Intent Specifications: An Approach to Building Human-Centered Specifications

Nancy G. Leveson, *Member*, *IEEE*

Abstract—This paper examines and proposes an approach to writing software specifications, based on research in systems theory, cognitive psychology, and human-machine interaction. The goal is to provide specifications that support human problem solving and the tasks that humans must perform in software development and evolution. A type of specification, called *intent specifications*, is constructed upon this underlying foundation.

Index Terms—Requirements, requirements specification, safety-critical software, software evolution, human-centered specifications, means-ends hierarchy, cognitive engineering.

1 The Problem

Software is a human product and specification languages are used to help humans perform the various problem-solving activities involved in requirements analysis, software design, review for correctness (verification and validation), debugging, maintenance and evolution, and reengineering. This paper describes an approach, called intent specifications, to designing system and software specifications that potentially enhances human processing and use by grounding specification design on psychological principles of how humans use specifications to solve problems, as well as on basic system engineering principles. Using such an approach allows us to design specification languages with some confidence that they will be usable and effective.

A second goal of intent specifications is to integrate formal and informal aspects of software development and enhance their interaction. While mathematical techniques are useful in some parts of the development process and are crucial in developing software for critical systems, informal techniques will always be a large part (if not most) of any complex software development effort: Our models have limits in that the actual system has properties beyond the model, and mathematical methods cannot handle all aspects of system development. To be used widely in industry, our approach to specification must be driven by the need 1) to systematically and realistically balance and integrate mathematical and nonmathematical aspects of software development and 2) to make the formal parts of the specification easily readable, understandable, and usable by everyone involved in the development and maintenance process.

Specifications should also enhance our ability to engineer for quality and to build evolvable and changeable systems. Essential system-level properties (such as safety and security) must be built into the design from the beginning; they cannot be added on or simply measured afterward. Up-front planning and changes to the development process are needed to achieve particular objectives. These changes include using notations and techniques for reasoning about particular properties, constructing the system and the software in it to achieve them, and validating (at each step, starting from the very beginning of system development) that the evolving system has the desired qualities. Our specifications must reflect and support this process. In addition, systems and software are continually changing and evolving; they must be designed to be changeable and the specifications must support evolution without compromising the confidence in the properties that were initially verified.

Many of the ideas in this paper are derived from attempts by cognitive psychologists, engineers, and human factors experts to design and specify human-machine interfaces. The human-machine interface provides a representation of the state of the system that the operator can use to solve problems and perform control, monitoring, and diagnosis tasks. Just as the control panel in a plant is the interface between the operator and the plant, system and software requirements and design specifications are the interface between the system designers and builders or builders and maintainers. The specifications help the designer, builder, tester, debugger, or maintainer understand the system well enough to create a physical form or to find problems in or change the physical form.

The paper is divided into two parts: The first part describes some basic ideas in systems theory and cognitive engineering.[1] The second part describes a type of specification method called *intent specifications*, built upon these basic ideas, that is designed to satisfy the goals listed

• *The author is with the Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Room 33-406, 77 Massachusetts Ave., Cambridge, MA 02139-4307. E-mail: leveson@.mit.edu.*

Manuscript received 13 Aug. 1997; revised 22 July 1998; accepted 17 Nov. 1998.
Recommended for acceptance by H.A. Muller.
For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number 105499.

1. *Cognitive engineering* is a term that has come to denote the combination of ideas from systems engineering, cognitive psychology, and human factors to cope with the challenges of building high-tech systems composed of humans and machines. These challenges have necessitated augmenting traditional human factors approaches to consider the capabilities and limitations of the human element in complex systems.

above, i.e., to enhance human processing and problem solving, to integrate formal and informal aspects of software development, and to enhance our ability to engineer for quality and to build evolvable and changeable systems.

2 Specifications and Human Problem Solving

To be useful to and usable by humans to solve problems, specification language and system design should be based on an understanding of the problem or task that the user is solving. The systems we design and the specifications we use impose demands on humans. We need to understand those demands and how humans use specifications to solve problems if we are to design specifications that reflect reasonable demands and that assist humans in carrying out their tasks.

Not only does the language in which we specify problems have an effect on our problem-solving ability, it also affects the errors we make while solving those problems. Our specification language design needs to reflect what is known about human limitations and capabilities.

A problem-solving activity involves achieving a goal by selecting and using strategies to move from the current state to the goal state. Success depends on selecting an effective strategy or set of strategies and obtaining the information necessary to carry out that strategy successfully. Specifications used in problem-solving tasks are constructed to provide assistance in this process. Cognitive psychology has firmly established that the representation of the problem provided to problem solvers can affect their performance (see Norman [30] for a survey of this research). In fact, Woods claims that there are no neutral representations [48]: The representations available to the problem solver either degrade or support performance. To provide assistance for problem solving, then, requires that we develop a theoretical basis for deciding which representations support effective problem-solving strategies. For example, problem-solving performance can be improved by providing representations that reduce the problem solver's memory load [21] and that display the critical attributes needed to solve the problem in a perceptually salient way [20].

A problem-solving strategy is an abstraction describing one consistent reasoning approach characterized by a particular mental representation and interpretation of observations [31]. Examples of strategies are hypothesis and test, pattern recognition, decision tree search, reasoning by analogy, and topological search.

Some computer science researchers have proposed theories about the mental models and strategies used in program understanding tasks (examples of such models are [4], [22], [32], [38], [40]). Although this approach seems useful, it may turn out to be more difficult than it appears on the surface. Each of the users of a specification may (and probably will) have different mental models of the system, depending on such factors as prior experience, the task for which the model is being used, and their role in the system [1], [13], [28], [36]. The same person may have multiple mental models of a system and even having two contradictory models of the same system does not seem to constitute a problem for people [28].

Strategies also seem to be highly variable. A study that used protocol analysis to determine the trouble-shooting strategies of professional technicians working on electronic equipment found that no two sequences of actions were identical, even though the technicians were performing the same task every time (i.e., finding a faulty electronic component) [34]. Not only do search strategies vary among individuals for the same problem, but a person may vary his or her strategy dynamically during a problem-solving activity: Effective problem solvers change strategies frequently to circumvent local difficulties encountered along the solution path and to respond to new information that changes the objectives and subgoals or the mental workload needed to achieve a particular subgoal.

It appears, therefore, that to allow for multiple users and for effective problem solving (including shifting among strategies), specifications should support all possible strategies that may be needed for a task to allow for multiple users of the representation, for shedding mental workload by shifting strategies during problem solving, and for different cognitive and problem-solving styles. We need to design specifications such that users can easily find or infer the information they need regardless of their mental model or preferred problem-solving strategies. That is, the specification design should be related to the general tasks users need to perform with the information, but not be limited to specific predefined ways of carrying out those tasks.

One reason why many software engineering tools and environments are not readily accepted or easily used is that they imply a particular mental model and force potential users to work through problems using only one or a very limited number of strategies, usually the strategy or strategies preferred by the designer of the tool. The goal of specification language design should be to make it easy for users to extract and focus on the important information for the specific task at hand without assuming particular mental models or limiting the problem-solving strategies employed by the users of the document. The rest of this paper describes an approach to achieve this goal.

3 Components of a Specification Methodology to Support Problem-Solving

Underlying any methodology is an assumed *process*. In our case, the process must support the basic system and software engineering tasks. A choice of an underlying system engineering process is the first component of a specification methodology. In addition, cognitive psychologists suggest that three aspects of interface design must be addressed if the interface is to serve as an effective medium:

1. *content* (what semantic information should be contained in the representation given the goals and tasks of the users,
2. *structure* (how to design the representation so that the user can extract the needed information), and
3. *form* (the notation or format of the interface) [46].

The next sections examine each of these four aspects of specification design in turn.

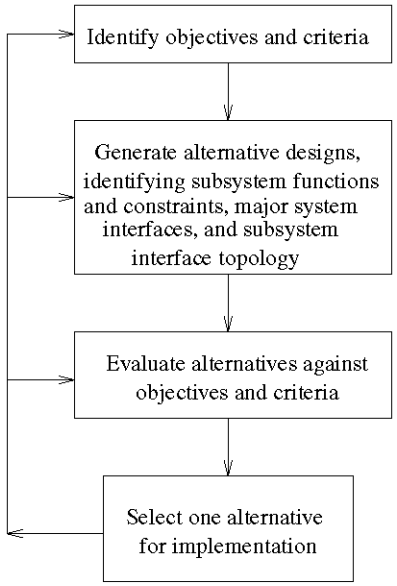

Fig. 1. The basic systems engineering process.

3.1 Process

Any system specification method should support the systems engineering process. This process provides a logical structure for problem solving (see Fig. 1). First, a need or problem is specified in terms of objectives that the system must satisfy and criteria that can be used to rank alternative designs. Then, a process of system synthesis takes place that results in a set of alternative designs. Each of these alternatives is analyzed and evaluated in terms of the stated objectives and design criteria and one alternative is selected to be implemented. In practice, the process is highly iterative: The results from later stages are fed back to early stages to modify objectives, criteria, design alternatives, and so on.

Design alternatives are generated through a process of system architecture development and analysis. The system engineers break down the system into a set of subsystems, together with the functions and constraints imposed upon the individual subsystem designs, the major system interfaces, and the subsystem interface topology. These aspects are analyzed with respect to desired system performance characteristics and constraints, and the process is iterated until an acceptable system design results. The preliminary design at the end of this process must be described in sufficient detail that subsystem implementation can proceed independently.

The software requirements and design process are simply subsets of the larger system engineering process. System engineering views each system as an integrated whole, even though it is composed of diverse, specialized components, which may be physical, logical (software), or human. The objective is to design subsystems that, when integrated into the whole, provide the most effective system possible to achieve the overall objectives. The most challenging problems in building complex systems today arise in the interfaces between components. One example is the new highly automated aircraft where most incidents and accidents have been blamed on human error, but more properly reflect difficulties in the collateral design of the aircraft, the avionics systems, the cockpit displays and controls, and the demands placed on the pilots.

What types of specifications are needed to support humans in this system engineering process and to specify the results? Design decisions at each stage must be mapped into the goals and constraints they are derived to satisfy, with earlier decisions mapped (traced) to later stages of the process, resulting in a seamless (gapless) record of the progression from high-level system requirements down to component requirements and designs. The specifications must also support the various types of formal and informal analysis used to decide between alternative designs and to verify the results of the design process. Finally, they must assist in the coordinated design of the components and the interfaces between them.

3.2 Content

The second component of a specification methodology is the content of the specifications. Determining appropriate *content* requires considering what the specifications will be used for, that is, the problems that humans are trying to solve when they use specifications. Previously, we looked at a narrow slice of this problem—what should be contained in blackbox requirements specifications for process control software to ensure that the resulting implementations are internally complete [19], [24]. This paper again considers the question of specification content, but within a larger context.

This question is critical because cognitive psychologists have determined that people tend to ignore information during problem solving that is not represented in the specification of the problem. In experiments where some problem solvers were given incomplete representations while others were not given any representation at all, those with no representation did better [14], [39]. An incomplete problem representation actually *impaired* performance because the subjects tended to rely on it as a comprehensive and truthful representation—they failed to consider important factors deliberately omitted from the representations. Thus, being provided with an incomplete problem representation (specification) can actually lead to worse performance than having no representation at all [46].

One possible explanation for these results is that some problem solvers did worse because they were unaware of important omitted information. However, both novices and experts failed to use information left out of the diagrams with which they were presented, even though the experts could be expected to be aware of this information. Fischoff et al., who did such an experiment involving fault tree diagrams, attributed it to an "out of sight, out of mind" phenomenon [14].

