text{sub}(G_0, G_1)) \longrightarrow \text{deniable}(G_0)$$

Axiom 2

$$\neg\text{sup: } \text{Propositions} \times \text{Propositions.} \\ \text{satisfied}(G_1) \wedge \text{satisfied}(\neg\text{sup}(G_0, G_1)) \longrightarrow \text{deniable}(G_0)$$

$$\neg\text{sub: } \text{Propositions} \times \text{Propositions.} \\ \text{denied}(G_1) \wedge \text{satisfied}(\neg\text{sub}(G_0, G_1)) \longrightarrow \text{satisficeable}(G_0) \\ \text{If } \neg\text{sub}(G_0, G_1) \text{ then there exists } G_2, \dots, G_n \text{ such that} \\ \neg(\text{satisfied}(G_2) \wedge \dots \wedge \text{satisfied}(G_n) \wedge \text{satisfied}(\neg\text{sub}(G_0, G_1))) \longrightarrow \text{deniable}(G_0)$$

but

$$\text{satisfied}(G_1) \wedge \text{satisfied}(G_2) \wedge \dots \wedge \text{satisfied}(G_n) \wedge \text{satisfied}(\neg\text{sub}(G_0, G_1)) \longrightarrow \text{deniable}(G_0)$$

positive or negative influence between G_0 and G_1 .

Likewise, $+und$ and $-und$ indicate, respectively, possible positive or negative influence between two propositions.

C. Methods

Goals may be refined by the designer, who is then responsible for satisficing not only the goal's offspring but also the refinement itself represented as a link. Alternatively, the framework provides goal refinement methods (methods for short) which represent generic procedures for refining a goal into one or more offspring, such as:

"To maintain accurately data about class \mathbf{x} , you need to maintain accurately data about all relevant subclasses of \mathbf{x} ."

Every such refinement is represented in terms of a link having one of the types of the previous section and which is considered satisfied.

Generally, a method has the form

$$\begin{aligned} x_1/C_1, x_2/C_2, \dots, x_n/C_n : SelP(x_1, x_2, \dots, x_n) \\ G_0(x_1, \dots, x_n) \xrightarrow{L} \{G'(x_1, \dots, x_n)\} \\ \text{For all } G' \text{ such that } Pred(G', x_1, \dots, x_n) \end{aligned}$$

Here G_0 represents the parent goal, predicate $Pred$ determines the set of offspring while L is the link type relating G_0 to its offspring. The refinement of G_0 through a method is subject to the method's selection criterion, $SelP$, consisting of a Boolean expression with free variables x_1, x_2, \dots, x_n . These are bound to objects of type c_1, \dots, c_2 , respectively, when the method is applied.

There are three types of goal refinement methods, corresponding to the three types of goals introduced earlier.

1) *Goal Decomposition Methods*: These are usually AND decomposition methods of the form $SelP : G \xrightarrow{AND} \{G_1, G_2, \dots, G_n\}$ used to decompose a goal G into an AND set of offspring G_1, G_2, \dots, G_n . For instance, the following method decomposes a goal having a class as argument into goals having as arguments its immediate specializations:

$$\begin{aligned} x/Class : G[x] \xrightarrow{AND} \left\{ G_{[x_i]} | (x_i \text{ isA } x) \wedge \right. \\ \left. \forall x_j \left[((x_i \text{ isA } x_j) \wedge (x_j \text{ isA } X)) \right. \right. \\ \left. \left. \Rightarrow ((x_j = x_i) \vee (x_j = x)) \right] \right\}. \end{aligned}$$

Since there are three specializations of **Employee** in our example, the accuracy goal *Accuracy* [**attributes(Employee)**] (abbreviated as $A[\mathbf{attributes(Employee)}]$) can be refined using the subclass goal decomposition method:

$$\begin{aligned} A[\mathbf{attributes(Employee)}] \xrightarrow{AND} \\ \{A[\mathbf{attributes(Researcher)}], \dots, \\ A[\mathbf{attributes(Secretary)}]\}. \end{aligned}$$

In Fig. 2, offspring are shown underneath the parent goal. Link types are sometimes omitted from figures. Now each of

these goals needs to be satisfied in turn. Likewise, satisficing the goal of $A[\mathbf{attributes(Researcher)}]$ requires that all attributes of **Researcher** be maintained accurately. This decomposition can be accomplished by a method of the form:

$$\begin{aligned} x/Class : A[\mathbf{attributes(x)}] \xrightarrow{AND} \\ \{A[\mathbf{attr(x)}] | attr \in \mathbf{attributes(x)}\}. \end{aligned}$$

This method leads to the following further decomposition of $A[\mathbf{attributes(Researcher)}]$. Assuming that research employees have attributes **degree** and **publ** (publications), in addition to those of **Employee**,

$$\begin{aligned} A[\mathbf{attributes(Researcher)}] \xrightarrow{AND} \\ \{A[\mathbf{Researcher.name}], \dots, \\ A[\mathbf{Researcher.degree}], \\ A[\mathbf{Researcher.publ}]\}. \end{aligned}$$

2) *Goal Satisficing Methods*: Such methods refine a goal into a set of satisficing goals, thereby committing the design that is being generated to particular design decisions. Returning to our example, there may be two satisficing methods offered for the goal $A[\mathbf{Researcher.publ}]$. If the publication record of each researcher is obtained from existing databases, the accuracy of this information might be ensured through periodic auditing of those databases. If, on the other hand, these data are fed directly by the employee in question, a method may call for the validation of the data by the employee's manager:

$$\begin{aligned} i/InformationItem : A[i] \xrightarrow{+und} Audit[i] \\ i/InformationItem : A[i] \xrightarrow{sup} Validation[i]. \end{aligned}$$

Using these methods, $A[\mathbf{Researcher.publ}]$ can be refined to $Audit[\mathbf{Researcher.publ}]$ or $Validation[\mathbf{Researcher.publ}]$ through $+und$ and sup links, respectively. Note that the designer may later change the type of the $+und$ link once the design has proceeded further and it can be determined that auditing indeed leads to more accurate publication data. Clearly, selection of one of the two alternatives leads to very different types of user interfaces for the system under development. In particular, if validation is selected, all publication information will have to be confirmed by another person, while auditing calls for the inclusion of an audit requirement on the database from which publication data are imported.

3) *Argumentations Methods*: These methods refine a goal or a link into an argumentation goal, thereby indicating evidence/counter-evidence, in terms of arguments, for the satisficing of a goal. For instance, a formal claim consisting of a conjunction could be refined into claims of each conjunct related to the parent through an AND link.

Fig. 2 illustrates the goal structure that might be generated by the simple example we have been introducing piecemeal. In the bigger picture of information system development, a source object, say a component of a requirements specification, is mapped into one (or possibly several) target object(s), say components of a design specification [13]. The dependencies among these objects are shown through dependency links

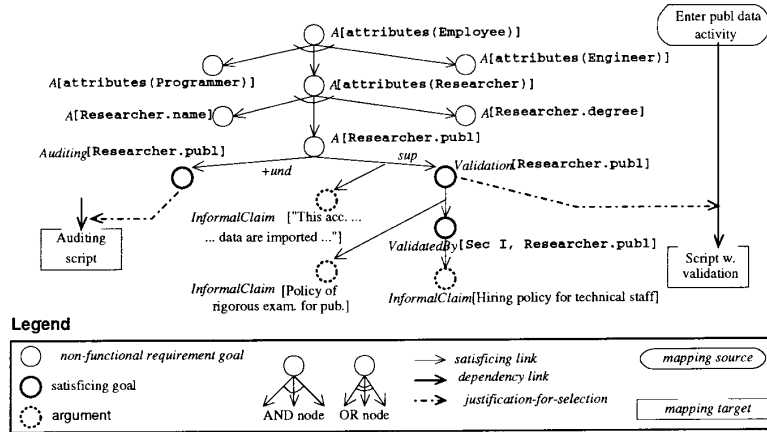


Fig. 2. Goal graph structure for accurate employee attributes.

on the left- and right-hand sides of Fig. 2. The use of the goal structure generated by the designer from nonfunctional requirements, possibly with the help of methods, is intended to help her *select* among alternatives and *justify* her design decisions. She can selectively focus attention, thus controlling goal structure expansion.

