Representing and Reasoning about Physical Systems from a Functional Viewpoint†

Mahmoud Pegah
Computer Science Dept
219 Wright Hall
University of Northern Iowa
Cedar Rapids, IA 50614
pegah@iscsvax.unl.edu

Jon Sticklen
AI/KBS Laboratory
CPS Dept, A714 Wells Hall
Michigan State University
East Lansing, MI 48824-1017
sticklen@cps.msu.edu

William Bond
Intelligent Systems Applications
Dept 313, PO Box 516
McDonnell Douglas Aerospace
St. Louis, MO 63166
bond@mdcgwy.mdc.com

James K. McDowell
Dept of Chemical Engineering
Research Complex - Engineering
Michigan State University
East Lansing, MI 48824-1326
mcdowelj@cps.msu.edu

Abstract: A pivotal challenge of the 90s is to develop engineering methodologies which robustly address issues such as design for manufacturability, design to requirements, and conceptual design in engineering domains. Over the last several years, we have explored ways in which a functional reasoning viewpoint can be utilized to address such issues. In this report, we describe our extension and application of the Functional Modeling (FM) approach to represent and qualitatively simulate a significant portion of the fuel system of the McDonnell Douglas F/A-18 aircraft. Our general goals have been two fold: to test the scalability of our approach against a formidable real world problem, and to extend our approach to include a library facility from which standard parts may be instantiated into an evolving engineering design. Results support attainment of these goals. Knowledge acquisition for the research reported here centered on reverse engineering from a technical manual for the F/A-18 fuel system. Our KA experience on this project indicates that a functional viewpoint in general provides a strong backbone for reverse engineering.

1 Introduction

Reasoning explicitly about physical systems offers a way to circumvent the brittleness of reasoning systems built solely on associational knowledge. In addition, MBR is attractive because it captures an intuition that is especially cogent in engineering areas: in order to troubleshoot a device, or redesign a device to new specifications, or ..., it is useful to know how the device "works" — i.e., to represent and reason with a model of the device.

There are two variations on the above theme. Each is involved in the larger picture of representing devices in the world and reasoning about them. One branch of research has focused on how models of behavior are derived. The naive physics work of deKleer (deKleer & Brown, 1984), Forbus (Forbus, 1984), Kuipers (Kuipers, 1984), and Bylander (Bylander, 1986) exemplify this research. The second variation focuses on how models of behavior are used, as typified by the circuit diagnosis work of deKleer and Williams (deKleer & Williams, 1987) and Davis (Davis & Hamscher, 1988), and the function-based thrust of Chandrasekaran (Sembugamoorthy & Chandrasekaran, 1986), Franke (Franke, 1989), Sticklen (Sticklen,Chandrasekaran, & Bond, 1989), and others.

The functional approach to device understanding begins with the intuition that if we know the purposes of a device then we have a very powerful basis for organizing our causal understanding of that device. The underlying thesis is that when we know the purposes of a device (i.e., what it will be used for), we enhance our abilities to organize our causal knowledge of the device. This organization uses the abstractly stated purposes to index causal behaviors of the device that achieve those purposes. Our variation on the functional approach, functional modeling (FM), also includes a strong commitment to simulation as a core reasoning strategy with which we can utilize a functional representation to reason about device performance. In FM we first decompose the complex causal knowledge of a device along functional lines, then, given a particular situation, we compose a particularized causal story for how the device will operate given stated boundary conditions. That is, we exercise a *duality* between representational decomposition for managing complexity, and situation-specific composition for simulation.

Although we have applied FM in a number of domains, in this report we center on an application in the aerospace domain. In an earlier preliminary report, we sketched our intuitions that FM would prove to be a leveraged technique in

† This article is based on "Representing and Reasoning about the Fuel System of the McDonnell Douglas F/A-18 from a Functional Perspective" by Pegah, Bond, and Sticklen, which was first published in **IEEE Expert**, [8:2, pp 65-71].

aerospace (Sticklen,Bond, & St. Clair, 1988). In this account, we report on the results of applying FM to a realistically sized aerospace application. In Section 2.0, we discuss our project goals. In Section 3.0, we review the rudiments of FM. In Section 4.0, we describe our testbed domain, the fuel system of the F/A-18 aircraft. In Section 5.0, we describe extensions to FM which are currently being completed to build a standard-parts library into our approach. In Section 6.0, we describe the major issues related to our research results and our future research targets.

2 Research Goals

Our research goals focused on two general areas. First, we hoped to test the scalability of FM in a realistic domain drawn from the aerospace industry. We use the term "scalability" in two senses. First, we use it to mean the characteristic of an existing knowledge-based system which supports expansion of its current domain coverage with minimal change in the existing device representation or in the reasoning techniques. Our initial work utilizing FM was in a well circumscribed, small domain: the human body complement system, a sub-system of the human body immune system. On theoretical grounds, we have argued that the FM approach would scale (in this first sense) to larger domains because of its inherent modularity. However, an empirical test of that theoretical position was needed to substantiate our claims.

Second, we use the term scalability to mean the characteristic of a knowledge-based approach to be applicable to a number of domain areas, in short, we mean the domain independence of the approach. One of our central long term goals is to demonstrate that the functional approach is applicable to capturing causal understanding across diverse problem domains. Again, on theoretical grounds, we have argued that the basic issues of device representation and reasoning cut across various types of domains, but to be convincing we required a practical test to substantiate these claims.

Our second major project goal was to extend the representational power of FM to include a "standard parts" library facility. In engineering domains, it is useful to have a repository of standard parts which can be instantiated to facilitate construction of a device model. We anticipated that the incorporation of a library facility into FM would be of little theoretical interest, although such an addition would provide a needed practical mechanism for model builders. On undertaking the extension, however, we developed a new understanding of the organizational power inherent in any functional approach. In short, the library facility pointed the way to another level of device organization which was not anticipated (the functional role), which we are currently developing as an added feature of the functional repertoire.

