

Functional Representation of Designs and Redesign Problem Solving

Ashok Goel and B. Chandrasekaran*
Laboratory for Artificial Intelligence Research
Department of Computer and Information Science
The Ohio State University
Columbus, Ohio

Abstract

The information processing task of redesign and its subtasks of diagnosis and repair are analyzed. Various kinds of knowledge required for redesign problem solving are identified, and a scheme for representing them is described. In this scheme, the functions of the device and its structural components are represented explicitly, and causal and anticipatory knowledge about its design is organized around these functions. This functional representation language also provides primitives for representing and accessing knowledge of domain principles such as Physics laws. The use of functional representation of designs in redesign problem solving is illustrated for the redesign of the reaction wheel assembly aboard the Hubble space telescope.

1 Design: Proposal, Verification, and Redesign

The design problem can be abstractly characterized as a constrained *function-to-structure* mapping. The design task takes as input the specifications of the desired functions of a device and the constraints on the design, and produces as output a specification of a structure that realizes the desired functions and satisfies the constraints. One way to analyze a complex task such as design is to identify the methods that can be applied to the task, the knowledge and control that these methods require, and the subtasks generated by them. This analysis produces a *task structure* [Chandrasekaran, 1989], i.e., a task-subtask decomposition of the problem, along with a specification of the knowledge required for each of the subtasks. For a given task in this task structure, the choice of the method can depend on the knowledge available to the problem solver and the computational efficiency of finding the solution by various methods applicable to the task.

One method for solving design problems is *propose, verify, and redesign* [Chandrasekaran, 1988]. This method identifies and orders three subtasks, each of which in turn can be performed in different domains by different methods. For instance, case-based methods

have recently become a subject of research for the *propose* subtask, and the *verify* subtask can be performed by a variety of methods, including actual testing of the device, analytic methods such as finite element analysis, and various simulation techniques. One goal of this paper is to perform an analysis of the task of *redesign* in terms of the subtasks into which it can be decomposed, and the methods applicable to them. The second goal of this paper is to explore the use of *function-structure models* for the redesign task, specifically, to investigate the utility of the *functional representation* scheme Sembugamoorthy and Chandrasekaran, 1986 which models the relationship between the structure of a device, the behaviors that arise from it, and the teleology of the device as a whole. This research builds on our earlier work [Coel and Chandrasekaran, 1988] in which we proposed the use of functional representation of designs in critiquing a proposed design, i.e., in localizing the failure to deliver a function to a part of the structure.

1.1 Redesign: Corrective and Compensatory

Redesign is triggered as a task whenever the *verify* subtask shows that the proposed design falls short of the desired, either because some of the desired functions are not realized or because some of the behaviors are undesirable. Once the proposed design has been modified, the *verify-redesign* cycle is repeated if the design is getting closer to the desired one, or a different candidate design is sought from the *propose* subtask. In this paper, we are particularly concerned with redesign problem solving when the *verify* subtask finds an undesirable device behavior. The redesign of a ball bearing assembly which generates excess heat due to large rotational loads, where the generation of excess heat is an unintended and undesirable device behavior, is an example of this generic class of redesign problems.

Solutions to this redesign problem can be *corrective*, or *compensatory*, or some combination of the two. The redesigner may diagnose and repair the structural fault responsible for an undesirable behavior, or it may propose additional structures that can compensate for the undesirable behavior. If, for instance, isolating the structural fault responsible for an undesirable behavior or fully correcting it is not feasible, or is computationally too expensive, then the redesigner may devise a compensatory solution to the problem. In the ball bearing example, the proposal for the use of a cooler to remove the excess heat generated is a compensatory redesign solution. In this paper, we are especially interested in corrective redesign

*This research has been supported by the Defense Advanced Research Projects Agency, RADC contract F30602-85-C-0010, the National Science Foundation, grant CBT-87-03745, and the McDonnell Douglas Corporation, research contract WS-MDRL-2931. Computer facilities were enhanced by gifts from the Xerox Corporation.

problem solving.

The method of correcting an undesirable behavior further decomposes the redesign task into two subtasks: *diagnosis* and *repair*. The diagnosis subtask takes the proposed structure and its undesirable behaviors as input, and gives the structural causes for the undesirable behaviors as the output. The repair subtask of redesign takes the desired functions, the proposed structure, the undesirable behaviors and their structural causes as input, and produces as output a modified structure that realizes the desired functions without the undesirable behavior. The diagnosis subtask can be performed by a variety of methods ranging from associative mapping of behavior to structure to techniques based on simulation of behavior from structure. Below we present a method for diagnosis and repair that makes use of functional representations of designs in the form of stored *structure-to-function maps*.

2 Reaction Wheel Assembly

In order to make the present discussion more concrete, let us consider the specific problem of redesigning the reaction wheel assembly (RWA) aboard the Hubble space telescope, a slice of which is shown in Figure 1. The desired function of RWA is to make the telescope point at a chosen area of the sky. The given structure of the RWA consists of a rapidly spinning rotor mounted on a shaft. The rotating shaft is connected to a stator at both ends via assemblies of anti-friction ball bearings. The power that drives the rotor comes from a motor that is remotely controlled from earth. The stator itself is mounted on the walls of the telescope bay. The constraint on the design of RWA is to keep its mass as small as possible.

The functioning of RWA is based on the law of conservation of angular momentum. When the telescope is to be oriented in a specific direction, a signal from earth is sent to the motor that results in a change in the power supplied to the rotor. This causes a change in the angular velocity of the rotor and a corresponding change in its angular momentum. Due to the conservation of angular momentum, the angular momentum of the telescope as a whole changes in the opposite direction. When the telescope nears its desired orientation, a change in the angular momentum of the telescope in the opposite direction is achieved in a similar manner, and the telescope angular velocity is reduced to zero.

A common problem in the operation of RWA arises due to friction in the bearing assemblies. The load on the bearings due to the rapid spin of the rotor causes deformation of the bearing balls which results in increased frictional forces in the bearing assembly. This causes generation of heat in the bearing assembly. The resultant increase in temperature is detected by temperature sensors located near the bearing assemblies. Since the increase in temperature depends on the load on the bearings, a typical redesign solution to this problem is to increase the load capacity of the bearings by increasing the size of the balls.