D. Correlation Rules

As indicated earlier, the nonfunctional requirements set down for a particular system may be contradictory. For instance, having built-in procedures for validating or auditing the data managed by the information system in general requires additional manpower thereby interfering with the operating cost requirement, *Operating Cost[manpower]*. Guidance is needed by the designer in discovering such implicit relationships and in selecting the satisfying goals that best meet a set of given NFR goals. This is achieved either through external input by the designer herself or through the representation of generic interactions between goals through *correlation rules*.

Consider a satisfying goal whereby the system under design will offer an interface for "casual" users, say, all company employees, who wish to query or update the system's database.

$$\text{OperatingCost}[\text{manpower}] \xrightarrow{\text{sub}} \text{CasualUserInterface}[\text{Employee, Database}].$$

Unfortunately, making the database readily available to all employees is likely to lead to data inaccuracy, thereby interfering with accuracy goals. This can be expressed, for example, by a rule such as:

$$\text{CasualUserInterface}[e, i] \wedge \text{cardinality}(e) > 5 \wedge A[i'] \\ \wedge i' \subset i \rightarrow -\text{sub}(A[i], \text{CasualUserInterface}[e, i]).$$

This rule can be used to infer a *-sub* link between the goals *Casual User Interface[Employee, Database]* and *A[attributes(Employee)]*, assuming that *attributes(Employee) ⊂ Database* and that there are more than five employees.

Likewise, consider a security goal (with sort *S*) discouraging secretaries from accessing research publication data. Now a

validation goal which is positive for *A[Researcher.publ]*, with a class I secretary as validator, would also contribute negatively to such a security goal, and vice versa:⁴

$$(S[i, e, \text{accessCond}] \wedge \text{ValidatedBy}[e', i'] \wedge (i \subseteq i') \\ \wedge \text{isA}(e, e') \wedge \text{HigherClassification}(e, e') \\ \wedge \text{accessCond}) \\ \rightarrow (-\text{sup}(S[i, e, \text{accessCond}], \text{Validation}[i']) \\ \wedge -\text{sup}(\text{Validation}[i'], S[i, e, \text{accessCond}])).$$

Finally, consider the case where two satisfying goals interfere with each other because of dependence on a critical resource. This competition may be *synergistic* or *antagonistic*, leading, respectively, to positive or negative argumentation. For instance, two unrelated goals calling for auditing and validation of information may influence each other positively (through *sub* links) if there is no one on staff assigned to either task, because they jointly suggest the hiring of personnel data but may not individually justify hiring additional staff. If, however, the argumentative structure indicates that they can share an agent, one new staff member may be hired for the two tasks.

$$\text{Validation}[i] \wedge \text{Audit}[i'] \wedge \\ \neg \exists e, e' / \text{Employee} : \text{ValidatedBy}[e, i] \\ \wedge \text{AuditedBy}[e', i'] \\ \rightarrow \text{sub}(\text{Validation}[i], \text{Audit}[i']) \\ \wedge \text{sub}(\text{Audit}[i'], \text{Validation}[i]).$$

It is now possible to describe the expansion procedure that starts with a set of NFR goals and iteratively expands them into a goal graph structure. Throughout the expansion, the system maintains a list of all propositions that are to be refined, called *Open*, while the list *Closed* includes all propositions that have been completely refined.

Once a proposition has been selected from *Open* for refinement, the designer chooses whether she wants to propose a refinement or apply one of the available methods. Carrying out a chosen refinement involves creating propositions for the

⁴The representation of security requirements is adopted from [20].

offspring and the newly-created link and adding each to *Open*. Correlation links are then introduced for the new propositions, using both the designer's judgement and correlation rules in the system. This process is repeated for the chosen proposition until there are no more refinements the system or the designer can offer. At this time, the proposition is placed on the *Closed* list and another open proposition is selected.

E. The Labelling Procedure

Given partially constructed goal graph structure, the labelling procedure determines the status of each node on the graph through the assignment of a label. A node or link of the graph is labelled *satisfied* if it is satisficeable and not deniable; *denied* if it is deniable but not satisficeable; *conflicting* if it is both satisficeable and deniable; and *undetermined* if it is neither. These labels are denoted, respectively, by S, D, C , and U . They are similar to those in [16] and generally ones used in qualitative reasoning frameworks [1]. The U label, in particular, is intended to represent situations where either there is both positive and negative support, albeit inconclusive, for a given goal, or there is neither positive nor negative support.

The labelling algorithm consists of two basic steps. For each proposition \mathbf{P} on a given goal graph, the algorithm first computes the individual effect of each satisfied outgoing link. Secondly, the individual effects of all outgoing links are combined into a single label taking one of the four possible values mentioned earlier.

Given the open-ended nature of the argumentation process (i.e. the premise built into this framework that only some of the relevant knowledge is formally represented, the rest remaining with designers) the framework calls for an interactive labelling procedure where the designer may be asked to step in and determine the appropriate label for a particular proposition having supporting but inconclusive evidence. For this reason, the labels characterizing the influence of one set of offspring towards a parent include S, D, C , and U , as mentioned before, but also U^- and U^+ representing, respectively, inconclusive positive or negative support for a parent. Moreover, $?$ indicates a situation where the designer is to determine the label that characterizes the contribution of a proposition toward another. Note that the labels U^-, U^+ , and $?$ are introduced by the first step of the labelling algorithm and are eliminated by the second when the set of all contributions from all outgoing links associated with a given proposition are combined into a single label, S, D, C , or U .

Table II shows the propagation rules along different link types (always from offspring to parent). According to these rules, *sup* propagates S while *sub* propagates D ; *-sup* inverts an S label into a D and *-sub* inverts a D label into a S one. The propagation rules for *AND* and *OR* links are based on the ordering of labels $S \geq U, C \geq D$ and is defined as follows:

$$\text{Assuming } \text{AND}(G_0, \{G_1, G_2, \dots, G_n\}) \\ \text{then } \text{label}(G_0) = \min_i(\text{label}(G_i))$$

$$\text{Assuming } \text{OR}(G_0, \{G_1, G_2, \dots, G_n\}) \\ \text{then } \text{label}(G_0) = \max_i(\text{label}(G_i))$$

TABLE II
THE INDIVIDUAL EFFECT OF SOURCE LABEL UPON ITS DESTINATION LABEL

label _{source}	link type				
	<i>sub</i>	<i>sup</i>	<i>-sub</i>	<i>-sup</i>	<i>und</i>
S	U^+	S	U^-	D	U
D	D	U^-	S	U^+	U
C	$?$	$?$	$?$	$?$	U
U	U	U	U	U	U

Once all contributed labels have been collected for a given proposition, the second step of the labelling procedure combines them into a single label. Assuming that \mathbf{L} is the bag⁵ contributed to a given proposition, consisting of labels from the set S, D, C, U, U^-, U^+ , the U^+ and U^- labels are first combined by the designer into one or more S, D, C , and U labels. The resulting set of labels is then combined into a single one, by choosing the minimal element, $\min_{l \in \mathbf{L}}(l)$, and assuming a label ordering $S, D \geq U \geq C$.