One place to start in deciding what should be in a system specification is with basic systems theory, which defines a *system* as a set of components that act together as a whole to achieve some common goal, objective, or end. The components are all interrelated and are either directly or indirectly connected to each other. This concept of a system relies on the assumptions that the system goals can be defined and that systems are atomistic, that is, capable of being separated into component entities such that their interactive behavior mechanisms can be described.

The system *state* at any point in time is the set of relevant properties describing the system at that time. The system *environment* is a set of components (and their properties) that are not part of the system but whose behavior can affect the system state. The existence of a boundary between the system and its environment implicitly defines as *inputs* or *outputs* anything that crosses that boundary.

It is very important to understand that a system is always a model—an abstraction *conceived by the analyst*. For the same man-made system, an observer may see a different purpose than the designer and may also focus on different relevant properties. Thus, there may be multiple "correct" system models or specifications. To ensure consistency and enhance communication, a common specification is required that defines the:

- system boundary,
- inputs and outputs,
- components,
- structure,
- relevant interactions between components and the means by which the system retains its integrity (the behavior of the components and their effect on the overall system state), and
- purpose or goals of the system that makes it reasonable to consider it to be a coherent entity [8].

All of these properties need to be included in a complete system model or specification along with a description of the aspects of the environment that can affect the system state. Most of these aspects are already included in our current specification languages. However, the last, information about purpose or intent, is often not.

One of the most important limitations of the models underlying most current specification languages, both formal and informal, is that they cannot allow us to infer what is not explicitly represented in the model, including the intention of doing something a particular way. This intentional information is critical in the design and evolution of software. As Harman has said, practical reasoning is concerned with what to intend, while formal reasoning with what to believe [18].

> Formal logic arguments are a priori true or false with reference to an explicitly defined model, whereas functional reasoning deals with relationships between models, and truth depends on correspondence with the state of affairs in the real world [18].

In the conclusions to our paper describing our experiences specifying the requirements for TCAS II (an aircraft collision avoidance system), we wrote:

> In reverse engineering TCAS, we found it impossible to derive the requirements specification strictly from the pseudocode and an accompanying English language description. Although the basic information was all there, the intent was largely missing and often the mapping from goals or constraints to specific design decisions. Therefore, distinguishing between requirements and artifacts of the implementation was not possible in all cases. As has been discovered by most people attempting to maintain such systems, an audit trail of the decisions and the reasons why decisions were made is absolutely essential. This was not done by TCAS over the 15 years of its development, and those responsible for the system today are currently attempting to reconstruct decision-making information from old memos and corporate memory. For the most part, only one person is able to explain why some decisions were made or why things were designed in a particular way [26].

There is widespread agreement about the need for design rationale (intent) information in order to understand complex software or to correctly and efficiently change or analyze the impact of changes to it. Without a record of intent, important decisions can be undone during maintenance: Many serious accidents and losses can be traced to the fact that a system did not operate as intended because of changes that were not fully coordinated or fully analyzed to determine their effects [24]. What is not so clear is the content and structure of the information that is needed.

Simply keeping an audit trail of decisions and the reasons behind them as they are made is not practical. The number of decisions made in any large project is enormous. Even if it were possible to write them all down, finding the proper information when needed seems to be a hopeless task if not structured appropriately. What is needed is a specification of the intent (goals, constraints, and design rationale) from the beginning, and it must be specified in a usable and perceptually salient manner. That is, we need a framework within which to select and specify the design decisions that are needed to develop and maintain software.

3.3 Structure

The third aspect of specifications, *structure*, is the basis for organizing information in the specification. The information may all be included somewhere, but it may be hard to find or to determine the relationship to information specified elsewhere.

Problem solving in technological systems takes place within the context of a complex causal network of relationships [12], [34], [35], [46], and those relationships need to be reflected in the specification. The information needed to solve a problem may all be included somewhere in the assorted documentation used in large projects, but it may be hard to find when needed or to determine the relationship to information specified elsewhere. Psychological experiments in problem solving find that people attend primarily to perceptually salient information [20]. The goal of specification language design should be to make it easy for users to extract and focus on the important information for the specific task at hand, which includes all potential tasks related to use of the specification.

Cognitive engineers speak of this problem as "information pickup" [48]. Just because the information is in the interface does not mean that the operator can find it easily. The same is true for specifications. The problem of information pickup is compounded by the fact that there is so much information in system and software specifications, while only a small subset of it may be relevant in any given context.

3.3.1 Complexity

The problems in building and interacting with systems correctly are rooted in complexity and intellectual manageability. A basic and often noted principle of engineering is to keep things simple. This principle, of course, is easier to state than to do. Ashby's Law of Requisite Variety [2] tells

us that there is a limit to how simple we can make control systems, including those designs represented in software, and still have them be effective. In addition, basic human ability is not changing. If humans want to build and operate increasingly complex systems, we need to increase what is intellectually manageable. That is, we will need to find ways to *augment* human ability.

The situation is not hopeless. As Rasmussen observes, the complexity of a system is not an objective feature of the system [33]. Observed complexity depends upon the level of resolution upon which the system is being considered. A simple object becomes complex if observed through a microscope. Complexity, therefore, can only be defined with reference to a particular representation of a system, and then can only be measured relative to other systems observed at the same level of abstraction.

Thus, a way to cope with complex systems is to structure the situation such that the observer can transfer the problem being solved to a level of abstraction with less resolution. The complexity faced by the builders or users of a system is determined by their *mental models* (representations) of the internal state of the system. We build such mental models and update them based on what we observe about the system, that is, by means of our interface to the system. Therefore, the apparent complexity of a system ultimately depends upon the technology of the interface system [33].

The solution to the complexity problem is to take advantage of the most powerful resources people have for dealing with complexity. Newman has noted,

> People don't mind dealing with complexity if they have some way of controlling or handling it... If a person is allowed to structure a complex situation according to his perceptual and conceptual needs, sheer complexity is no bar to effective performance [29], [33].

Thus, complexity itself is not a problem if humans are presented with meaningful information in a coherent, structured context.

3.3.2 Hierarchy Theory

Two ways humans cope with complexity is to use top-down reasoning and stratified hierarchies. Building systems bottom-up works for relatively simple systems. But, as the number of cases and objects that must be considered increases, this approach becomes unworkable—we go beyond the limits of human memory and logical ability to cope with the complexity. Top-down reasoning is a way of managing that complexity. At the same time, we have found that pure top-down reasoning is not adequate alone; humans need to combine top-down with bottom-up reasoning. Thus, the structure of the information must allow reasoning in both directions.

In addition, humans cope with complexity by building stratified hierarchies. Models of complex systems can be expressed in terms of a *hierarchy* of levels of organization, each more complex than the one below, where a level is characterized by having *emergent* properties. The concept of emergence is the idea that, at any given level of complexity, some properties characteristic of that level (emergent at that level) are irreducible. Such properties do not exist at lower levels in the sense that they are meaningless in the language appropriate to those levels. For example, the shape of an apple, although eventually explainable in terms of the cells of the apple, has no meaning at that lower level of description.

Regulatory or *control* action involves imposing constraints upon the activity at one level of a hierarchy. Those constraints define the "laws of behavior" at that level that yield activity meaningful at a higher level (emergent behavior). Hierarchies are characterized by control processes operating at the interfaces between levels. Checkland explains it:

> Any description of a control process entails an upper level imposing constraints upon the lower. The upper level is a source of an alternative (simpler) description of the lower level in terms of specific functions that are emergent as a result of the imposition of constraints [8 p, 87].

Hierarchy theory deals with the fundamental differences between one level of complexity and another. Its ultimate aim is to explain the relationships between different levels: What generates the levels, what separates them, and what links them. Emergent properties associated with a set of components at one level in a hierarchy are related to constraints upon the degree of freedom of those components. In the context of this paper, it is important to note that describing the emergent properties resulting from the imposition of constraints requires a language at a higher level (a metalevel) *different* than that describing the components themselves. Thus, different description languages are required at each hierarchical level.

The problem then comes down to determining appropriate types of hierarchical abstraction that allow both top-down and bottom-up reasoning. In computer science, we have made much use of part-whole abstractions where each level of a hierarchy represents an aggregation of the components at a lower level and of information-hiding abstractions where each level contains the same conceptual information but hides some details about the concepts, that is, each level is a refinement of the information at a higher level. Each level of our software specifications can be thought of as providing *what* information, while the next lower level describes *how*.

Such hierarchies, however, do not provide information about *why*. Higher-level emergent information about purpose or intent cannot be inferred from what we normally include in such specifications. Design errors may result when we either guess incorrectly about higher-level intent or omit it from our decision-making process. For example, while specifying the system requirements for TCAS II [26], we learned from experts that crossing maneuvers are avoided in the design for safety reasons. The analysis on which this decision is based comes partly from experience during TCAS system testing on real aircraft and partly as a result of an extensive safety analysis performed on the system. This design constraint would not be apparent in most design or code specifications unless it were added in the form of comments, and it could easily be violated during system modification unless it was recorded and easily located.

But, there are abstractions that can be used in stratified hierarchies other than part-whole abstraction. While investigating the design of safe human-machine interaction,

Rasmussen studied protocols recorded by people working on complex systems (process plant operators and computer maintainers) and found that they structured the system along two dimensions: 1) a part-whole abstraction in which the system is viewed as a group of related components at several levels of physical aggregation and 2) a means-ends abstraction [34].

3.3.3 Means-Ends Hierarchies

In a means-end abstraction, each level represents a different model of the same system. At any point in the hierarchy, the information at one level acts as the goals (the ends) with respect to the model at the next lower level (the means). Thus, in a means-ends abstraction, the current level specifies *what*, the level below *how*, and the level above *why* [34]. In essence, this intent information is emergent in the sense of system theory:

> When moving from one level to the next higher level, the change in system properties represented is not merely removal of details of information on the physical or material properties. More fundamentally, information is added on higher-level principles governing the coordination of the various functions or elements at the lower level. In man-made systems, these higher-level principles are naturally derived from the purpose of the system, i.e., from the reasons for the configurations at the level considered [34].

A change of level involves both a shift in concepts and in the representation structure, as well as a change in the information suitable to characterize the state of the function or operation at the various levels [34].

Each level in a means-ends hierarchy describes the system in terms of a different set of attributes or "language." Models at the lower levels are related to a specific physical implementation that can serve several purposes, while those at higher levels are related to a specific purpose that can be realized by several physical implementations. Changes in goals will propagate downward through the levels, while changes in the physical resources (such as faults or failures) will propagate upward. In other words, states can only be described as errors or faults with reference to their intended functional purpose. Thus, reasons for proper function are derived "top-down." In contrast, causes of improper function depend upon changes in the physical world (i.e., the implementation) and, thus, they are explained "bottom up" [46].

Mappings between levels are many-to-many: Components of the lower levels can serve several purposes, while purposes at a higher level may be realized using several components of the lower-level model. These goal-oriented links between levels can be followed in either direction, reflecting either the means by which a function or goal can be accomplished (a link to the level below) or the goals or functions an object can affect (a link to the level above). So, the means-ends hierarchy can be traversed in either a top-down (from ends to means) or a bottom-up (from means to ends) direction.

As stated earlier, our representations of problems have an important effect on our problem-solving ability and the strategies we use, and there is good reason to believe that representing the problem space as a means-ends mapping provides useful context and support for decision making and problem solving. Consideration of purpose or reason (top-down analysis in a means-ends hierarchy) has been shown to play a major role in understanding the operation of complex systems [33].

Rubin's analysis of his attempts to understand the function of a camera's shutter (as cited in [35]) provides an example of the role of intent or purpose in understanding a system. Rubin describes his mental efforts in terms of conceiving all the elements of the shutter in terms of their function in the whole rather than explaining how the individual parts worked: How they worked was immediately clear when their function was known. Rasmussen argues that this approach has the advantage that solutions of subproblems are identifiable with respect to their place in the whole picture and it is immediately possible to judge whether a solution is correct or not. In contrast, arguing from the parts to the way they work is much more difficult because it requires synthesis: Solutions of subproblems must be remembered in isolation and their correctness is not immediately apparent.