It is important to view the motivation for the research reported here in proper perspective. We did not seek to implement and field an "industrial strength" computer system which would be used in the current mainstream of engineering. Our goals rather were to enhance the current FM representational power (with the "functional library") and to provide a proof-of-principle working prototype to support claims of domain independence of the FM approach, and to support claims of ease-of-scalability of the FM approach.

3 Basics of Functional Modeling

Our approach springs from the broad framework of the Generic Task theory of knowledge based systems (Chandrasekaran, 1983; Chandrasekaran, 1986), extends the earlier framework of Sembugamoorthy and Chandrasekaran (Sembugamoorthy & Chandrasekaran, 1986), and builds on our initial conception of how a compiled level problem solver and a deep level problem can interact (Sticklen & Chandrasekaran, 1985).

Understanding FM is best done in two stages: understanding the principles of device representation in FM, and understanding the reasoning mechanism that uses such a representation. From the representational perspective, FM largely adopts the original formalism for functional representation of Sembugamoorthy and Chandrasekaran. This formalism centers on the organization of causal device knowledge. To represent a device functionally the device is first recursively decomposed into its constituent subdevices. In engineered artifacts, this decomposition typically parallels the major structural systems of the device.

The second step in representing a device functionally is to enumerate the functions of each of the subdevices. A function is composed of three elements:

- a *Provided* clause which states the conditions under which the function will be applicable. This amounts to a precondition for the function.
- a *ToMake* clause which states the result which will be achieved after the function completes. The ToMake clause may be thought of as a postcondition.
- a *By* clause which points to the causal description of how the function is implemented. We have so far limited our functional representations to implement functions by behaviors, as described below.

Functions provide a means of abstractly knowing what can be achieved (ToMake), what must be true for a given function to be applicable (Provided), and a pointer to a causal description of how the function is implemented (By). Below, a fourth element for function description is described: the functional role.

To complete a functional representation for a device, the behaviors which implement functions (pointed to *via* the "By" clauses in functions) should be described. Behaviors are directed graph structures in which the start nodes of the graph are tests of state variables of the device, and other nodes are descriptions of changes in state variables. Behaviors resemble fragments of causal nets. However, unlike causal nets, the edges of the directed graph are annotated and point to an elaboration of why each node transition takes place. These annotations are either pointers to "world knowledge" or to other parts of the functional description itself; i.e., to lower level functions or behaviors.

To summarize, there are four central facets of the FM approach to device representation. First, the functional representation is a conceptual abstraction of what a device is and how it works. The "what it is" part is represented as a collection of sub-devices related by a "ComponentOf" relation. The "how it works" is represented as the functionality of which it is capable and the behaviors that accomplish those functions. Second, a functional description exhibits a natural modularity. A sub-device of the overall device may be replaced with another totally different sub-device which accomplishes the same functions.

Third, in understanding from the top level the device functionality, we are normally led via a chain of

device => function => behavior
=> sub-device ...

to lower and lower levels of sub-devices. However, this path of understanding may be terminated before the lowest levels of the device are reached. Once a level is reached at which a particular functionality of some underlying sub-device may be "assumed true," then further probing along the current path is unnecessary. This ability to probe only as far as needed follows directly from the modularity of representation adopted. Put another way, in the functional approach to device understanding, there is a implicit natural "layering of understanding" from the most abstract levels of device description to the most detailed. Finally, and related to the last point, each behavior in a functional representation can be thought of as a fragment of a complete causal net. Each of the fragments carries with it (in its start nodes) predicates which indicate when the fragment is applicable.

The points above are not unrelated. Overall, FM manages the complexity involved in comprehending a complex device by a divide and conquer strategy; i.e., by decomposition. The decomposition is two fold: the device-subdevice dimension, and the device causality dimension. In the device causality dimension, fragments of causal knowledge are "behaviors" which are indexed by abstractly stated functions.

From a representational perspective, the methodology of FM parallels the original work of Sembugamoorthy and Chandrasekaran (Sembugamoorthy & Chandrasekaran, 1986). However, the computational goals differ. We set a consequence finding information processing task for our function-based problem solving. More specifically, consequence finding is undertaken in response to a particular set of boundary conditions and amounts to building up a full state change diagram from the fragments that exist in the behaviors of the functional representation. In other words, we will build a specialized causal net for particular boundary conditions. Note the symmetry between the foundation of our representational viewpoint (decomposition to handle complexity) and the core of our computational process (composition tailored to a particular context).

The consequence finding algorithm is as follows.

1. Specify the initial conditions. The device variables have "default values" so only device variable bindings which are not the normal state of the device need be set. Likewise, only missing functions or altered functions need to be explicitly input.

2. Determine the functions/behaviors that should be used as a starting point. Once the initial conditions of the device are specified, it is possible to use those initial conditions to index behaviors and functions of the device that would be applicable under those starting conditions. After a round of "filtering" to remove redundant functions/behaviors, we are left with a set of functions/behaviors we will call the "invocable functions/behaviors." If there are no invocable functions/behaviors, the functional reasoner halts. An example of the redundancy we need to "filter" would be two functions, say F_a and F_b, such that both have the same Provided clause (i.e., same precondition), and such that F_a contains a knowledge pointer to the function F_b. In this case, we want to filter out F_b.

3. Starting with the invocable functions/behaviors from the previous step, construct a new state-change graph structure particularized for the current situation. This new structure is termed a "Particularized State Diagram" (PSD). Each node in the PSD will be a partial state description in the same sense as before, as a pointer to a variable of the device and a statement about how that variable is altered.