It is interesting to compare our labelling procedure with those of truth maintenance systems (TMS's) [15], [17]. They record and maintain beliefs, their justifications and assumptions, while distinguishing facts from defeasible beliefs, which are either accepted or rejected. As with TMS's, our graph labelling procedure recursively propagates values of offspring to parents. However, our procedure is not automatic, but interactively allows the designer to deal with inconclusive evidence. While we have *AND* and *OR*, comparable to TMS conjunction and disjunction, our link types have additional values, all of which are inputs to computing *individual effect* in our first step. In applying the propagation rules of Table II, *links*, which are not included in TMS beliefs, must be *satisfied*. Unlike TMS's, we then combine individual effects of label values including qualitative (*conflicting*) and open-ended (*undetermined*) ones, using a label ordering in the second step.

III. DEALING WITH ACCURACY REQUIREMENTS

A major consideration in building an information system is the degree to which its design encourages accuracy of the information being managed. For example, a system that allows users to update information in their own files may be user-friendly, but one will not have confidence in the information it contains. There are many ways to promote accuracy requirements for an information system. Restricting access to resources is only one such technique.

Within our framework treating accuracy requirements as goals offers directional guidance for the overall design process. In particular, accuracy requirements are used below as criteria for selecting a particular design in order to address elements of a given functional requirement.

A. Goals of Accuracy Requirements

The goals of accuracy requirements have *Accuracy* as the

⁵ \mathbf{L} is a bag because duplicate labels are useful; for instance, several positive supporting links indicated by several U^+ 's may be combined into an S label by the designer.

sort and **InformationItem** (abbreviated **Info**) as the parameter. They are expressed as $Accuracy[i]$ (abbreviated as $A[i]$), where i is a collection of information items. Information items may be categorized into three types of propositions: i) that an entity in the system has the *property* of some class during some time interval; ii) that an entity in the system has an *attribute* with a certain value during some time interval; iii) that an object in the system, say a record, has one and only one corresponding *entity* in the application domain, say an employee. Accuracy requirements can then be expressed on collections of such information items, such as the employee *attributes* of Section II (see [12] or [13] for this). In general, satisficing accuracy goals is understood in terms of the degree of confidence in the accuracy of information items maintained by the project system.

B. Goal Refinement Methods

1) *Goal Decomposition Methods*: We present below some examples of accuracy decomposition methods, to be used in the illustration of Section III-D.

- *Subclass* method: In order to establish the accuracy of a class c of information items, establish the accuracy of each immediate specialization, c_i , of c . This is a special case of the goal decomposition methods mentioned in Section II-C.
- *Subset* method: To establish the accuracy of a set of information items, establish the accuracy of each subset of information items. Similarly, a *superset* method can be provided.
- *IndividualAttributes* method: To establish the accuracy of the attributes of a class of information items, establish the accuracy of each attribute of the class.
- *DerivedInfo* method: To establish the accuracy of a set of information items, establish that the function that derives them is correctly designed and that all of the function's source parameters, currently in the system, are accurate.
- *AttributeSelection* method: To establish the accuracy of an information item obtained by a sequence of attribute selections (e.g., **Joe.project.budget**), establish the accuracy of each information item obtained in the sequence (e.g., **Joe.project**, **Project.budget**).
- *Conservation* method: To establish the accuracy of a collection of information items that can no longer be decomposed into information items currently in the system, establish i) their accuracy, when received by the system from some external agent, and ii) their correct internal manipulation by the system.
- *CorrectExternalManipulation* method: To establish the accuracy of information items upon receipt, establish *CorrectInfoFlow*, i.e., they were accurate when they were first transmitted by the original sender, and have subsequently been correctly manipulated until receipt by the system. *CorrectInfoFlow* is a sub-sort of *Correctness* goals which, unlike accuracy goals, are related to actions that induce certain results.

2) *Goal Satisficing Methods*: Taking the premise that the

accuracy of information items depends entirely on the process in which they are manipulated within the system and its environment, accuracy satisficing goals alter *that* process.⁶ Accuracy satisficing goals include *preventive*, *curative*, and *precautionary techniques*. They affect the level of our confidence in the accuracy of information items.

Preventive accuracy satisficing goals detect and disallow inaccuracies, when information items are received by the system. Most of them require direct interaction between the system and agents in the application domain. They can be specialized by varying the agent who performs the needed task, the volume of information items, evidences attached, the time of processing and output, etc.

- *Confirmation*: The informant, either a machine or a person, double-checks the previously submitted information item. This technique can be specialized: to *confirmation-via-identical-channel* if the confirmation and first transmission use the channel; otherwise to *confirmation-via-distinct-channel* (e.g., via a daisy-channel).
- *Verification*: A *verifier*, who is a co-worker of the sender of information item makes a duplicate entry of the item onto some medium in the system (e.g., via duplicate IBM key-entry operation). As with confirmation, verification can be specialized to *verification-via-identical-channel* or *verification-via-distinct-channel*.
- *Validation*: A *validator* performs checking in the application domain, using certain records or procedural guidelines to ensure that the information item meets predetermined standards. The type and thoroughness of the checking can be reflected in specialized methods: *creation-validation* for directly contacting the information source, *experimentation* for re-testing the information item, etc.
- *Audit*: An accuracy *auditor* uses procedures to periodically go through suspicious sampled information items.
- *Consistency-checking*: To prevent frequently-occurring errors, the system enforces certain integrity constraints (e.g., check-sums incorporated into ISBN's).

Curative satisficing goals trace inaccuracies to their source, and provide for recovery from inaccuracies. *Precautionary* satisficing goals make information flow more reliable in terms of what is involved, such as senders, receivers, and communication channels.

3) *Goal Argumentation Methods*: These methods support or deny the use of accuracy satisficing goals and various refinements in terms of arguments. Examples include:

- *resource-assignment*: In performing a task for a satisficing goal, assign resources in the application domain. For example, one can support a refinement from a goal of validating expense summaries to one assigning a staff member to the task, by claiming that class I secretaries

⁶Martin [30], for instance, offers a glossary of techniques for improving accuracy.

will perform the validation.

$$\text{Validation}[\text{Summary}] \xrightarrow{\text{sup}} \text{FormalClaim}[\exists e : \text{ValidatedBy}[\text{e, Summary}] \wedge \text{EmpStatus}(\text{e, Sec I})]$$

- *policy-manual-consultation*: When a question arises about the applicability of various types of methods, consult policy manuals in the application domain.
- *priority-based-selection*: Select a method among alternatives according to their relative priority. For example, for a satisficing goal which is good for high-priority accuracy goal A but bad for goal B , the priority would be a positive argument for A but negative for B .