Support for this argument can be found in the difficulties AI researchers have encountered when modeling the function of mechanical devices "bottom-up" from the function of the components. DeKleer and Brown found that determining the function of an electric buzzer solely from the structure and behavior of the parts requires complex reasoning [10]. Rasmussen suggests that the resulting inference process is very artificial compared to the top-down inference process guided by functional considerations as described by Ruben.

> In the DeKleer-Brown model, it will be difficult to see the woods for the trees, while Rubin's description appears to be guided by a birds-eye perspective [35].

Glaser and Chi suggest that experts and successful problem solvers tend to focus first on analyzing the functional structure of the problem at a high level of abstraction and then narrow their search for a solution by focusing on more concrete details [16]. Representations that constrain search in a way that is explicitly related to the purpose or intent for which the system is designed have been shown to be more effective than those that do not because they facilitate the type of goal-directed behavior exhibited by experts [44]. Therefore, we should be able to improve the problem solving required in software development and evolution tasks by providing a representation (i.e., specification) of the system that facilitates goal-oriented search by making explicit the goals related to each component.

Viewing a system from a high level of abstraction is not limited to a means-ends hierarchy, of course. Most hierarchies allow one to observe systems at a less detailed level. The difference is that the means-ends hierarchy is explicitly *goal oriented* and, thus, assists goal-oriented problem solving. With other hierarchies (such as the part-whole hierarchies often used in computer science), the links between levels are not necessarily related to goals. So, although it is possible to use higher-levels of abstraction to select a subsystem of interest and to constrain search, the subtree of the hierarchy connected to a particular subsystem does not necessarily contain system components relevant to the goals the problem solver is considering.

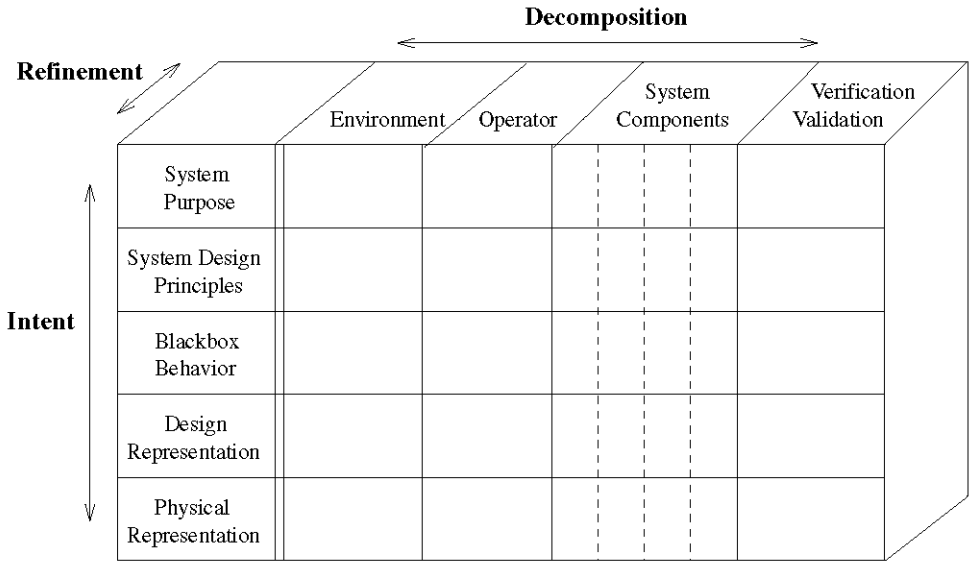

Fig. 2. The structure of an intent specification for software systems.

3.4 Form (Notation)

The final aspect of specification design is the actual form of the specification. Although this is often where we start when designing languages, the four aspects actually have to be examined in order, first defining the process to be supported, then determining what the content should be, then how the content will be structured to make the information easily located and used, and, finally, the form the language should take. All four aspects need to be addressed, not only in terms of the analysis to be performed on the specification, but also with respect to human perceptual and cognitive capabilities.

Note that the form itself must also be considered from a psychological standpoint: The usability of the language will depend on human perceptual and cognitive strategies. For example, Fitter and Green describe the attributes of a good notation with respect to human perception and understanding [15]. Casner [7] and others have argued that the utility of any information presentation is a function of the *task* that the presentation is being used to support. For example, a symbolic representation might be better than a graphic for a particular task, but worse for others.

No particular specification language is being proposed here. We first must clarify what needs to be expressed before we can design languages that express that information appropriately and effectively. In addition, different types of systems require different types of languages. All specifications are abstractions—they leave out unimportant details. What is important will depend on the problem being solved. For different types of systems, the important and difficult aspects differ. For example, specifications for embedded controllers may emphasize control flow over data flow (which is relatively trivial for these systems), while data transformation or information management systems might place more emphasis on the specification of data flow than control flow. Attempts to include everything in the specification are not only impractical, but involve wasted effort and are unlikely to fit the budgets and schedules of industry projects. Because of the importance of completeness, as argued earlier, determining exactly what needs to be included becomes the most important problem in specification design.

This paper deals with process, content and structure, but not form (notation). We are defining specification languages built upon the foundation laid in this paper and on other psychological principles, but they will be described in future papers.

4 Intent Specifications

These basic ideas provide the foundation for what can be called *intent specifications*. They have been developed and successfully used in cognitive engineering by Vicente and Dinadis for the design of operator interfaces, a process they call *ecological interface design* [11], [43].

The exact number and content of the means-ends hierarchy levels may differ from domain to domain. Here, a structure is presented for process systems with shared software and human control. In order to determine the feasibility and scalability of this approach when specifying a complex system, we extended the formal TCAS II aircraft collision avoidance system requirements specification previously written [26] to include intent information and other information that cannot be expressed formally, but is needed in a complete system requirements specification. We are currently applying the approach to other examples, including a NASA robot and part of the U.S. Air Traffic Control System. The TCAS II specification is used as an example in this paper.[2] The table of contents for the example TCAS II System Requirements Specification (shown in Fig. 3) may be helpful in following the description of intent specifications. Note that the only part of TCAS that we specified previously is Section 3.4 and parts of Section 3.3.

2. Our TCAS II Intent Specification (complete system specification) is over 800 pages long. Obviously, the entire specification cannot be included in this paper. It can be accessed from http://sunnyday.mit.edu.

In the intent specifications we have built for real systems, we have found the approach to be practical; in fact, most of the information in an intent specification is already located somewhere in the often voluminous documentation for large systems. The problem in these systems usually lies in finding specific information when it is needed, in tracing the relationships between information, and in understanding the system design and why it was designed that way. Intent specifications are meant to assist with these tasks.

System and software specifications of the type being proposed (see Fig. 2) are organized along a vertical dimension using intent abstraction and two horizontal dimensions using two types of part-whole abstraction (refinement and decomposition). These three dimensions constitute the problem space in which the human navigates. The horizontal refinement and decomposition dimensions allow users to change their focus of attention to more or less detailed views within each intent level or model. The vertical dimension (based on intent abstraction) specifies the level of intent at which the problem is being considered, i.e., the language or model that is currently being used.

4.1 Part-Whole Dimension

Computer science commonly uses two types of part-whole abstractions. *Parallel decomposition* (or its opposite, aggregation) separates units into (perhaps interacting) components of the same type. In Statecharts, for example, these components are called orthogonal components and the process of aggregation results in an orthogonal product. Each of the pieces of the parallel decomposition of Statecharts is a state machine, although each state machine will in general be different.

The second type of part-whole abstraction—*refinement*—takes a function and breaks it down into more detailed steps. An example is the combining of a set of states into a superstate in Statecharts. In Petri-nets, such abstractions have been applied both to states and to transitions—they provide a higher-level name for a piece of the net. In programming, refinement abstractions are represented by procedures or subprograms.

Note that neither of these types of abstraction is an "emergent-property" or means-ends abstraction—the whole is simply broken up into a more detailed description. Additional information, such as intent, is not provided at the higher level.

Along these horizontal dimensions, intent specifications are broken up into four parts. The first part (the first column in Fig. 2) contains information about characteristics of the environment that affects the ability to achieve the system goals and design constraints. For example, in TCAS, the designers need information about the operation of the ground-based ATC system in order to fulfill the system-level constraint of not interfering with it. Information about the environment is also needed for some types of hazard analysis and for normal system design. For example, the design of the surveillance logic in TCAS depends on the characteristics of the transponders carried on the aircraft with which the surveillance logic interacts.

The second column of the horizontal dimension is information about human operators or users. Too often human factors design and software design is done independently. Many accidents and incidents in aircraft with advanced automation have been blamed on human-error that has been induced by the design of the automation. For example, Weiner introduced the term *clumsy automation* to describe automation that places additional and unevenly distributed workload, communication, and coordination demands on pilots without adequate support [47]). Sarter et al. [37] describe additional problems associated with new attentional and knowledge demands and breakdowns in mode awareness and "automation surprises," which they attribute to *technology-centered automation*: Too often, the designers of the automation focus exclusively on technical aspects, such as the mapping from software inputs to outputs, on mathematical models of requirements functionality, and on the technical details and problems internal to the computer; they do not devote enough attention to the cognitive and other demands of the automation design on the operator.

One goal of intent specifications is to integrate the information needed to design "human-centered automation" into the system requirements specification. We are also working on analysis techniques to identify problematic system and software design features in order to predict where human errors are likely to occur [27]. This information can be used in both the automation design and in the design of the operator procedures, tasks, interface, and training.

The third part of the horizontal dimension is the system itself and its decomposition into subsystems or components. Finally, each level also includes information about the verification and validation activities and results appropriate for that specification level.

4.2 Intent Dimension

The Intent (vertical) dimension has five hierarchical levels, each providing intent ("why") information about the level below. Each level is mapped to the appropriate parts of the intent levels above and below it, providing *traceability* of high-level system requirements and constraints down to code (or physical form) and vice versa.

Each level also supports a different type of reasoning about the system, with the highest level assisting systems engineers in their reasoning about system-level goals, constraints, priorities, and trade-offs. The second level, System Design Principles, allows engineers to reason about the system in terms of the physical principles and laws upon which the design is based. The Blackbox Behavior level enhances reasoning about the logical design of the system as a whole and the interactions between the components, as well as the functional state without being distracted by implementation issues. The lowest two levels provide the information necessary to reason about individual component design and implementation issues. The mappings between levels provide the relational information that allows reasoning across hierarchical levels.