The increased temperature in the bearing assembly is an example of an unintended and undesirable behavior.




Figure 1: Reaction Wheel Assembly

The designer of HWA *anticipated* the potential for this undesirable behavior and included sensors in the design specifically to detect its presence. Note that because of the constraint of keeping the mass as small as possible, increasing the size of the bearing balls by an arbitrarily large amount is not an acceptable redesign solution. Also, since the effects of rotational loads on bearing assemblies are not known analytically, it is not possible to exactly compute the smallest size ball bearings that can support a given rotational load.

In using the redesign of RWA as an illustration of our analysis of redesign problem solving, we assume that the angular momentum of the telescope as a whole is initially zero, and that the angular momentum of the rotor is in the anticlockwise direction. We also assume that the command from earth is to increase the angular velocity of the rotor so that the telescope acquires an angular momentum in the opposite direction, and that the desired change in the magnitude of angular momentum is proportional to the magnitude of the command signal. While these assumptions reduce the size of the problem, they do not entail any loss of generality.

3 Functional Organization of Design Knowledge

3.1 Knowledge for Redesign Problem Solving

Efficient and effective redesign problem solving requires knowledge specific to the proposed design for a device.

Design-specific knowledge can be of several kinds, including knowledge of the *structure* of the device, its components and the relations between them; knowledge of the desired *functions* of the device, and the functional abstractions of the structural components; knowledge of the *composition* of the device function from the functions of its structural components; knowledge of the *justifications* for the choice of various structural components; knowledge of the device states, the state variables characterizing them, and the *causal* dependencies between them; and knowledge of the *anticipated* side effects of the functioning of the device.

In addition, redesign problem solving requires knowledge of the design domain that goes beyond any specific design or particular device. This includes knowledge of *primitive components* and primitive relations available in the design domain; knowledge of *primitive substances* in the domain including abstract substances such as heat; knowledge of *primitive processes* in the domain such as friction and their effects on components and substances; knowledge of *generic engineering mechanisms* such as rolling friction and generic engineering devices such as bearing assemblies, and knowledge of *general domain principles* such as the law of conservation of angular momentum and general domain relations such as the momentum of a rotating object is proportional to its angular velocity.

3.2 Representation of Structure-to-Function Maps

We now describe the functional representation scheme for representing and organizing knowledge of the structure, function, and structure-to-function maps of a design. The structure-to-function maps, called *causal behaviors*, explicitly represent design-specific knowledge, and contain pointers to more general domain knowledge¹.

Let us begin with representation of function. The functions of the device and its components are represented as schemas; the schema for the function of RWA is shown in Figure 2.² The underlined expressions in the figure are the primitives of the functional representation language. The schema specifies the device state the function takes as input and the device state that the function gives as output. It also specifies the causal behavior *BehaviorChange.Momentum* that results in transforming the given input state into the desired output state, and the conditions under which the transformation is possible. Finally, the schema specifies the anticipated side-effects of achieving the function in the form of *Behavior Generate Heat*.

¹Note that, the term behavior is being used in two different contexts, to refer to the device outputs as in *undesirable behavior*, and to refer to the sequences of device states as in *causal behavior*.

²The arrow on top of a variable, such as $L_{\text{telescope}}$, indicates that the variable is a vector quantity, i.e. it has both a magnitude and a direction associated with it. The vertical bars on the sides of a vector variable indicate that only the magnitude of the variable is being used. The symbol Δ denotes a change in the value of the variable. The symbol / denotes proportionality, while the positive sign indicates direct proportionality.

Functions:
Given: Control Signal c
To-Make: Change Angular Momentum of Telescope
 $L_{\text{telescope}}$
 $|\Delta L_{\text{telescope}}| = f(+c)$
Provided: Large Angular Velocity of Rotor
rotor
By: BehaviorChangeMomentum
Side-Effect: Generation of Heat in Bearing
 Q_{bearing}
 $Q_{\text{bearing}} = f(-|\vec{\omega}_{\text{shaft}}|)$
By: BehaviorGenerateHeat
End Functions

Figure 2: Function of RWA

Knowledge of structure in the functional representation scheme is organized in a structure-substructure hierarchy. Each substructure in this hierarchy is represented as a schema. A part of the schema for the ball bearing assembly is shown in Figure 3; not all the primitives shown are used in redesign problem solving below. The schema specifies the functional abstraction of the device, the domain principles and relations underlying its operation along with the operating range, its structural relations with other components, and the justification for its choice. The schema also contains pointers to the state transitions in causal behaviors in which the component plays some role, specifically to the transitions $state2 \rightarrow state3$, $state3 \rightarrow state4$, and $stateA \rightarrow state5$ in *Behavior Generate Heat*.

Causal behaviors compose the functions of the structural components into the device functions, and are represented as acyclic directed graphs. A node in such a causal graph represents a causal state of the device characterized by its state variables. An edge between two nodes in the causal graph represents a causal state transition. The causal graphs for *BehaviorChange Momentum* and *Behavior Generate Heat* are shown in Figures 4 and 5, respectively. Note the causal dependencies between the state variables characterizing the causal states in the two figures.³ A state transition in a causal graph can be one of several types. For instance, a transition could be due to the function of some component, e.g., the transition $state1 \rightarrow state2$ in *BehaviorChangeMomentum* shown in Figure 4, or it could be based on some domain principle, e.g., the transition $state2 \rightarrow state6$ also in Figure 4. Often, domain principles are applicable only in the context of some structural component or relation, e.g., the transitions $state2 \rightarrow state3$ in Figure 5 and $state's \rightarrow stateA$ in Figure 4; sometimes they may require additional assumptions, e.g., the transition $state \rightarrow state7$ in Figure 4. Also, a state transition may point to a more detailed sequence of state transitions, e.g., the transition $state2 \rightarrow stateA$ in *Behavior Change Momentum* points

³The circular arrows adjacent to some of the vector variables in Figures 4 and 5 indicate the direction of rotational motion about the rotor axis, clockwise or anticlockwise.