   The PSD is constructed by traversing each of the applicable behaviors.

   a. When at a partial state, place into the PSD a corresponding node to mark a partial state change, and update the associated state variable database accordingly.

   b. When at an annotation which is non-decomposable, remember that succeeding partial states assume whatever the annotation points to but make no changes in the PSD that is being built.

   c. When at an annotation which is decomposable (that is, another function or behavior), remember that succeeding partial states assume the function/behavior pointed to (as in b) and expand the function/behavior pointed to whenever possible. To determine whether a given function/behavior may be expanded, check its starting predicates.

4. The process of expanding the annotation links continues until there are no more links that are decomposable.

Summarizing, the PSD is built by following all decomposable annotations that were in the starting behaviors and expanding them recursively until what is left is a PSD which includes only partial state transitions. Each node in the PSD contains knowledge of the state variable it alters and the nature of the alteration. In addition, each node contains a listing of the assumptions under which this state change takes place. Once the PSD has been constructed, it is easy to determine what the effect on the device will be by traversing the PSD and noting cumulative changes that take place in the state description variables of the device.

4 Functional Representation of the F/A-18 Fuel System

Our problem domain is the fuel system of the F/A-18C and F/A-18D Navy model aircraft built by the McDonnell Douglas Corporation (MDC), as shown schematically in Figure 1. In the F/A-18, fuel is carried internally in six tanks (shown in Figure 1) and up to three external tanks (not shown). The two internal wing tanks and Tanks 1 and 4 are transfer tanks. The internal fuselage tanks, Tanks 2 and 3, are engine feed tanks. The fuselage fuel transfer system pumps enough fuel from the transfer tanks to the feed tanks to insure that the feed tanks are full at all times. The fuel in either transfer tank circuit can be transferred to either of the feed tanks. The wing fuel transfer system pumps fuel from the wing tanks to the feed tanks. The right wing transfers only to Tank 3 and the left wing transfers only to Tank 2.

The engine fuel supply system is powered by motive flow fuel pressure generated by two driven motive flow/boost pumps. This "motive flow" is basically a hydraulic subsystem which uses the fuel of the aircraft as the hydraulic fluid. Each engine is supplied fuel by separate feed systems for redundancy.

Tank 2 and the left motive flow/boost pump supplies fuel to the left engine. Tank 3 and the right motive flow/boost pump supplies fuel to the right engine. The motive flow/boost pumps provide high pressure fuel (motive flow) for operation of engine fuel turbine pumps, the fuel dump system, and internal fuel transfer.

4.1 Knowledge Acquisition: Reverse Engineering

The detailed knowledge which we gathered and organized in our project was obtained from a technical manual for the F/A-18 fuel system. This manual consisted of schematics for the fuel system and information about the operation of components, it and included no direct information about intended engineering "uses" for various components or subsystems.

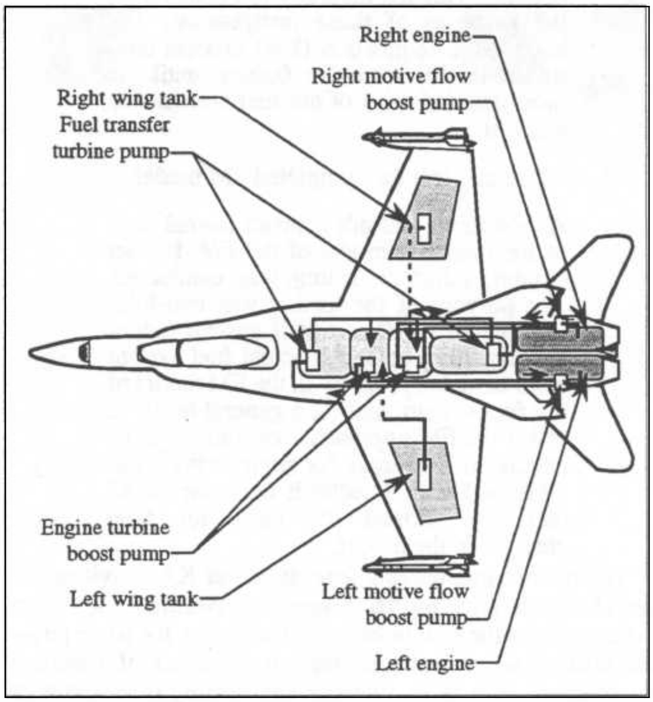

Figure 1: Schematic of F/A-18 fuel system

Our knowledge gathering consisted of three phases as follows...

1. obtain a top level understanding of the fuel delivery system from MDC engineers.

   This phase was relatively short, consisting of several interviews.

2. use the (putative) understanding of the engineering intentionality to be achieved by the F/A-18 fuel system as a start to reverse engineer a full FM model of the fuel system from the technical manual, and recursively continue until an FM model was completed

   This was by far the most time consuming part of the project, requiring approximately a two year effort.[1] This phase consisted of using (the evolving) purpose-oriented understanding of the fuel system to guide the development of deep and deeper levels of understanding. Thus we started with the top level, purpose-anchored understanding of the F/A-18 fuel system (from #1 just

1. This is a two year effort by the calendar. Major implementation responsibility for the final software system was in the hands of the first author, a graduate student at the time. Thus the two calendar year duration is not indicative of the span that would have been required in industry.

above), and used it to help organize the causal understanding of lower level components. This led naturally to enumerating the purposes of those components. The knowledge acquisition (KA) process continued in this recursive fashion until the most detailed level of the fuel system was reached.

3. informally test the completed FM model

   Although we did not conduct formal tests of the final FM model of the F/A-18 fuel system, informal testing was conducted. The purpose of the testing was two-fold: (at a detail level) to test if system redundancies inherent in the actual fuel system were modeled properly in the FM model of the fuel system, and (at a general level) to test if the FM approach were judged to be of potential leverage for engineering modeling at MDC. Feedback from the MDC engineers formed the basis for final changes in the model.