4.2.1 System Purpose

Along the vertical dimension, the highest specification level, *System Purpose*, contains:

1. System Purpose

- 1.1 Introduction
- 1.2 Historical Perspective
- 1.3 Environment
  - 1.3.1 Environmental Assumptions
  - 1.3.2 Environmental Constraints
- 1.4 Operator
  - 1.4.1 Tasks and Procedures
  - 1.4.2 Pilot-TCAS Interface Requirements
- 1.5 TCAS System Goals
- 1.6 High-Level Functional Requirements
- 1.7 System Limitations
- 1.8 System Constraints
  - 1.8.1 General Constraints
  - 1.8.2 Safety-Related Constraints
- 1.9 Hazard Analysis

2. System Design Principles

- 2.1 General Description
- 2.2 TCAS System Components
- 2.3 Surveillance and Collision Avoidance Logic
  - 2.3.1 General Concepts
  - 2.3.2 Surveillance
  - 2.3.3 Tracking
  - 2.3.4 Traffic Advisories
  - 2.3.5 Resolution Advisories
  - 2.3.6 TCAS/TCAS Coordination
- 2.4 Performance Monitoring
- 2.5 Pilot-TCAS Interface
  - 2.5.1 Controls
  - 2.5.2 Displays and Aural Annunciations
- 2.6 Testing and Validation
  - 2.6.1 Simulations
  - 2.6.2 Experiments
  - 2.6.3 Other Validation Procedures and Results

3. Blackbox Behavior

- 3.1 Environment
- 3.2 Flight Crew Requirements
  - 3.2.1 Tasks
  - 3.2.2 Operational Procedures
- 3.3 Communication and Interfaces
  - 3.3.1 Pilot-TCAS Interface
  - 3.3.2 Message Formats
  - 3.3.3 Input Interfaces
  - 3.3.4 Output Interfaces
  - 3.3.5 Receiver, Transmitter, Antennas
- 3.4 Behavioral Requirements
  - 3.4.1 Surveillance
  - 3.4.2 Collision Avoidance
  - 3.4.3 Performance Monitoring
- 3.5 Testing Requirements

4. Physical and Logical Function

- 4.1 Human-Computer Interface Design
- 4.2 Pilot Operations (Flight) Manual
- 4.3 Software Design
- 4.4 Physical Requirements
  - 4.4.1 Definition of Standard Conditions
  - 4.4.2 Performance Capability of Own Aircraft's Mode S Transponder
  - 4.4.3 Receiver Characteristics
  - 4.4.4 TCAS Transmitter Characteristics
  - 4.4.5 TCAS Transmitter Pulse Characteristics
  - 4.4.6 TCAS Pulse Decoder Characteristics
  - 4.4.7 Interfence Limiting
  - 4.4.8 Aircraft Suppression Bus
  - 4.4.9 TCAS Data Handling and Interfaces
  - 4.4.10 Bearing Estimation
  - 4.4.11 High-Density Techniques
- 4.5 Hardware Design Specifications
- 4.6 Verification Requirements

5. Physical Realization

- 5.1 Software
- 5.2 Hardware Assembly Instructions
- 5.3 Training Requirements (Plan)
- 5.4 Maintenance Requirements

A. Constant Definitions

B. Table Definitions

C. Reference Algorithms

D. Physical Measurement Conventions

E. Performance Requirements on Equipment that Interacts with TCAS

F. Glossary

G. Notation Guide

H. Index

Fig. 3. The contents of the sample TCAS Intent Specification.

- system goals,
- design constraints,
- assumptions,
- limitations,
- design evaluation criteria and priorities, and
- results of analyses for system level qualities.

Examples of high-level *goals* (purpose) for TCAS II are to:

G1. *Provide affordable and compatible collision avoidance system options for a broad spectrum of National Airspace System users.*
G2. *Detect potential midair collisions with other aircraft in all meteorological conditions.*

Usually, in the early stages of a project, goals are stated in very general terms. One of the first steps in defining system requirements is to refine the goals into testable and achievable high-level requirements. For G1 above, a refined subgoal is:

R1. *Provide collision avoidance protection for any two aircraft closing horizontally at any rate up to 1,200 knots and vertically up to 10,000 feet per minute.*

This type of refinement and reasoning is done at the System Purpose level, using an appropriate specification language (most likely English).

Requirements (and constraints) are also included for the human operator, for the human-computer interface, and for the environment in which TCAS will operate. Requirements on the operator (in this case, the pilot) are used to guide the design of the TCAS-pilot interface, flightcrew tasks and procedures, aircraft flight manuals, and training plans and program. Links are provided to show the relationships. Example TCAS II operator requirements are:

O1. *After the threat is resolved, the pilot shall return promptly and smoothly to his/her previously assigned flight path.*
O2. *The pilot must not maneuver on the basis of a Traffic Advisory only.*

Design constraints are restrictions on how the system can achieve its purpose. For example, TCAS is not allowed to interfere with the ground-level air traffic control system while it is trying to maintain adequate separation between aircraft. Avoiding interference is not a goal or purpose of TCAS—the best way to achieve it is not to build the system at all. It is, instead, a constraint on how the system can achieve its purpose, i.e., a constraint on the potential system designs. Because of the need to evaluate and clarify trade-offs among alternative designs, separating these two types of intent information (goals and design constraints) is important.

For safety-critical systems, constraints should be further separated into normal and safety-related. Examples of *nonsafety constraints* for TCAS II are:

C1. *The system must use the transponders routinely carried by aircraft for ground ATC purposes.*
C2. *No deviations from current FAA policies and philosophies must be required.*

Safety-related constraints should have two-way links to the system hazard log and, perhaps, links to any analysis results that led to that constraint being identified. Hazard analyses specified on this level are linked to Level 1 requirements and constraints on this level, to design features on Level 2, and to system limitations (or accepted risks). Example safety constraints are:

SC1.
The system must generate advisories that require as little deviation as possible from ATC clearances.
SC2.
The system must not disrupt the pilot and ATC operations during critical phases of flight.

Note that *refinement* occurs at the same level of the intent specification (see Fig. 2). For example, the safety-constraint SC3 can be refined

SC3.
The system must not interfere with the ground ATC system or other aircraft transmissions to the ground ATC system.

SC3.1.
The system design must limit interference with ground-based secondary surveillance radar, distance-measuring equipment channels, and with other radio services that operate in the 1030/1090 MHz frequency band.

SC3.1.1.
The design of the Mode S waveforms used by TCAS must provide compatibility with Modes A and C of the ground-based secondary surveillance radar system.
SC3.1.1.
The frequency spectrum of Mode S transmissions must be controlled to protect adjacent distance-measuring equipment channels.
SC3.1.1.
The design must ensure electromagnetic compatibility between TCAS and...

SC3.2.
Multiple TCAS units within detection range of one another (approximately 30 nmi) must be designed to limit their own transmissions. As the number of such TCAS units within this region increases, the interrogation rate and power allocation for each of them must decrease in order to prevent undesired interference with ATC.

Environment requirements and constraints may lead to restrictions on the use of the system or to the need for system safety and other analyses to determine that the requirements hold for the larger system in which the system being designed is to be used. Examples for TCAS include:

E1. *Among the aircraft environmental alerts, the hierarchy shall be: Windshear has first priority, then the Ground Proximity Warning System (GPWS), then TCAS.*
E2. *The behavior or interaction of non-TCAS equipment with TCAS must not degrade the performance of the TCAS equipment or the performance of the equipment with which TCAS interacts.*
E3. *The TCAS alerts and advisories must be independent of those using the master caution and warning system.*

Assumptions are specified, when appropriate, at all levels of the intent specification to explain a decision or to record fundamental information on which the design is based.

These assumptions are often used in the safety or other analyses or in making lower level design decisions. For example, operational safety depends on the accuracy of the assumptions and models underlying the design and hazard analysis processes. The operational system should be monitored to ensure:

1. that it is constructed, operated, and maintained in the manner assumed by the designers,
2. that the models and assumptions used during initial decision making and design were correct, and
3. that the models and assumptions are not violated by changes in the system, such as workarounds or unauthorized changes in procedures, or by changes in the environment [24].

Operational feedback on trends, incidents, and accidents should trigger re-analysis when appropriate. Linking the assumptions throughout the document with the hazard analysis (for example, to particular boxes in the system fault trees) will assist in performing safety maintenance activities.

Examples of assumptions associated with requirements on the first level of the TCAS intent specification:

R1. *Provide collision avoidance protection for any two aircraft closing horizontally at any rate up to 1,200 knots and vertically up to 10,000 feet per minute.*

Assumption.
This requirement is derived from the assumption that commercial aircraft can operate up to 600 knots and 5,000 fpm during vertical climb or controlled descent (and, therefore, two planes can close horizontally up to 1,200 knots and vertically up to 10,000 fpm).

R3. *TCAS shall operate in enroute and terminal areas with traffic densities up to 0.3 aircraft per square nautical miles (i.e., 24 aircraft within 5 nmi).*

Assumption.
Traffic density may increase to this level by 1990, and this will be the maximum density over the next 20 years.

An example of an assumption associated with a safety constraint is:

SC5.
The system must not disrupt the pilot and ATC operations during critical phases of flight nor disrupt aircraft operation.

SC5.1.
The pilot of a TCAS-equipped aircraft must have the option to switch to the Traffic-Advisory-Only mode, where TAs are displayed but display of resolution advisories is inhibited.

Assumption.
This feature will be used during final approach to parallel runways, when two aircraft are projected to come close to each other and TCAS would call for an evasive maneuver.

Assumptions may also apply to features of the environment. Examples of environment assumptions for TCAS are that:

EA1.
All aircraft have legal identification numbers.
EA2.
All aircraft carry transponders.
EA3.
The TCAS-equipped aircraft carries a Mode-S air traffic control transponder whose replies include encoded altitude when appropriately interrogated.
EA4.
Altitude information is available from intruding targets with a minimum precision of 100 feet.
EA5.
Threat aircraft will not make an abrupt maneuver that thwarts the TCAS escape maneuver.

System *limitations* are also specified at Level 1 of an intent specification. Some may be related to the basic functional requirements, such as:

L1. *TCAS does not currently indicate horizontal escape maneuvers and, therefore, does not (and is not intended to) increase horizontal separation.*

Limitations may also relate to environment assumptions. For example, system limitations related to the environment assumptions above include:

L2. *TCAS provides no protection against aircraft with nonoperational transponders.*
L3. *Aircraft performance limitations constrain the magnitude of the escape maneuver that the flight crew can safely execute in response to a resolution advisory. It is possible for these limitations to preclude a successful resolution of the conflict.*
L4. *TCAS is dependent on the accuracy of the threat aircraft's reported altitude. Separation assurance may be degraded by errors in intruder pressure altitude as reported by the transponder of the intruder aircraft.*

Assumption.
This limitation holds for existing airspace, where many aircraft use pressure altimeters rather than GPS. As more aircraft install GPS systems with greater accuracy than current pressure altimeters, this limitation will be reduced or eliminated.

Limitations are often associated with hazards or hazard causal factors that could not be completely eliminated or controlled in the design. Thus, they represent accepted risks. For example:

L5. *TCAS will not issue an advisory if it is turned on or enabled to issue resolution advisories in the middle of a conflict (→ FTA-405)*[3]
L6. *If only one of two aircraft is TCAS equipped while the other has only ATCRBS altitude-reporting capability, the assurance of safe separation may be reduced (→ FTA-290).*

In our TCAS intent specification, both of these system limitations have pointers to boxes in the fault tree generated during the hazard analysis of TCAS II.

Finally, limitations may be related to problems encountered or trade-offs made during the system design process

3. The pointer to FTA-405 denotes the box labeled 405 in the Level-1 fault tree analysis.

(recorded on lower levels of the intent specification). For example, TCAS has a Level 1 performance monitoring requirement that led to the inclusion of a self-test function in the system design to determine whether TCAS is operating correctly. The following system limitation relates to this self-test facility:

L7. *Use by the pilot of the self-test function in flight will inhibit TCAS operation for up to 20 seconds depending upon the number of targets being tracked. The ATC transponder will not function during some portion of the self-test sequence.*

Most of these system limitations will be traced down in the intent specification levels to the user documentation. In the case of an avionics system like TCAS, this specification includes the Pilot Operations (Flight) Manual on Level 4 of our TCAS intent specification. An example is shown in Section 4.2.2.

Evaluation criteria and *priorities* are used to resolve conflicts among goals and design constraints and to guide design choices at lower levels. This information has not been included in the TCAS example specification as I was unable to find out how these decisions were made during the TCAS design process.