Structure:
Functional Abstraction:
 Support Rotation with Minimal Friction
Operation:
Range:
 Low: <min. value ($L_{d_{bearing}}$)>
 High: <max. value ($L_{d_{bearing}}$)>
Relations:
 Frictional Forces Proportional to Load on Bearing
 $F_{r_{bearing}} = f(+L_{d_{bearing}})$
 Maximum Load Supported Proportional to Size of Bearing Balls
 $max. (L_{d_{bearing}}) = R_{ball}$
Pointers to Behaviors:
 BehaviorGenerateHeat
 Transitions: State 2 -> State 3
 State 3 -> State 4
 State 4 -> State 5
Structural Relations:
 Connected to Shaft
 Connected to Stator
 Connected to Temperature Sensor
 Contained in Reaction Wheel Assembly
 Contains Bearing Balls
Justification: Large max. ($L_{d_{bearing}}$)
End Structure

Figure 3: Structure of Bearing Assembly

to Behavior Generate Heat.

Note the reference to the law of conservation of angular momentum in transition *stateC* --> *stateI* in Behavior Change Momentum. Knowledge of such Physics laws can be represented as behavioral templates. Knowledge of the law of conservation of angular momentum, for example, can be represented as a small set of behavioral templates corresponding to the prototypical situations governed by the law. In one prototypical situation, for instance, if one object is contained in another, and the angular momentum of the first object changes then, on account of the conservation law, the angular momentum of the second will also change with an equal magnitude but in the opposite direction. In fact, the functioning of RWA is based on this use of the law. Thus, the behavioral template representing this prototypical application of the law of conservation of angular momentum (not shown here) is instantiated in Behavior Change Momentum that results in the achievement of the function of RWA (transition *state6* --> *stateI* in Figure 4).

The ball bearing assembly, which has been treated as a generic device in functional representation of RWA, can be similarly represented in terms of its structural components and their functional abstractions, making available finer grained design knowledge. At a larger grain size, the telescope as a whole can be represented. Thus the design knowledge is organized in two hierarchies: the classical structure-substructure hierarchy, and the function-behavior hierarchy. Since the causal state transitions in a behavior contain pointers to the substructures, and the schemas representing substructures contain back-pointers to the causal state transitions in which they play some role, knowledge in one hierarchy is accessible from the other. This function-structure model of the design is generated by the propose subtask of design.