C. Correlation Rules

Accuracy satisficing goals, such as verification, usually contribute positively to accuracy goals (such as $A[\text{attribute}(\text{Researcher})]$) provided that the (verification) process is rapid. Otherwise, information items will become less timely. This perturbation is an example of a satisficing goal becoming negative. For example:⁷

$$\text{VerifiedBy}[\text{e, i, t}] \wedge \text{Excessive}(t) \wedge A[i'] \wedge i' \subseteq i \rightarrow -\text{sub}(A[i'], \text{Verification}[i])$$

Verification may be negative for a security goal if the verifier is not allowed to access the information item to be verified.

A security satisficing goal (such as *Mutual-ID*) or a user-friendliness satisficing goal (such as Casual User Interface) can be positive or negative for an accuracy goal. Consider *Mutual-ID*[\mathbf{a} :Agent, \mathbf{i} :Info, \mathbf{p} :Procedure, \mathbf{t} :Time]. To mutually ensure the identity of the agent \mathbf{a} , attempting to access certain information items \mathbf{i} , and the identity of the system process, both the agent and the system, during time interval \mathbf{t} , go through a test procedure, \mathbf{p} , which requires alternating queries and answers by the two (this is similar to the *challenge response* process [37]). This would be positive for accuracy goals if a malicious user, in the absence of mutual identification, would penetrate the system and falsify the information item.

Table III summarizes some of the correlations. Entries of the form: $\langle \text{condition, orientation} \rangle$ mean “if the *condition* holds, then the relationship between the requirements goal and satisficing goal is given by the *orientation*.”

The table is similar in spirit to the “relationship matrix” [21], which indicates, informally and without correlation rules, how much each engineering characteristic affects each customer quality requirement in terms of four types of values: strong positive, medium positive, medium negative, or strong negative.

An accuracy satisficing goal can be synergistic or antagonistic with respect to another satisficing goal, for one or more types of nonfunctional requirements goals. Suppose a single channel can sometimes be shared for confirmation and verification. Now *confirmation-via-distinct-channel* and *verification-via-distinct-channel* are mutually synergistic if a

⁷The time parameter (\mathbf{t}) is omitted when not needed.

TABLE III
CORRELATION OF NFR GOAL (NFR GOAL)
WITH SATISFICING GOALS (SATGOAL)

SatGoal	NFR Goal	Accuracy	Security
Verification		$\langle R_1, \text{sub} \rangle, \langle R_2, -\text{sub} \rangle$	$\langle R_4, -\text{sub} \rangle$
Mutual-ID		$\langle R_3, \text{sub} \rangle$	$\langle \text{True}, \text{sub} \rangle$
CasualUser Interface		$\langle \text{True} - \text{sub} \rangle$	

$$R_0 : A[i'] \wedge i' \subseteq i$$

$$R_1 : \text{VerifiedBy}[\text{e, i, t}] \wedge R_0$$

$$R_2 : R_1 \wedge \text{Excessive}[\text{t}]$$

$$R_3 : \text{Mutual-ID}[\text{e, i, p, t}] \wedge R_0 \wedge$$

$$\text{Informant-ID} - \text{established}[\text{c}]$$

$$R_4 : S[\text{i, e, AccessCond}] \wedge \text{VerifiedBy}[\text{e}', \text{i}', \text{t}] \wedge \text{i} \subseteq \text{i}'$$

$$\wedge \text{is.A}(\text{e, e}') \wedge \text{HigherClassification}(\text{e, e}') \wedge \text{AccessCond}$$

new channel can be installed for shared use by the two, but mutually antagonistic if the channel is unshareable.

D. Illustration

Consider the example of research expense management system in Section I. Now assume that $A[\text{attributes}(\text{Rpt})]$ is the root node of the goal tree representing an accuracy requirement, “all the attributes of expense reports should be accurate”. The root goal can be refined with the *subclass* method into three offspring corresponding to the subclasses of Rpt , specified as part of functional requirements (See Fig. 3):

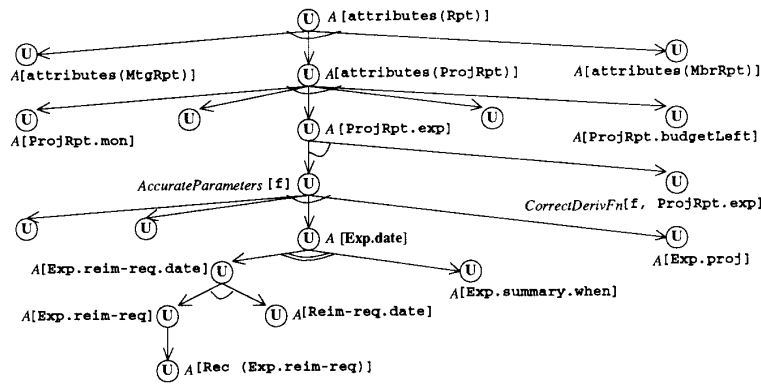
$$A[\text{attributes}(\text{Rpt})] \xrightarrow{\text{AND}} \{A[\text{attributes}(\text{ProjRpt})], A[\text{attributes}(\text{MtgRpt})], A[\text{attributes}(\text{MbrRpt})]\}.$$

Now each of these offspring needs to be satisfied. Focusing on the subgoal of $A[\text{attributes}(\text{ProjRpt})]$, the goal of $A[\text{attributes}(\text{ProjRpt})]$ is decomposed by the *individualAttributes* method in terms of the accuracy of the attributes.

$$A[\text{attributes}(\text{ProjRpt})] \xrightarrow{\text{AND}} \{A[\text{ProjRpt.mon}], \dots, A[\text{ProjRpt.budgetLeft}]\}$$

The legend or the symbols are given in Fig. 2. (When omitted, assume that the link type for satisficing and argumentation methods is *sup* in the remainder of this paper.)

Focusing on $A[\text{ProjRpt.exp}]$, the designer indicates that ProjRpt.exp is a derived information item, where the derivation function, \mathbf{f} , is shown in Fig. 3. Thus, the *derivedInfo* decomposition method is instantiated: the function needs to be correctly designed and the parameters of the function should be accurate. Next the *subset* method is instantiated for the



where $f = \text{ComputeAmount}(\text{Exp}, \text{Exp.date}, \text{Exp.proj}, \text{ProjRpt.mon}, \text{ProjRpt.proj})$

Fig. 3. Goal graph structure with decompositions for accurate expense-reports attributes.

decomposition of $\text{AccurateParameters}[f]$:

$$A[(\text{ProjRpt.exp})] \xrightarrow{\text{AND}} \{ \text{CorrectDerivFn} \\ [f, \text{ProjRpt.exp}], \\ \text{AccurateParameters}[f] \}$$

$$\text{AccurateParameters}[f] \xrightarrow{\text{AND}} \{ A[\text{Exp.date}], \\ \dots, A[\text{Exp.proj}] \}.$$

Two competing alternatives (i.e., disjunctive refinements) are foreseen by the designer for the date the expense was incurred: it may come from either the expense reimbursement requests (by requiring the members to send their reimbursement request forms to the central management office), or the expense summary (by requiring the secretary to submit it directly):

$$A[\text{Exp.date}] \xrightarrow{\text{OR}} \\ \{ A[\text{Exp.reim-req.date}], A[\text{Exp.summary.when}] \}.$$

To explore the first alternative, the designer applies the *attributeSelection* method:

$$A[\text{Exp.reim-req.date}] \xrightarrow{\text{AND}} \\ \{ A[\text{Exp.reim-req.date}], A[\text{Reim-req.date}] \}.$$

To illustrate manipulation of information items, we introduce some method applications which were not shown in Fig. 2. The designer indicates that **Exp.reim-req** should be received from an external agent. According to the *conservation* method, **Exp.reim-req** should be received from an external agent. According to the conservation method, **Exp.reim-req** should be both accurate when received and correct when processed by the system:

$$A[(\text{Exp.reim-req})] \xrightarrow{\text{AND}} \{ A[\text{Rec}(\text{Exp.reim-req})], \\ \text{CorrectProcessing} \\ [\text{Reim-req.date}] \}.$$

When invoked by the designer, the labelling procedure assigns *U* (undetermined) to all goals, since there are no closed leaves.