Finally, Level 1 contains the analysis results for system-level (emergent) properties such as safety or security. For the TCAS specification, a hazard analysis (including fault tree analysis and failure modes and effects analysis) was performed and is included and linked to the safety-critical design constraints on this level and to lower-level design decisions based on the hazard analysis. Whenever changes are made in safety-critical systems or software (during development or during maintenance and evolution), the safety of the change needs to be evaluated. This process can be difficult and expensive. By providing links throughout the levels of the intent specification, it should be easy to assess whether a particular design decision or piece of code was based on the original safety analysis or safety-related design constraint.

4.2.2 System Design Principles

The second level of the specification contains *System Design Principles*—the basic system design and scientific and engineering principles needed to achieve the behavior specified in the top level. The horizontal dimension again allows abstraction and refinement of the basic system principles upon which the design is predicated.

For TCAS, this level includes such general principles as the basic *tau* concept, which is related to all the high-level alerting goals and constraints:

PR1.
Each TCAS-equipped aircraft is surrounded by a protected volume of airspace. The boundaries of this volume are shaped by the tau *and DMOD criteria.*

PR1.1.
TAU: In collision avoidance, time-to-go to the closest point of approach (CPA) is more important than distance-to-go to the CPA. Tau is an approximation of the time in seconds to CPA. Tau equals 3,600 times the slant range in nmi, divided by the closing speed in knots.

PR1.2.
DMOD: If the rate of closure is very low, a target could slip in very close without crossing the tau boundaries and triggering an advisory. In order to provide added protection against a possible maneuver or speed change by either aircraft, the tau boundaries are modified (called DMOD). DMOD varies depending on own aircraft's altitude regime. See Table 2.

The principles are linked to the related higher level requirements, constraints, assumptions, limitations, and hazard analysis, as well as linked to lower-level system design and documentation. Assumptions used in the formulation of the design principles may also be specified at this level. For example, the TCAS design has a built-in bias against generating advisories that would result in the aircraft crossing paths (called *altitude crossing advisories*).

PR36.2
A bias against altitude crossing RAs is also used in situations involving intruder level-offs at least 600 feet above or below the TCAS aircraft. In such a situation, an altitude-crossing advisory is deferred if an intruder aircraft that is projected to cross own aircraft's altitude is more than 600 feet away vertically($\downarrow Alt\_Separation\_Test_{m-351}$).

Assumption.
In most cases, the intruder will begin a level-off maneuver when it is more than 600 feet away and, so, should have a greatly reduced vertical rate by the time it is within 200 feet of its altitude clearance (thereby, either not requiring an RA if it levels off more than ZTHR[4] *feet away or requiring a noncrossing advisory for level-offs begun after* ZTHR *is crossed, but before the 600 foot threshold is reached).*

The example above includes a pointer down to the part of the black box requirements specification ($Alt\_Separation\_Test$) that embodies the design principle. As another example of the type of links that may be found between Level 2 and the levels above and below it, consider the following: TCAS II advisories may need to be inhibited because of an inadequate climb performance for the particular aircraft on which TCAS II is installed. The collision avoidance maneuvers posted as advisories (called RAs or Resolution Advisories) by TCAS II assume an aircraft's ability to safely achieve them. If it is likely they are beyond the capability of the aircraft, then TCAS II must know beforehand so it can change its strategy and issue an alternative advisory. The performance characteristics are provided to TCAS II through the aircraft interface. An example design principle (related to this problem) found on Level 2 of the intent specification is:

PR39.
Because of the limited number of inputs to TCAS for aircraft performance inhibits, in some instances, where inhibiting RAs would be appropriate it is not possible to do so (↑ L3). In these cases, TCAS may command maneuvers that may significantly

4. The vertical dimension, called ZTHR, used to determine whether advisories should be issued varies from 750 to 950 feet, depending on the TCAS aircraft's altitude.

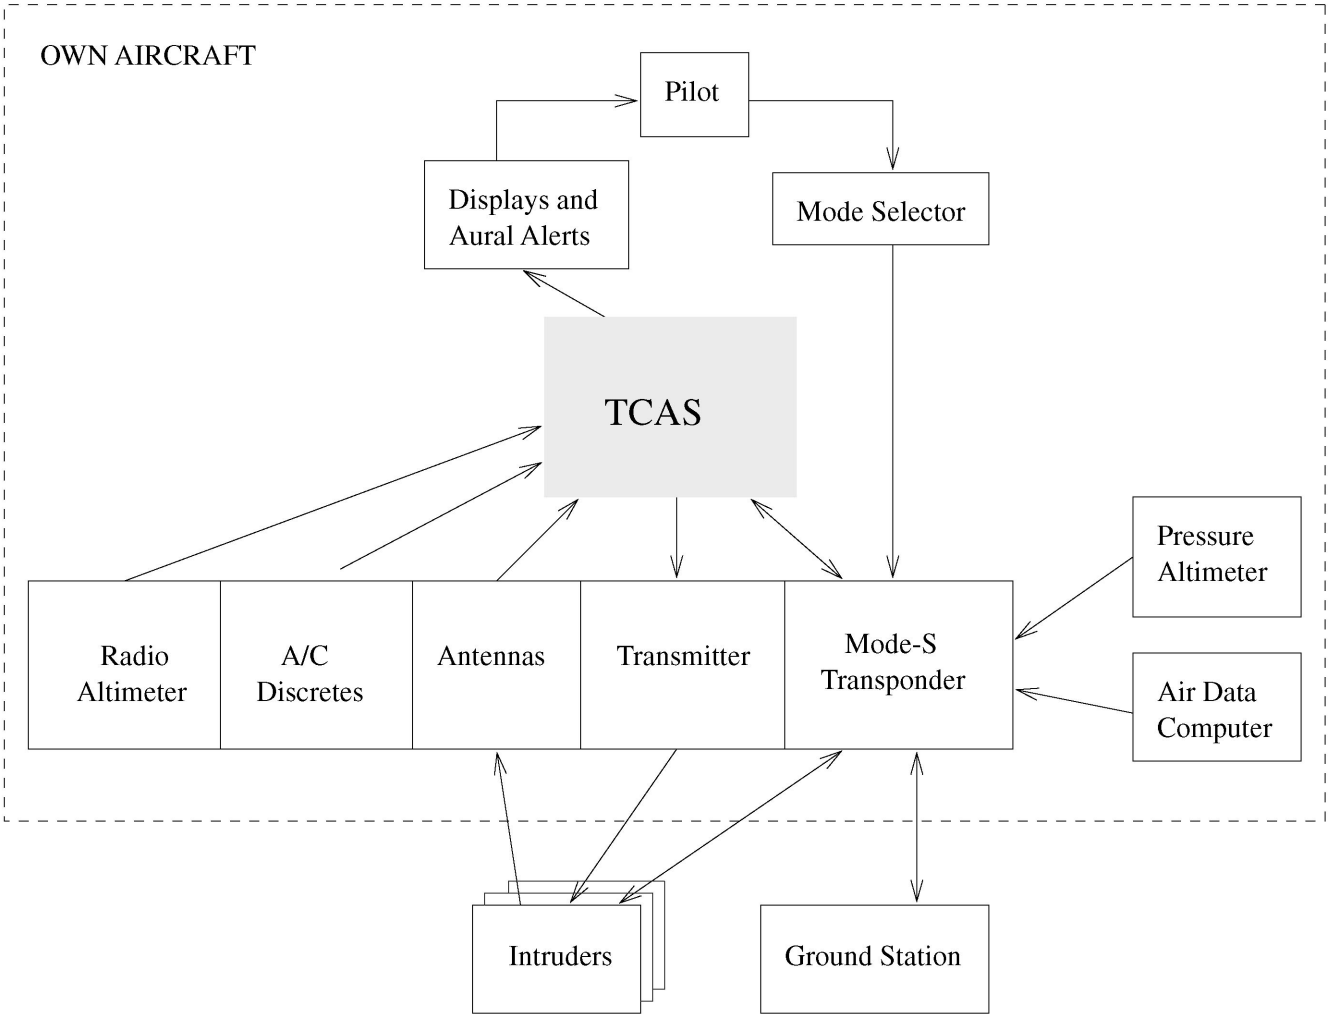

Fig. 4. System viewpoint showing the system interface topology for the Blackbox Behavior level of the TCAS specification.

reduce stall margins or result in stall warning (↑ SC9.1). Conditions where this may occur include... The aircraft flight manual or flight manual supplement should provide information concerning this aspect of TCAS so that flight crews may take appropriate action (↓ [Pilot procedures on Level 3 and Aircraft Flight Manual on Level 4]).

Finally, principles may reflect trade-offs between higher-level goals and constraints. As examples:

PR3.
Trade-offs must be made between necessary protection (G1) and unnecessary advisories (SC5). This is accomplished by controlling the sensitivity level, which controls the tau and, therefore, the dimensions of the protected airspace around each TCAS-equipped aircraft. The greater the sensitivity level, the more protection is provided but the higher is the incidence of unnecessary alerts. Sensitivity level is determined by...

PR38.
The need to inhibit CLIMB *RAs because of inadequate aircraft climb performance will increase the likelihood of TCAS II 1) issuing crossing maneuvers, which in turn increases the possibility that an RA may be thwarted by the intruder maneuvering (↑ SC7.1, FTA-1150), 2) causing an increase in* DESCEND *RAs at low altitude (↑ SC8.1), and 3) providing no RAs if below the descend inhibit level (1,200 feet above ground level on takeoff and 1,000 feet above ground level on approach).*

4.2.3 Blackbox Behavior

Beginning at the third level, or *Blackbox Behavior* level, the specification starts to contain information more familiar to software engineers. Above this level, much of the information, if located anywhere, is found in system engineering specifications. The Blackbox Behavior model at the whole system viewpoint specifies the system components and their interfaces, including the human components (operators). Fig. 4 shows a system-level view of TCAS II and its environment. Each system component behavioral description and each interface is refined in the normal way along the horizontal dimensions.

The environment description includes the assumed behavior of the external components (such as the altimeters and transponders for TCAS), including, perhaps, failure

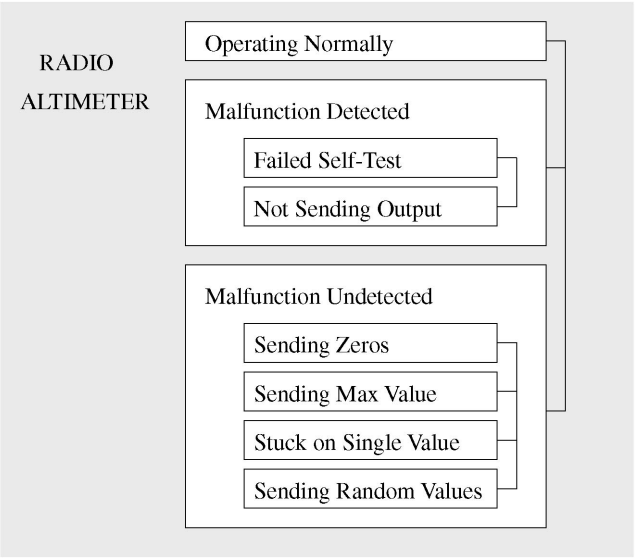

Fig. 5. Part of the SpecTRM-RL description of an environment component (a radio altimeter). Modeling failure behavior is especially important for safety analyses. In this example, 1) the altimeter may be operating correctly, 2) it may have failed in a way that the failure can be detected by TCAS II (i.e., it fails a self-test and sends a status message to TCAS or it is not sending any output at all), or 3) the malfunctioning is undetected and it sends an incorrect radio altitude.

behavior, upon which the correctness of the system design is predicated, along with a description of the interfaces between the TCAS system and its environment. Fig. 5 shows part of a state-machine description of an environment component, in this case an altimeter.