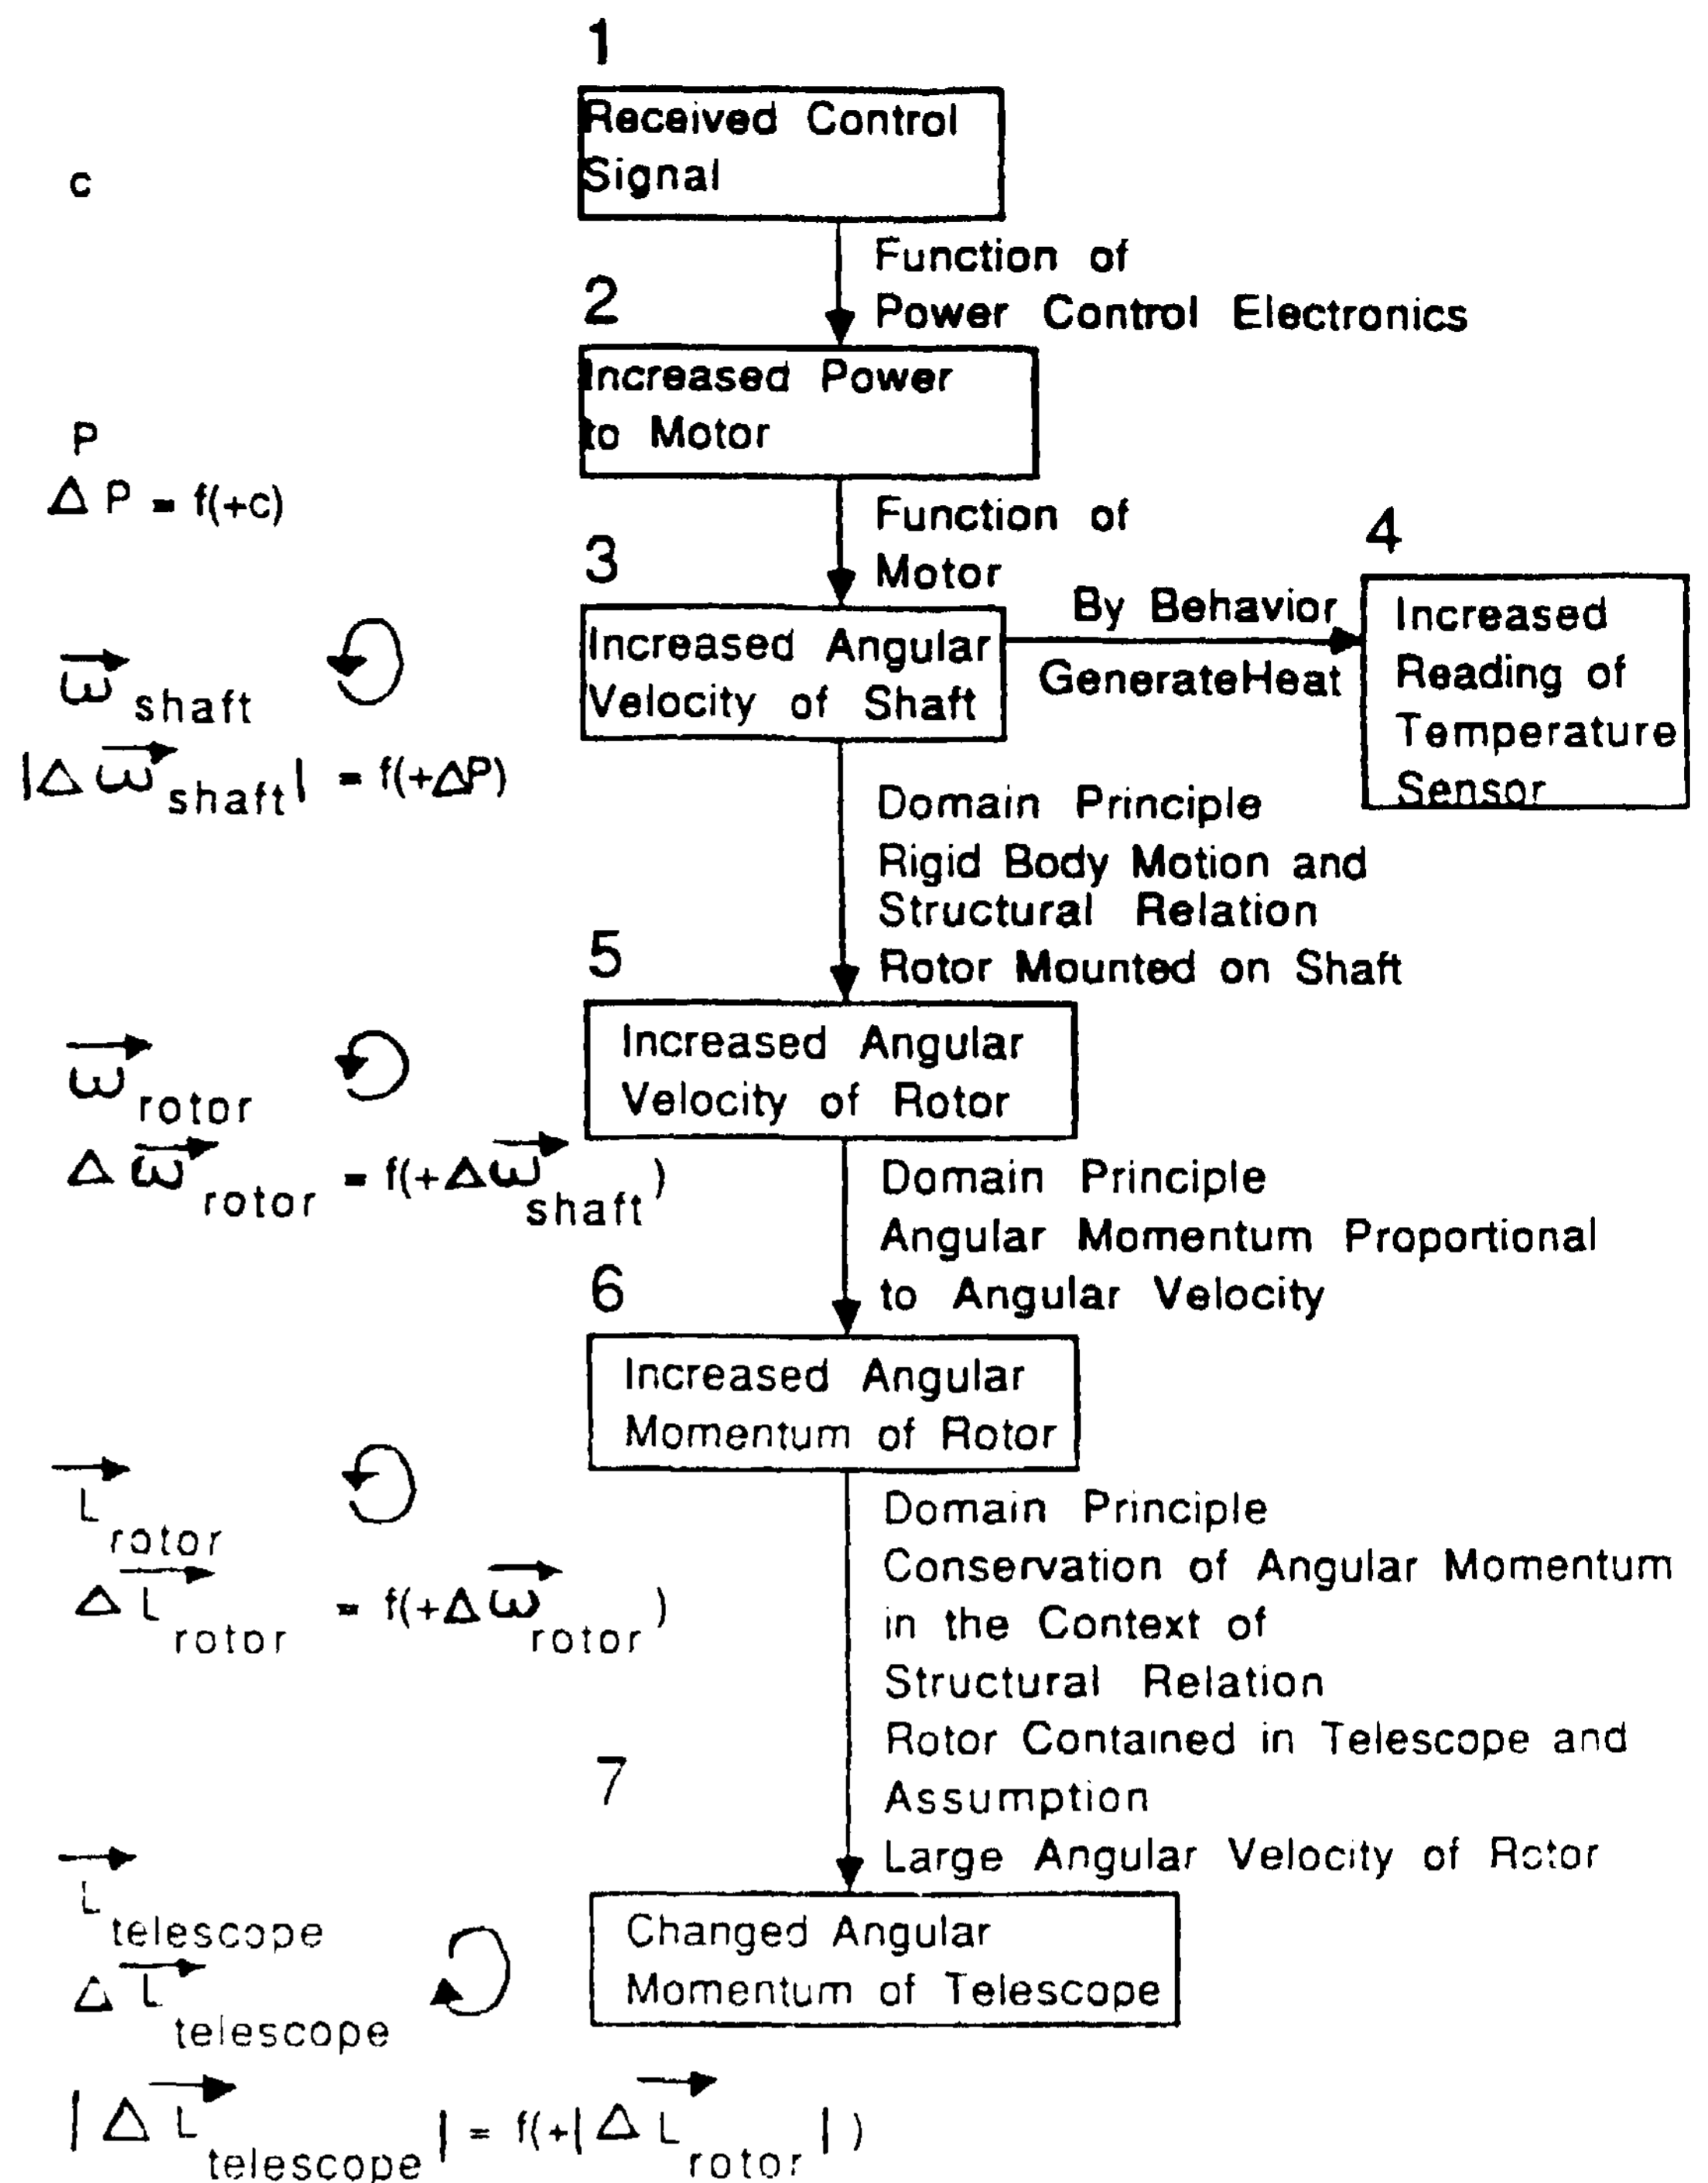


Figure 4: BehaviorChangeMomentum of RWA

4 Corrective Redesign: Diagnosis and Repair

4.1 Diagnosis of Undesirable Behavior

Let us now consider how the functional representation of designs helps in solving redesign problems in which an undesirable behavior is to be corrected. As mentioned above, the method for corrective redesign decomposes the redesign task into the subtasks of diagnosis of the structural faults(s) responsible for the undesirable behavior and repair of the structural faults. The method of functional reasoning further decomposes the diagnostic task into three subtasks: identification of the causal behavior(s) in which the sensor that detects the undesirable behavior plays a functional role; identification of the malfunction responsible for the undesirable behavior; and identification of the structural fault responsible for the undesirable behavior⁴.

⁴ We assume that the sensor itself is functioning properly. Sensor validation is a related but different problem which requires functional representation of the sensor.