The reverse engineering aspects of our KA activities (#2 just above) require further discussion. Normally, a design engineer faces the task of creating the design for some physical artifact which will accomplish some set of specified requirements. The task of reverse engineering is to start with a (typically structural) description of an existing artifact, and from it develop an understanding of how the artifact "works" sufficient, e.g., to redesign part or all of the artifact to altered specifications. The central goal of reverse engineering is to re-capture the functional understanding of the artifact.

Typically, artifact descriptions are conceptually similar to blueprints, which represent physical structure, but not artifact subsystem function/purpose. Although in principle such blueprint representations contain all knowledge necessary to understand how the artifact "works," assimilating that knowledge typically involves assigning purpose to the various subsystems of the artifact, and that task from structure alone would be a very formidable undertaking. In fact, that task would involve the typical task of Qualitative Physics (determining large scale behavior from structure and small scale component behavior), followed by the selection from those possible behaviors the small subset of behaviors which the original design engineer intended.

This is in fact why reverse engineering from descriptions of physical structure are difficult. It is also why some version of artifact purpose/goals (such as is used in VHDL for digital electronics) is an important aspect of device representation. Given a device description such as VHDL, or as that employed in FM, artifact purpose/goals are included from the start.

4.2 Major Subsystems

We decomposed the fuel system of the F/A-18 into the major subsystems shown in Figure 2. The subsystems on which we concentrated were the internal fuel transfer system and the motive flow system. In terms of complexity, these two subsystems constitute most of the fuel system. On a base component by base component count, the current model represents approximately 70% of the entire fuel system.

The job of the internal fuel transfer system is to deliver fuel from the tanks to the engines, while the job of the motive flow system is to provide the hydraulic power which operates the pumps to allow the fuel transfer system to operate.

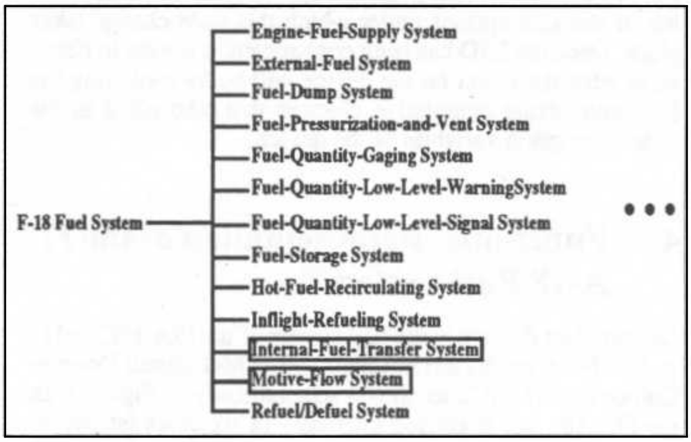

Figure 2: Top Level Components of the F/A-18 Fuel System

4.3 Representing The F/A-18 Fuel Transfer Control System

The following section describes a portion of the F/A-18 aircraft fuel system. We will discuss only a portion of the total representation following from one particular starting point; i.e., we will look at a vertical slice through the representation. We start with Component - Function - Behavior of the internal fuel transfer system. Figure 3 shows the functions each of the components is capable of, and the behaviors which implement these functions.

One of the functions of the Internal-Fuel-Transfer System is shown in Figure 4: the function which accomplishes a movement of fuel from the wing tanks to the feed tanks. One

FUNCTION :	***TransferFuelToFeedTanks***
Provided :	**right motive flow at tube restrictor side side of Tank 3 wash filter present?**
ToMake :	**fuel transfer to feed tank enabled**
By:	**to-Tank2, to-Tank3.**

Figure 3: Top Level Function

of the implementing behaviors for the TransferFuelToFeedTanks function is the "to-tank-3" behavior, shown in Figure

5. This behavior results in enabling fuel transfer to Tank3. The link annotation shows how this is achieved.

Consider the annotation on the first link in Figure 5. This annotation is a pointer to another behavior of the representation. Thus, if more detail is desired about how the transition from the presence of motive flow (i.e., hydraulic driving force) at the tube restriction side of the Tank3 wash filter leads to motive flow pressure at the inlet side of the Tank3 cutoff valve, we can examine the behavior pointed to by the link annotation: "control motive flow pressure to Tank3 transfer shutoff valve."

This behavior is shown in Figure 6. Note that the link annotations for the behavior of Figure 6 all begin "by knowledge." These are annotations to world knowledge rather than annotations pointing to deeper parts of our functional representation. The behavior of Figure 6 demonstrates how a functional representation can naturally "bottom out" when a causal transition need not be further explained for the reasoning task that will be required. The behavior of Figure 6 basically is a statement of connectivity between four points in the fuel system. Clearly, there is more detail that can be used to explain the causality involved in the concept of "incompress-