Remember that the boundaries of a system are purely an abstraction and can be set anywhere convenient for the purposes of the specifier. In this case, any component that was already on the aircraft or in the airspace control system and was not newly designed or built as part of the TCAS effort was included as environment.

Going along this level to the right, each arrow in Fig. 4 represents a communication and needs to be described in more detail. Each box (component) also needs to be refined. What is included in the decomposition of the component will depend on whether the component is part of the environment or part of the system being constructed. The language used to describe the components may also vary. State-machine language called SpecTRM-RL (Specification Tools and Requirements Methodology—Requirements Language), which is a successor to the language (RSML) used in our official TCAS II specification [36] was used. Fig. 6 shows part of the SpecTRM-RL description of the behavior of the CAS (collision avoidance system) subcomponent. SpecTRM-RL specifications are intended to be both easily readable with minimum instruction and formally analyzable (we have a set of analysis tools that work on these specifications).

Note that the behavioral descriptions at this level are purely blackbox: They describe the inputs and outputs of each component and their relationships *only* in terms of externally visible variables, objects, and mathematical functions. Any of these components (except the humans, of course) could be implemented either in hardware or software (and, in fact, some of the TCAS surveillance functions are implemented using analog devices by some vendors). Decisions about physical implementation, software design, internal variables, and so on are limited to levels of the specification below this one.

Other information at this level might include flight crew requirements such as description of tasks and operational procedures, interface requirements, and the testing requirements for the functionality described on this level. We have developed a visual operator, task modeling language that can be translated to SpecTRM-RL and, thus, permits integrated simulation and analysis of the entire system, including human-computer interactions [5].

4.2.4 Design Representation

The two lowest levels of an intent specification provide the information necessary to reason about component design and implementation. The fourth level, *Design Representation*, contains design information. Its content will depend on whether the particular function is being implemented using analog or digital devices or both. In any case, this level is the first place where the specification should include information about the physical or logical implementation of the components.

For functions implemented on digital computers, the fourth level might contain the usual software design documents or it might contain information different from that normally specified. Again, this level is linked to the higher level specification.

The design intent information may not all be completely linked and traceable upward to the levels above the Design Representation—for example, design decisions based on performance or other issues unrelated to requirements or constraints, such as the use of a particular graphics package because the programmers are familiar with it or it is easy to learn. Knowing that these decisions are *not* linked to higher level purpose is important during software maintenance and evolution activities.

The fourth level of the example TCAS intent specification simply contains the official pseudocode design specification. But, this level might contain information different than we usually include in design specifications. For example, Soloway et al. [41] describe the problem of modifying code containing delocalized plans (plans or schemas with pieces spread throughout the software). They recommend using pointers to chain the pieces together, but a more effective approach might be to put the plan or schema at the higher design representation level and point to the localized pieces in the lower level Code or Physical representation. The practicality of this approach, of course, needs to be determined.

Soloway et al. [41] also note that reviewers have difficulty reviewing and understanding code that has been optimized. To assist in code reviews and walkthroughs, the unoptimized code sections might be shown in the refinement of the Design Representation along with mappings to the actual optimized code at the lower implementation level.

The possibilities for new types of information and representations at this level of the intent hierarchy is the subject of long-term research.

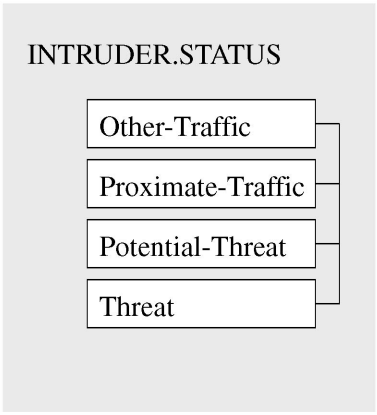

Threat → Other-Traffic

				OR	
	Alt-Reporting **in-state** Lost	T	T	T	.
	Bearing-Valid $_{m\text{-}478}$	F	.	T	.
A	Range-Valid $_{v\text{-}398}$	.	F	T	.
N	Proximate-Traffic-Condition $_{m\text{-}498}$	.	.	F	.
D	Potential-Threat-Condition $_{m\text{-}494}$	.	.	F	.
	Other-Aircraft **in-state** On-Ground	.	.	.	T

Description: A threat is reclassified as other traffic if its altitude reporting has been lost (⩓PR13) and either the bearing or range inputs are invalid; if its altitude reporting has been lost and both the range and bearing are valid but neither the proximate nor potential threat classification criteria are satisfied; or the aircraft is on the ground (⩓PR12).

Mapping to Level 2: ⩓PR23, ⩓PR29

Mapping to Level 4: ⩔Section 7.1, Traffic-Advisory

Fig. 6. Part of a SpecTRM-RL Blackbox Behavior level description of the criteria for downgrading the status of an intruder (into our protected volume) from being labeled a threat to being considered simply as other traffic. Intruders can be classified in decreasing order of importance as a threat, a potential threat, proximate traffic, and other traffic. In the example, the criterion for taking the transition from state *Threat* to state *Other Traffic* is represented by an AND/OR table, which evaluates to TRUE if any of its columns evaluates to TRUE. A column is TRUE if all of its rows that have a "T" are TRUE and all of its rows with an "F" are FALSE. Rows containing a dot represent "don't care" conditions. The subscripts denote the type of expression (e.g., v for input variable, m for macro, t for table, and f for function) as well as the page in the document on which the expression is defined. A macro is simply an AND/OR table used to implement an abstraction that simplifies another table.

Other information at this level might include hardware design descriptions, the human-computer interface design specification, the pilot operations (flight) manual, and verification requirements for the requirements and design specified on this level.

4.2.5 Physical Representation

The lowest level includes a description of the physical implementation of the levels above. It might include the software itself, hardware assembly instructions, training requirements (plan), etc.

4.2.6 Example

To illustrate this approach to structuring specifications, a small example is used related to generating resolution advisories. TCAS selects a resolution advisory (vertical escape maneuver) against other aircraft that are considered a threat to the aircraft on which the TCAS system resides. A resolution advisory (RA) has both a sense (upward or downward) and a strength (vertical rate), and it can be positive (e.g., CLIMB) or negative (e.g., DON'T CLIMB). In the software to evaluate the sense to be chosen against a particular threat, there is a procedure to compute what is called a "Don't-Care-Test." The software itself (Level 5) would contain comments about implementation decisions and also a pointer up to the Level 4 design documentation and from there up to the Level 3 black-box description of this test, shown in Fig. 7.

In turn, the blackbox (Level 3) description of the Dont-Care-Test would be linked to Level 2 explanations of the intent of the test and the reason behind (why) the design of the test. For example, our Level 2 TCAS intent specification contains the following:

> PR35.
> **Don't-Care-Test.** *When TCAS is displaying an RA against one threat and then attempts to choose a sense against a second threat, it is often desirable to choose the same sense against it as was chosen against the first threat, even if this sense is not optimal for the new threat. One advantage is display continuity (↑ SC6). Another advantage is that the pilot may maneuver more sharply to increase separation against both threats. If a dual sense advisory is given, such as* DON'T CLIMB AND DON'T DESCEND, *a vertical maneuver to increase separation against one threat reduces separation against the other threat. The most important advantage, however, is to avoid sacrificing separation inappropriately against the first threat in order to gain a marginal advantage against the second threat.*
>
> *The don't-care test determines the relative advantages of optimizing the sense against the new threat versus selecting the same sense for both threats. When the former outweighs the latter, the threat is called a do-care threat; otherwise, the threat is a don't-care threat.*

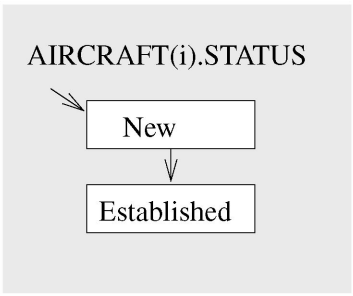

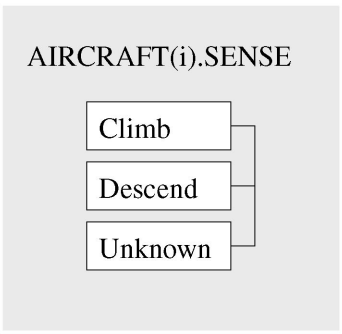

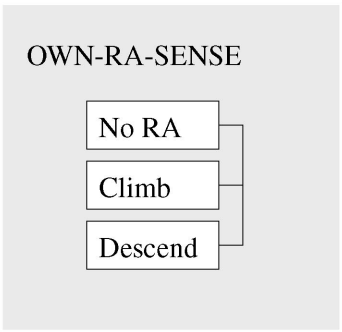

Macro: Don't-Care-Test(i)

					OR	
AND	Aircraft(i).Capability $_{v\text{-}392}$ = TCAS-TA/RA	F	F	F	F
	Aircraft(i).Status **in-state** New	T	T	T	T
	Aircraft(i).Sense **in-state** Climb	T	.	T	.
	Aircraft(i).Sense **in-state** Descend	.	T	.	T
	Down-Separation $_{f\text{-}517}$ < ALIM [⩚PR1]	F	.	.	.
	Up-Separation $_{f\text{-}542}$ < ALIM [⩚PR1]	.	F	.	.
	Own-RA-Sense **in-state** Descend	.	.	T	.
	Own-RA-Sense **in-state** Climb	.	.	.	T
	Some Aircraft(j).Sense **not-in-same-state-as** Aircraft(i).Sense	.	.	T	T
	Some Aircraft(j).Vertical-Miss-Distance $_{f\text{-}543}$ (RELALT, TAUM, TRTRU, TVPE) < Separation-Second-Choice(i) $_{f\text{-}538}$	.	.	T	T

Comment: The last two entries in the AND/OR table ensure that there exists at least one other aircraft that is a threat and has selected a sense opposite that of the current aircraft, and that the modeled separation for that aircraft following a leveloff is worse than the modeled separation for the current aircraft in the opposite (second choice) sense.

Mapping to Level 2: ⩚PR35

Mapping to Level 4: ⩛Sense.Dont-care-test

Abbreviations:

ALIM = Positive-RA-Altitude-Limit-Threshold $_{t\text{-}545}$ [Alt-Layer-Value $_{f\text{-}510}$]

RELALT = Own-Tracked-Alt $_{f\text{-}529}$ + (4 s x Own-Tracked-Alt-Rate $_{f\text{-}528}$) - Other-Tracked-Alt $_{f\text{-}524}$

TAUM = Min (Max (Modified-Tau-Capped $_{f\text{-}522}$, 10 s, True-Tau-Uncapped $_{f\text{-}542}$)

TRTRU = True-tau-Capped $_{f\text{-}542}$

TVPE = XTVPETBLX $_{t\text{-}552}$ [Other-Sensitivity-Level $_{v\text{-}391}$]

Fig. 7. This macro is used in defining which resolution advisory will be chosen when multiple aircraft (threats) are involved, among the most complicated aspects of the collision avoidance logic. Abbreviations are used to enhance readability.

This Level 2 description in turn points up to high-level goals to maintain separation between aircraft and constraints (both safety-related and nonsafety-related) on how this can be achieved. We found while constructing the TCAS intent specification that having to provide these links identified goals and constraints that did not seem to be documented anywhere but were implied by the design and some of the design documentation.

```
PROCESS Sense.Don't_care_test;

        (↑)Don't_Care_Test_{m-357}, (↑)Climb_Desc._Inhibit_{m-317}

    {WL threat = threat whose WL entry is input to task}
    {TF threat = threat examined in loop below}
    IF (either sense provides adequate separation)
       THEN SET Don't_care flag for WL threat;
       ELSE CLEAR Don't_care flag for WL threat;
            IF (own resolution advisories show a Positive in second-choice sense)
               THEN calculate own altitude following a leveloff;
                    REPEAT WHILE (more entries in threat file AND don't_care flag
                                  for WL threat not set);
                                  IF (resolution against TF threat shows a Positive in same sense
                                      as second choice for WL threat)
                                      THEN calculate altitude relative to TF threat and
                                           time for leveloff;
                                           {result of 'do care' for WL threat}
                                           CALL vertical_miss_distance_calculation
                                                IN (rel alt, rel vert rate, start time (WL threat)
                                                    end time (WL threat), clip time (WL threat));
                                           IF (sep with leveloff vs. TR threat less than that
                                                for second choice maneuver vs. WL threat)
                                                THEN SET Don't_care flag for the WL threat;
                                                     {allow second choice sense}
                                  Select next threat file entry;
                    ENDREPEAT;
END Don't_care_test;
```

Fig. 8. The pseudocode for the Don't-Care-Test.