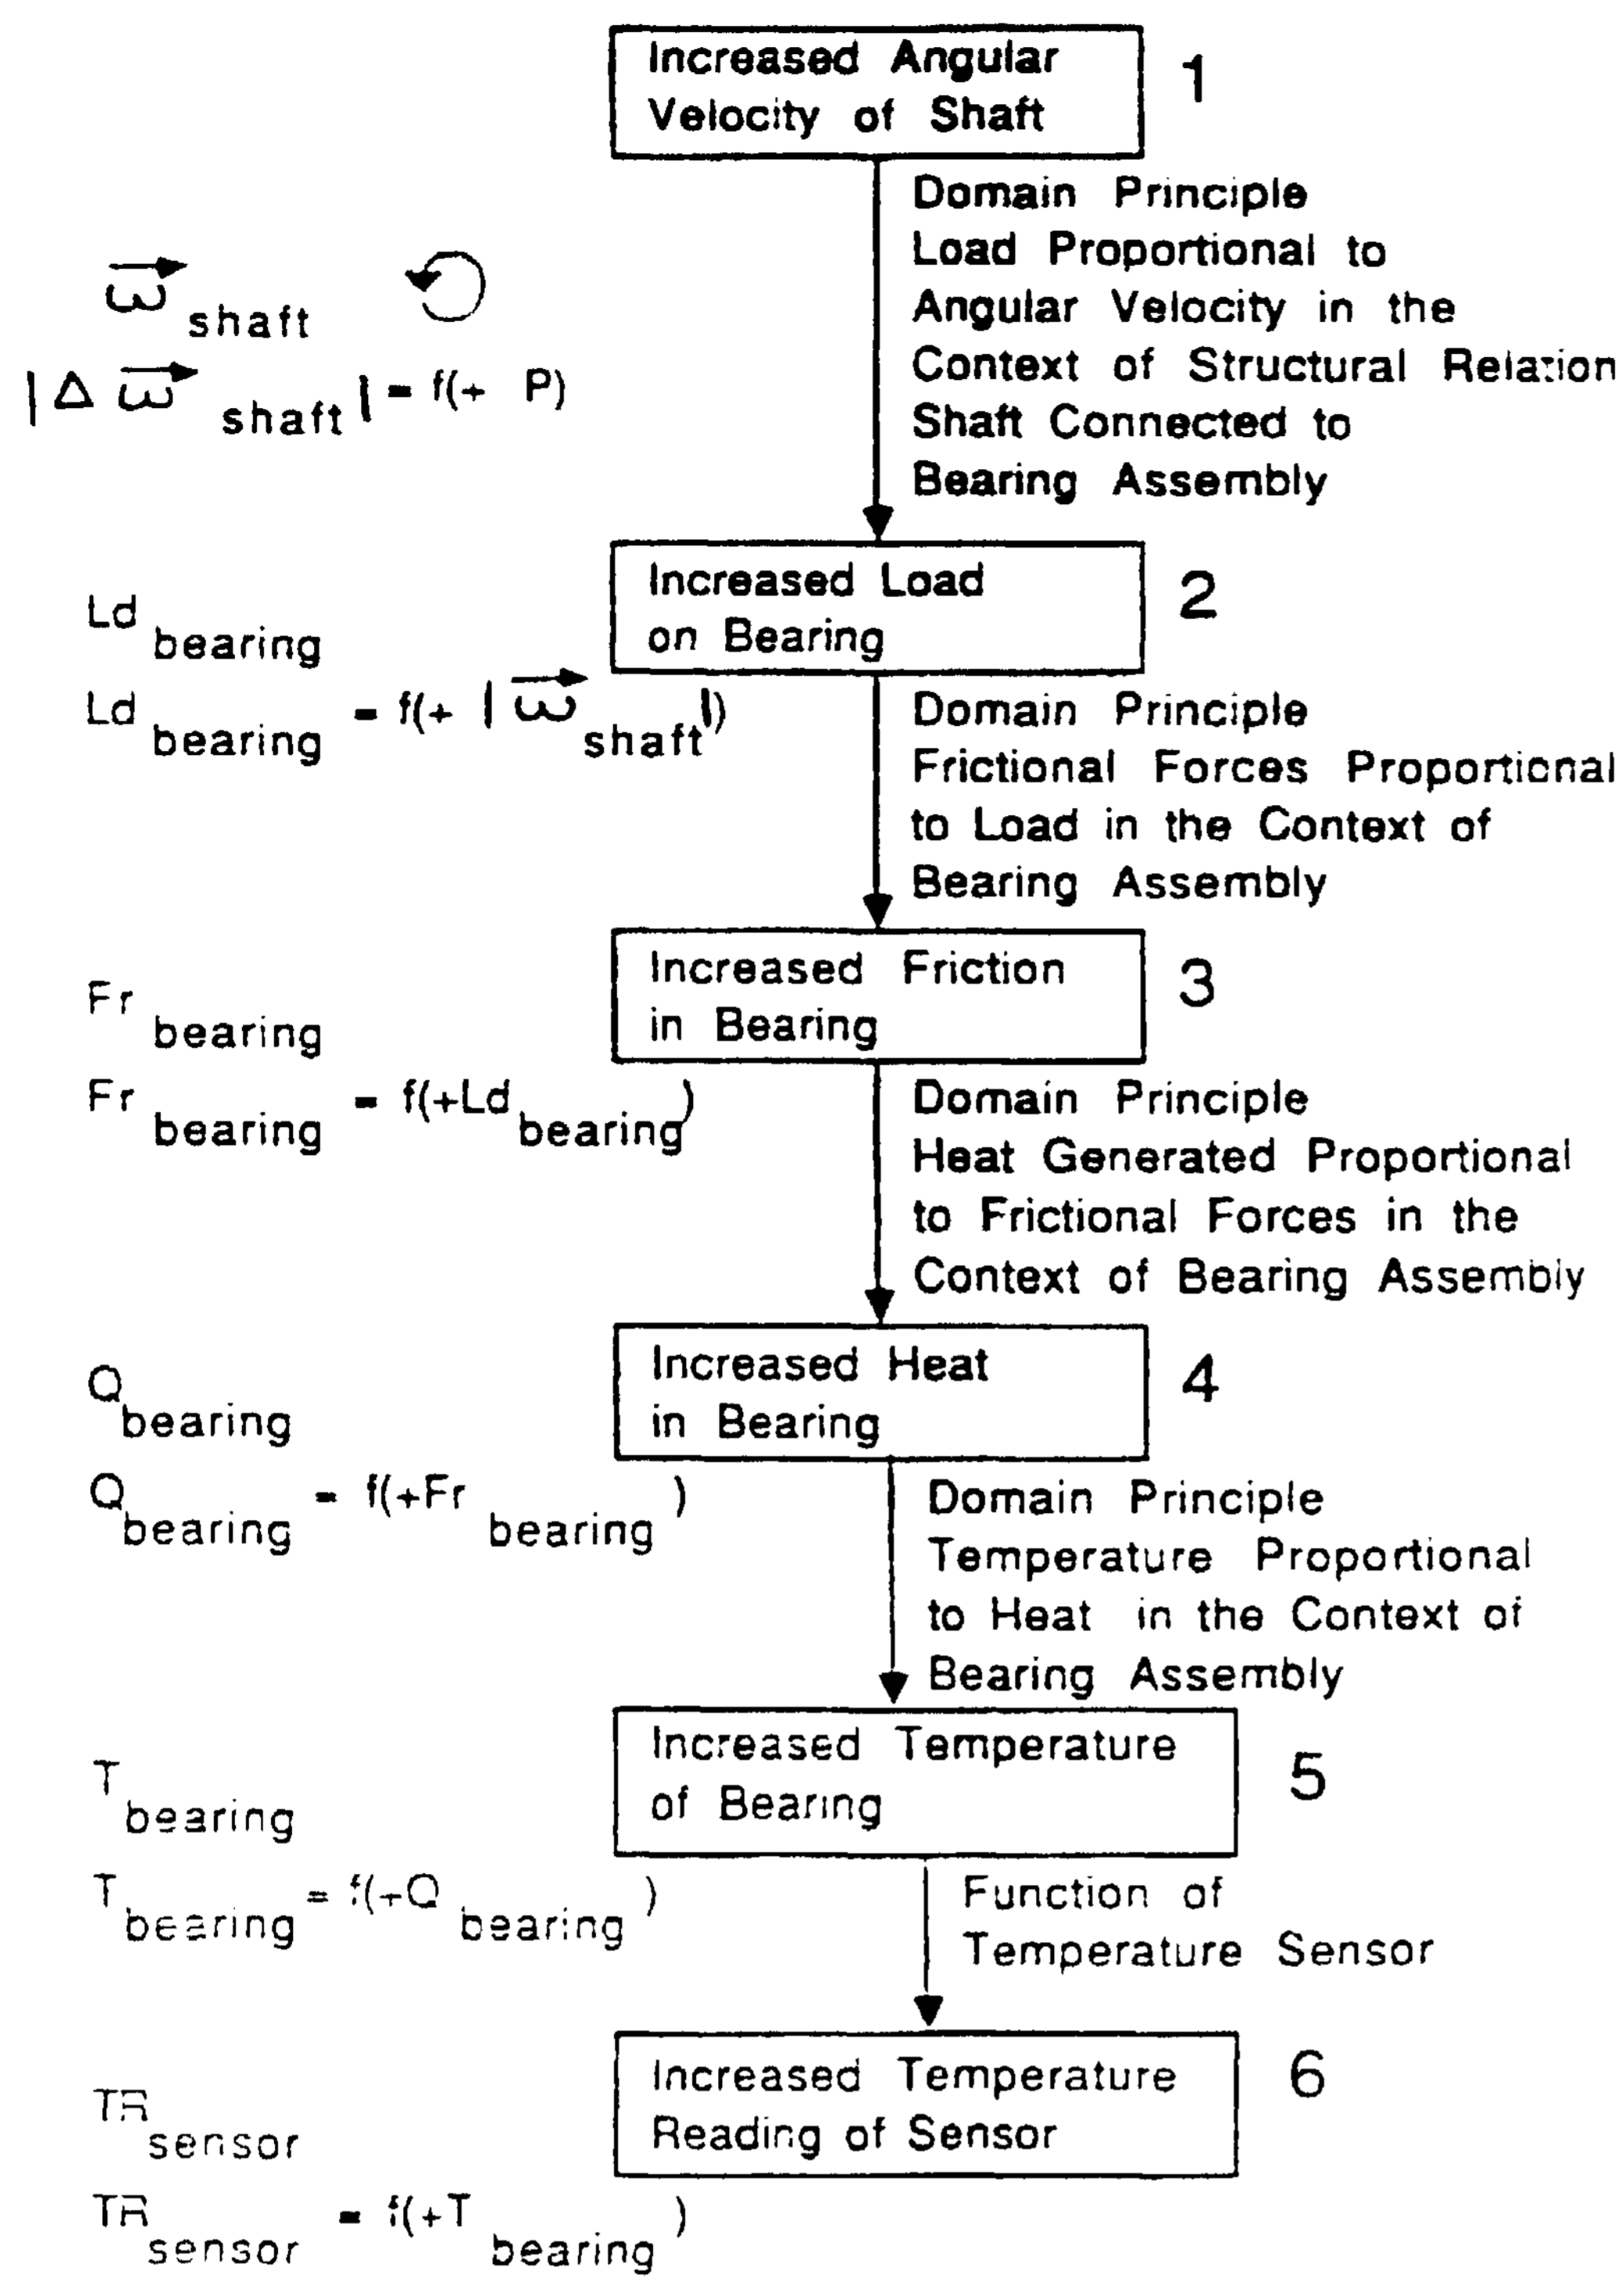


Figure 5: BehaviorGenerateHeat of RWA

Let us begin by analyzing the task of identifying the causal behavior in which the sensor that detects the undesirable behavior plays a functional role. Since the undesirable behavior does not realize any explicit function of the design, searching the function-behavior hierarchy for the causal state transition in which the sensor plays a role is likely to be computationally very expensive. However, because of the specification of the structural relations in the schema for each substructure, the structure-substructure hierarchy can be searched more efficiently to locate the schema for the sensor. The schema for the sensor specifies the causal behavior(s) in which it plays a functional role. In the RWA problem, for example, the schema for the temperature sensor (not shown here) specifies that the transition $state5 \rightarrow state6$ in *BehaviorGenerateHeat* is based on the function of the sensor (see Figure 5).

The second subtask of diagnosis is to identify the malfunction responsible for the undesirable behavior. This is accomplished by backward tracing of the causal behavior(s) determined in the first subtask, starting from the causal state transition in which the sensor plays a functional role. The values of the state variables characterizing the preceding causal states are com-

puted and checked against the operating ranges of the structural components specified in the state transitions. In the RWA problem, this backward tracing of *BehaviorGenerateHeat* leads to $state2$ characterized by the variable $Ld_{bearing}$ (Figure 5). On comparing the value of $Ld_{bearing}$ with the operating range specified in the schema for the bearing assembly (see Figure 3), it is determined that the load on the bearing is beyond what the bearing assembly can support. This identifies the malfunctioning of the bearing assembly as the cause of the given undesirable behavior.

The last subtask of diagnosis is to identify the structural cause of the undesirable behavior. This is accomplished by using knowledge of the relations underlying the operation of the malfunctioning substructure identified in the second subtask. In the RWA problem, for example, the schema for the bearing assembly (Figure 3) shows that the load capacity of the bearing assembly, $max.Ld_{bearing}$ depends on the size of the bearing balls, R_{ball} . This enables the identification of the structural cause of given undesirable behavior, namely, the size of the bearing balls is too small for the load.

4.2 Repair of Faulty Structure

Once the structural cause for the undesirable behavior has been determined, the redesigner has to repair the structure to correct the behavior. The method of functional reasoning decomposes the repair task into three subtasks: selection of a repair strategy for correcting the structural fault; proposal of a repair solution; and testing whether the proposed solution necessitates additional structural modifications.