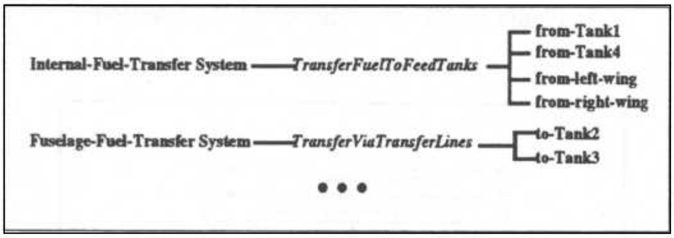

Figure 4: Device-Function-Behavior for a slice of the F/A-18 Fuel System

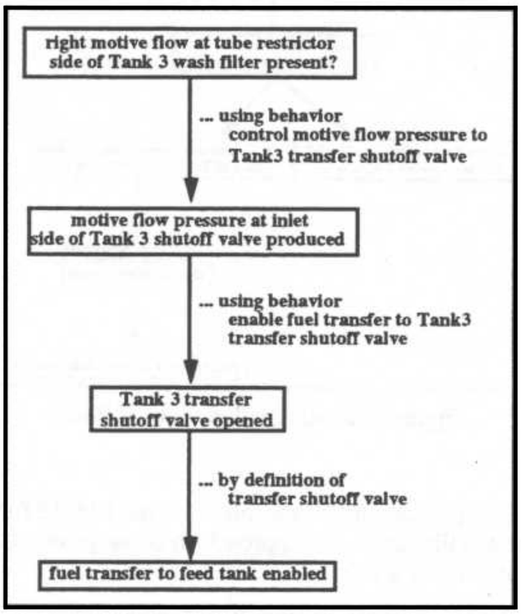

Figure 5: "to-tank-3" behavior

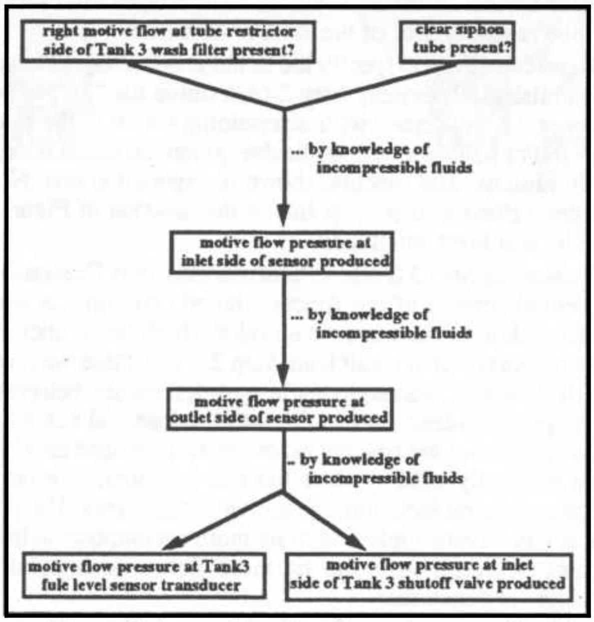

Figure 6: "control motive flow pressure to Tank3 transfer shutoff valve" behavior

ible fluids." But for the purposes we had, it was sufficient to say that because two points are connected, and because the fuel can be treated as an incompressible fluid, we know that when upstream points have pressure, downstream points will also have pressure. We can say this by making reference to our store of world knowledge.

We started examining our representation of the F/A-18 fuel system at the highest level of the control system for fuselage fuel transfer (Figure 4). At that level, it was easy to see what the overall functionality of the control system was. Given the modular nature of a functional representation, it was also easy to follow the link annotations to deeper and deeper levels of detail in the system (Figure 5, Figure 6).

To this point, we have followed a "browsing" mode of exposition. We have shown how our representation of the F/A-18 fuel system is naturally modular, and how one might use it as a kind of automated textbook to understand various aspects of the overall system. In the next section, we will show how the functional representation of the F/A-18 fuel system is used as the basis for consequence finding.

4.4 Reasoning About The Internal Fuel Transfer System

In this section, we will step through the application of the reasoning algorithm outlined in Figure 3.0. Assume the following starting condition:

right motive flow at tube restrictor side
of Tank 3 wash filter

We initiate the consequence finding algorithm described above. Our goal is to determine the effect on the fuselage transfer system of the presence of right motive flow fuel at the tube restrictor side of the Tank 3 wash filter.

Reasoning Step 1 (specify the initial conditions) is already accomplished. Reasoning Step 2 (determine the "invocable" behaviors) is concerned with determining which of the functions and/or behaviors are applicable, given the stated boundary conditions. The function shown in Figure 4 is one. Note that this behavior directly points to the function of Figure 5 from its first level annotation.

Reasoning Step 3 (build a Particularized State Diagram) is the central activity of our function-based consequence finding algorithm. In this step, we utilize the high level functions (or behaviors) which result from Step 2. From these we index (via the link annotations) lower level functions and behaviors whose preconditions are met. Those functions and behaviors whose predicates are true are expanded and "spliced into" the place originally held by their links in a process similar to macro expansion techniques in software languages. The process is recursively applied until no more decomposable links remain; i.e., until there are no more links which point to behaviors or functions.

Following the rules of Step 3 of our algorithm, we would start with the function of Figure 4. At this level of detail, we would see the simple causal net-like structure of Figure 7. At each iteration in the consequence finding algorithm, we expand all function/behavior links currently visible. After one round of expansion, the PSD has grown by expanding the behavior of Figure 5 to yield the diagram of Figure 8. Finally, the behavior of Figure 6 is expanded to yield the final, most detailed PSD, as shown in Figure 9. Note that the view of Figure 9 is not fully expanded; the links showing dots are not expanded in this diagram for pedagogical purposes.

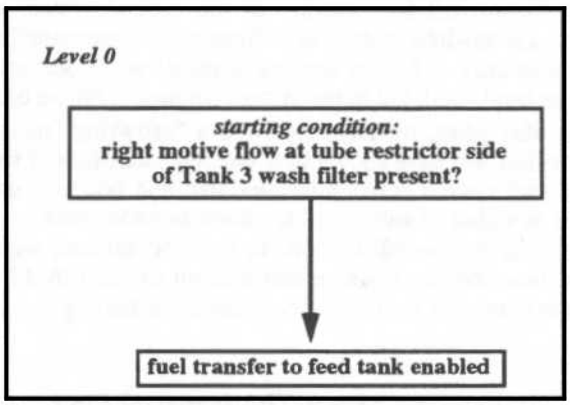

Figure 7: Coarse grain view of PSD

Once a complete PSD is produced, it is straightforward to determine the cumulative effects (i.e., consequences) resulting from the given initial conditions. The PSD graph structure is traversed, keeping a running tally of all changes made to state variables. The cumulative effects are then read from this "tally sheet."