Understanding the design of the Don't-Care-Test also requires understanding other concepts of sense selection and aircraft separation requirements that are used in the blackbox description (and in the implementation) of the Don't-Care-Test procedure. For example, the separation between aircraft in Fig. 7 is defined in terms of ALIM. The concept is used in the Level 3 documentation, but the meaning and intent behind using the concept is defined in the basic TCAS design principles at Level 2:

PR2.
ALIM *is the desired or "adequate" amount of separation between aircraft that TCAS is designed to meet. This amount varies from 400 to 700 feet, depending on own aircraft's altitude.* ALIM *includes allowances to account for intruder and own altimetry errors and vertical tracking uncertainties that affect track projections (see PR22.3). The value of ALIM increases with altitude to reflect increased altimetry error (↑ SC4.5) and the need to increase tracked separation at higher altitudes.*

The blackbox behavioral specification shown in Fig. 7 also points to the module that implements this required behavior in the design specification on Level 4. For TCAS II, pseudocode was used for the design specification. Fig. 8 shows the pseudocode provided by MITRE for the Don't-Care-Test.

The structure of intent specifications has advantages in solving various software engineering problems—such as changing requirements, program understanding, maintaining and changing code, and validation—as discussed in the next section.

5 Intent Specification Support for Software Engineering Problem Solving

As stated earlier, our representations of problems have an important effect on our problem-solving ability and the strategies we use. A basic hypothesis of this paper is that intent specifications will support the problem solving required to perform software engineering tasks. This hypothesis seems particularly relevant with respect to tasks involving education and program understanding, search, design, changing requirements, fault tolerance, safety assurance, maintenance, and evolution.

5.1 Education and Program Understanding

Curtis et. al. [9] did a field study of the requirements and design process for 17 large systems. They found that substantial design effort in projects was spent coordinating a common understanding among the staff of both the application domain and of how the system should perform within it. The most successful designers understood the application domain and were adept at identifying unstated requirements, constraints, or exception conditions and mapping between these and the

computational structures. This is exactly the information that is included in the higher levels of intent specifications and the mappings to the software. Thus, using intent specifications should help with education in the most crucial aspects of the system design for both developers and maintainers and augment the abilities of both, i.e., increase the intellectual manageability of the task.

5.2 Search Strategies

Vicente and Rasmussen have noted that means-ends hierarchies constrain search in a useful way by providing traceability from the highest level goal statements down to implementations of the components [46]. By starting the search at a high level of abstraction and then deciding which part of the system is relevant to the current goals, the user can concentrate on the subtree of the hierarchy connected to the goal of interest: The parts of the system not pertinent to the function of interest can easily be ignored. This type of "zooming-in" behavior has been observed in a large number of psychological studies of expert problem solvers. Recent research on problem-solving behavior consistently shows that experts spend a great deal of their time analyzing the functional structure of a problem at a high level of abstraction before narrowing in on more concrete details [3], [6], [16], [34], [42].

With other hierarchies, the links between levels are not necessarily related to goals. So, although it is possible to use higher levels of abstraction in a standard decomposition or refinement hierarchy to select a subsystem of interest and to constrain search, the subtree of the hierarchy connected to a particular subsystem does not necessarily contain system components that are relevant to the goals and constraints that the problem solver is considering.

Upward search in the hierarchy, such as that required for debugging, is also supported by intent specifications. Vicente and Rasmussen claim (and have experimental evidence to support) that, in order for operators to correctly and consistently diagnose faults, they must have access to higher-order functional information since this information provides a reference point defining how the system *should* be operating. States can only be described as errors or faults with reference to the intended purpose. Additionally, causes of improper functioning depend upon aspects of the implementation. Thus, they are explained bottom up. The same argument seems to apply to software debugging. There is evidence to support this hypothesis. Using protocol analysis, Vessey found that the most successful debuggers had a "system" view of the software [42].

5.3 Design Criteria and Evaluation

An interesting implication of intent specifications is their potential effect on system and software design. Such specifications might not only be used to understand and validate designs, but also to guide them.

An example of a design criterion appropriate to intent specifications might be to be minimize the number of one-to-many mappings between levels in order to constrain downward search and limit the effects of change in higher levels upon the lower levels. Minimizing many-to-many (or many-to-one) mappings, would in addition, ease activities that require following upward links and minimize the side effects of lower-level changes.

Intent specifications assist in identifying intent-related structural dependencies (many-to-many mappings across hierarchical levels) to allow minimizing them during design, and they clarify the tradeoffs being made between conflicting goals. Software engineering attempts to define coupling between modules have been limited primarily to the design level. Perhaps an intent specification can provide a usable definition of coupling with respect to emergent properties and to assist in making design tradeoffs between various types of high-level coupling.

5.4 Minimizing the Effects of Requirements Changes

Hopefully, the highest levels of the specification will not change, but sometimes they do, especially during development, as system requirements become better understood. Functional and intent aspects are represented throughout an intent specification, but in increasingly abstract and global terms at the higher levels. The highest levels represent more stable design goals that are less likely to change (such as detecting potential threats in TCAS), but, when they do, they have the most important (and costly) repercussions on the system and software design and development, and they may require analysis and changes at all the lower levels. We need to be able to determine the potential effects of changes and, proactively, to design to minimize them.

Reversals in TCAS are an example of this. About four years after the original TCAS specification was written, experts discovered that it did not adequately cover requirements involving the case where the pilot of an intruder aircraft does not follow his or her TCAS advisory and, thus, TCAS must change the advisory to its own pilot. This change in basic requirements caused extensive changes in the TCAS design, some of which introduced additional subtle problems and errors that took years to discover and rectify.

Anticipating exactly what changes will occur and designing to minimize the effects of those changes is difficult and the penalties for being wrong are great. Intent specifications theoretically provide the flexibility and information necessary to design to ease high-level requirements changes without having to predict exactly which changes will occur: The abstraction and design are based on intent (system requirements) rather than on part-whole relationships (which are the least likely to change with respect to requirement or environment changes).

5.5 Design of Run-Time Assertions

Intent specifications may assist software engineers in designing effective fault tolerance mechanisms. Detecting unanticipated faults during execution has turned out to be a very difficult problem. For example, in one of our empirical studies, we found that programmers had difficulty writing effective assertions for detecting errors in executing software [25]. We have suggested that using results from safety analyses might help in determining which assertions are required and where to detect the most important errors [23].

The information in intent specifications tracing intent from requirements, design constraints, and hazard analyses through the system and software design process to the software module (and back) might assist with writing effective and useful assertions to detect general violations of system goals and constraints.

5.6 Safety Assurance

A complete safety analysis and methodology for building safety-critical systems requires identifying the system-level safety requirements and constraints and then tracing them down to the components [24]. After the safety-critical behavior of each component has been determined (including the implications of its behavior when the components interact with each other), verification is required that the components do not violate the identified safety-related behavioral constraints. In addition, whenever any change is made to the system or when new information is obtained that brings the safety of the design into doubt, revalidation is required to ensure that the change does not degrade system safety. To make this verification (and reverification) easier, safety-critical parts of the software should be isolated and minimized.

This analysis cannot be performed efficiently unless those making decisions about changes and those actually making the changes know which parts of the system affect a particular safety design constraint. Specifications need to include a record of the design decisions related to basic safety-related system goals, constraints, and hazards (including both general design principles and criteria and detailed design decisions), the assumptions underlying these decisions, and why the decisions were made and particular design features included. Intent specifications capture this information and provide the ability to trace design features upward to specific high-level system goals and constraints.

5.7 Software Maintenance and Evolution

Although intent specifications provide support for a top-down, rational design process, they may be even more important for the maintenance and evolution process than for the original designer, especially of smaller or less complex systems. Software evolution is challenging because it involves many complex cognitive processes—such as understanding the system's structure and function, understanding the code and documentation and the mapping between the two, and locating inconsistencies and errors—that require complex problem-solving strategies.

Intent specifications provide the structure required for recording the most important design rationale information, i.e., that related to the purpose and intent of the system, and locating it when needed. They, therefore, can assist in the software change process.

While trying to build a model of TCAS, we discovered that the original conceptual model of the TCAS system design had degraded over the years as changes were made to the pseudocode to respond to errors found, new requirements, better understanding of the problem being solved, enhancements of various kinds, and errors introduced during previous changes. The specific changes made often simplified the process of making the change or minimized the amount of code that needed to be changed, but complicated or degraded the original model. Not having any clear representation of the model also contributed to its degradation over the 10 years of changes to the pseudocode.

By the time we tried to build a representation of the underlying conceptual model, we found that the system design was unnecessarily complex and lacked conceptual coherency in many respects, but we had to match what was actually flying on aircraft. I believe that making changes without introducing errors or unnecessarily complicating the resulting conceptual model would have been simplified if the TCAS staff had had a blackbox requirements specification of the system. Evolution of the pseudocode would have been enhanced even more if the extra intent information had been specified or organized in such a way that it could easily be found and traced to the code.

Tools for restructuring code have been developed to cope with this common problem of increasing complexity and decreasing coherency of maintained code [17]. Using intent specifications will not eliminate this need, but we hope it will be reduced by providing specifications that assist in the evolution process and, more important, assist in building software that is more easily evolved and maintained. Such specifications may allow for backing up and making changes in a way that will not degrade the underlying conceptual model because the model is explicitly described and its implications traced from level to level. Intent specifications may also allow controlled changes to the higher levels of the model if they become necessary.

Maintenance and evolution research has focused on ways to identify and capture information from legacy code. While useful for solving important short-term problems, our long-term goal should be to specify and design systems that lend themselves to change easily—that is, evolvable systems. Achieving this goal requires devising methodologies that support change throughout the entire system life cycle—from requirements and specification to design, implementation, and maintenance. For example, we may be able to organize code in a way that will minimize the amount of code that needs to be changed or that needs to be evaluated when deciding if a change is safe or reasonable.

In summary, the author believes that effective support for such evolvable systems will require a new paradigm for specification and design and hypothesize that such a paradigm might be rooted in abstractions based on intent. Intent specifications provide the framework to include the information maintainers need in the specification. They increase the information content so that less inferencing (and guessing) is required. Intent specifications not only support evolution and maintenance, but they may be more evolvable themselves, which would ease the problem of keeping documentation and implementation consistent. In addition, they also provide the possibility of designing for evolution so that the systems we build are more easily maintained and evolved.

6 Conclusions

Specifications are constructed to help us solve problems. Any theory of specification design, then, should be based on fundamental concepts of problem-solving behavior. It should also support the basic systems engineering process. This paper has presented one such approach to system and software specifications based on underlying ideas from psychology, systems theory, human factors, system engineering, and cognitive engineering.