Although omitted, all the links in Fig. 3 have *S* as their labels, since they are the results of generic-method applications.

The designer uses the correct External Manipulation method to refine the accuracy of the received information item (Fig. 4):

$$A[\text{Rec}(\text{Exp.reim-req})] \xrightarrow{\text{AND}} \{ \text{CorrectCreation} \\ [\text{Rec}(\text{Exp.reim-req})], \\ \text{CorrectInfoFlow} \\ [\text{Rec}(\text{Exp.reim-req})] \}.$$

Unfortunately, ensuring the correct creation and subsequent transmissions of the item from the creator to the system is in many cases costly and impractical. Accordingly, the designer may resign himself to using some satisficing methods for $A[\text{Rec}(\text{Exp.reim-req})]$. In selecting a method, the designer uses the argumentation method of *policy-manual-consultation, Designer's Consultation Guidelines* (DCG). The designer regard the validation method to be appropriate:

$$\text{CorrectInfoFlow}[\text{Rec}(\text{Exp.reim-req})] \xrightarrow{\text{sup}} \\ \text{Validation}[\text{Exp.reim-req}].$$

Note how the method above (call it validation_e) is supported by a designer-supplied argument:

$$\text{validation}_e \xrightarrow{\text{sup}} \text{InformalClaim}["\text{DCG} : \text{Careful examination is preferred for those materials that are directly related to issuing a check}"].$$

To satisfy the goal of validation, the designer again consults the DCG and discovers that class I secretary is one, but not the only, good class of candidate for carrying out the validation. Thus, a class I secretary is assigned (call the assignment, assign_v) and the assignment is supported by:

$$\text{Validation}[(\text{Exp.reim-req})] \xrightarrow{\text{eqI}} \text{FormalClaim} \\ [\text{ValidatedBy}[\text{Sec I}, \dots] \wedge \dots] \\ \text{assign}_v \xrightarrow{\text{sup}} \text{InformalClaim} \\ ["\text{DCG} : \text{For } \dots, \text{consider class Isecretary.} "].$$

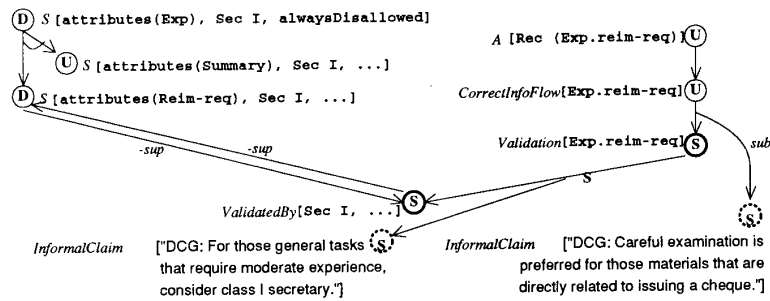


Fig. 4. Mutual exclusion in satisfying accuracy of received reimbursement information.

Now suppose, as in Fig. 4, that a security requirement was considered earlier: reimbursements should not be revealed to secretaries with a job classification below II. However, this is in direct conflict (i.e., mutually exclusive or sufficiently negative) with using a secretary of class I as the validator. Now the system uses the correlation rules to propose two new links with type *-sup*:

$$\begin{aligned}
 &S[\text{attributes}(\text{Reim} - \text{Req}), \\
 &\quad \text{SecI, AlwaysDisallowed}] \xrightarrow{-sup} \\
 &\quad \text{ValidatedBy}[\text{SecI, Exp.reim} - \text{req}, \dots] \\
 &\text{ValidatedBy}[\text{SecI, Exp.reim} - \text{req}, \dots] \xrightarrow{-sup} \\
 &\quad S[\text{attributes}(\text{Reim} - \text{req}), \text{SecI}, \dots]
 \end{aligned}$$

Suppose that the designer assigns *FormalClaim*[**ValidatedBy**(**Sec I**, ...)] the label *S* (satisfied), and labels all the other leaves.⁸

The labelling procedure of Section II propagates the labels upwards. Some of the results are shown in Fig. 4. Since the validation by a class I secretary is sufficient counter-evidence, the security goal (see left-hand side of figure) is labelled *D* (denied). This value and the *U* value in the AND link in the upper-left corner are further propagated; the minimum value of the two is selected, resulting in *D*.

Note that the denial of the root security goal is not final. Instead of a class I secretary, the designer may see if a higher-ranking staff member can do the validation. Other satisficing methods may be considered as well. The designer will choose one alternative and provide an argument for later use in justifying the final design; then the labelling procedure will update the labels which reflect the current status of the process.

The success, or lack thereof, of goal satisficing methods relies on the cooperation between the system and agents in the environment, which is described in the *user's procedure manual*,⁹ which is initially drafted during the design process. The manual indicates policies that the agents in the environment should obey when interacting with the system in order to satisfy the methods selected. For instance, if a *verification* method is selected, the manual indicates that a member must

transfer his expense information to the system and to the project office which will enter the same information into the system.

At the design stage, the choice of method (related to requirements for accuracy, security, and the like) results in selection among design alternatives.¹⁰ In the next section, we consider how performance goals are dealt with in the implementation stage.

IV. DEALING WITH PERFORMANCE REQUIREMENTS¹¹

The previous section illustrates the dynamic *process* aspect of design. This section focuses on *performance requirements*, as a second example of how a class of nonfunctional requirements can be treated within our proposed framework. Unlike accuracy requirements, which were treated in the context of system design, performance requirements will be treated during the implementation phase when *designs* are mapped on to *implementations*.

A starting point for understanding good system performance is the set of standard definitions from computer systems theory (e.g., [27]), such as achieving low response time and suitable device utilizations. In practice,¹² performance goals often focus on response time and throughput, and are developed for particular applications systems. They are often stated briefly, yet users expect the system to somehow meet their (implicit) performance concerns. And as we will see, performance goals can result in very complex goal-graph structures.

When implementing an information system using performance as a main criterion, the implementor has to abandon generic implementation algorithms and structures. Instead, implementation techniques have to be selected on a case-by-case basis from a number of alternatives. Inputs to this mapping process are: 1) a given set of *implementation alternatives*; 2) the *source schema* (some portion of the design specification); 3) a *workload characterization* for the particular system (e.g., an estimate of the number of researchers to be handled by the expense management system); 4) *performance goals*, specified for a particular system. As examples of performance goals, one

¹⁰ See the description of dependency types in Section II.C.

¹¹ An earlier version [36] of portions of this section appears in the *Proc. 3rd Int. Workshop on Database Programming Languages*, Nafplion, Greece, August 1991.

¹² Many thanks to Michael Brodie for his insight on the use of performance goals in industry.

⁸ To resolve conflicts, a negotiation-based approach may be taken (e.g., [24], [40]). We use argumentation methods to record how conflicts are resolved, e.g., by attachment of priorities.