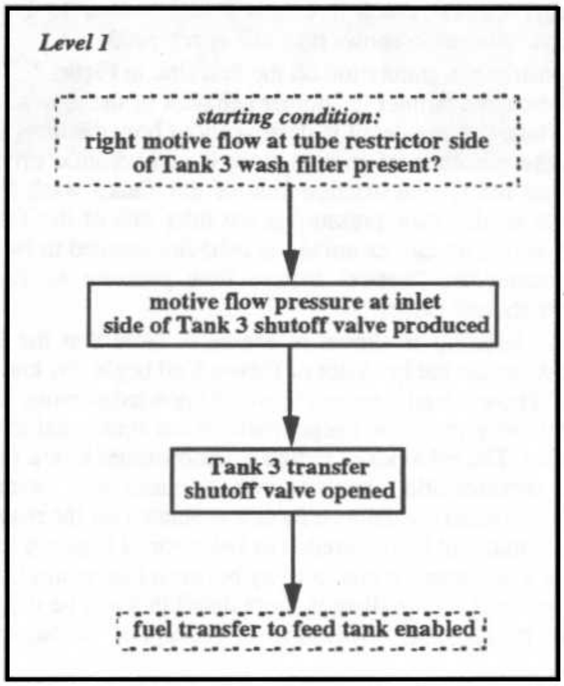

Figure 8: PSD after one round of expansion

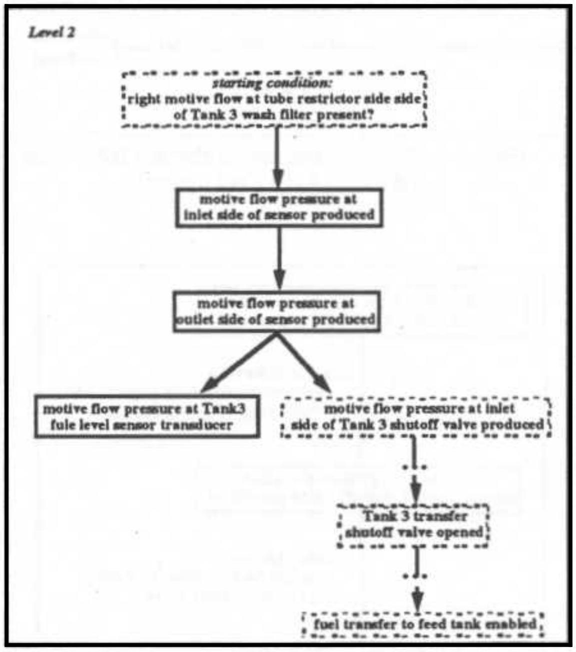

Figure 9: Most detailed view of PSD

This simple reasoning example from the F/A-18 fuel system domain illustrates our approach to consequence finding. The central ideas are that

- we decompose a complex device making strong use of its known functionality to organize the decomposition, then

- for consequence finding, we compose a situation specific causal net by a process resembling macro expansion.

Space does not allow us to present more complex reasoning examples in this report. Our results have supported the idea that the FM representational methodology scale acceptably to larger systems. Our representation of the F/A-18 fuel system currently includes 89 component devices, 92 functions, 118 behaviors, and 181 state variables. Although still not very large, the F/A-18 system is an order of magnitude more complex than any system yet represented functionally. Moreover, as we have argued previously, the reason for the good scaling characteristics, in the first sense of scaling discussed above, is the natural modularity offered by a functional approach to representation.

Informal testing of the FM model of the F/A-18 fuel system was carried out in St. Louis under the direction of Dr. Bond. This testing centered on MacAIR engineers familiar with the F/A-18 fuel system acquainting themselves with the representational and simulation features of the model, and running representative (to them) simulation tests on the model. The informal testing did not involve the standard library facility as described below in Section 5.0.

The informal testing revealed two items. First, the MacAIR engineers were in general satisfied with the ability of the FM model to produce indicative results. Second, the testing revealed the specific ability of the FM simulation package to indicate engineered redundancies in the fuel system design. As a high level comment, the MacAIR engineers deemed the FM technique as promising.

5 A Standard Library Facility - The Functional "Role"

The F/A-18 was chosen to exercise the supposed scalability of Functional Modeling. The model developed was basically an application of the techniques developed in other, less complex, domains. However, as part of our F/A-18 work, the representational framework was extended to include a standard library facility.

One of the most tedious and error prone parts of design is the need to copy the same type of component into a design many times. Most modern CAD systems provide a standard library facility which contains templates for parts which may be instantiated into a design. Similarly, we have developed a device library for the F/A-18 fuel system, which is partially shown in the Type hierarchy of Figure 10.

There are over 90 different valves in the entire fuel system of the F/A-18 aircraft. In general, any valve has two functions; enabling flow (open function) and disabling flow (close function). However, the manner in which those functions are carried out will depend on the specific type of valve under scrutiny. The Type hierarchy organization for the standard library (partially shown in Figure 10), provides inheritance support for modeling low level objects. For example, the device Valve has a function which is to "shut down flow."

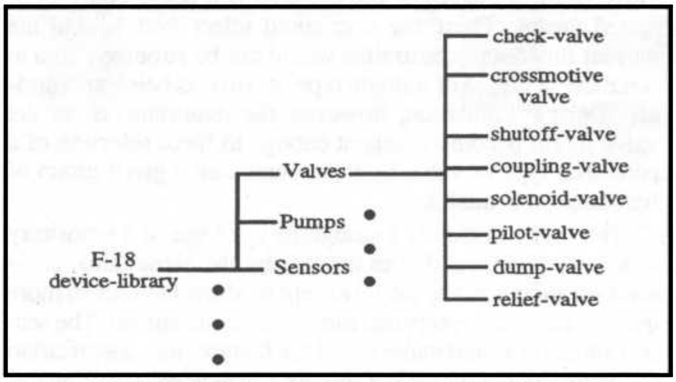

Figure 10: Standard parts library for F/A-18 Fuel System

Each type of valve (check-valve, crossmotive-valve, etc) inherits this function. But for each type of valve we can specialize the behavioral implementation that achieves these functions.