The subtask of selecting a repair strategy requires a memory of repair strategies indexed by the type of repair tasks for which they are appropriate. For instance, one common repair strategy is *component replacement*, i.e., to replace the component responsible for the undesirable behavior with a functionally equivalent component that meets the design requirements. This repair strategy is useful for repair tasks in which the parameter of some component is responsible for the undesirable behavior. The functional representation of design helps in identifying the type of repair task, which can then be used to select the appropriate repair strategy. In the RWA problem, for example, the diagnostic task showed that the small size of the bearing balls is responsible for the abnormally high reading of the temperature sensor. This leads to the selection of the repair strategy of component replacement.

The second subtask of repair is proposal of a repair solution. Functional representation is of limited help in performing this task; the repair solution is produced by the application of the strategy selected in the first subtask. In the RWA problem, the repair strategy of component replacement uses the relation between $Ld_{bearing}$ and R_{ball} to propose the solution of replacing the bearing balls with larger ones.

The third subtask of repair is to check whether the proposed solution necessitates additional structural modifications. This is accomplished by causally propa-

gating the effects of the structural modification. Starting from the state transition where the structural modification is proposed, the causal behavior is traced forward. New values of the state variables characterizing the succeeding causal states are calculated using the causal dependencies between them, and compared against the operating ranges of the structural components specified in state transitions. If the value of some state variable is beyond the range of the corresponding component, then another structural modification is made along the lines indicated above. This process is repeated until the behavior is traced fully, and the values of the state variables show that the undesirable behavior has been corrected.

The modified structure produced by the repair task can now be verified, and if needed, redesigned again. Note that the causal propagation of the effects of a structural modification in the repair task helps to *locally* verify that the undesirable behavior has been corrected. However, this does not constitute verification of the design as a whole.

4.3 Limitations of the Method

There is of course no guarantee that, the method for corrective redesign described above would succeed in solving an arbitrary redesign problem. Its success or failure depends on whether or not the needed knowledge is available to it. For instance, if, in the RWA problem, knowledge of the causal relation between the load capacity of the bearing assembly and the size of the bearing balls, or the repair strategy of component replacement were not available, the redesigner could not have reached the solution of replacing the bearing balls with larger ones. Instead it would have continued to trace the causal behavior backwards, until, using the knowledge of the causal relation between the angular velocity of the shaft and the load on the bearing assembly, it decided that the angular velocity of the shaft was too high, and proposed reducing the angular velocity of the shaft as the solution to the redesign problem. The designer could have failed altogether if even this knowledge was unavailable. However, the problem solving does terminate, even if in failure, once the causal behavior has been traced fully.