⁹ In acquiring formal requirements [39] recognizes the need for generating documents which are in spirit similar to our manuals.

could require that a researcher registering for a meeting should get from the system under design fast response time, and that storage requirements for information on all researchers be minimized. The framework detailed in section II is then applied for the satisficing of these qualitative goals. Outputs of the process are the target implementation, goal graphs, and a prediction of performance [36] calculated in terms of a performance model.

It is interesting to contrast the treatment offered in this section with other research based on the transformational approach, such as the TI system [2]. TI, like its transformation-based peers, focuses on correctness requirements, i.e. making sure that the generated implementation is consistent with the original specification. Performance, if treated at all, is treated as a selection criterion among alternative transformations. Kant's early work [25], on the other hand, does address performance goals. Her framework, however, focuses on conventional programming-in-the-small rather than information system development, relies on quantitative performance measures (which are available for her chosen domain but are, unfortunately, not available for information systems because of their complexity) and assumes an automatic programming setting rather than the dialectical software development process adopted here.

A. Layered Goal Structures

Since generating efficient implementations is better understood than some of the other phases of information system development, we can impose additional structure in the representation of performance goals. This is accomplished through a series of language layers, which account for potentially interacting data model features, implementation techniques, and performance characteristics of design languages. This layered approach is inspired by a framework for prediction of performance of relational databases [22]. As design decisions are made at higher layers, corresponding to higher levels of abstraction, they are reflected in lower layers that describe the system in more detail. The layering shows where to introduce inputs related to design components, thus providing the information needed to make implementation decisions, while controlling the number of concepts to consider at a time.

We apply this layering approach to performance-based selection among implementation alternatives for conceptual design specification languages.¹³ Our layering organizes some recent work on performance and implementation from the areas of semantic data models and object oriented systems.¹⁴ For each layer, there are goal graphs whose refinements have an impact on graphs at lower layers. Fig.5 shows a series of linguistic subsets, where higher-level languages include more features supported by semantic data models: 0) The target relational data model, such as the database system facilities offered by the DBPL language [6]; 1) entities, both persistent data entities (such as **John**, an instance of **Researcher**), and finite entities (e.g., integers), arranged in classes; 2) attributes,

¹³See [36] for more on performance prediction for conceptual design languages.

¹⁴This includes results on record layout for entities and attributes [4], [9], [35], [46], enforcement of constraints [8], [44], and process scheduling [11].




Fig. 5. Layers arranged in a grid.




Fig. 6. The performance sort.

defined on entity classes, roughly corresponding to the Entity-relationship Model [10]; 3) transactions, modeled as classes with attributes and instance entities; 4) entities and transactions with attributes, and classes arranged in *IsA hierarchies*, roughly corresponding to the Taxis subset described in [35]; 5) the above Taxis subset, extended with constraints; 6) the source conceptual design specification language, including constraints and long-term processes (whose nature has aspects of entities and activities, as well as constraints), comparable to Taxis [11] or TDL [6].

B. Performance Goal Refinement Methods

Performance goals drive selection of implementation alternatives, and are stated in terms of concepts applicable to information systems, such as response time. Many of our methods are based on features specific to information systems, and their implementation. All performance goals use the *Performance* sort. There are several *sub-sorts*, some of which are shown in Fig. 6.

1) *Goal Decomposition Methods*: One aspect of goal decomposition involves the selection of an appropriate *sub-sort*. For example, we can use the *time-space* goal decomposition method to decompose the goal of “good performance for the Researcher class at layer 4—in into the goals of good time performance for Researcher at layer 4 and good space performance for Researcher at layer 4:¹⁵

$$P[\text{Researcher}, 4] \xrightarrow{\text{AND}} \{ \text{Time}[\text{Researcher}, 4], \text{Space}[\text{Researcher}, 4] \}$$

Likewise, a goal involving time can be decomposed by the *throughput-response time* method, and a goal involving space can be decomposed by the *main memory–secondary storage* method.

Another aspect of goal decomposition involves the decomposition of goal parameters. The *subclass* and *individualAttributes* performance goal decomposition methods are similar

¹⁵Here *P* stands for the *Performance* sort.

to the structural methods with the same names described in Section III.

a) *Operational method*: A performance goal on an information item *i* (such as a class, or an attribute of a class) can be decomposed into the corresponding goal for the *operations* o_j on the item.

$$P[i, \text{Layer}] \xrightarrow{\text{sub}} \{P[o_j(i), \text{Layer}] \mid o_j(i) \text{ is an operation on } i\}$$

This method can be specialized. The *individual-bulk operations* method decomposes a goal on the basis of whether an operation manipulates one or many items. By the *implementation components* method, a goal for an operation is decomposed into lower-layer components of the operation.

b) *Static-dynamic schema method*: While the conceptual design (or schema) of an information system may remain constant, in some cases it may be expected to change. For example, new specializations of **Researcher** might be added over time with relative efficiency, without requiring the entire system to be shut down and restarted. This method decomposes a performance goal for an information item, on the basis of whether the schema is expected to change.

2) *Goal Satisficing Methods*: Some performance goal satisficing methods are available from systems performance engineering and semantic data model implementation techniques. *Indexing* is positive for time but negative for space. By *earlyFixing*, early connection is made between an action and the instructions that achieve it [43]. A specialization of *earlyFixing* is *staticOffsetDetermination*, which determines offsets statically, rather than at execution time. Using *accessManyAttributesPerTuple*, if many of the attributes in a tuple will frequently be accessed, time goals can be positively satisfied.

3) *Goal Argumentation Methods*: Expected or actual usage statistics, and predictions of performance of implementation alternatives, can be used as arguments for a choice of satisficing methods. Suppose we know that all references to information item *i* in a segment of code can be uniquely determined statically, rather than being expressions with several possible values. We write: *ExplicitReferences*[*i*, **Layer**]. An argument that information item is subject to frequent changes in the schema can be written: *FrequentSchemaChanges*[*i*, **Layer**].

C. Illustration

Returning to our research expense management system example, we will illustrate how a designer builds a goal graph for a few layers starting at Layer 4 (IsA hierarchies), showing some goal refinement methods and the impact of higher-layer goals upon lower ones. Fig. 1 shows part of the IsA hierarchy for the example. Of the 12 attributes (not shown) of the **Researcher** class, ten, including **Name**, are inherited from **Employee**, while two others, including **Meeting**, are not inherited. Additional input information, such as the distribution of class populations, is required to characterize the workload. Of the 2000 employees, for instance, 1000 are researchers, including 700 computer researchers and 300 mathematicians. Of the two noninherited attributes of **Researcher**, the **Meeting**

attribute is very frequently accessed. This information can be included in argumentation structures.

Layer 4 selects implementations for attributes of entity classes; in the presence of IsA hierarchies, there are several possible implementation techniques. Inheritance hierarchies result in collections of attribute values whose appearance is more like a “staircase” than a relational table:¹⁶

	Name	Meeting	OperatingSystems
ComputerResearcher			
Researcher			
Employee			

As a result, a simple relational representation may waste space. Options include using one relation per class: storing either all attributes (newly defined or inherited) of a particular class in the corresponding relation (*horizontal splitting*), or only the newly defined attributes (*vertical splitting*).