The implementation of the standard library is conceptually straightforward. One point that is not straightforward, however, lies in making sure that the "connections" between the standard part and the rest of the model are properly made when instantiating the standard part. In a functional model, these connections are of two different types. First, the state variables of the standard part must be mapped into the overall model. Second, any physical connections (pipes in our case) must be properly attached from the standard part to the overall model. Our initial strategy to produce solutions to both these problems is to rely on the user of the system to drive the instantiation of a standard part and, more particularly, to make the necessary connections between the instantiated part and the larger model. We rely on the computer system only for bookkeeping functions. Research efforts in compositional modeling (Falkenhainer, 1991) are expected to have an impact on this issue.

The initial reason for developing the library facility was to facilitate more rapid development of an FM model. However, once conceptually completed, we noted a strong potential synergism between our library utility for FM models and research reported by Nayak, Addanki, and Joscowicz (Nayak,Addanki, & Joscowicz, 1990). Nayak et al suggest that a way to automatically select appropriate high level models is to represent the primitive behaviors of the models in a context dependent manner. Our framework for FM, in particular the standard library, enables a natural extension to Nayak's notion of context dependent behavior by allowing selection of functional components based on context dependent information.

There are three ways that a modeler might utilize the standard library in constructing an FM model. First, he may use the standard library as a static repository of parts. For example, when describing a portion of the fuel system model, a user may directly instantiate a "relief value." Second, a user may select "valve" and then, because of constraints imposed by the functional requirements, an appropriate version of valve could be selected, and the selection could be by automated means. Third, the user could select "valve," but the current functional constraints would not be strong enough to force the selection of a single type of valve as being appropriate. During simulation, however, the requirements on the valve might become stringent enough to force selection of a particular type of valve by that demanding a given group of functions be available.

The first situation is a straightforward use of a repository of standard parts, and does not require the hierarchical organization shown in Figure 10 except to allow the user to more quickly find the appropriate subdevice to instantiate. The second situation would make use of the hierarchical organization of Figure 10, but would not require any new epistemic analysis; the objects under the subdevice "valve" are simply types-of "valve."

The third situation is conceptually of more interest. In FM, we have to date thought of devices as being made up of components (the listing of subdevices), functions, and the behaviors that implement the functions. Individual functions are context dependent because of their preconditions. In addition to that level of context dependency, we now add the idea of the "functional role." Within a device, a number of functions may be grouped together, and the applicability of this organization of functions may be context dependent. So, for example, when we find a thermal ballast in a hydraulic system, using the new notion of functional roles, it would be possible to represent the device as having one set of functions which act to stabilize the temperature of a system, while a second set of functions acts to store hydraulic fluid (i.e., as a reservoir). The set of functions to which attention should be directed at a particular point in reasoning are context dependent. From the framework provided by FM, the functional role is a natural extension which provides an additional and higher level indexing capability to causal understanding of a device. This third utility of functional roles has not been worked out in detail as yet, nor have we incorporated the idea into an existing model.

Nayak et al proposed their context dependent behaviors as an aid to model selection. Given components that could act in different modes, their proposal attacks the problem of selecting the appropriate mode. Our application of their idea within the FM framework, and specifically to the standard library, is also aimed at the problem of selecting an appropriate component for instantiation into a model; i.e., our proposal is aimed at problems of engineering design.

6 Discussion

For the research described above, our most important result was to demonstrate that the Functional Approach to represen0ting and reasoning about devices can scale to (at least some) real world problems in the aerospace domain. Prior to this research project, the Functional Approach has been applied in a number of problem areas, but the largest domain was a project to model the human body complement system (a subsystem of the immune system). This physiological model contained approximately 20 components. The model of the fuel system contains approximately 100 components, and was constructed using the same representational and reasoning approach as the earlier work.

Our results support the view that the FM approach is applicable to realistically sized systems in the aerospace domain. Early in this report, we noted two senses of "scalability" that we hoped to address. The second sense was that of being useful over diverse domains. The F/A-18 fuel system is a distinctly engineered domain. By successfully representing it, and providing a mechanism for reasoning about it, we have shown that the Functional Approach is applicable to (at least some) engineered systems.

A second outcome of our research has been to support the idea that a functional viewpoint is useful in projects involving reverse engineering. This is not a surprising result. Reverse engineering in practice typically means determining from structural schematics how an engineered artifact is meant (by the original design engineer) to function. The entire functional view centers on capturing such knowledge. Thus FM and the other functional techniques form a natural template for capturing knowledge in reverse engineering situations.

In his survey of Model Based Reasoning, Davis lists three crux research issues that MBR must deal with: issues of domain independence, issues of scalability, and issues of model selection (Davis & Hamscher, 1988). Although the survey is explicitly for the area of troubleshooting, the same three issues may be raised for the entire area of reasoning about physical systems. Suppose we examine some answers offered by the Functional Approach in the light of these issues.

The issue of domain independence centers on whether a particular technique is applicable to a limited domain only, or is more generally useful. We have described our experience in two very diverse domains in which we have applied the Functional Approach. Although the representational primitives were augmented for our engineering domain, the basic strategy for reasoning and the basic primitives of representation were the same in the engineering domain and in the medical diagnosis domain. Successful results have been obtained in both domains.