This analysis also provides focus to the issue of spatial and geometrical reasoning in redesign problem solving, since the realization of a redesign solution often involves reasoning about the shapes and contours of structural components. In the RWA problem, for instance, once the redesign solution of increasing the size of the ball bearings is reached, it still remains to be decided how this solution is going to be realized. The functional representation of RWA makes the knowledge of the structural relations of the bearing assembly available to the redesigner. However, what changes have to be made to the shaft and the stator so that larger sized ball bearings can be used is not clear. This requires the capabilities of spatial and geometrical reasoning about the shapes and contours of ball bearings, the rotor, and the stator.

5 Concluding Discussion

We have presented an analysis of the redesign problem and shown how function-structure models of designs can be used for solving a generic class of redesign problems. This work follows a rich literature on redesign problem solving. Stallman and Sussman [1977; introduced dependency-directed backtracking to decide what structural component to modify when a design failed to achieve the desired functions. The causal behaviors of the functional representation scheme serve a similar purpose. These behaviors capture the causal dependencies between the device states which enables the redesigner to trace the structural cause of an undesirable behavior. The REDESIGN system [Steinberg and Mitchell, 1985] makes use of the purposes of structural components in a design which is similar to the notion of functional abstractions of structural components in the functional representation scheme. REDESIGN's redesign knowledge, however, is largely associative rather than in the form of function-structure models. The PROMPT system [Murthy and Addanki, 1987] uses modification operators, and decides on their applicability by testing their preconditions. The functional reasoning method for redesign seeks to identify various types of modification tasks, and calls for a functionally organized memory of modification strategies indexed by the tasks. The CHEF system [Hammond, 1989] makes use of anticipatory knowledge about potential problems with a plan for retrieving the best-matching plan from memory. In our framework, the redesigner uses knowledge of anticipated side effects for correcting undesirable behaviors.

In a different line of research, Rieger [1976] has used functional models of devices for problem solving as well as for natural language understanding. The function-structure model described in this paper can be similarly viewed as providing both a partial theory of comprehension of the functioning of devices as well as a language for capturing design knowledge useful in redesign and diagnostic problem solving. However, while Rieger's models focus on the identification of various types of causality, the functional representation scheme emphasizes the organization of causal knowledge.

The main contributions of the present research are two-fold. First, it provides a partial task structure for the redesign problem. That is, it identifies a task-subtask decomposition for the redesign problem, some of the methods applicable to the subtasks, and the knowledge required by these methods. This analysis begins to provide a framework for capturing the interactions between the tasks, methods, and knowledge for redesign problems. The second main contribution of this work is to show how the functional representations of designs can be used for solving a class of redesign problems. It also specifies constructs in the language for representing and organizing anticipatory knowledge of undesirable behaviors, and for accessing this knowledge from the specification of the structural components in the design.

The decomposition of the redesign problem into a task structure and the functional organization of design knowledge provides a method for managing the com-

plexity of the problem. Both the diagnosis and repair subtasks of corrective redesign problem solving can be computationally very complex. In diagnosis, every component and every relation between components in the structure can potentially be the cause of an undesirable behavior. In repair, every substructure in the design is modifiable in potentially very large number of ways. Moreover, each structural modification can potentially affect the entire design. The decomposition of the redesign task into a number of smaller subtasks and the functional organization of design knowledge helps in focusing the attention of the redesigner and localizing the search at each step in redesign problem solving. Finally, we note that the functional representation of a device also provides a causal explanation of the functioning of the device as well as justification for the design decisions.

Acknowledgments

The problem domain of redesign of the reaction wheel assembly aboard the Hubble space telescope was suggested to us by Professor Edward Feigenbaum. Our description of the reaction wheel assembly has been adapted from [Feigenbaum *et al*, 1988]. S. Prabhakar has made significant contributions to the ongoing implementation of the KRITIK system, which embodies the general framework for adaptive design problem solving presented in this paper. This paper has benefited from helpful comments by the reviewers.

References

[Chandrasekaran, 1988] B. Chandrasekaran. Design: An Information Processing Level Analysis. In D. Brown and B. Chandrasekaran, *Design Problem Solving: Knowledge Structures and Control Strategies*, London, UK: Pitman; in press.

[Chandrasekaran, 1989] B. Chandrasekaran. Task-structures, Knowledge Acquisition, and Learning. Guest editorial, *Machine Learning*; in press.

[Feigenbaum *et al*, 1988] E. Feigenbaum, R. Keller, Y. Iwasaki, and R. Engelmores. Proposal for Research on Multi-Use Knowledge Bases. Knowledge Systems Laboratory, Stanford University, September 1988.

[Goel and Chandrasekaran, 1988] A. Goel and B. Chandrasekaran. Integrating Model-Based Reasoning with Case-Based Reasoning for Design Problem Solving. In *Proceedings of the AAAI-88 workshop on AI and Design*. St. Paul, 1988; in press.

[Hammond, 1989] K. Hammond. *Case-Based Planning: Viewing Planning as a Memory Task*. San Diego, California: Academic Press, 1989.

[Murthy and Addanki, 1987] S. Murthy and S. Addanki. PROMPT: An Innovative Design Tool. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 637:642, Seattle, 1987.

[Rieger, 1976] C. Rieger. An organization of knowledge for problem solving and language comprehension. *Artificial Intelligence*, 7(2):89-127, 1976.

[Sembugamoorthy and Chandrasekaran, 1986] V. Sembugamoorthy and B. Chandrasekaran. Functional Representation of Devices and Compilation of Diagnostic Problem-Solving Systems. In *Experience, Memory and Reasoning*. J. Kolodner and C. Riesbeck: editors, Hillsdale, NJ: Lawrence Erlbaum, 1986, pp. 47-73.

[Stallman and Sussman, 1977] R. Stallman and G. Sussman. Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis. *Artificial Intelligence*, 9:135-196, 1977.

[Steinberg and Mitchell, 1985] L.I. Steinberg and T.M. Mitchell. The REDESIGN System: A Knowledge-based Approach to VLSI CAD. *IEEE Design and Test of Computers*, 2(1):45-54, 1985.