Turning to the top of the goal graph (See Fig. 7 and the legend for symbols in Fig. 2), the implementor’s Layer 4 goal is good performance for the attributes of the **Researcher** entity class. First, the implementor decides to use the *time-space* method to decompose the goal into good time performance and good space performance for the attributes. The implementor can then use the *individual-bulk operations* method to decompose the time goal based on whether operations affect many entities, or just an individual entity. The goal of good time performance for individual operations on attributes of the **Researcher** class can now be decomposed by the *individualAttributes* method, resulting in goals for individual operations on the **Name** attribute, the **Meeting** attribute, etc. The implementor then focuses on the **Meeting** attribute, and observes that while most of the attributes of **Researcher** are inherited, **Meeting** is one of the two that is not. The implementor also recalls that **Meeting** is frequently accessed. By storing only the non-inherited attributes together, we have a small tuple size; moreover, of the attributes which are stored in the tuple, a high proportion will be frequently accessed. The actual value of this ratio (50%) is recorded as an argument for selecting the satisficing goal of improving time performance by accessing many attributes per tuple. This satisficing goal lead to selection of an implementation using vertical splitting for the attributes of **Researcher** and its specializations. Another implementation alternative is horizontal splitting, which can offer better space performance.

At Layer 4, the designer dealt with the **Researcher** class in an IsA hierarchy, leaving the mapping target at Layer 3 being the **Researcher** class using vertical splitting. The satisficing goal *AccessManyAttributesPerTuple*[...] is refined to the Layer 3 (transactions) goal of good time performance for individual operations on the **Meeting** attribute of **Researcher**. Thus the implementor continues addressing the goal of good time performance, but at Layer 3, which deals with operations without inheritance.

¹⁶In the illustration, not all attributes are shown.

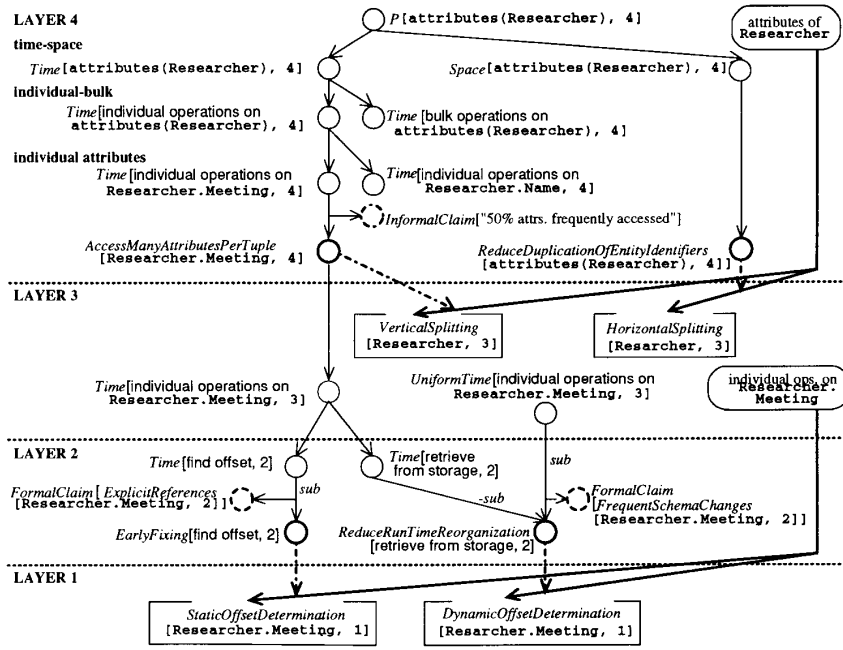


Fig. 7. A Performance goal graph.

The implementor decomposes the Layer 3 *Time* Goal (see middle left of Fig. 7) according to the implementation components of the operation (only some of which are shown). The result is a set of Layer 2 (attributes) time goals — for finding the offset for the *Meeting* attribute field within a relational tuple, retrieving the value from secondary storage, etc. The implementor focuses on finding the offset quickly; *earlyFixing* is positive. The implementor reviews the source schema and observes that *Meeting* is always referenced explicitly in the code. *ExplicitReferences[Researcher.Meeting, 2]* is recorded as an argument for the *sub* link, and static offset determination for the *Meeting* attribute is chosen as an implementation technique. Thus the implementor has dealt with a Layer 2 issue, resulting in a mapping target at Layer 1.

An alternative implementation is dynamic offset determination. The *-sub* link records its negative impact on the goal of minimizing time. However, this would have a positive impact on another goal—offering *uniform* time performance. As shown in the lower right-hand side of Fig. 7, when dealing with frequent schema changes, a structure that reduces expensive run-time reorganization can offer less variation in response time.

V. CONCLUSIONS

The main contribution of this research is that it offers a concrete framework for integrating nonfunctional requirements into the software development process, at least for information systems. In tackling this task, our research extends earlier work by Lee [14], [28], [29] and [38]. The framework is still under refinement and a prototype implementation is underway, intended to provide a vehicle for more thorough testing and for gaining experience with the framework's strengths and

weaknesses.

Much remains to be done with this work. Firstly, the framework needs to be applied to other types of nonfunctional requirements and life-size examples. Secondly, the framework needs a theoretical foundation for representing and reasoning with nonfunctional requirements. This foundation needs to include a semantics for nonfunctional requirements. For example, what does it really mean to claim that a particular design decision enhances system accuracy concerning employee data? Moreover, a proof theory based on this semantics is required, including efficient algorithms for special classes of inferences related to nonfunctional requirements. The whole framework we have offered here can then be justified on formal semantic grounds rather than informal, intuitive ones.

Unfortunately, it seems that such a formal semantic treatment of nonfunctional requirements would need to be done individually for different types of requirements and is therefore a long-term research project. In the meantime, an experimental approach such as the one adopted here can offer solutions that may find immediate use in an area of computer practice that is in great need of concepts, methodologies, and tools.

ACKNOWLEDGMENT

We would like to thank the referees for their constructive and detailed comments, as well as E. Yu and S. McIlraith for providing helpful suggestions.

REFERENCES

- [1] *Artif. Intell. J.*, vol. 24, Dec. 1984.
- [2] R. Balzer, "A 15 year perspective on automatic programming," *IEEE Trans. Software Eng.*, vol. SE-11, pp. 1257–1268, Nov. 1985.
- [3] V. R. Basili and J. D. Musa, "The future engineering of software: A management perspective," *IEEE Computer*, vol. 24, pp. 90–96, Sept.