The issue that Davis raises of scalability is a central concern for representation and reasoning approaches. One way to argue for a scalable approach is to point out ways in which the approach modularizes a domain. The Functional Approach deals with this issue very directly; behaviors in the Functional Approach are causal net fragments. The organization of the fragments is by the known functionality of the device that is being modeled. Because a Functional Representation of a device is inherently compartmentalized, it is straightforward from a representation viewpoint to add new subdevices.

The final issue identified by Davis, model selection, is both the most interesting of his three issues, and the hardest to pin down. One of the reasons for the difficulty is that the selection of a model is a multidimensional task. Along one dimension, we must select the level at which we want to represent our model. As Davis points out, no model is complete. The Functional Representation deals straightforwardly with this fact by including the ability to point to "world knowledge" as the reason for a given state variable transition (in a behavior). This allows the modeler to construct a model that "bottoms out" at whatever level is appropriate. The level at which the bottoming is legitimate is determined by whether or not the world knowledge can be treated as a monolithic entity for purposes of the current model. A full discussion of this issue is beyond the scope of this paper.

The determination of the type of model we want to construct should be based on (a) the representational primitives offered by a particular type of model, and (b) the reasoning that a particular type of model enables. If the knowledge we have of a device to be modeled can be expressed in the primitives of a particular approach, and if the output of reasoning with that approach matches what we need to have in terms of output, then that particular type of modeling approach would be a good candidate. This statement may seem self-evident. Yet for the most part, model-based reasoning and qualitative physcis has not dealt explicitly with issues of types of models in these terms. We believe that one of the strongest arguments supporting the Functional Approach is the relative clarity of statement of the representational primitives of the approach, and of the reasoning methods that come bundled with the approach.

Our plans for future research on engineering applications of FM include full elaboration of our "standard library" facility to implement the concept of the functional role. To date, we have experimented with the relatively easy ways to leverage the functional role: as a CAD aid for building functional models, and as a constraint on the choice of particular types of devices (e.g., a particular type of valve). However, we have yet to develop the use of the functional role as a way of organizing alternative groups of behaviors in a device, nor the selection one or the other of these groups on a run time basis. We intend to accomplish both as a last extension our work on the F/A-18 fuel system.

In a parallel track, we have begun initial examination of how a functional representation of a device may be used to directly support troubleshooting. We are exploring this issue in the context of a new project in collaboration with McDonnell Douglas Space Systems Company: trouble shooting the external thermal control system of Spacestation FREEDOM.

The F/A-18 fuel system model has provided an important milestone in the verification of the validity of the Functional Approach to device modeling, and to its evolution. Our long term goal remains to explore the generality of the functional point of view for representing and reasoning about complex devices across diverse domains.

7 Acknowledgments

We are indebted to Ahmed Kamel (MSU AI/KBS) and Barry Flachsbart (McDonnell Douglas) who read earlier drafts of this text and whose comments were most helpful. Gene Wallingford played an important role in the FM-based initial implementation of the F/A-18 fuel system at MSU, from which the current model evolved.

Research reported here is supported in part by funding from the McDonnell Douglas Research Laboratories' Independent Research and Development funds. It is also supported by DARPA (ARPA 8673) and the NSF Center for High Speed Low Cost Polymer Composite Processing at MSU. Generous equipment support for the MSU AI/KBS Laboratory has been provided by Apple Computer.

8 References

Bylander, T. (1986) Consolidation: A Method for Reasoning about the Behavior of Devices. The Ohio State University.

Chandrasekaran, B. (1983). Towards a Taxonomy of Problem-Solving Types. AI Magazine, 4(1), 9-17.

Chandrasekaran, B. (1986). Generic Tasks in Knowledge-Based Reasoning: High-Level Building Blocks for Expert System Design. IEEE Expert, 23-30.

Davis, R., & Hamscher, W. (1988). Model Based Trouble Shooting. In H. Shrobe (Eds.), Exploring Artificial Intelligence Morgan Kaufmann.

deKleer, J., & Brown, J. S. (1984). A Qualitative Physics Based on Confluences. Artificial Intelligence, 24, 7-83.

deKleer, J., & Williams, B. (1987). Diagnosing Multiple Faults. Artificial Intelligence, 32, 97-130.

Falkenhainer, B. a. K. F. (1991). Compositional modeling: Finding the right model for the job. Artificial Intelligence, 51(1-3), 95-143.

Forbus, K. D. (1984). Qualitative Process Theory. Artificial Intelligence, 24, 85-168.

Franke, D. (1989). Representing and Acquiring Teleological Descriptions. In Proceedings of the 1989 Workshop on Model Based Reasoning (11th IJCAI), (pp. 62-68).

Kuipers, B. (1984). Commonsense Reasoning about Causality: Deriving Behavior from Structure. Artificial Intelligence, 24, 169-203.

Nayak, P. P.,Addanki, S., & Joscowicz, L. (1990). Modeling with Context Dependent Behaviors. In Working Papers for The Second AAAI Workshop on Model Based Reasoning, (pp. 1-8).

Sembugamoorthy, V., & Chandrasekaran, B. (1986). Functional Representation of Devices and Compilation of Diagnostic Problem-Solving Systems. In J. Kolodner & C. Reisbeck (Eds.), Experience, Memory, and Learning Lawrence Erlbaum Associates.

Sticklen, J.,Bond, W. E., & St. Clair, D. C. (1988). Introducing Functional Reasoning into the Aerospace Domain. In Proceedings of the 4th Conference on AI Applications for Space, NASA, .

Sticklen, J., & Chandrasekaran, B. (1985). Use of Deep Level Reasoning in Medical Diagnosis. In Proceedings of Expert Systems in Government Symposium, .

Sticklen, J.,Chandrasekaran, B., & Bond, W. (1989). Distributed Causal Reasoning. Knowledge Acquisition, 1, 139-162.