The choice of content, structure, and form of specifications have a profound effect on the kind of cognitive processing that the user must bring to bear to use a specification for the tasks involved in system and software design and construction, maintenance, and evolution. Intent specifications provide a way of coping with the complexity of the cognitive demands on the builders and maintainers of automated systems by basing our specifications on means-ends as well as part-whole abstractions. The author believes that the levels of the means-ends hierarchy reflect a rational design philosophy for the systems engineering of complex systems and, thus, a rational way to specify the results of the process. They provide mapping (tracing) of decisions made earlier into the later stages of the process. Design decisions at each level are linked to the goals and constraints they are derived to satisfy. A seamless (gapless) progression is recorded from high-level system requirements down to component requirements, design, and implementation.

In addition, intent specifications provide a way of integrating formal and informal aspects of specifications. Completely informal specifications of complex systems tend to be unwieldy and difficult to validate. Completely formal specifications provide the potential for mathematical analysis and proofs, but omit necessary information that cannot be specified formally. Some formal approaches require building special models in addition to the regular system specifications. The author believes that the widespread use of formal specifications in industry will require the development of formal specifications that are readable with minimal training requirements and that are integrated with informal specifications. Ideally, formal analysis should not require building special models that duplicate the information included in the specification or it is unlikely that industry will find the use of formal methods to be cost effective.

An example intent specification for TCAS II has been constructed and was used as an example in this paper. The reader is cautioned, however, that intent specifications are a logical abstraction that can be realized in many different physical ways. That is, the particular organization used for the TCAS specification is simply one possible physical realization of the general logical organization inherent in intent specifications.

Acknowledgments

This work was partially supported by NASA Grant NAG-1-1495 and by U.S. National Science Foundation Grant CCR-9396181.

References

[1] *Mental Models and Human-Computer Interaction,* D. Ackermann and M.J. Tauber, eds. Amsterdam: North-Holland, 1990.
[2] W.R. Ashby, "Principles of the Self-Organizing System," *Principles of Self-Organization,* H. Von Foerster and G.W. Zopf, eds., Pergamon, 1962.
[3] M. Beveridge and E. Parkins, "Visual Representation in Analogical Program Solving," *Memory and Cognition,* vol. 15, 1987.
[4] R. Brooks, "Towards a Theory of Comprehension of Computer Programs," *Int'l J. Man-Machine Studies,* vol. 18, pp. 543-554, 1983.
[5] M. Brown and N.G. Leveson, "Modeling Controller Tasks for Safety Analysis," *Proc. Second Workshop Human Error and System Development,* Apr. 1998.
[6] M.A. Buttigieg and P.M. Sanderson, "Emergent Features in Visual Display Design for Two Types of Failure Detection Tasks," *Human Factors,* vol. 33, 1991.
[7] S.M. Casner, "A Task Analytic Approach to the Automated Design of Graphic Presentations," *ACM Trans. Graphics,* vol. 10, no. 2, Apr. 1991.
[8] P. Checkland, *Systems Thinking, Systems Practice.* John Wiley & Sons, 1981.
[9] B. Curtis, H. Krasner, and N. Iscoe, "A Field Study of the Software Design Process for Large Systems," *Comm. ACM,* vol. 31, no. 2, pp. 1,268-1,287, 1988.
[10] J. DeKleer and J.S. Brown, "Assumptions and Ambiguities in Mechanistic Mental Models," *Mental Models,* D. Gentner and A.L. Stevens, eds., Lawrence Erlbaum, 1983.
[11] N. Dinadis and K.J. Vicente, "Ecological Interface Design for a Power Plant Feedwater Subsystem," *IEEE Trans. Nuclear Science,* 1996.
[12] D. Dorner, "On the Difficulties People Have in Dealing with Complexity," *New Technology and Human Error,* J. Rasmussen, K. Duncan, and J. Leplat, eds., pp. 97-109, New York: John Wiley & Sons, 1987.
[13] K.D. Duncan, "Reflections on Fault Diagnostic Expertise," *New Technology and Human Error,* J. Rasmussen, K. Duncan, and J. Leplat, eds., pp. 261-269, New York: John Wiley & Sons, 1987.
[14] B. Fischoff, P. Slovic, and S. Lichtenstein, "Fault Trees: Sensitivity of Estimated Failure Probabilities to problem Representation," *J. Experimental Psychology: Human Perception and Performance,* vol. 4, 1978.
[15] M.J. Fitter and T.R.G. Green, "When Do Diagrams Make Good Programming Languages?" *Int'l J. Man-Machine Studies,* vol. 11, pp. 235-261, 1979.
[16] R. Glaser and M.T.H. Chi, "Overview," *The Nature of Expertise,* R. Glaser, M.T.H. Chi, and M.J. Farr, eds., Hillsdale, N.J.: Erlbaum, 1988.
[17] W. Griswold and D. Notkin, "Architectural Tradeoffs for a Meaning-Preserving Program Restructuring Tool," *IEEE Trans. Software Eng.,* vol. 21, no. 3, pp. 275-287, Mar. 1995.
[18] G. Harman, "Logic, Reasoning, and Logic Form," *Language, Mind, and Brain,* T.W. Simon and R.J. Scholes, eds., Lawrence Erlbaum, 1982.
[19] M.S. Jaffe, N.G. Leveson, M.P.E. Heimdahl, and B. Melhart, "Software Requirements Analysis for Real-Time Process-Control Systems," *IEEE Trans. Software Eng.,* vol. 17, no. 3, Mar. 1991.
[20] C.A. Kaplan and H.A. Simon, "In Search of Insight," *Cognitive Psychology,* vol. 22, 1990.
[21] K. Kotovsky, J.R. Hayes, and H.A. Simon, "Why Are Some Problems Hard? Evidence from Tower of Hanoi," *Cognitive Psychology,* vol. 17, 1985.
[22] S. Letovsky, "Cognitive Processes in Program Comprehension," *Proc. First Workshop Empirical Studies of Programmers,* pp. 58-79, 1986.
[23] N.G. Leveson, "Software Safety in Embedded Computer Systems," *Comm. ACM,* vol. 34, no. 2, Feb. 1991.
[24] N.G. Leveson, *Safeware: System Safety and Computers.* Addison-Wesley, 1995.
[25] N.G. Leveson, S.S. Cha, J.C. Knight, and T.J. Shimeal, "The Use of Self-Checks and Voting in Software Error Detection: An Empirical Study," *IEEE Trans. Software Eng.,* vol. 16, no. 4, Apr. 1990.
[26] N.G. Leveson, M.P.E. Heimdahl, H. Hildreth, and J.D. Reese, "Requirements Specification for Process-Control Systems," *IEEE Trans. Software Eng.,* vol. 20, no. 9, Sept. 1994.
[27] N.G. Leveson, L.D. Pinnel, S.D. Sandys, S. Koga, and J.D. Reese, "Analyzing Software Specifications for Mode Confusion Potential," *Proc. Workshop Human Error and System Development,* 1977.

[28] D.A. Lucas, "Mental Models and New Technology," *New Technology and Human Error,* J. Rasmussen, K. Duncan, and J. Leplat, eds., pp. 321-325, New York: John Wiley & Sons, 1987.
[29] J.R. Newman, "Extension of Human Capability through Information Processing and Display Systems," Technical Report SP-2560, System Development Corp., 1966.
[30] D.A. Norman, *Things that Make Us Smart.* Addison-Wesley, 1993.
[31] J. Rasmussen and A. Pejtersen, "Virtual Ecology of Work," *An Ecological Approach to Human Machine Systems I: A Global Perspective,* J.M. Flach, P.A. Hancock, K. Caird and K.J. Vicente, eds., Hillsdale, N.J.: Erlbaum, 1995.
[32] N. Pennington "Stimulus Structures and Mental Representations in Expert Comprehension of Computer Programs," *Cognitive Psychology* vol. 19, pp. 295-341, 1987.
[33] J. Rasmussen "The Role of Hierarchical Knowledge Representation in Decision Making and System Management," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 15, no. 2, Mar./Apr. 1985.
[34] J. Rasmussen, "Information Processing and Human-Machine," *Interaction: An Approach to Cognitive Engineering,* North-Holland, 1986.
[35] J. Rasmussen, "Mental Models and the Control of Action in Complex Environments," *Mental Models and Human-Computer Interaction,* D. Ackermann and M.J. Tauber, eds., North-Holland: Elsevier, pp. 41-69, 1990.
[36] J. Reason, *Human Error.* Cambridge Univ. Press, 1990.
[37] N.D. Sarter, D.D. Woods, and C.E. Billings, "Automation Surprises," *Handbook of Human Factors/Ergonomics,* second ed., G. Salvendy ed., New York: Wiley, 1995.
[38] B. Shneiderman and R. Mayer, "Syntactic/Semantic Interactions in Programmer Behavior: A Model and Experimental Results," *Computer and Information Sciences,* vol. 8, no. 3, pp. 219-238, 1979.
[39] G.F. Smith, "Representational Effects on the Solving of an Unstructured Decision Problem," *IEEE Trans. Systems, Man, and Cybernetics,* pp. 1,083-1,090, vol. 19, 1989.
[40] E. Soloway and K. Ehrlich, "Empirical Studies of Programming Knowledge," *IEEE Trans. Software Eng.,* vol. 10, no. 5, pp. 595-609, 1984.
[41] E. Soloway, J. Pinto, S. Letovsky, D. Littman, and R. Lampert, "Designing Documentation to Compensate for Delocalized Plans," *Comm. ACM,* vol. 31, no. 2, pp. 1,259-1,267, 1988.
[42] I. Vessey, "Expertise in Debugging Computer Programs: A Process Analysis," *Int'l J. Man-Machine Studies,* vol. 23, 1985.
[43] K.J. Vicente, "Supporting Knowledge-Based Behavior through Ecological Interface Design," PhD thesis, Univ. of Illinois at Urbana-Champaign, 1991.
[44] K.J. Vicente, K. Christoffersen, and A. Pereklit, "Supporting Operator Problem Solving through Ecological Interface Design," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 25, no. 4, pp. 529-545, 1995.
[45] K.J. Vicente and J. Rasmussen, "The Ecology of Human-Machine Systems II: Mediating Direct Perception in Complex Work Domains," *Ecological Psychology,* vol. 2, no. 3, pp. 207-249, 1990.
[46] K.J. Vicente and J. Rasmussen, "Ecological Interface Design: Theoretical Foundations," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 22, no. 4, July/Aug. 1992.
[47] E.L. Wiener, "Human Factors of Advanced Technology ("Glass Cockpit") Transport Aircraft," NASA Contractor Report 177528, NASA Ames Research Center, June 1989.
[48] D.D. Woods, "Toward a Theoretical Base for Representation Design in the Computer Medium: Ecological Perception and Aiding Human Cognition," *An Ecological Approach to Human Machine Systems I: A Global Perspective,* J.M. Flach, P.A. Hancock, K. Caird and K.J. Vicente, eds., Hillsdale, N.J.: Erlbaum, 1995.

Nancy G. Leveson is a professor of aerospace information systems in the Aeronautics and Astronautics Department at the Massachusetts Institute of Technology, Cambridge. Previously, she was Boeing Professor of computer science and engineering at the University of Washington, Seattle. She has served as editor-in-chief of the *IEEE Transactions on Software Engineering* and on the board of directors of the International Council on Systems Engineering. Dr. Leveson is a fellow of the ACM and is currently an elected member of the board of directors of the Computing Research Association, a member of the U.S. National Research Council (NRC) Commission on Engineering and Technical Systems, as well as liaison to the U.S. NRC Aeronautics and Space Engineering board. She received the 1995 AIAA Information Systems award for "developing the field of software safety and for promoting responsible software and system engineering practices where life and property are at stake" and also the 1999 ACM Alan Newell Award. Recently, Dr. Leveson was elected to the National Academy of Engineering. She is author of the book, *Safeware: System Safety and Computers,* published by Addison-Wesley, and publishes, speaks, and consults widely.