- 1991.
- [4] V. Benzaken, "An evaluation model for clustering strategies in the O_2 object-oriented database system," in *Proc. 3rd Int. Conf. Database Theory*, pp. 126-140, 1990.
 - [5] B. W. Boehm *et al.*, *Characteristics of Software Quality*. Amsterdam: North-Holland, 1978.
 - [6] A. Borgida *et al.*, "Support for data-intensive applications: Conceptual design and software development," in *Proc. 2nd Int. Workshop on Database Programming Languages*, pp. 258-280, 1990.
 - [7] T. P. Bowen *et al.*, "Specification of software quality attributes," *Rep. RADC-TR-85-37*, Rome Air Development Center, Griffiss Air Force Base, NY, Feb. 1985.
 - [8] S. Ceri and J. Widom, "Deriving production rules for constraint management," in *Proc. 16th Int. Conf. Very Large Data Bases*, pp. 566-577, Aug. 1990.
 - [9] A. Chan *et al.*, "Storage and access structures to support a semantic data model," in *Proc. 8th Int. Conf. Very Large Data Bases*, pp. 122-130, Sept. 1982.
 - [10] P. P.-S. Chen, "The entity-relationship model—toward a unified view of data," *ACM Trans. Database Systems*, vol. 1, pp. 9-36, Mar. 1976.
 - [11] K. L. Chung *et al.*, "Process management and assertion enforcement for a semantic data model," in *Proc. EDBT '88, Int. Conf. Extending Database Technology*, pp. 469-487, Mar. 1988.
 - [12] L. Chung, "Representation and utilization of nonfunctional requirements for information system design," in *Proc. CAISE '91* pp. 5-30, 1991.
 - [13] K. L. Chung, "From information system requirements to designs: A mapping framework," *Information Systems*, vol. 16, pp. 429-461, 1991.
 - [14] J. Conklin and M. L. Begeman, "gIBIS: A hypertext tool for explanatory policy discussions," *ACM Trans. Office Information Systems*, vol. 6, pp. 303-331, 1988.
 - [15] J. de Kleer, "Problem solving with the ATMS," *Artif. Intell. J.*, vol. 28, pp. 127-162, 1986.
 - [16] C. DiMarco, "Computational stylistics for natural language translation," Ph.D. dissertation, Dept. Computer Science, Univ. Toronto, 1990.
 - [17] J. Doyle, "A truth maintenance system," *Artif. Intell. J.*, vol. 12, pp. 231-272, 1979.
 - [18] S. F. Fickas, "Automating the transformational development of software," *IEEE Trans. Software Eng.*, vol. SE-11, pp. 1268-1277, 1985.
 - [19] U. Hahn *et al.*, "Teamwork support in a knowledge-based information systems environment," *IEEE Trans. Software Eng.*, vol. 17, pp. 467-482, May 1991.
 - [20] H. R. Hartson and D. K. Hsiao, "Full protection specifications in the semantic model for database protection languages," in *Proc. ACM Annual Conf.*, pp. 90-95, Oct. 1976.
 - [21] J. R. Hauser and D. Clausing, "The house of quality," *Harvard Business Review*, pp. 63-73, May-June 1988.
 - [22] W. F. Hyslop, "Performance prediction of relational database management systems," Ph.D. dissertation, Dept. Computer Science, Univ. Toronto, 1991.
 - [23] M. Jarke *et al.*, "DAIDA: An environment for evolving information systems," *ACM Trans. Information Systems*, vol. 10, Jan. 1992.
 - [24] W. L. Johnson *et al.*, "Representation and presentation of requirements knowledge," USC/Information Sciences Institute, Oct. 1991.
 - [25] E. Kant, "On the efficient synthesis of efficient programs," *Artif. Intell. J.*, vol. 20, pp. 253-305, May 1983.
 - [26] S. E. Keller *et al.*, "Specifying software quality requirements with metrics," in *Tutorial: System and Software Requirements Engineering*, R. H. Thayer and M. Dorfman, Eds. IEEE Computer Society Press, 1990, pp. 145-163.
 - [27] E. D. Lazowska *et al.*, *Quantitative System Performance*. Englewood Cliffs, NJ: Prentice-Hall, 1984.
 - [28] J. Lee, "SIBYL: A qualitative decision management system," in *Artificial Intelligence at MIT: Expanding Frontiers*, vol. 1, P. H. Winston and S. A. Shellard, Eds. Cambridge, MA: The MIT Press, 1990, pp. 105-133.
 - [29] ———, "Extending the Potts and Bruns model for recording design rationale," in *Proc. 13th Int. Conf. Software Eng.*, pp. 114-125, May 1991.
 - [30] J. Martin, *Security, Accuracy, and Privacy in Computer Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
 - [31] J. Mostow, "Towards better models of the design process," *AI Magazine*, vol. 6, pp. 44-57, 1985.
 - [32] J. Mylopoulos *et al.*, "A language facility for designing database-intensive applications," *ACM Trans. Database Systems*, vol. 5, pp. 185-207, June 1980.
 - [33] J. Mylopoulos *et al.*, "Telos: representing knowledge about information systems," *ACM Trans. Information Systems*, vol. 8, pp. 325-362, Oct. 1990.
 - [34] N. Nilsson, *Problem-Solving Methods in Artificial Intelligence*. New York, McGraw-Hill, 1971.
 - [35] B. Nixon *et al.*, "Implementation of a compiler for a semantic data model: Experiences with taxis," in *Proc. ACM SIGMOD 1987 Annual Conf.*, pp. 188-131, Dec. 1987.
 - [36] B. Nixon, "Implementation of information system design Specifications: A performance perspective," in *Database Programming Languages: Bulk Types & Persistent Data*. San Mateo, CA: Morgan Kaufmann, 1992, pp. 149-168.
 - [37] C. P. Pfleeger, *Security in Computing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
 - [38] C. Potts and G. Bruns, "Recording the reasons for design decisions," in *Proc. 10th Int. Conf. Software Eng.*, pp. 418-427, 1988.
 - [39] H. Reubenstein, "Automated acquisition of evolving informal descriptions," Ph.D. dissertation; also *Tech. Rep. 1205*, MIT Artif. Intell. Lab., 1990.
 - [40] W. N. Robinson, "Negotiation behavior during requirement specification," in *Proc. 12th Int. Conf. Software Eng.*, pp. 268-276, Mar. 1990.
 - [41] G.-C. Roman, "A taxonomy of current issues in requirements engineering," *IEEE Computer*, vol. 18, pp. 14-23, Apr. 1985.
 - [42] H. A. Simon, *The Sciences of the Artificial, 2nd ed.* Cambridge, MA: MIT Press, 1981.
 - [43] C. U. Smith, *Performance Engineering of Software Systems*. Reading, MA: Addison-Wesley, 1990.
 - [44] M. Stonebraker, "Triggers and inference in database systems," in *On Knowledge Base Management Systems*, M. L. Brodie and J. Mylopoulos, Eds. New York: Springer-Verlag, 1986, pp. 297-314.
 - [45] R. H. Thayer and M. C. Thayer, "Glossary," in *Tutorial: System and Software Requirements Engineering*, Richard H. Thayer and Merlin Dorfman, Eds. IEEE Computer Society Press, 1990, pp. 605-676.
 - [46] G. E. Weddell, "Selection of indexes to memory-resident entities for semantic data models," *IEEE Trans. Knowledge and Data Eng.*, vol. 1, pp. 274-284, June 1989.



John Mylopoulos received the Ph.D. degree from Princeton University, Princeton, NJ, in 1970.

He is currently a Professor of computer science at the University of Toronto, and is also a principal investigator of a project funded by the Information Technology Research Centre of Ontario, the aim of which is the design of a knowledge-based management system. He also leads a project funded by Canada's Networks of Excellence Programme. His research interests in the past included the design of Taxis and Telos; his current research interests include knowledge representation systems, knowledge-based systems, and their applications in building information systems. He has published approximately 100 refereed journal and conference proceedings and has edited three books; he is also a member of several editorial boards.



Lawrence Chung received the B.Sc. and M. Sc. degrees in computer science from the University of Toronto in 1981 and 1984, respectively, and is currently a Ph.D. student there.

He has participated in the Taxis implementation project, and his interests include implementation of semantic data models, development of information systems, application of artificial intelligence to software engineering, and representation of non-functional requirements.



Brian Nixon received the B.Sc. and M.Sc. degrees in computer science from the University of Toronto in 1980 and 1983, respectively, and is now a Ph.D. student there.

He has participated in the Taxis implementation project, and his interests include the implementation of semantic data models and performance of information systems.